

BRNO UNIVERSITY OF TECHNOLOGY
FACULTY OF ELECTRICAL ENGINEERING AND COMMUNICATION

PROCEEDINGS II OF THE 32ND STUDENT EEICT 2026

ELECTRICAL ENGINEERING, INFORMATION SCIENCE,
AND COMMUNICATION TECHNOLOGIES

SELECTED PAPERS

April 28, 2026
Brno, Czech Republic

BRNO UNIVERSITY OF TECHNOLOGY
FACULTY OF ELECTRICAL ENGINEERING AND COMMUNICATION



Title Proceedings II of the 32nd Conference STUDENT EEICT 2026

Editor Assoc. Prof. Vítězslav Novák

Publisher Brno University of Technology, Faculty of Electrical Engineering and Communication

Year released 2026

1st Edition

Content and language issues are the responsibility of the authors.

ISBN 978-80-214-6418-6

ISSN 2788-1334

BRNO UNIVERSITY OF TECHNOLOGY
FACULTY OF ELECTRICAL ENGINEERING AND COMMUNICATION



SPONSORS

Diamond tier

Honeywell

Golden tier

ABB

e·on

ROHDE&SCHWARZ
Make ideas real

SKODA

ThermoFisher
SCIENTIFIC



SPONSORS

Silver tier

Garrett
ADVANCING MOTION

MESIT

NXP

RENESAS

Valeo

Honorary tier

SKUPINA ČEZ

TECHNICAL SPONSORS & MEDIA PARTNERS

IEEE
IEEE Czechoslovakia Section

VERONI
kafestánek

Contents

Foreword

Ladislav Duraja
Real-Time Detection and Visualization of the Superficial Venous Systém 9
Veronika Šromová, Dinara Sobola
Development of an Optimized Workflow for the Fabrication and Quantitative Analysis of Microstructured Bio-Substrates 13
Helena Géryková, Radovan Smíšek
Analysis of High-Frequency Components of the P Wave 17
Vojtech Smejkal
Opportunistic Rainfall Detection Using Microwave Backhaul Links of Regional ISPs 21
Vojtěch Krehan, David Kuřátko, Ondřej Jaroš
Reflector Dipole Antenna for Secondary Radar Band 24
Katarína Kúdelová, Willi Lazarov, Jakub Vostoupal
Online Radicalisation and Extremism: An Integrated OSINT and Legal Analysis 28
Jan Kalina, Michal Ptacek
Centralized Dispatch of Photovoltaic Plants, AVEVA Implementation and Operating Data Analysis 32
Tomáš Rojík, Jan Koudelka
Feed-Forward Curves Calculation for a Control Valve and Circulation Pump Performance Improvement During Ancillary Services Provision 36
Dominik Heral, Martin Světlík
Finite Element Analysis of Magnetic Fields in a High-Voltage Synchronous Machine 40
Marek Ďurďa, Šimon Sloboda
Simple Active Attitude Determination and Control System for Cubesat 44
Filip Majchrák, Jakub Arm
Web-Based 3D Simulator for Industrial Automation Processes 48
Vojtěch Kropáč, Dominik Říčánek
Prediction of contamination level in a water treatment plant using neural networks 52
Dominik Chalupka, Ladislav Polák
Litemate: An Assistive Tool for Video Image Color Correction 56
Jakub Stastny, Karel Jurik
A Compact RF Power Meter with 40 dB Dynamic Range from DC to 4 GHz 60
Martin Zatloukal, Jan Kufa
Temperature-Controlled Chamber for Comparing Temperature Sensors 64
Adéla Fialová, Karel Sedlář
Functional Profiling of Bacterial Genomes Using Constrained Ordination and Dimensionality Reduction 68
Borys Aleksandrenko, Martin Súkeník, Karel Sedlář
Hardware and Software Augmentation of Bioreactors for Culturing Bacteria 72
Dominika Kubová, Martin Králík, Martin Mézl
Interactive rehabilitation device for hand and wrist 76
Matej Grega
Design of a Pulsed Laser Diode Driver for Optical Fiber Sensor Measurement 80
Miloslav Kužela, Tomáš Frýza
Augmented reality for assisted testing of radar sensors 84
Yonatan Agafari, Petr Kriz
Classification of Normal and Anomalous Air Using BME688 Sensor Data 88
Martin Kopka
Adaptive Digital Control Techniques for Wide-Range Efficiency Optimization in a Buck Converter 92
Tomáš Paulysko, Jiří Veverka
Design of MIBICO Payload Controller for the CIMER Mission 96

Jan Bažil, Jiří Veverka
Spacecraft electrical passivation switch providing a physical gap on the power line 99
Jan Jilčík
Comparative Study on the Electrochemical Performance of Lead Foil in Different Electrolytes for Sodium-Ion Batteries 103
Martin Poul
Apparent Power Time Aggregation Algorithms 107
Martin Kyseľ, Kamil Jaššo
Compact Redesign and Safety Upgrade of a Bidirectional Power Supply for Flexible Automotive Testing 111
Jan Gadas, Aleš Povalač
Modular High-Power Electronic Load 115
Hernán A. Bravo Urrea
Influence of Rotor Flux Barrier Geometry on Torque Ripple in a Synchronous Reluctance Motor for Electric Vehicle Applications 119
Radim Vancl, Michal Husák
Thermal System Modeling and Validation 123
Róbert Ruman, Viet Anh Phan
Design and Implementation of a Framework for Automatic Generation of Attack Graphs Using Open Sources 128
Jan Havlín, Adam Fridrich, Tomáš Lieskovan, Patrik Dobiáš, Jan Hajný
Post-Quantum Encryption on FPGA Platform: Czech-Estonian Pilot 132
Jiří Hrachovina, Sara Ricci
Post-Quantum Migration of the EST Protocol: Evaluation on Constrained Microcontrollers 139
Inna Zumberg, Johannes Jun Wei Low, Anna Taubenberger, Vratislav Cmiel
A PDMS stopper-based cell migration assay as an alternative to the Scratch assay 143
Romana Ščotková, Matyáš Dostalík, Zdenka Fohlerová
Electroactive Double-Layer Janus Biohybrid Hydrogel for Cardiac Tissue Engineering 148
Kateřina Krejčí, Michal Nohel, Jiří Chmelík, Roman Jakubíček, Jan Valošek
Segmentation of Cervical Spinal Rootlets in MRI Scans of Patients with Cord Compression 152
Michal Švento
Audio Inpainting with Endless Sounds 157
David Leitgeb
Gray-Box Modeling of Flanger Effect 162
Matej Liska
Defect-Aware Approach to Validation of MEMS Microphones in Spherical Arrays 167
Pavel Palurik, Antonin Cerny, Juraj Goryl, Dmitri Moltchanov, Radim Dvorak, Pavel Masek
Design and Implementation of a Modular IoT Gateway for Heterogeneous Networks 172
Martin Kousal, Roman Maršálek
Cognitive Radio implementation using GNU Radio 177
Radim Dvořák, Pavel Paluřík, Thao Le Dinh, Dmitri Moltchanov
The Impact of Over-the-Air Updates on Cellular IoT Systems Performance 182
Viet Anh Phan, Jan Jerabek
ML-Driven Real-Time In-Network Anomaly Detection on P4-Programmable Switches 187
Willi Lazarov, Roman Klampar, Zdenek Martinasek
Interactive Graph-based Environment for Open Source Intelligence 192
Patrik Dobias, Lukas Malina
Message Recovery Attack on ML-KEM 197
Karolína Fůsková
Emission optical spectrometry of plasma 202
Jan Schilhab, Alexandr Matějka, Petr Mrázek, Adam Šlechta
A Novel Self-Biased Current Reference for Ultra Low-Voltage Operation in 22 nm CMOS 207
Alexandr Matějka, Jan Schilhab, Adam Šlechta, Petr Mrázek
Low Voltage Low Noise Amplifier Design for an Ultrasonic Receiver in Automotive Applications 212
Antonín Šimek, Jaromír Kotlán, Ondřej Čech, David Trochta, Pavel Čudek, Tomáš Kazda
Preparation of hard carbon derived from waste biomass for use in sodium-ion batteries 218
Jiří Báňa, David Trochta, Tomáš Kazda
Design and Optimization of a Pre-treatment Line for Li-ion Battery Recycling 223

Jaromír Jarušek, Imrich Gablech
Fabrication of electrically conductive thick films using masked stereolithography 228
Marek Sedlařík, Petr Vyroubal, Martin Mačák
Development of a Reduced Order Thermal Model for Real-Time Lithium-ion Battery Storage Simulation 233
Miroslav Kozumplík
Semi-Automated Parameter Identification and Rapid Control Prototyping of a DC Motor Drive Using an STM32 Microcontroller 238
Vít Šimeček
Fatigue Life Analysis Methods for Vibration Shakers 243
Jan Čtvrtníček
Comparative Analysis of Fitness Functions in Coverage Path Planning 248
Stepan Jezek, Daniel Duraj, Radim Burget
Zero-Shot Traffic Accident Detection in Urban Surveillance Video Using Large Vision-Language Models 253
Miloš Cihlář, Petr Šopák, Stanislav Svědiroh, Tomáš Horeličan
Orpheus XAI: Autonomous Navigation and Immersive VR Telepresence for a Military-Grade Unmanned Ground Vehicle 258



OPENING WORD OF THE DEAN

This Proceeding contains papers presented during the 32nd annual STUDENT EEICT conference, held at the Faculty of Electrical Engineering and Communication, Brno University of Technology, on April 28, 2026. The fruitful tradition of joining together creative students and seasoned science or research specialists and industry-based experts gave again a valuable opportunity to exchange information and experience.

The EEICT involves multiple corporate partners, collaborators, and evaluators whose intensive support is highly appreciated. Importantly, the competitive, motivating features of the conference are associated with a practical impact: In addition to encouraging students to further develop their knowledge, interests, and employability potential, the forum directly offers career opportunities through the affiliated PerFEKT JobFair, a yearly job-related workshop and exhibition complementing the actual EEICT sessions. In this context, the organizers acknowledge the long-term assistance from the Ministry of Education, Youth and Sports of the Czech Republic, which has proved essential for refining the scope and impact of the symposium.

In total, 164 peer-reviewed full papers distributed between 20 sessions were submitted and evaluated by examining boards with industry and academic specialists. The presenting authors exhibited a very high standard of knowledge and communication skills, and the best competitors received prize money and/or valuable gifts. This Proceeding comprises 57 award winning full papers, all selected by the conference's evaluation boards.

Our sincere thanks go to the sponsors, experts, students, and collaborators who participated in, contributed to, and made the conference a continued success.

Considering all the efforts and work invested, I believe that the 32nd STUDENT EEICT (2026) has been beneficial for all the participants.

I am convinced that the inspiration gathered during the event will contribute to the further advancement of open science and research, giving all the attendees a chance to freely discuss their achievements and views.

prof. Ing. Jaroslav Koton, Ph.D.
Dean of the Faculty of Electrical Engineering and Communication

Real-Time Detection and Visualization of the Superficial Venous System

Ladislav Duraja
Department of Biomedical Engineering
Brno University of Technology, FEEC
Brno, Czech Republic
246994@vutbr.cz

***Abstract*— This paper deals with the non-contact detection and real-time visualization of the superficial venous system. Based on a literature review of current near-infrared (NIR) imaging methods, a device was designed using a Raspberry Pi platform with an additional NIR illumination unit and an infrared-sensitive camera. The proposed system allows for the detection of subcutaneous vein structures with immediate visualization of results, which is crucial for clinical practices such as intravenous injections. Pilot measurements confirmed the functionality of the IR illumination and the effectiveness of the chosen image processing algorithms.**

Keywords*— *venous system, near-infrared imaging, vein detection, real-time processing, image processing, Raspberry Pi, NIR camera, visualization

I. Introduction

The establishment of peripheral intravenous (IV) access is a critical procedure in modern medicine, required for the administration of fluids, medications, and blood products. Despite its frequency, the procedure is often complicated by "Difficult Intravenous Access" (DIVA) factors, such as obesity, dehydration, or dark skin pigmentation, which obscure the visual and tactile identification of veins. Failure to secure access on the first attempt leads to increased patient anxiety, physical trauma, and delays in treatment [1].

While several commercial visualization tools exist, they often present limitations. Devices like the AccuVein AV500 use laser projection (830 nm) which, although effective for superficial veins, can be costly for widespread deployment [2][3]. Ultrasound-based systems like the Butterfly iQ+ offer deep vessel visualization but require specialized training and direct skin contact [4].

This paper proposes a low-cost, non-contact alternative based on the Raspberry Pi 4 platform. The system utilizes the unique absorption properties of deoxygenated hemoglobin in the near-infrared (NIR) spectrum to isolate venous structures from the surrounding tissue [5][6].

The primary objective is to develop a real-time processing pipeline that enhances vessel contrast while maintaining a high frame rate suitable for clinical assistance.

II. Hardware Design and modifications

The hardware setup was optimized for high-speed acquisition and mechanical stability. The system is based on the Raspberry Pi 4B platform.

A. Computational Platform and Sensor

The system is built around the Raspberry Pi 4 Model B (4GB RAM) [7]. This single-board computer was selected for its native Camera Serial Interface (CSI), which allows for high-bandwidth data transfer directly from the sensor to the GPU, bypassing the latency associated with USB interfaces. The image is captured by an OV5647 CMOS sensor featuring a 5-megapixel resolution and a lens with an f/1.8 aperture. To enable NIR detection, the internal IR-cut filter was removed, making the sensor sensitive to wavelengths up to 1000 nm. The system maintains a native 480x800 vertical orientation for the 4.3" DSI display without performance degradation.



Fig. 1. Experimental hardware setup featuring the Raspberry Pi 4 Model B and NIR-sensitive camera module with integrated illumination unit

B. Custom NIR Illumination Control

Proper illumination is paramount for achieving sufficient contrast. The area of interest is illuminated by a dual-LED NIR array operating at a wavelength of approximately 850 nm. In the original module configuration, the LED intensity was controlled by a photoresistor (LDR). However, this led to insufficient brightness in well-lit clinical environments. To resolve this, a 50k Ω linear potentiometer was integrated into the circuit to

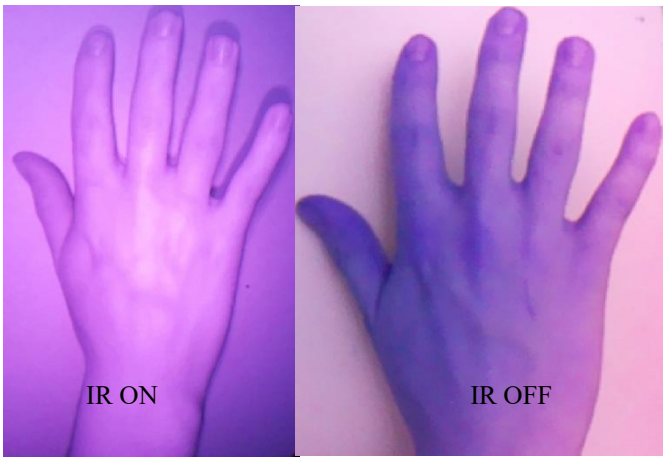

Fig. 2. Raw IR images with IR diodes ON/OFF

allow manual control over the drive current. This modification enables the operator to fine-tune the illumination to the specific skin type and depth of the patient's veins.

III. Optical Principles and Image Processing

To achieve maximum visualization of the venous network, a multi-stage digital processing pipeline was developed.

A. NIR Interaction with Biological Tissue

The physical basis for this device is the "optical window" of biological tissue in the NIR range (760–900 nm) [8][9]. In this region, skin and subcutaneous fat exhibit relatively low absorption and high scattering. In contrast, deoxygenated hemoglobin (Hb) acts as a strong absorber [5][6]. The attenuation of light through the tissue can be modeled by the Beer-Lambert law:

$$I = I_0 \cdot e^{-\mu_a \cdot d} \quad (1)$$

where I is the detected intensity, $I_0$is the incident intensity, μ_a is the absorption coefficient of the hemoglobin, and d is the path length. This difference in absorption coefficients causes the veins to appear as low-intensity (darker) regions compared to the surrounding tissue [8][9].

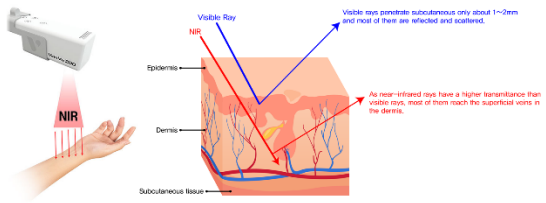

Fig. 3. Principle of vein detection [9]

B. Spectral Channel Extraction and CLAHE

The processing begins with the extraction of the green channel from the RGB input. In NIR-sensitive sensors, the

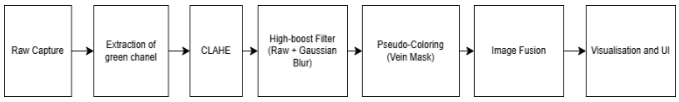

Fig. 4. Block diagram of image processing

green channel often exhibits the highest contrast for hemoglobin absorption. This is followed by Contrast Limited Adaptive Histogram Equalization (CLAHE) applied in 8x8 blocks to enhance local details without over-amplifying noise.

C. High-boost Filtering and Color Fusion

To further sharpen the vessel edges, a High-boost filter is implemented by subtracting a Gaussian-blurred version of the image from the original, effectively emphasizing high-frequency components:

$$Image_{sharp} = A \cdot Image_{orig} - Image_{blurred} \quad (2)$$

Finally, a color fusion technique is used. The original grayscale NIR image (70% weight) is blended with the filtered vein mask (30% weight) using the *addWeighted* function. The mask is tinted cyan to provide clear visual separation between the vessels and the surrounding tissue.

IV. Practical Implementation and Comparative Evaluation

To evaluate the efficacy of the proposed system, a series of pilot measurements were conducted under controlled laboratory conditions. The testing focused on the system's ability to isolate venous structures across different enhancement stages and its performance in a real-time environment.

A. Methodology and Setup

The experimental setup involved positioning the NIR camera at a constant distance of 35 cm from the dorsal side of the subject's hand. The camera was oriented perpendicularly to the skin surface to minimize perspective distortion and ensure uniform light distribution from the dual-LED NIR array. All recordings were captured at a resolution of 480 x 800 pixels, providing a balance between sufficient spatial resolution, low computational overhead and fitting to 4.3" touch display.

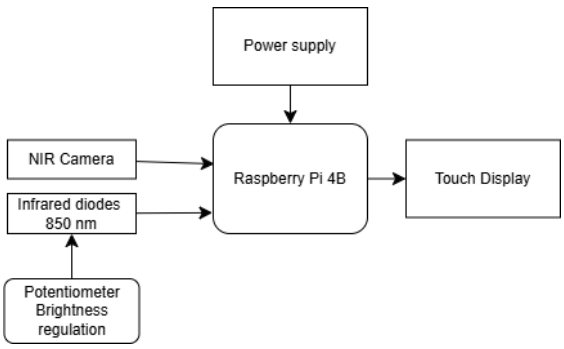

Fig. 5. Block diagram of the device

B. *Touch-friendly User Interface (GUI)*

A custom graphical user interface was developed to facilitate real-time interaction. The system includes a touch handler for the DSI display, featuring intuitive "FOTO" (screenshot) and "STOP" (stops program) buttons. A real-time FPS counter is integrated to monitor system stability.

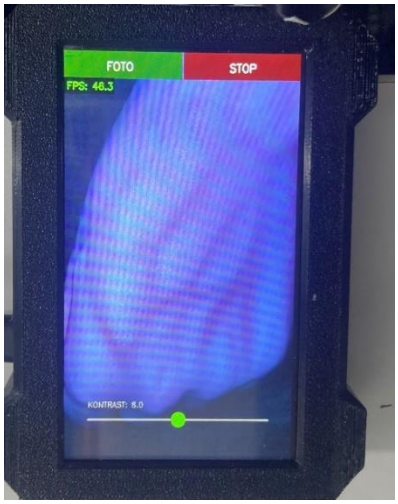

Fig. 6. Display with interaction buttons.

C. *Adaptive Contrast Control*

An interactive graphical slider allows the operator to dynamically adjust the CLAHE *clipLimit* in Fig. 6 a parameter labeled "*Kontrast"* from 1.1 to 15.0. Pilot testing identified three distinct operating modes:

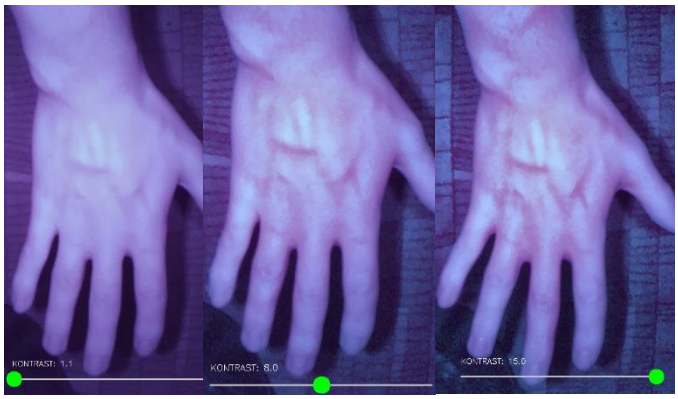

Fig. 7. Dorsum of the hand images with different clipLimit.

D. *Evaluation on Diverse Subjects*

Beyond basic verification, the system was tested on a group of 3 subjects representing diverse physiological profiles to analyze the impact of BMI, age, and dermal features. The characteristics of the experimental group are summarized in Table I and examples of individual comparisons are shown in Figures 8, 9, 10, and 11.

TABLE I.

PHYSIOLOGICAL CHARACTERISTICS AND FEATURES OF TESTED SUBJECTS

Subject ID	AGE	Gender	BMI CATEGORY	SPECIFIC FEATURES	Region of Interest	Figure number
01	12	Female	Healthy Weight	-	Dorsum of hand	8
02	24	Male	Healthy Weight	Forearm tattoo	Antecubital	9
03	57	Male	Overweight	Pigment spots	Dorsum of hand / Antecubital	10/11



Fig. 8. Comparison between standard visual capture (left) and the enhanced NIR visualization (right) for Subject 01. (Pediatric).



Fig. 9. NIR visualization near dermal tattooing.

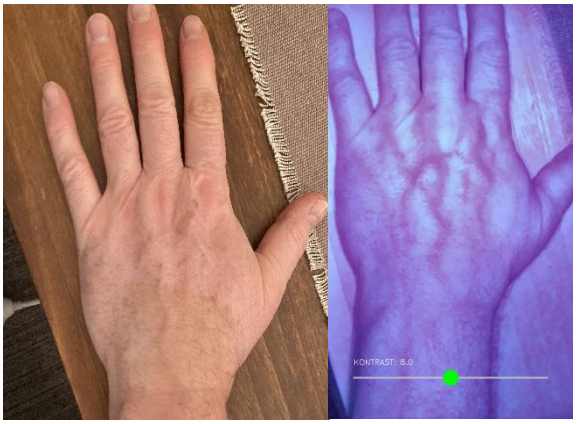

Fig. 10. NIR visualization on the dorsum of the hand for Subject 03 (Overweight).

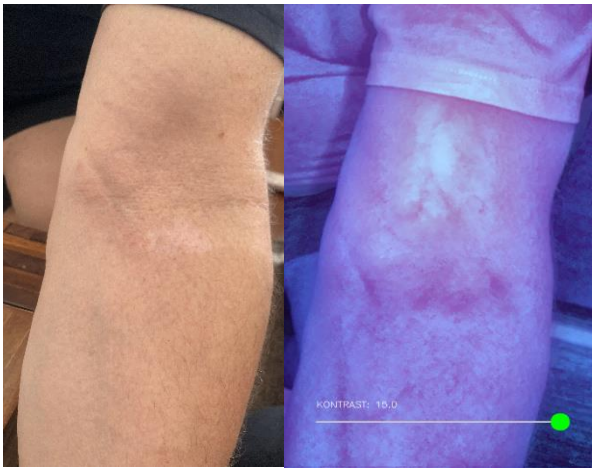

Fig. 11. Comparative view of the antecubital fossa with small pigment spot.

V. Results and Discussion

The experimental evaluation confirms the functional viability of the proposed NIR imaging system for real-time superficial vein visualization. However, to fully assess its clinical applicability, it is necessary to critically analyze the method's inherent limitations and the environmental factors that influence its performance.

A. Image Enhancement Results

As shown in Fig. 2. the raw NIR image provided only basic low contrast of the venous network. The application of green-channel extraction and CLAHE processing noticeably improved the local contrast. By utilizing a high-boost filter, the edges of the superficial veins became clearly defined. The final color fusion (70% original, 30% cyan-tinted mask) allowed for an intuitive visual separation of the vascular structures shown on Fig. 7.

B. Influence of Physiological Factors, Tattoos and Ambient Light Conditions

The main limitation of optical vein detection is the thickness of subcutaneous adipose tissue. Adipose tissue exhibits high scattering properties at a wavelength of 850 nm. As observed in individuals with a higher BMI, this scattering acts as a physical low-pass filter that attenuates the reflected signal from deeper-lying vessels, for example in the *antecubital fossa* (see Fig. 11), and significantly reduces the contrast of the resulting image.

Unlike age-related pigment spots, which do not significantly affect NIR imaging, extensive dermal tattoos represent a significant limitation. As observed during the evaluation of Subject 02, tattooing presents a critical challenge because the ink consists primarily of carbon-based compounds. Since this ink is deposited directly within the *dermis*, it completely blocks the optical contrast between the superficial veins and the surrounding tissue in this case.

Reducing the influence of ambient lighting can improve image quality by minimizing interference caused by unwanted light components. Ideally, the scene is illuminated primarily by controlled infrared lighting, which allows for more stable and reproducible results.

C. System Performance

An important result of this work is the achieved real-time performance. The system maintains a stable frame rate of more than 30 FPS, which is essential for dynamic medical procedures such as catheter insertion. Despite this, rapid limb movements may introduce temporary motion blur and affect the stability of CLAHE processing, particularly at tile boundaries. This limitation indicates that patient motion can reduce image quality during initial vein localization.

References

[1] I PUTU ADI SURYA GUNAWAN; SIGIT, Riyanto a GUNAWAN, Agus. Vein Visualization System Using Camera and Projector Based on Distance Sensor. Online. In:2018. Available from: https://doi.org/10.1109/elecsym.2018.8615501. [cit. 2025-10-01].

[2] ACCUVEIN. *Vein visualization system* [online]. Medford, NY: AccuVein Inc., ©2025 [cit. 2025-10-20]. Available from: https://www.accuvein.com/

[3] J ACCUVEIN INC. *AccuVein AV500 User Manual* [online]. Medford, NY: AccuVein Inc., 2023 [cit. 2025-10-20]. Available from: https://www.accuvein.com/wp-content/uploads/2023/08/AccuVein_AV500_User_Manual.pdf.

[4] BUTTERFLY NETWORK INC. *Butterfly iQ+ Portable Ultrasound System* Burlington, MA: Butterfly Network Inc., 2024 [cit. 2025-10-20]. Available from: https://www.butterflynetwork.com/iq-plus

[5] SCOTT A. PRAHL, 1999. Optical absorption of hemoglobin. Online. 1999-01-01. Available from: https://omlc.org/spectra/hemoglobin/ [cit. 2025-10-01].

[6] TANG, Chaoying; QIAN, Mengen a WANG, Biao. *Deep learning-based vein* visualization and its mobile implementation. Online. Biomedical Signal Processing and Control. 2024, vol. 102, s. 107272-107272. ISSN 1746-8094. Available from: https://doi.org/10.1016/j.bspc.2024.107272. [cit. 2025-10-01].

[7] DATACAPTURECONTROL**.** *Raspberry Pi 4B (Top View)*. [online obrázek]. 2023. In: *Raspberry Pi 4B Board Layout, Pinout, I/O, Power, and Specifications*. [Cit. 2025-12-03]. Available from: https://datacapturecontrol.com/articles/io-devices/single-board-computers/raspberry-pi-4b

[8] ZHOU, Xinkai; XIA, Yunjia; UCHITEL, Julie; COLLINS-JONES, Liam; YANG, Shufan et al. Review of recent advances in frequency-domain near-infrared spectroscopy technologies [Invited]. Online. *Biomedical Optics Express*. 2023, vol. 14, no. 7, s. 3234. ISSN 2156-7085. Available from: https://doi.org/10.1364/BOE.484044. [cit. 2025-11-13].

[9] VEINVU, "Blood Vein Scanner: Portable Vein Finder Device," 2026. [Online]].[cit. 2026-03-15].Available from: https://veinvu.com/vein-finder-device/blood-vein-scanner/

[10] SANDELL, Julia a TIMOTHY C. ZHU. A review of *in-vivo* optical properties of human tissues and its impact on PDT. Online. *Journal of Biophotonics*. 2011, vol. 4, no. 11-12, s. 773-787. ISSN 1864-063X. Available from: https://doi.org/10.1002/jbio.201100062. [cit. 2025-11-10].

Development of an Optimized Workflow for the Fabrication and Quantitative Analysis of Microstructured Bio-Substrates

1st Veronika Šromová
Department of Physics
Brno University of Technology
Technická 2848/8, 61600 Brno, the Czech Republic
237011@vutbr.cz

2nd doc. Mgr. Dinara Sobola, Ph.D.
Institute of Physics of Materials
Czech Academy of Sciences
Žižkova 513/22, Brno, 61662
sobola@vutbr.cz

***Abstract*—Reliable evaluation of newly fabricated biomaterial substrates requires reproducible fabrication protocols and objective methods for quantifying cellular responses. This study presents an integrated workflow for the design, fabrication, and quantitative validation of microstructured biomaterials produced by resin-based additive manufacturing.**

Microtopographical substrates were designed using CAD software and fabricated via resin 3D printing using a biocompatible photopolymer. Printability limits were assessed through scanning electron microscopy. To objectively quantify biological response, a custom MATLAB-based image analysis programme was developed to measure cell confluency from phase-contrast microscopy images. This pipeline works on the principle of automated segmentation and systematic spatial sampling.

Indirect cytotoxicity testing was performed using LW13K2 cells exposed to extracts from printed samples subjected to different UV post-curing intervals and sterilization protocols. Statistical analysis using Two-Way ANOVA demonstrated that autoclave sterilization significantly improved cell proliferation compared with ethanol sterilization ($p < 0.001$), while extended UV curing showed no statistically significant effect on cell viability.

The presented workflow combines additive manufacturing, automated image quantification, and statistical modeling to provide a reproducible framework for evaluating the biocompatibility of 3D-printed biomaterial substrates.

***Index Terms*—3D printing, cell viability, image analysis, sterilization methods**

I. Introduction

Additive manufacturing has become an important tool in tissue engineering and biomaterials research, as it enables the fabrication of substrates with precisely controlled surface architectures [1]. Such microstructured interfaces are then investigated for their ability to influence cellular processes, such as adhesion, proliferation, and migration [2]–[4]. However, the reliable evaluation of newly fabricated biomaterial surfaces requires not only precise manufacturing but also robust and objective methods for assessing biological responses.

While traditional manual counting approaches or qualitative microscopic observation remain a standard practice, they are often time-consuming and inherently subjective [5]. Variability between observers, differences in expertise, and fatigue during prolonged counting sessions can all lead to inconsistent results. Furthermore, such approaches are not well suited for large-scale analyses [6]. To address these limitations, this research introduces an automated image analysis technique designed to support and complement traditional biological observation.

The present study proposes such a workflow by combining CAD-based design, 3D printing, automated MATLAB image segmentation, and statistical analysis. The objective was to establish a reproducible framework for evaluating the cytocompatibility of printed biomaterial substrates and for optimizing post-processing conditions through quantitative biological measurements.

II. Materials and Methods

A. CAD Design of Microstructured Substrates

The microstructured substrates investigated in this study were designed using the computer-aided design (CAD) software SolidWorks (Dassault Systèmes, France) [7]. The design process focused on generating well-defined microtopographical features, such as spine and microwell arrays, that could later be used to investigate interactions between cells and the material's surface. At the same time, it was necessary to design the individual models so that they met the resolution limits of the printer. Particular attention was therefore paid to feature dimensions and spacing in order to ensure successful fabrication without printing artifacts.

B. Additive Manufacturing and Post-Processing

Fabrication of the designed substrates was performed using an Elegoo Saturn 3D resin printer (Elegoo Inc, China). The printer utilizes UV light to start a chemical reaction that bonds monomers and oligomers into a polymer network. Printing was performed with a layer thickness of 0.05 mm and an exposure time of 20 s per layer. The bottom layer count was set to 5, followed by 10 transition layers with a linear exposure profile. The actual fabrication time was approximately 1 hour.

A commercially available biocompatible photopolymer resin Liqcreate Bio-Med Clear was used as a printing material [8]. This resin is specifically developed for medical and

laboratory applications and is compatible with sterilization procedures commonly used in biological experiments.

Following fabrication, the printed samples underwent a multi-step post-processing procedure to remove residual uncured resin and ensure complete polymerization. First, the samples were rinsed in Liqcreate Resin Cleaner to eliminate excess resin from the surface. Subsequently, they were subjected to UV post-curing with a specific wavelength of 405 nm for defined time intervals using a UW-01 Washing/Curing Machine (Creality) [9] in order to further crosslink the polymer network and reduce the presence of residual monomers.

To evaluate the influence of curing conditions on biological compatibility, selected samples were exposed to multiple post-curing durations ranging from 10 to 150 minutes. These samples were later used in indirect cytotoxicity experiments.

C. Morphological Characterization

The resolution limits of the printer were characterized using LYRA3 scanning electron microscope (Tescan, Brno, Czech Republic) to define optimal parameters and prevent structural artifacts.

Prior to imaging, the samples were sputter-coated using a coater EM ACE600 (Leica, Wetzlar, DEU) with a thin conductive layer to minimize charging effects during electron beam exposure. SEM imaging then allowed detailed visualization of the printed surface morphology and potential printing artifacts.

This analysis therefore provided valuable feedback for refining design parameters. Also, this step was important for ensuring that the fabricated microstructures remained within the effective resolution capabilities of the printing system.

D. Cell Culture

Cell viability was evaluated using the LW13K2 fibroblast cell line, which is originally isolated from spontaneously transformed rat embryo fibroblasts of inbred strain Lewis. These cells exhibit spindle-shaped morphology and they grow in monolayers, with doubling time of approximately 24-30 hours.

LW13K2 cells were cultured in standard growth medium (MEM; Minimum Essential Medium) under controlled conditions at 37 °C in a humidified atmosphere containing 5% CO_2.

E. Microscopy and Image Acquisition

Cell proliferation and viability was monitored using phase-contrast microscopy. This non-invasive technique allowed for the direct visualization of living cells within their culture dishes. Observations were conducted using an Eclipse TS100 inverted microscope (Nikon, Tokyo, Japan) equipped with a CFI Achromat ADL 10× objective for routine morphological and confluency assessments.

Digital images were captured using a 1/2" CCD camera (1280 × 950 px resolution). All image acquisition and primary metadata analysis were managed using NIS-Elements D software (Version 4.10, Nikon).

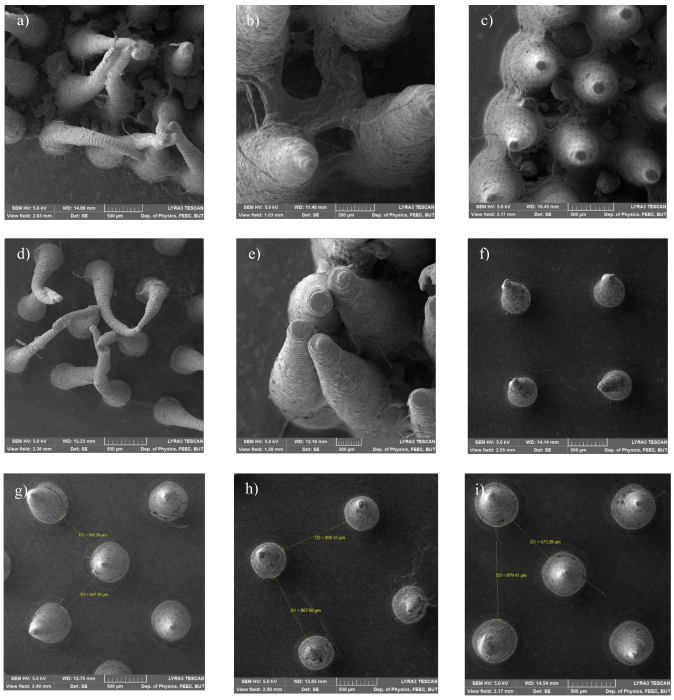

Fig. 1: SEM revealing some of the morphological outcomes of the designed samples; reducing the spacing between the spines to less than 1 mm increased unwanted artifacts.

F. Indirect Cytotoxicity Protocol

The indirect cytotoxicity assessment was performed to evaluate the compatibility of the printed material independently of surface topography. Cubic test specimens ($1 \times 1 \times 1$ cm^3) were fabricated using the same material and printing parameters as the microstructured substrates. These samples were subjected to ten distinct UV post-curing intervals (10-150 min). Following curing, two sterilization methods were compared:

- Autoclave: 121 °C at 15 psi for 60 minutes.
- Ethanol: 70% ethanol immersion for 30 minutes followed by rinsing in phosphate-buffered saline.

Sterilized samples were incubated in 2.5 ml of MEM culture medium for 24 hours at 37 °C in order to obtain material extracts that could potentially contain residual leachables. This corresponded to an extraction ratio of 2.4 cm^2/ml based on the sample surface area, which is appropriate for solid, non-porous materials. The extracts were subsequently applied to LW13K2 fibroblast cells cultured under standard conditions. Cell viability was subsequently evaluated after 72 hours of incubation with the extracts.

A negative control group was included and it consisted of pure, untreated MEM culture medium that was incubated under the same conditions (37 °C, 24 hours) but without any contact with the printed material.

G. MATLAB-Based Quantitative Analysis

Quantification of cellular proliferation was performed using a custom image-processing workflow developed in MATLAB

(MathWorks, USA) [10]. The primary objective of the analysis pipeline was to provide an objective and reproducible measurement of cell coverage from phase-contract microscopy images.

To ensure representative sampling across each culture dish, a systematic spatial sampling strategy was employed. For every sample, twenty-five independent fields of view (FOVs) were captured across the culture surface. This sampling grid included one central region and multiple peripheral locations distributed across the dish in order to reduce spatial bias. Based on the dimensions of the 1/2" CCD sensor (6.4 × 4.8 mm) and the magnification of the 10× objective lens, each FOV corresponded to a physical area of approximately 0.64 × 0.48 mm. The analysis therefore covered a total non-overlapping area of approximately 7.68 mm^2 per culture dish.

Image processing consisted of several sequential steps. First, grayscale images were preprocessed to improve contrast between cells and background. Adaptive thresholding was then applied to generate a binary mask representing cell-covered regions.

Cell proliferation was quantified using the confluency metric, defined as the percentage of pixels corresponding to detected cellular regions relative to the total number of pixels in the image:

$$\text{Confluency}(\%) = \left(\frac{\sum \text{pixel}_{\text{cell}}}{\sum \text{pixel}_{\text{total}}} \right) \times 100 \qquad (1)$$

The confluency values obtained from the individual FOVs were averaged to produce a single quantitative value representing each biological replicate.

H. Statistical Analysis

Statistical evaluation of the quantitative data was performed using a two-factor analysis of variance (Two-Way ANOVA). The analysis was designed to assess the influence of two independent variables on cellular proliferation:

- Sterilization method (autoclave or ethanol),
- UV post-curing duration.

Table I shows results from the statistical evaluation. The dependent variable in the model was the measured cell confluency obtained from the MATLAB-based image analysis. The ANOVA model allowed evaluation of the individual effects of both factors as well as potential interactions between them. Statistical significance w as a ssessed a t a t hreshold l evel of $p < 0.05$.

TABLE I: Two-Way ANOVA results evaluating the main effects of the Sterilization Method and UV Curing Duration on LW13K2 cell confluency at the 72-hour mark.

Source	Sum Sq.	d.f.	Mean Sq.	F	*p*-value
Sterilization	28895.8	1	28895.8	93.62	< 0.001
Curing Group	2451.0	9	272.3	0.88	0.542
Interactions	2426.4	9	269.6	0.87	0.550
Error	55554.3	180	308.6		
Total	91911.0	199			

III. Results and Discussion

A. Fabrication Fidelity and Structural Integrity

The morphological assessment via LYRA3 scanning electron microscope (Fig. 1) confirmed that while the printer is capable of producing suitable substrates for subsequent biological testing, the inter-feature spacing is the critical limiting factor. Reducing the spacing between spines to less than 1 mm resulted in significant unwanted artifacts such as incomplete drainage of uncured material between features.

Furthermore, minor surface roughness typical for layer-by-layer photopolymerization was observed, but no major structural defects such as feature collapse or incomplete polymerization were detected when the specified design limits were respected.

It should be noted that while this resolution is sufficient for large-scale micro-architectures, applications requiring submicron or nanostructured topography would necessitate the use of higher-resolution systems.

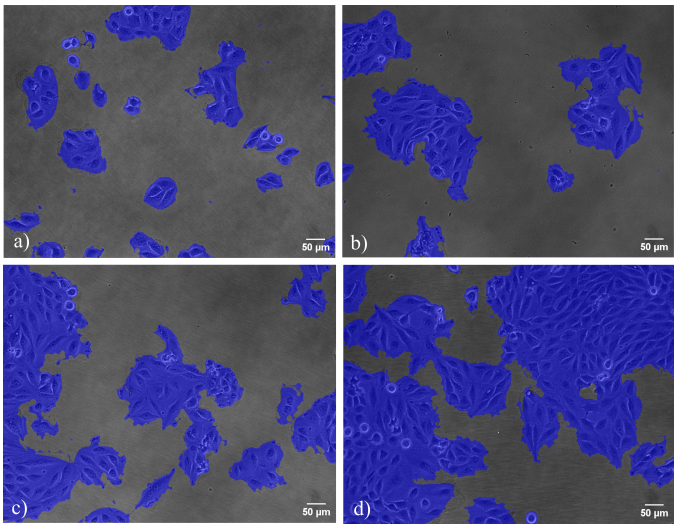

Fig. 2: Calculation of confluence using implemented code in the MATLAB programming language. The figure shows some of the final outputs. The pictures were taken from cultures in the early stages before reaching a confluent monolayer. The calculated confluency in these Fields of View are: a) 17.11%; b) 22.29%; c) 36.91%; d) 86.54%

B. Performance of the Image Analysis Pipeline

The developed MATLAB image-processing pipeline successfully identified cellular regions across a wide range of cell densities (see Fig. 2). The segmentation algorithm was able to distinguish adherent cell clusters from the background in phase-contrast microscopy images. The use of adaptive thresholding allowed the algorithm to remain robust against variations in illumination and image contrast.

Together, these elements enabled reproducible estimation of cell confluency suitable for statistical evaluation. However, it should be noted that confluency represents an indirect measure of cell proliferation, as cell spreading and morphology may

vary. Therefore, complementary assays could further refine the assessment of proliferation.

C. Quantitative Cytotoxicity Evaluation

Quantitative analysis of the microscopy images revealed clear differences in cell proliferation depending on the applied sterilization protocol. The Two-Way ANOVA indicated that the sterilization method had a statistically significant effect on the measured confluency values ($F = 93.62$, $p < 0.001$).

Samples sterilized using autoclaving exhibited high confluency values ranging from approximately 90% to 97%, comparable to the negative control group. In contrast, samples sterilized with ethanol showed significantly reduced cell proliferation (see Fig. 3). However, the marked reduction in cell proliferation observed in this ethanol-treated group warrants further analysis. The 30-minute ethanol immersion likely exceeded the resin's chemical resistance threshold, potentially causing polymer swelling or micro-fractures that facilitated the release of unreacted monomers into the extract.

The duration of UV post-curing did not produce a statistically significant effect on cell confluency ($p = 0.542$), indicating that extended curing beyond the shortest tested interval did not substantially improve cytocompatibility. These findings suggest that the elevated temperature and pressure of the autoclave may promote additional polymer crosslinking and facilitate the removal or degradation of residual monomers, thereby improving the biological compatibility of the printed material.

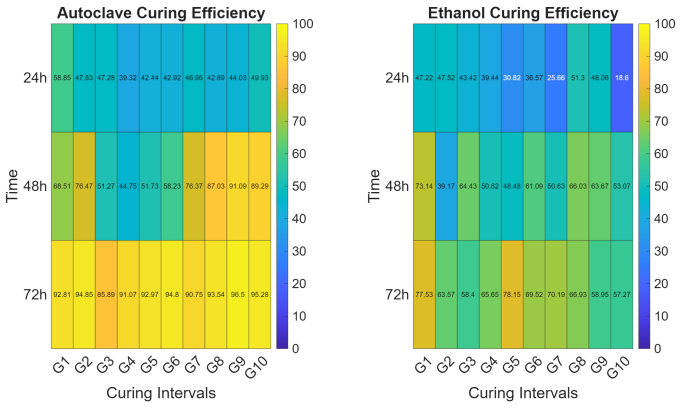

Fig. 3: This heatmap visualizes the progression of average LW13K2 cell confluency over 24, 48, and 72 hours. The color gradient indicates cell confluency, with darker/cooler colors representing lower cell proliferation, and brighter/warmer col-ors representing higher confluency rates.

IV. Conclusion

This study presents an integrated methodological workflow for the evaluation of resin-based 3D printed biomaterial substrates.

The developed image-processing pipeline enabled objective quantification of cell proliferation through automated segmentation and systematic spatial sampling of microscopy images. This approach reduced observer-related variability and provided reliable quantitative metrics suitable for statistical comparison between experimental groups.

Statistical evaluation of the indirect cytotoxicity assay demonstrated that the sterilization method significantly influences the biological compatibility of the printed material. Autoclave sterilization produced confluency levels comparable to the negative control, whereas ethanol sterilization resulted in reduced cell proliferation. In contrast, extended UV post-curing did not significantly affect cell viability within the tested range.

Furthermore, the methodology may serve as a foundation for future studies investigating how engineered surface architectures and fabrication parameters influence cellular responses.

Acknowledgement

This publication was supported by the project "Mechanical Engineering of Biological and Bioinspired Systems", funded as project No. CZ.02.01.01/00/22_008/0004634 by Programme Johannes Amos Commenius. This work was conducted as a bachelor's thesis under the supervision of doc. Mgr. Dinara Sobola, Ph.D. I would also like to thank Ing. Pavel Kaspar, Ph.D. for guidance during the preparation of this paper.

References

[1] A. Borah and D. S. Kumar, "7 - Overcoming the barriers of two-dimensional cell culture systems with three-dimensional cell culture systems: techniques, drug discovery, and biomedical applications," in Biomedical Product and Materials Evaluation, Woodhead Publishing Series in Biomaterials, P. V. Mohanan, Ed. Woodhead Publishing, 2022, pp. 179–229. doi: 10.1016/B978-0-12-823966-7.00003-7.

[2] N. Gui, W. Xu, D. Myers, R. Shukla, H. Tang, and M. Qian, "The effect of ordered and partially ordered surface topography on bone cell responses: A review," *Biomaterials Science*, vol. 6, pp. 250–264, Dec. 2017. doi: 10.1039/C7BM01016H.

[3] R. Tudureanu, I. Handrea-Dragan, S. Boca, and I. Botiz, "Insight and recent advances into the role of topography on cell differentiation and proliferation on biopolymeric surfaces," *International Journal of Molecular Sciences*, vol. 23, 2022. doi: 10.3390/ijms23147731.

[4] D. Hamilton, S. Ghrebi, H. Kim, B. Chehroudi, and D. Brunette, "Surface Topography and Cell Behavior," pp. 1–15, May 2008. doi: 10.1201/b18990-246.

[5] D. S. J. Ting, G. S. L. Peh, D. J. H. Neo, X. Y. Ng, B. Y. L. Tan, R. C. B. Wong, H. S. Ong, and J. S. Mehta, "Cell counting and cell cycle analysis of simple non-cultured endothelial cell injection (SNEC-I) therapy: Characterization for clinical translation," *Cells*, vol. 14, p. 986, 2025. doi: 10.3390/cells14130986.

[6] D. Cadena-Herrera, J. E. Esparza-De Lara, N. D. Ramírez-Ibañez, C. A. López-Morales, N. O. Pérez, L. F. Flores-Ortiz, and E. Medina-Rivero, "Validation of three viable-cell counting methods: Manual, semi-automated, and automated," *Biotechnology Reports*, vol. 7, pp. 9–16, 2015. doi: 10.1016/j.btre.2015.04.004.

[7] Dassault Systèmes SolidWorks Corp., "SOLIDWORKS," version 2023, Waltham, MA: Dassault Systèmes, 2023. [Online]. Available: https://www.solidworks.com.

[8] Liqcreate, "Bio-Med Clear," 2025. [Online]. Available: https://www.liqcreate.com/product/bio-med-clear-biocompatible-resin/.

[9] Creality, "UW-01 Washing/Curing Machine," Creality - Official Website. [Online]. Available: https://www.creality.com/products/creality-uw-01-washing-curing-machine.

[10] The MathWorks, Inc., "MATLAB," version 2023b, Natick, MA: The MathWorks, Inc., 2023. [Online]. Available: https://www.mathworks.com.

Analysis of High-Frequency Components of the P Wave

Helena Géryková
Department of Biomedical Engineering
FEEC, Brno University of Technology
Brno, Czech Republic
256982@vutbr.cz

Radovan Smíšek
Department of Biomedical Engineering
FEEC, Brno University of Technology
Institute of Scientific Instruments of the CAS
Brno, Czech Republic
smisek@vutbr.cz

***Abstract*—This paper focuses on the high-frequency components of the P wave in electrocardiogram (ECG) recordings. The study aims to determine the optimal frequency range for reliable analysis and to assess the high-frequency characteristics of healthy and pathological patients.**

Two complementary methods were employed: the spectrogram method and frequency band signal-to-noise (SNR) analysis. The spectrogram method was used to determine the highest frequency at which the P wave can be reliably analysed. Results showed that the high-frequency components of the P wave were reliably detected up to 140 Hz in recordings across three patient groups.

The frequency band analysis observed the SNR in fourteen predefined frequency bands from 60 to 1500 Hz. The highest SNR values were detected in the first and lowest frequency band (60–150 Hz). A clear difference was observed between all three groups, with the highest value present in the healthy group, lower values in the group with paroxysmal fibrillation, and the lowest SNR in the group with atrial pacemakers.

These results demonstrate that spectrogram and band-based analyses provide complementary information about the spectral properties and quality of the P wave signal.

***Index Terms*—P wave, ECG, high-frequency analysis, signal averaging, signal-to-noise ratio**

I. Introduction

Electrocardiography (ECG) remains a valuable tool used to assess the electrical activity of the myocardium. The ECG recording is composed of several components, one of which is the P wave. Small abnormalities in P wave morphology can indicate pathological cases involving atrial arrhythmias.

Despite its diagnostic potential, the P wave remains susceptible to electromagnetic interference, baseline wander, and physiological artifacts. Signal averaging methods may be applied to improve the quality of P wave analysis [1].

High-frequency (HF) analysis of the ECG has gained considerable attention in recent years, particularly for studying ventricular depolarization. Studies have demonstrated that HF components of the QRS complex contain critical information regarding myocardial conduction and arrhythmias that may be obscured in conventional low-frequency analysis. However, the application of HF analysis to the P wave has received considerably less attention, likely due to the aforementioned amplitude limitations [2], [3].

The aim of this paper is to explore signal processing methods for HF analysis of the P wave and to determine the optimal frequency range based on SNR evaluation. This work may also provide methodological insights for future applications of HF P wave analysis in both clinical and research settings.

II. Theoretical Background

A. P wave and Atrial Depolarization

The P wave represents atrial depolarization and reflects the propagation of electrical activation through the myocardium. Under physiological conditions, excitation originates in the sinoatrial node and spreads across both atria [4]. Compared to ventricular activity, atrial conduction is slower and involves a lower myocardial mass, resulting in a P wave with relatively low amplitude and a longer duration. Typically, the length of the P wave ranges between 80 and 120 ms, and the amplitude is usually in the order of tens of microvolts [5].

B. Signal-Averaged P Wave

To enhance low-amplitude ECG components, signal averaging techniques are employed. In the context of P wave analysis, extracted wave segments are aligned and averaged, which leads to the suppression of uncorrelated noise while preserving consistent atrial activity. This approach is labeled as signal-averaged P wave (SAPW) and has been used previously in the assessment of atrial conduction pathologies and atrial fibrillation. This improves the signal-to-noise ratio and enables the analysis of microvolt-level features that may not be observed in the raw signal [1]. Accurate alignment of P wave segments is a crucial step in effective averaging and subsequent frequency-domain analysis.

C. Signal-to-Noise Ratio

The signal-to-noise ratio is a fundamental metric used for the evaluation of biosignals and the reliability of extracted features. SNR is defined as the ratio of the desired signal power to the power of background noise. It can also be calculated using the ratio of signal and noise amplitudes. Low SNR values indicate that relevant signal features are polluted by noise, whereas higher values suggest a sufficient quality of the signal. For high-frequency analysis of the P wave, SNR-based evaluation is necessary, as higher frequencies can be accompanied by increased sensitivity to noise. Assessment of SNR across different frequency ranges provides an objective

upper limit of frequencies at which atrial activity can be reliably analysed [2], [6].

III. Methods

A. Data

The analysis was performed on multi-lead ECG recordings obtained from clinical datasets. The datasets were captured at St. Anne's University Hospital in Brno and provided by VDI Technologies. The sampling frequency was 5 kHz, and six precordial leads were used for evaluation. A high sampling frequency was necessary to preserve high-frequency components. Patients were split into three groups: healthy patients (n=87), patients with paroxysmal atrial fibrillation (PAF)(n=4), and patients with implanted atrial pacemaker (n=6). Recordings of the patients with PAF were obtained during their time in sinus rhythm. P waves were detected using the reference points of the QRS complex, obtained in the accompanying annotation files.

B. P wave detection

For every QRS complex, a predefined window of 220–120 ms before the centre of the QRS was used to determine the P wave peak. The algorithm accounted for both positive and negative peaks, as well as the minimum peak distance to the QRS complex. P wave candidates were independently identified in each of the six precordial leads. Candidate peaks were grouped based on proximity using a window of 25 ms. The final P wave peak position was calculated as the weighted mean of the grouped peaks. To ensure successful detection, a P to R peak distance check was employed in a 120–220 ms window before the R peak. Finally, a median based filter with a tolerance of 35 ms was applied to remove any outlier detections. An example of P wave detection is shown below in Figure 1.

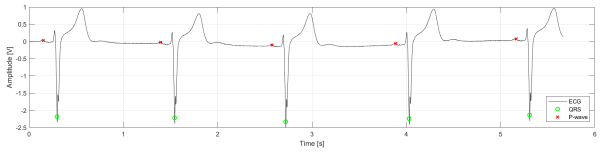

Fig. 1: ECG recording, with annotated QRS complex positions, and newly detected P waves. From lead V3.

C. Extraction, Alignment and Outlier Rejection

P wave segments were extracted using a window of 120 ms before and 60 ms after the P wave peak. Baseline wander was corrected by subtracting the mean of the first 20 ms of every segment. To improve averaging and consistency, P wave segments were aligned using delay estimation based on the median wave template. The final delay was calculated by obtaining the median delay across the evaluated leads. Morphological outliers were rejected based on the Pearson correlation coefficient (threshold 0.95) between each segment and the template segment. When fewer than ten segments remained, thirty segments with the highest correlation scores were retained. Aligned and normal segments are illustrated in Figure 2.

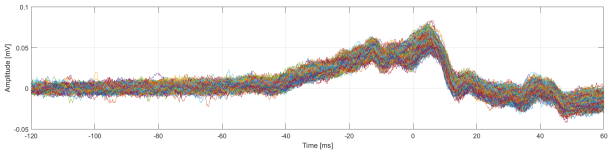

Fig. 2: P wave segments after alignment and removal of outliers. Lead V3.

D. Spectral Analysis

Two independent approaches were employed to assess the high frequency of the P wave. The first of these was a time-frequency approach using short-time Fourier transformation (STFT) with a 64-sample window, 75% overlap, and $NFFT = 2048$. The second method was based on predefined frequency bands.

The second method relied on the transformation of the signal using the Fast Fourier Transform (FFT), and specific frequency components corresponding to predefined bands were retained. The filtered signal was reconstructed using inverse FFT, and the envelope of the signal was extracted. In total, fourteen overlapping bands, ranging from 60 to 1500 Hz, were evaluated. Lastly, P waves were extracted and averaged, which can be seen in Figure 3.

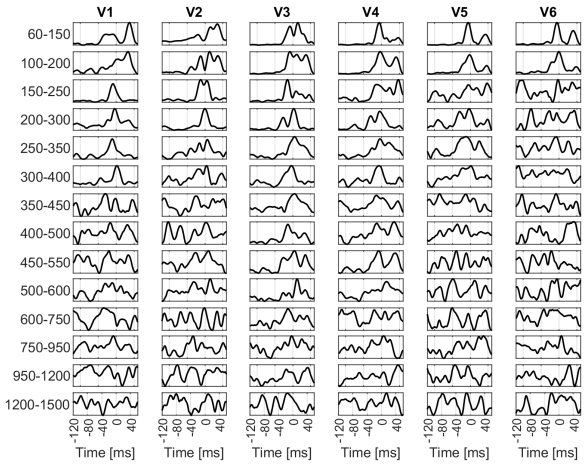

Fig. 3: Average P waves from a healthy patient across 14 frequency bands (rows), and six precordial leads (columns). Time 0 corresponds to the detected P wave peak.

E. Evaluation of SNR

For both approaches, SNR was computed by comparing the signal and noise within predefined time intervals relative to the P wave peak. In both cases, two temporal windows were defined: the signal window included -45 ms before and +40 ms after the P wave peak, while the noise window was defined from -120 ms to the -70 ms before the P wave peak. In the pacemaker group, the noise window was shortened to -100 ms to avoid the artifact from atrial stimulation.

In the spectrogram analysis, the SNR was evaluated for each frequency bin for each detected P wave as the ratio of the mean power of the signal window to the median power of the noise window. Cutoff frequency was then defined as the highest frequency within 40—600 Hz at which the SNR remained continuously above a predefined threshold of 8.

In the frequency band analysis, SNR was computed separately for each band as the RMS ratio between the signal and noise window of the averaged P wave. In this method, the signal was deconstructed into predefined frequency bands and later reconstructed, ensuring that the signal contained only specific bands.

Both methods served an important role in the analysis of high-frequency components. The first method (spectrogram) was primarily used to determine the optimal frequency range and to evaluate the presence of said HF components. The second method (frequency bands) served as both a visualization tool and an independent confirmation of the range set by the spectrogram.

F. Creation of High-Frequency Maps

Along with the SNR evaluation, high-frequency maps were computed for visualization purposes. To create these maps, the P waves from the first four frequency bands (60–300 Hz) were averaged. For each patient, a matrix containing the averaged P waves for all leads was then normalized so that the maximum value was 1. The signal from the first lead was inserted at the top of the image to enable interpolation of the remaining leads. Interpolation was subsequently performed between each pair of leads using a predefined interpolation factor of 80. The result is a high-frequency map displaying the propagation of the P wave across leads.

IV. Results

A. Spectrogram Analysis

Analysis was conducted across three patient groups and precordial leads. In the healthy group, the median cutoff frequency was 141.6 Hz across all leads, with a standard deviation of 29.0 Hz. The lowest median was observed in lead V6 at 139.2 Hz.

In the patient group with paroxysmal atrial fibrillation (PAF), the median cutoff frequency was lower at 135.5 Hz. The lowest median cutoff was detected in V1 at 128.2 Hz, and the highest was detected in V3 at 140.4 Hz. The standard deviation was also low at 22.5 Hz.

Finally, in the patient group with atrial pacemakers, the median cutoff frequency across all leads was 142.2 Hz, nearly identical to that of the healthy group. Nevertheless, the standard deviation was higher (45.0 Hz) with a wider range of 39 to 221 Hz, resulting in a group mean of 123.3 Hz.

B. Frequency Band Method

This secondary method allowed for analysis across fourteen bands, ranging from 60 to 1500 Hz. Signal-to-noise ratio was calculated in each band. In the healthy group, the highest SNR values were observed in the lowest frequency band of 60–150 Hz. Here, the median reached 11.0, while the values gradually decreased with higher frequencies. In the 1200–1500 Hz band, the SNR value was only 1.17.

Compared to the healthy subjects, patients with PAF showed lower SNR values, with the highest median of 7.01 in the 60–150 Hz band.

Lastly, the pacemaker group showed the lowest SNR values among the three groups, with the highest value of 2.84 in the lowest frequency band. The distribution of SNR values across frequency bands is shown in Figure 4. Table I summarizes the differences between groups. Along with the band-based SNR,

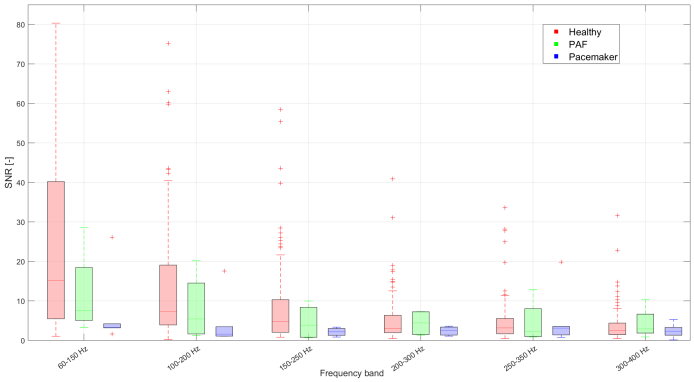

Fig. 4: Boxplot displaying the distribution of SNR values across frequency bands in lead V3. Separate boxplots were used for the healthy patients, patients with PAF and pace-maker patients. Outliers above 85 are not displayed for clarity purposes (n=5).

lead-dependent SNR was also observed. While the difference between leads was not too great, higher SNR values were observed in the septal (V1,V2) and anterior (V3,V4) leads, while lower values were found in the lateral (V5,V6) leads. For example, in the first band, the median SNR value ranged from 15.19 in V3 to 8.33 in V6. This trend can also be seen in the HF maps, where the healthy patient (Figure 5a) shows focused energy around the peak. On the contrary, the patient from the PAF and pacemaker group shows that the P wave energy is more scattered and broader (Figure 5b,5c).

TABLE I: Summary of HF analysis results across groups

Group	Median cutoff freq (Hz)	Median SNR
Healthy	141.6	11.0
PAF	135.5	7.01
Pacemaker	142.2	2.84

V. Discussion

The paper focused on the evaluation of high-frequency components of the P wave using two methods: a spectrogram based cutoff frequency approach and a frequency band SNR analysis. Each of these methods provides a slightly different view of the ECG recording, together offering a more complete characterization of the signal.

The spectrogram analysis indicated that the high-frequency components of the P wave extended up to 140 Hz in most

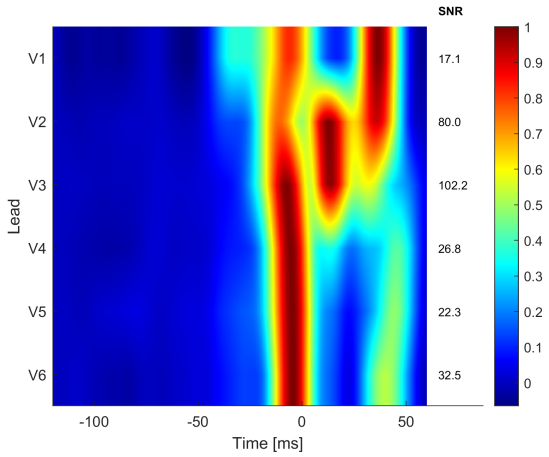

(a) Healthy

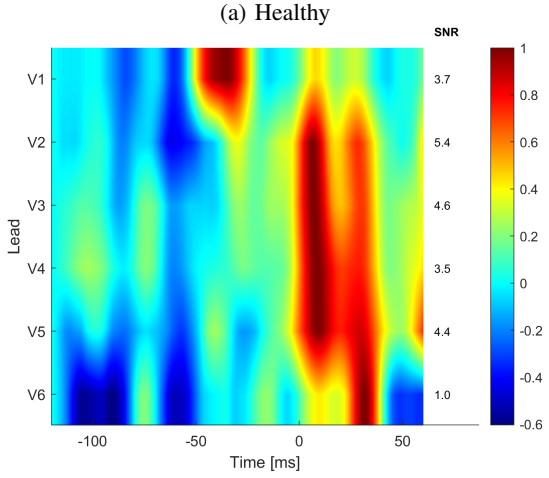

(b) Paroxysmal atrial fibrillation

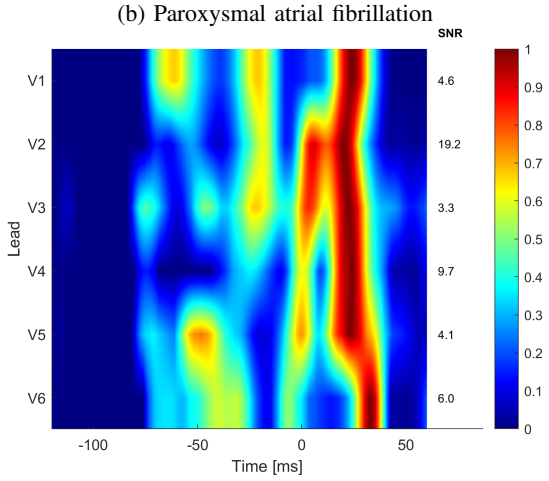

(c) Atrial pacemaker

Fig. 5: Illustrative high-frequency maps showing the distri-bution of energy (60–300 Hz) for one representative subject from each group. SNR values for each lead were computed as the RMS ratio of the signal window to the noise window across the averaged bands. Time 0 corresponds to the detected P wave peak.

recordings. This value was consistent across leads and patient groups, which suggests that high-frequency components are indeed present in the P wave, however their amplitude relative to noise is low.

The frequency band method provided a complementary approach. By decomposing the signal into specific bands, a more detailed view of these components was obtained. This method also showed a clear difference between the three groups. Healthy subjects showed the highest signal-to-noise ratio, while patients with PAF showed a reduced SNR. The lowest SNR values were observed in the group of patients with pacemakers. This finding suggests that even though high frequencies are present, their amplitude and quality may be reduced in a pathological setting. Lower SNR values in PAF patients may reflect structural remodeling linked to the condition, which might reduce the coherence of the P wave. In the pacemaker group, the artificial impulse results in higher variability, accompanied by broader, earlier energy distribution, visible in Figure 5.

The spectrogram method is a reliable tool to determine the presence of small but consistent frequency components, whereas the band method evaluates their actual strength within a given range.

VI. Conclusion

This study analysed high-frequency components of the P wave using spectrogram and frequency band analysis. The spectrogram method demonstrated that detectable P wave components extend up to 140 Hz. The frequency based SNR analysis revealed clear differences between patient groups, with the highest signal quality in the healthy group and the lowest in patients with atrial pacemakers. In addition, higher SNR values were also observed in septal and anterior leads compared to lateral leads. These results demonstrate that the spectrogram and frequency band approaches provide complementary information about the high-frequency components of the P wave. Future work will focus on expanding the analysis to a larger dataset and optimizing different processing parameters.

References

[1] D. Darbar, A. Jahangir, S. C. Hammill, and B. J. Gersh, “P wave signal-averaged electrocardiography to identify risk for atrial fibrillation,” *Pacing and Clinical Electrophysiology*, vol. 25, no. 10, pp. 1447–1453, 2002.

[2] S. Qiu et al.,“Revisiting the diagnostic and prognostic significance of high-frequency QRS analysis in cardiovascular diseases: A comprehensive review,” *Postgraduate Medical Journal*, vol. 100, no. 1189, pp. 785–795, Nov. 2024.

[3] L. G. Tereshchenko and M. E. Josephson, “Frequency content and characteristics of ventricular conduction,” *Journal of Electrocardiology*, vol. 48, no. 6, pp. 933–937, Nov.–Dec. 2015.

[4] S. Douedi and H. Douedi, “P wave,” in *StatPearls*. Treasure Island, FL, USA: StatPearls Publishing, 2025.

[5] P. G. Platonov, “P wave morphology: Underlying mechanisms and clinical implications,” *Annals of Noninvasive Electrocardiology*, vol. 17, no. 3, pp. 161–169, 2012.

[6] R. N. Cavalieri and P. Bertemes Filho, “Determination of maximum noise level in an ECG channel under SURE wavelet filtering for HRV extraction,” *Revista Mexicana de Ingeniería Biomédica*, vol. 41, no. 2, pp. 66–72, 2020.

Opportunistic Rainfall Detection Using Microwave Backhaul Links of Regional ISPs

Vojtěch Smejkal
Department of Telecommunications
Brno University of Technology
Brno, Czech Republic

***Abstract*—Commercial microwave links operated within regional Internet service provider networks generate continuous telemetry that can be reused for opportunistic environmental sensing. This paper presents the design and field deployment of a real-time microsystem that transforms operational backhaul measurements into rainfall estimates. The solution targets heterogeneous point-to-point links in the 5–80 GHz range and integrates automated acquisition, preprocessing, storage, rainfall retrieval, and visualization into a single workflow. Telemetry is collected using the Simple Network Management Protocol and HTTP-based application programming interfaces and stored in InfluxDB, while the TelcoRain framework is used to convert path attenuation into rainfall intensity estimates. The presented deployment was validated on a small regional ISP network near Polná in the Vysočina region of the Czech Republic, containing 11 usable links from three vendors. Results confirm that the system is technically deployable in small ISP environments and can detect rainfall events in real time, but they also show that limited network density and wet-antenna effects reduce the reliability of area-wide precipitation mapping and absolute intensity estimation.**

***Index Terms*—commercial microwave links, opportunistic sensing, rainfall monitoring, precipitation estimation, regional ISPs, backhaul telemetry, TelcoRain, InfluxDB**

I. Introduction

Accurate rainfall monitoring is essential for hydrology, urban drainage, and early warning systems. Conventional rain gauges provide accurate point measurements, while weather radar offers spatial coverage at the cost of indirect retrieval and terrain-related blind areas [1], [2]. For this reason, interest has increased in opportunistic sensing approaches that exploit already deployed communication infrastructure [3], [4].

Commercial microwave links (CMLs) are a promising source of such data because rain-induced attenuation is directly observable in operational telemetry. The relationship between specific attenuation and rainfall intensity is commonly modeled as

$$\gamma_R = kR^\alpha, \tag{1}$$

where γ_R is the specific attenuation in dB/km, R is rainfall intensity in mm/h, and k and α are frequency-dependent coefficients defined by ITU-R P.838-3 [5]. In practical terms, rain sensitivity is not uniform across the spectrum. Links below roughly 10 GHz usually react only to stronger rainfall or longer paths, mid-band links around 17–26 GHz provide a useful compromise between robustness and sensitivity, and E-band links near 80 GHz are highly sensitive even to weak precipitation but also more prone to wet-antenna attenuation and weather-related service degradation [8], [9]. Previous work has shown that dense CML networks can complement standard meteorological observations and support operational rainfall monitoring [6]–[8].

This paper focuses on a more practical question: can a regional ISP with a relatively small and heterogeneous backhaul network deploy a functional real-time rainfall monitoring microsystem using existing devices and open interfaces? The contribution is fourfold: 1) a compact architecture for multi-vendor telemetry collection, 2) a description of a real deployment in Polná, Vysočina, 3) qualitative validation against local weather stations, and 4) a discussion of reliability and accuracy limits in small-scale deployments.

II. Deployment in Polná, Vysočina

The evaluated deployment uses an operational regional ISP network near Polná, a small town in the Vysočina region of the Czech Republic. The locality is suitable for a feasibility study because it combines a compact urban center with surrounding rural access points and moderate terrain undulation. From the sensing perspective, this creates a mixed topology with short urban hops and longer inter-site backhaul paths. The network was not designed for meteorology; it was built for broadband service delivery, which makes it representative of the type of infrastructure that could realistically be reused by regional operators.

The dataset contains 11 usable point-to-point microwave links from Alcoma, Mimosa, and Ubiquiti. The monitored frequencies span 5, 10, 17, 24, 26, and 80 GHz. Link lengths range from 0.1 to 4.3 km and antenna diameters from 0.2 to 0.6 m. Seven links are provided by Alcoma, two by Mimosa, and two by Ubiquiti. This diversity is useful for testing because the network includes both low-sensitivity links in the lower microwave bands and highly rain-sensitive millimeter-wave links at 80 GHz.

Fig. 1 shows the deployed topology as visualized in the web application. The links are concentrated around Polná and its nearby access points, which is sufficient for event detection and cross-checking between bands, but still sparse for reconstructing a spatial rainfall field with high confidence.

TABLE I
SUMMARY OF THE DEPLOYED CML DATASET

Parameter	Value
Usable links	11
Vendors	Alcoma (7), Mimosa (2), Ubiquiti (2)
Frequencies	5, 10, 17, 24, 26, 80 GHz
Link length	0.1–4.3 km
Antenna diameter	0.2–0.6 m
Sampling period	10 s
Historical import window	11/2024–11/2025
Reference sensors	Warioweather ME13, EcoWitt GW2001

TABLE II
EXPECTED RAIN SENSITIVITY ACROSS DEPLOYED FREQUENCY BANDS

Band	Sensitivity	Practical role in this deployment
5–10 GHz	Low	Stable reference-like links; useful mainly for stronger rain or longer paths
17–26 GHz	Medium	Best compromise between path robustness and measurable rain attenuation
80 GHz	High	Strong response to weak rain on short links; more affected by wet antennas

III. ARCHITECTURE AND PROCESSING WORKFLOW

The architecture is organized as a lightweight processing chain suitable for small operators. Native Simple Network Management Protocol (SNMP) was used for Alcoma and Mimosa devices, while Ubiquiti links were accessed through the UISP HTTP application programming interface (API) because SNMP support was unavailable on the monitored units. Operational metrics, especially received signal level (RSL), transmitted signal level (TSL), and selected device parameters such as temperature, were acquired every 10 s and stored in an InfluxDB time-series database.

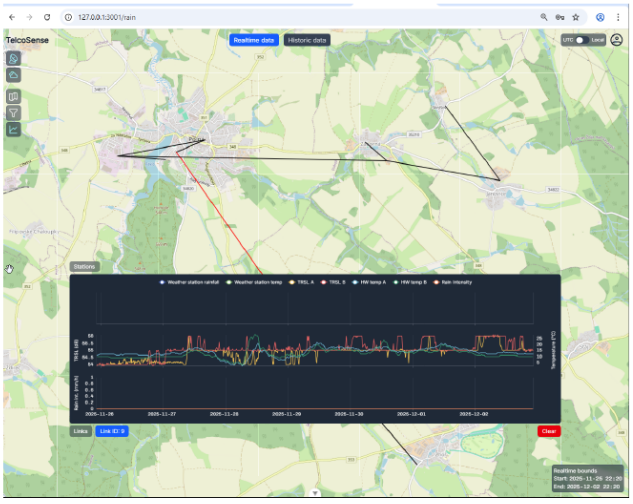

Fig. 1. Map-based visualization of the Polná deployment in the TelcoRain web interface. The view combines CML topology, local reference stations, and historical time series.

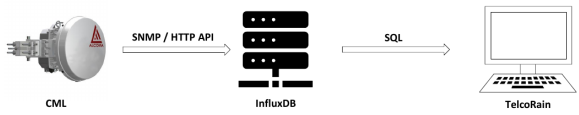

Fig. 2. Simplified processing chain from microwave link telemetry to rainfall estimation.

Telegraf was used for SNMP polling, whereas a custom Python script handled authenticated HTTP requests, JSON parsing, and InfluxDB writes for UISP-based devices. TelcoRain itself requires two data layers: static link metadata stored in MariaDB and time-series telemetry available in a format compatible with its internal queries. A periodic InfluxDB task was therefore introduced to normalize field names and move raw measurements from the acquisition bucket into the TelcoRain input bucket. The complete workflow is summarized in Fig. 2.

The implementation supports both real-time and retrospective analysis. Besides live polling, historical data exported from PandoraFMS were imported from CSV files into InfluxDB. This allowed the same pipeline to be applied to archived measurements collected over approximately one year. Such dual operation is important in practice because live monitoring alone is insufficient for debugging event detection, seasonal effects, and link-specific artifacts.

A key practical issue is that operational microwave links are not designed as meteorological sensors. Their measurements are affected by adaptive transmit power control, hardware differences, quantization, and wet-antenna attenuation (WAA). WAA is especially problematic because water remaining on the antenna radome can produce attenuation unrelated to path-averaged rainfall and therefore extend the detected event beyond the actual precipitation period [10], [11].

IV. VALIDATION, ACCURACY, AND RELIABILITY

The deployed microsystem successfully produced real-time rainfall estimates for all 11 usable links and exposed them through the TelcoRain web interface. Historical data import also worked correctly, enabling retrospective event analysis with the same processing stack used in the online mode. From an engineering perspective, the main result is that a heterogeneous small-ISP network can be integrated into a unified rainfall sensing workflow with modest additional software components.

A qualitative validation was performed using two local weather stations with 10 s sampling. Fig. 3 compares a TelcoRain-derived rainfall estimate from one CML with precipitation measured by a nearby Warioweather ME13 station during a rainfall event in early November 2025. The main onset of precipitation is captured by both sources in the same time window, which confirms good event-detection capability. The CML output, however, exhibits stronger short-term fluctuations and a slower decay after the main event. This behavior is

consistent with WAA and with the physical difference between point rain-gauge measurements and path-averaged microwave attenuation.



Fig. 3. Qualitative validation of a rainfall event. Top: rainfall intensity derived from a CML using TelcoRain. Bottom: precipitation from a nearby Warioweather station.

The present dataset supports event-level validation more reliably than absolute rain-rate accuracy assessment. In other words, the paper can defend that the system detected when rain occurred, but not yet that it estimated mm/h values with gauge-grade accuracy over the whole event. This limitation is caused by sparse link geometry, limited reference instrumentation, and the fact that a single gauge cannot fully represent the path-averaged quantity observed by a microwave link. A more rigorous accuracy study would require repeated events, denser ground truth, and explicit calibration of wet-antenna effects.

Reliability can nevertheless be discussed at the system level. During the illustrated rainfall episode, the acquisition, normalization, TelcoRain command-line interface processing, and visualization chain remained operational without loss of the displayed time series. This indicates that the proposed microsystem is sufficiently robust for unattended operation during ordinary adverse weather. At the same time, the study did not quantify availability under extreme conditions such as icing, prolonged power outages, or radio link failure. Therefore, the current contribution should be interpreted as a demonstrated operational prototype rather than a certified meteorological instrument.

The most important limitation was not algorithmic but geometric. Although the network covers a wide frequency range, it contains only 11 usable links and spans a limited area. This is sufficient for feasibility testing and link-scale event detection, but it is inadequate for reliable areal precipitation mapping. Longer links at 10, 17, and 26 GHz provided useful cumulative attenuation, while the 80 GHz links were highly sensitive even to weak events; nevertheless, the spatial density remained too low for robust rainfall-field reconstruction.

V. Conclusion

This paper presented a real-time microsystem for opportunistic rainfall detection based on commercial microwave links operated by a regional ISP near Polná, Czech Republic. The deployment combined SNMP and HTTP telemetry collection, InfluxDB-based storage and normalization, TelcoRain processing, historical data import, and web-based visualization in a unified workflow.

From the experimental deployment, several practical conclusions can be drawn. First, the system processed telemetry from 11 heterogeneous links with a 10 s sampling period and remained operational during the documented rainfall event, which supports the feasibility of unattended real-time deployment. Second, comparison with local weather stations showed good agreement in event timing, meaning that the system is reliable for rainfall occurrence detection. Third, the agreement in absolute rainfall intensity was only qualitative because wet-antenna attenuation, path averaging, and limited reference measurements still affect the retrieved magnitudes. Finally, the mixed 5–80 GHz dataset confirmed the expected spectral behavior: lower bands were more robust but less sensitive, while 80 GHz links were very responsive yet more vulnerable to residual attenuation.

Future work should therefore focus on three directions: broader topology coverage through federation of multiple regional ISPs, improved compensation of wet-antenna effects, and a quantitative accuracy study with denser meteorological ground truth and formally reported error metrics.

References

[1] World Meteorological Organization, *Guide to Meteorological Instruments and Methods of Observation*, WMO-No. 8, 8th ed. Geneva, Switzerland: WMO, 2018.

[2] E. Saltikoff, K. Friedrich, J. Soderholm, *et al.*, "An overview of using weather radar for climatological studies: Successes, challenges, and potential," *Bull. Amer. Meteorol. Soc.*, vol. 100, no. 9, pp. 1739–1752, 2019.

[3] H. Messer, A. Zinevich, and P. Alpert, "Environmental monitoring by wireless communication networks," *Science*, vol. 312, no. 5774, p. 713, 2006.

[4] R. Nebuloni, F. Giannetti, F. Sapienza, *et al.*, "A review of technical aspects and challenges in opportunistic rainfall estimation using satellite and terrestrial microwave links," *IEEE Geosci. Remote Sens. Mag.*, 2025, early access.

[5] International Telecommunication Union, "Recommendation ITU-R P.838-3: Specific attenuation model for rain for use in prediction methods," 2005.

[6] A. Overeem, H. Leijnse, and R. Uijlenhoet, "Rainfall monitoring using microwave links from cellular communication networks: The Dutch experience," in *Proc. IEEE Stat. Signal Process. Workshop*, 2018, pp. 110–114.

[7] G. Roversi, F. C. P. de Macedo, F. Porcù, *et al.*, "Commercial microwave links as a tool for operational rainfall monitoring in northern Italy," *Atmos. Meas. Tech.*, vol. 13, pp. 5779–5797, 2020.

[8] H. Leijnse, R. Uijlenhoet, and J. N. M. Stricker, "Microwave link rainfall estimation: Effects of link length and frequency, temporal sampling, power resolution, and wet antenna attenuation," *Adv. Water Resour.*, vol. 31, no. 11, pp. 1481–1493, 2008.

[9] International Telecommunication Union, "Recommendation ITU-R P.530: Propagation data and prediction methods required for the design of terrestrial line-of-sight systems," 2021.

[10] M. Schleiss, J. Rieckermann, and A. Berne, "Quantification and modeling of wet-antenna attenuation for commercial microwave links," *IEEE Geosci. Remote Sens. Lett.*, vol. 10, no. 5, pp. 1195–1199, 2013.

[11] M. Bubniak and P. Musil, "Quantification of a wet antenna attenuation for a measurement of rainfall intensity via the metropolitan network of microwave links in 10 GHz band," in *Microwave and Radio Electronics Week*, 2022.

Reflector Dipole Antenna for Secondary Radar Band

Vojtěch Krehan
Department of Radio Electronics
FEEC, Brno University of Technology
Brno, Czech Republic
256513@vutbr.cz

David Kuřátko
Department of Radio Electronics
FEEC, Brno University of Technology
Brno, Czech Republic
xkurat01@vut.cz

Ondřej Jaroš
Eldis Pardubice, s.r.o.
Pardubice, Czech Republic
ondrej.jaros@eldis.cz

***Abstract*—This paper presents the design, theoretical analysis, and step-by-step simulation of a reflector dipole antenna for the Secondary Surveillance Radar (SSR) system operating in the 1030/1090 MHz band. The proposed antenna employs a broadband dipole driven by an aperture-coupled microstrip balun. To ensure mechanical stability and environmental protection, an ABS dielectric radome with supporting structures is integrated into the design. Furthermore, a rod reflector is utilized to suppress back-lobe radiation and increase the forward gain. Simulated in CST Studio Suite using an IS400 substrate, the step-by-step optimization process highlights the impact of the radome and reflector on the antenna's resonance frequency. The final compensated model achieves a reflection coefficient better than -17 dB across the operating band and a realized gain exceeding 7 dBi, with a beamwidth of approximately 84° in the azimuthal plane.**

***Index Terms*—Antenna, dipole, reflector, SSR, aperture coupling, CST Studio Suite, radome.**

I. Introduction

Secondary Surveillance Radars (SSR) are a crucial technology in air traffic control and airspace safety. Unlike primary radars, SSR systems cooperate directly with transponders on board the aircraft to enable precise identification and data transmission. According to international standards, these systems operate in the L-band, strictly utilizing 1030 MHz for interrogation and 1090 MHz for reply [1]. For reliable communication, the antenna must exhibit excellent impedance matching, typically requiring a reflection coefficient $|S_{11}| < -10$ dB at both operating frequencies.

To replace traditional bulky structures, recent SSR antenna designs explore planar technologies, such as series-fed printed dipole arrays [2] and miniaturized cavity-backed slot antennas [3]. Standard industry solutions also provide robust baselines [4]. However, integrating protective radomes often significantly degrades impedance matching, requiring careful co-design.

This paper focuses on the design and optimization of a slot-fed dipole antenna tailored for the SSR band [5]. The design utilizes the aperture-coupled feeding technique, which prevents direct galvanic connection and allows independent optimization of the microstrip feed and the radiating elements.

The rest of this paper is organized as follows. Section II describes the theoretical design and methodology, including the substrate, feeding mechanism, and calculations for the radiating arms. Section III details the step-by-step optimization process, covering the dipole arms modification, radome integration, and reflector design. Section IV presents the final optimization along with the simulated impedance and radiation characteristics. Finally, Section V concludes the paper.

II. Theoretical Design and Methodology

The initial antenna model was designed and simulated using CST Studio Suite® HF Time Domain Solver [6][1] . The theoretical design is divided into the substrate and feeding mechanism, followed by the calculation of the dipole's radiating arms.

A. Substrate and Feeding Mechanism

The choice of the dielectric substrate is a critical first step for the Printed Circuit Board (PCB). An IS400 substrate (an equivalent to FR-4) with a thickness of 1 mm was selected. Although materials like Rogers offer significantly lower dielectric losses, IS400 was chosen due to its low cost and the availability of large panel sizes (up to 2 meters), which is essential for the future development of multi-element antenna arrays. The substrate parameters used in the simulation are relative permittivity $\varepsilon_r = 5.1149$ and loss tangent $\tan\delta = 0.019$. To achieve a characteristic impedance of approximately 50 Ω, the microstrip line width was optimized to 8 mm.

The aperture-coupled feed acts primarily as a series impedance with an inductive character. To compensate for this parasitic inductance and ensure resonance, an open-ended microstrip stub (balun) is utilized [7]. The theoretical length of this balun should be a quarter of the guided wavelength ($\lambda_g/4$) [9]. The guided wavelength and the theoretical balun length l_{balun} are calculated as:

$$\lambda_g = \frac{2\pi}{\beta} = 264.4 \text{ mm} \tag{1}$$

$$l_{balun} = \frac{\lambda_g}{4} = 66.15 \text{ mm} \tag{2}$$

where β is the phase constant obtained from the simulation.

To maintain symmetry and confine the electromagnetic fields, vias were strategically placed along the microstrip

[1]https://www.3ds.com/products/simulia/cst-studio-suite

feedline. The theoretical spacing s and diameter d of the vias must satisfy the following conditions [9]:

$$s \leq \frac{\lambda_g}{20} \implies s \leq 13.22 \text{ mm} \tag{3}$$

$$d \geq \frac{s}{10} \implies d \geq 1.32 \text{ mm} \tag{4}$$

Based on these theoretical constraints, the via spacing was set to $s = 11$ mm and the diameter to $d = 2$ mm.

B. Dipole Dimensions and Radiating Arms

The radiating structure consists of two metal arms separated by a coupling slot in the center. The theoretical length of each arm is designed to be $\lambda_0/4$ [9], making the total theoretical length of the dipole $\lambda_0/2$. The coupling slot length is theoretically set to $\lambda_0/4$.

At the chosen center frequency of 1060 MHz, the free-space wavelength λ_0 is:

$$\lambda_0 = \frac{c}{f} = \frac{299792458}{1060 \cdot 10^6} \approx 282 \text{ mm} \tag{5}$$

Applying this to the structure yields the theoretical dimensions for the radiating elements:

$$l_{arms} = \frac{\lambda_0}{2} = 141 \text{ mm} \tag{6}$$

$$l_{slot} = \frac{\lambda_0}{4} = 70.5 \text{ mm} \tag{7}$$

where l_{arms} is total length of the dipole arms and l_{slot} is total length of the slot between dipole arms.

For this initial theoretical model, the width of the dipole arms was set to 8 mm, and the width of the coupling slot was set to 6 mm. The theoretical model of the aperture-coupled dipole, including the folded microstrip balun and the via configuration, is illustrated in Fig. 1.

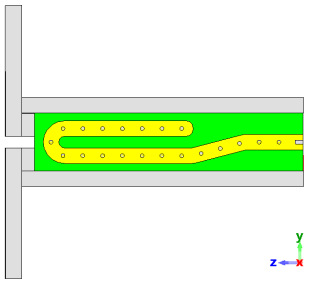

Fig. 1. Designed theoretical model of the aperture-coupled dipole with the slot and folded microstrip balun.

III. Step-by-Step Optimization

The theoretical dimensions provided a solid baseline, but to meet the strict SSR band requirements and prepare the antenna for physical integration into an array, the structure required iterative optimization.

A. Dipole Arms Optimization

The initial theoretical dipole exhibited suboptimal impedance matching above 1040 MHz. Furthermore, the original arm length (141 mm) is impractical for dense placement in a linear antenna array. Therefore, the dipole arms were significantly reduced and optimized. The final optimized dimensions are 59 mm in length and 9 mm in width for the arms, and 66 mm in length and 6 mm in width for the coupling slot. An N-type connector was also integrated into the model to simulate real-world feeding conditions accurately.

Fig. 2 illustrates the simulated surface current distribution on the optimized feeding structure. The current properly flows along the edges of the microstrip line and is symmetrically distributed across the slot into both dipole arms. The correct placement of the via at the bend of the microstrip is crucial for maintaining this symmetry.

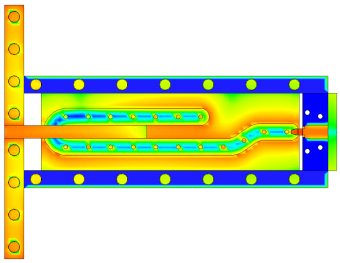

Fig. 2. Simulated surface current distribution on the optimized microstrip feed and coupling slot.

This dimensional optimization shifted the resonance frequency to approximately 1.09 GHz, where the reflection coefficient $|S_{11}|$ dropped below -28 dB (see Fig. 3). This upward frequency shift was intentionally introduced to pre-compensate for the expected capacitive loading from the radome in the subsequent design stage. The difference between the theoretical and optimized reflection coefficient is primarily caused by the significant reduction in the length of the dipole arms and the coupling slot.

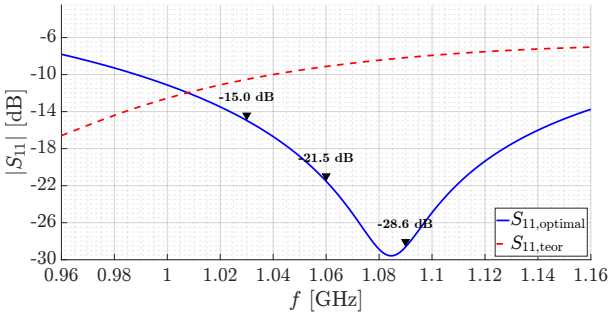

Fig. 3. Comparison of the reflection coefficient $|S_{11}|$ between the theoretical and optimized dipole.

B. Radome Integration

To provide mechanical stability and protect the antenna from environmental factors, a radome was designed using ABS material. ABS was selected for its good mechanical strength and suitability for 3D printing. The material parameters were set to $\varepsilon_r = 2.8$ and $\tan\delta = 0.004$ at 1 GHz [8][2]. The wall thickness was kept thin relative to the wavelength to minimize internal reflections [10], [11]. The radome (see Fig. 4) features a 3 mm wall thickness for the inserted U-shaped section and 2 mm for the fixed cover.

Electromagnetically, the radome acts as a dielectric load, effectively increasing the electrical length of the dipole [9]. Simulation results indicated that the radome shifted the resonance frequency downwards by approximately 50 MHz to near 1040 MHz. Additionally, four fixation posts made of the same ABS material were implemented to secure the antenna. While these posts had a negligible effect on impedance matching, they slightly altered the radiation pattern by accentuating the side lobes.

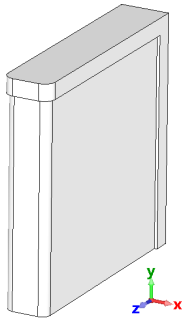

Fig. 4. Designed model of the radome.

C. Reflector Design

To achieve high directivity and suppress the back-lobe radiation, a reflector was added. To minimize the antenna's weight and aerodynamic drag, a rod reflector was utilized instead of a solid metal plate. The theoretical optimal distance r between the dipole and the reflector is $r = \lambda_0/4 = 70.5$ mm [9].

The reflector consists of vertically aligned metal rods, each 12 mm wide, 2 mm thick, and 168 mm long. The gap g between the rods was set to 28 mm, which satisfies the theoretical condition $g \leq \lambda_0/10$ (28.2 mm), ensuring the structure behaves electromagnetically like a solid plane [10].

The introduction of this large conductive structure in the near-field region drastically altered the impedance profile, causing a massive downward frequency shift of approximately 120 MHz. Consequently, the reflection coefficient above 1060 MHz no longer met the required -10 dB threshold.

[2]https://www.matweb.com

IV. Final Optimization and Results

To compensate for the severe resonance shift caused by the reflector without altering the carefully tuned dipole arms, the microstrip balun was optimized.

A. Impedance Matching

The balun length was significantly shortened from the theoretical 66.15 mm to 42.15 mm. Shortening the feed stub shifts the input impedance into the capacitive region, which successfully compensated for the accumulated parasitic inductance of the slot and the loading effects of the radome and reflector [7]. This tuning re-centered the resonance precisely at the 1060 MHz center frequency.

As a result, the final reflection coefficient $|S_{11}|$ is exceptionally well-matched across the operating band (see Fig. 5). The simulated values are presented in Table I.

TABLE I
Final Impedance Matching Parameters

Frequency [MHz]	$\lvert S_{11}\rvert$ [dB]
1030	-17.3
1060	-18.5
1090	-18.0

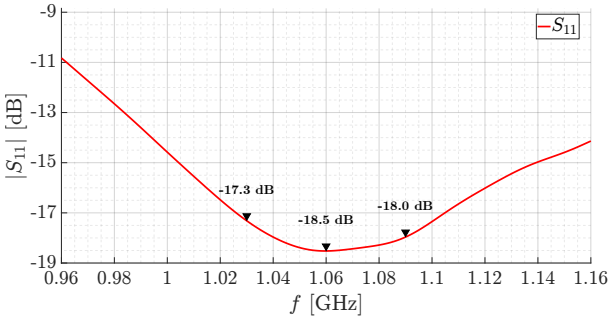

Fig. 5. Reflection coefficient $|S_{11}|$ of the final optimized antenna system.

B. Radiation Characteristics

The final mechanical structure (see Fig. 7), integrating all components and structural supports, successfully achieved the targeted directional pattern (see Fig. 6). The values of realized gain and beamwidth in the azimuthal plane are presented in Table II. Furthermore, the rod reflector effectively suppresses back radiation, yielding a front-to-back ratio of approximately 14 dB at 1090 MHz.

TABLE II
Final Radiation Characteristics Parameters

Frequency [MHz]	Gain [dBi]	Beamwidth [deg]
1030	8.01	84.3
1060	7.86	83.8
1090	7.78	83.0

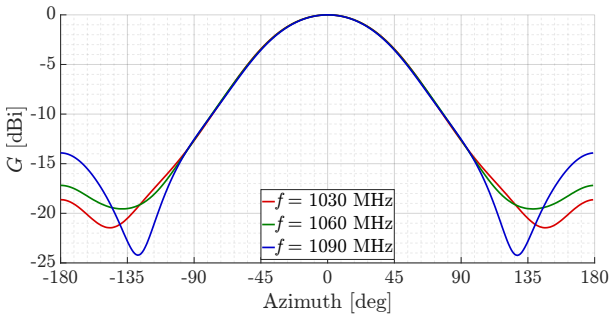

Fig. 6. Normalized realized gain of the final antenna in the azimuthal plane.

V. Conclusion

This paper detailed the design and optimization of a reflector dipole antenna for the SSR system operating at 1030 and 1090 MHz. The iterative simulation process in CST Studio Suite demonstrated how integrating a protective ABS radome and a weight-saving rod reflector significantly alters the antenna's resonance. By tuning the aperture-coupled microstrip balun, these effects were mitigated. The final design successfully meets the operational criteria, exhibiting an $|S_{11}|$ better than -17 dB and a realized gain exceeding 7.7 dBi. As summarized in Table III, compared to recent planar SSR solutions [2], [3] and industry baselines [4], the proposed reflector dipole achieves highly competitive forward gain and excellent broadband impedance matching, while successfully addressing the practical challenges of radome integration. Future work will focus on fabricating the physical prototype, experimentally verifying the radome's material properties, and measuring the radiation pattern and reflection coefficient in an anechoic chamber.

TABLE III
Comparison with SOTA SSR Antennas

Ref.	Antenna Type	Gain [dBi]	$\lvert S_{11}\rvert$ [dB]	Radome
[2]	Series-fed dipole array	20.0	< -12.0	No
[3]	Cavity Backed Slot	5.66	< -15.0	Yes (PTFE/E-glass)
[4]	Patch array (Baseline)	N/A	< -20.0	No
Proposed	Reflector Dipole	8.01	< -17.3	Yes (ABS)

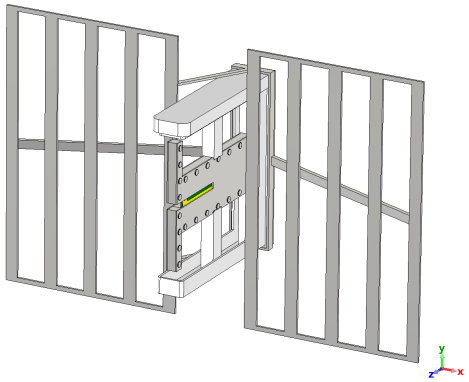

Fig. 7. Final mechanical model of the antenna including the ABS radome, structural supports, and the rod reflector.

References

[1] International Civil Aviation Organization, *Aeronautical Telecommunications: International Standards and Recommended Practices. Volume IV, Surveillance Radar and Collision Avoidance Systems*, 5th ed. Montreal, QC, Canada: ICAO, 2014.

[2] M. Abdolahi, G. Askari, and H. M. Sadeghi, "A new Chebyshev tapered series-fed printed dipole antenna comprising rhombus radiators at L band for SSR application," *2017 Iranian Conference on Electrical Engineering (ICEE)*, Tehran, Iran, 2017. doi: 10.1109/IranianCEE.2017.7985366.

[3] P. L. Vijayvergiya, R. Mishra, U. S., H. Naik, C. Suresh, J. Kumar, and A. K. Singh, "Miniaturized cavity backed H-slot antenna for SSR application in Airborne systems with different low loss radome materials," *2025 IEEE Space, Aerospace and Defence Conference (SPACE)*, Bangalore, India, 2025. doi: 10.1109/SPACE65882.2025.11171226.

[4] V. Schejbal, P. Bezousek, J. Pidanic, and M. Chyba, "Secondary Surveillance Radar Antenna," *IEEE Antennas and Propagation Magazine*, vol. 55, no. 2, pp. 164-170, April 2013. doi: 10.1109/MAP.2013.6578015.

[5] M. K. Leino, J. Ala-Laurinaho, A. Höök, B. Svensson, and V. Viikari, "Low-scattering chopped dipole for secondary surveillance radar," *IEEE Trans. Antennas Propag.*, vol. 70, no. 3, pp. 1694–1705, Mar. 2022.

[6] CST Studio Suite 2026, Dassault Systèmes.

[7] D. M. Pozar, "A review of aperture coupled microstrip antennas: History, operation, development, and applications," Univ. Massachusetts at Amherst, MA, USA.

[8] MatWeb, "Overview of materials for Acrylonitrile Butadiene Styrene (ABS), Extruded," [Online].

[9] C. A. Balanis, *Antenna Theory: Analysis and Design*, 4th ed. Hoboken, NJ, USA: John Wiley & Sons, 2016.

[10] S. Silver, *Microwave Antenna Theory and Design*. London, U.K.: Institution of Engineering and Technology, 2008.

[11] H. U. Tahseen, L. Yang, and X. Zhou, "Design of FSS-antenna-radome system for airborne and ground applications," *IET Commun.*, vol. 15, no. 13, pp. 1691–1699, 2021.

Online Radicalisation and Extremism: An Integrated OSINT and Legal Analysis

1st Katarina Kudelova
Department of Telecommunications
FEEC, Brno University of Technology
Brno, Czechia
257128@vut.cz

2nd Willi Lazarov
Department of Telecommunications
FEEC, Brno University of Technology
Brno, Czechia
lazarov@vut.cz

3rd Jakub Vostoupal
Dept. of International and European Law
Faculty of Law, Masaryk University
Brno, Czechia
jakub.vostoupal@law.muni.cz

***Abstract*—Online radicalisation has become one of the most significant security challenges in the digital age. Social media platforms, particularly TikTok, amplify this process through personalised recommendation algorithms that create narrow content streams and expose users to increasingly extreme narratives, a phenomenon known as the "rabbit-hole effect". At the same time, legal frameworks such as the EU Digital Services Act (DSA) attempt to regulate platform behaviour, yet face substantial enforcement challenges. This paper presents an integrated analysis combining theoretical background on radicalisation and disinformation with a legal assessment of EU and Czech regulatory tools, and a practical case study on TikTok. Using browser automation and metadata extraction, we found that TikTok's algorithm narrowed content toward politically sensitive and emotionally charged material within approximately 60 minutes of targeted interaction. Results confirm rabbit-hole dynamics in the Czech information space and cross-topic contamination across disinformation, pro-Russian narratives, and radicalisation content – revealing a persistent gap between DSA obligations and actual platform behaviour, and underscoring the need for stronger enforcement and algorithmic transparency.**

***Index Terms*—OSINT, cyberspace, radicalisation, extremism, disinformation, legal analysis**

I. Introduction

Radicalisation in the online environment represents one of the most discussed security challenges of the last decade. What has changed is not the existence of radicalisation itself, but the conditions under which it spreads. Today, communication takes place in fast, global, and highly personalised online environments. Research shows that digital spaces enable the unprecedentedly rapid spread of emotions, symbols, and shared meanings, creating qualitatively different conditions for the formation of radical attitudes [1], [5]. Algorithms and social feedback mechanisms can speed up radicalisation much faster than traditional face-to-face interactions, creating what researchers describe as "unprecedented online journeys", where users can move from simple curiosity to extreme views in a very short time [1].

Among contemporary platforms, TikTok stands out due to its highly responsive recommendation system, which adapts to user behaviour within minutes and can create narrow content streams exposing users to emotionally charged or one-sided narratives. Although these systems were designed to maximise engagement, they can unintentionally amplify polarising or radicalising content – making TikTok a suitable environment for studying how algorithmic dynamics interact with radicalisation in practice.

These challenges are particularly relevant for the European Union (EU), which has introduced the Digital Services Act (DSA) to ensure greater platform accountability in addressing systemic risks including disinformation and algorithmic manipulation. At the same time, the Czech Republic, a country whose information space has repeatedly been the target of foreign disinformation campaigns, illustrates how sensitive the debate on the limits of state intervention remains.

This paper addresses these challenges through an integrated approach that combines theory, law, and practice. It first establishes the theoretical foundations of online radicalisation and the role of algorithmic systems, then examines how the EU and Czech legal frameworks attempt to respond to these risks, and finally tests these dynamics in practice through a controlled OSINT-based case study on TikTok.

II. Theoretical Background

Radicalisation is generally understood as a gradual process, not a single event. It describes how individuals or groups move towards more extreme beliefs, emotions, or behaviours that can, in some cases, lead to violence [2]. Researchers approach this process from three main perspectives: psychology focuses on internal factors such as loss of significance, moral disengagement, or a search for identity; sociology looks at external influences like peer pressure, group polarisation, and social identity; and security studies combine both to understand how radicalisation can escalate towards violent extremism [3].

Research shows that radicalisation is a multi-level process shaped by psychology, social dynamics, and, increasingly, technology [2]. Digital platforms have fundamentally accelerated this process by lowering the threshold for participation and enabling individuals to anonymously engage without fear of social or legal consequences [5].

Two key phenomena explain how this works in practice. Echo chambers are online spaces where users mostly encounter opinions that match their own, creating a false sense of consensus [6]. Filter bubbles are created by algorithms that

The research described in this paper was financially supported by the Technology Agency of the Czech Republic, project No. OZ01020006.

show users content based on past behaviour, limiting exposure to alternative viewpoints [7]. Together, these effects can trap users in a closed worldview and reinforce an "us versus them" mindset that polarises societies [8].

A closely related mechanism is the so-called rabbit-hole effect – a process where users are gradually guided toward more extreme content through algorithmic recommendations. Platforms like TikTok use complex algorithms that analyse watch time, likes, and interactions to keep users engaged. Zhou (2024) [9] shows that TikTok's algorithm heavily weights metrics like like-ratios and trending hashtags, which can narrow the visible content diversity and push users toward more homogenous or extreme material [9]. Autoplay further reduces user control, making it easier to drift into increasingly radical content streams [10]. However, researchers stress that the rabbit-hole effect is not automatic or uniform – personal factors such as critical thinking and social context play an important role in how users respond [11].

To analyse such a comprehensive environment, Open Source Intelligence (OSINT) can be employed as an intelligence discipline for systematic collection, processing, and interpretation of publicly available data. It relies on legally obtained public information that is validated and analysed to fulfil a specific intelligence requirement [12]. In this context, OSINT offers a valuable approach for monitoring public social media content, including threads and user profiles. OSINT techniques enable researchers to track the spread of targeted content, map networks of radical actors, and examine how algorithmic mechanisms foster ideological echo chambers [13].

III. Legal Analysis

Addressing online radicalisation and the spread of extremist content requires not only technical tools but also effective legal frameworks. This section traces the evolution of platform regulation in the EU, examines the core obligations introduced by the DSA, and assesses where legal tools face their practical limits, with particular attention to the Czech context.

A. The Evolution of Platform Regulation

The limitations of voluntary cooperation became especially visible during major crises. Following the 2016 Brexit referendum and the Cambridge Analytica scandal, the EU established the High-Level Expert Group on Fake News and Disinformation in January 2018, which led to the first Code of Practice on Disinformation signed in October 2018 [14]. However, this Code remained voluntary and platforms could choose how to implement it. Independent assessments soon revealed that compliance was inconsistent and transparency was limited [15]. Russia's full-scale invasion of Ukraine in 2022 further exposed how easily coordinated influence campaigns could destabilise European societies [16]. Because national responses were fragmented and self-regulation had proven insufficient, the EU moved toward binding legislation. The DSA, adopted in October 2022, is the EU's first comprehensive and enforceable framework for online platforms [17]. One of its key innovations is the category of Very Large Online Platforms (VLOPs) – services with more than 45 million EU users that are subject to the strictest obligations. TikTok was formally designated as a VLOP in April 2023 [18].

B. The DSA and Platform Accountability

The DSA introduces a layered system of duties. For VLOPs like TikTok, the most important obligations include systemic risk assessments, transparency reporting, and data access for researchers [19]. Platforms must identify risks connected to the design and functioning of their services – including the spread of illegal content, threats to fundamental rights, impacts on electoral processes, and risks for minors – and take concrete measures to mitigate them, such as adjusting recommendation systems or improving content moderation [19].

A key feature of the DSA is the distinction between illegal content and harmful but legal content. The DSA does not define what is illegal online – this is left to other EU or national laws. Instead, it focuses on process: platforms must offer ways to report illegal content, cooperate with trusted flaggers, and comply with removal orders from national authorities. At the same time, the DSA prohibits general monitoring obligations to protect freedom of expression [19]. For harmful but legal content, such as borderline radicalising material that does not yet violate any law, the DSA requires risk assessments and mitigation measures, but stops short of mandating removal. This distinction is particularly relevant for online radicalisation, where much of the problematic content operates in a grey zone between protected speech and genuine harm.

C. Limits of Legal Protection

Despite its ambitions, the DSA faces significant enforcement challenges. The regulatory architecture is partially decentralised. While the European Commission oversees VLOPs directly, national Digital Services Coordinators enforce the rules for smaller platforms. Many of these bodies are underresourced, which can lead to case overload and an inability to investigate new violations [20]. In February 2024, the European Commission opened formal proceedings against TikTok under the DSA, examining potential breaches related to advertising transparency, algorithmic risk, and design features affecting minors [21]. The fact that these investigations remain ongoing without final decisions illustrates the time lag between rule-making and enforcement, during which harmful practices may continue and persist.

IV. Case Study on TikTok

TikTok is a short-form video platform owned by ByteDance whose recommendation system curates personalised content based on behavioural signals such as watch time, replays, and interactions. Its high virality and emotionally charged content environment make it particularly suitable for studying how algorithmic dynamics interact with radicalisation processes [22]. Even minimal interactions can shift recommendations toward narrower and more extreme content streams, directly illustrating the rabbit-hole effect. The following case study, whose phases are illustrated in Figure 1 and described in the sections below, demonstrates this dynamic in practice.

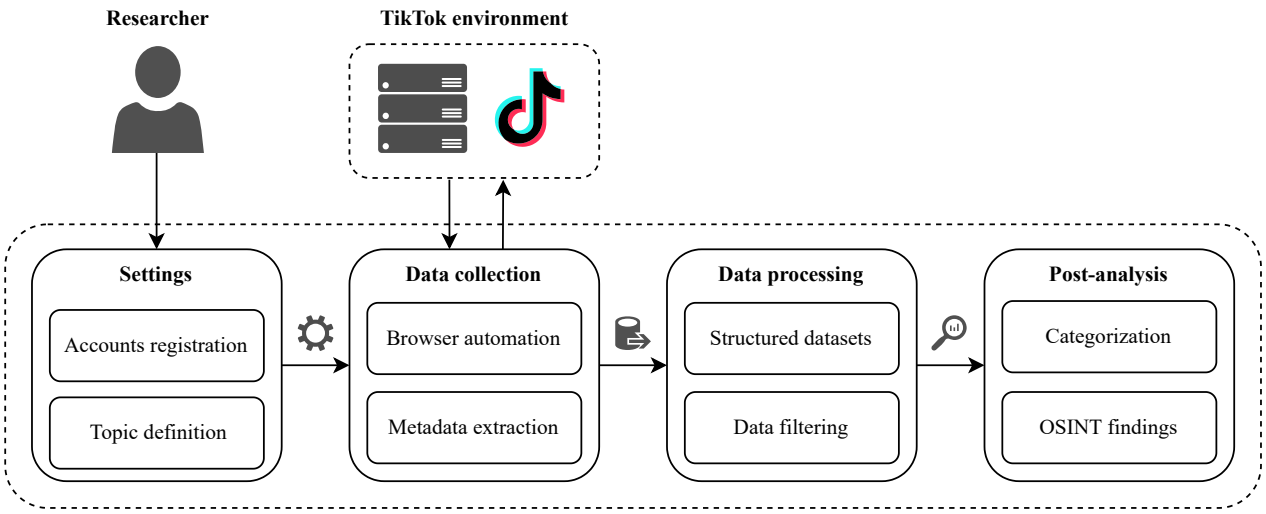

Fig. 1. Phases of the data collection and analysis in the TikTok environment

A. Scope and Parameters

The case study focused on three thematic areas relevant to hybrid threats and radicalisation:

- political disinformation and anti-institutional narratives;
- pro-Russian propaganda and war-related content;
- and radicalisation and extremist narratives.

Three separate TikTok accounts were created with the Czech language and IP address settings to simulate new users entering the Czech information space. Each account was trained on one specific theme through consistent interactions – watching relevant videos to completion, liking selected content, following thematic accounts, and using related hashtags and trending audio. The study was conducted in 2025 using publicly available content only, without accessing private data or bypassing platform security measures.

B. Data Collection

Data collection combined two complementary approaches. Manual observation provided qualitative insight into content framing, tone, and comment sections – context that automated tools cannot easily capture. Automated collection was implemented using Python with Playwright and Selenium, which simulate realistic user behaviour in a clean browser environment with no prior TikTok activity.

During automated browsing, visible attributes of each video were captured (including captions, hashtags, authors, and engagement metrics) and stored in a unified JSON format. Two separate datasets were created: a baseline dataset representing the feed before any interactions, and a post-interaction dataset recorded after the automated behavioural signals had been applied. These datasets were then compared side by side, focusing on shifts in topic distribution, emotional tone, narrative framing, and content homogeneity. The results indicate whether the algorithmic trajectory narrows or reinforces certain patterns over time, which is directly relevant for understanding the rabbit-hole effect.

C. Results and Discussion

The results confirm the presence of rabbit-hole dynamics in the Czech information space. Across all three profiles, TikTok's algorithm narrowed the content feed toward the trained theme within approximately 60 minutes of targeted scrolling, after which the feed of a completely new account became visibly focused on a specific topic through consistent interaction alone.

A common pattern across all three profiles was the high proportion of emotionally charged content. Videos were frequently built around simple polarising narratives, such as questioning the legitimacy of government institutions, framing the war in Ukraine as a Western provocation, or portraying ethnic and religious minorities as threats to social order. This was reinforced by recurring stylistic features: dramatic music, provocative captions, AI-generated text-to-speech narration, and rapid editing. Comment sections were often more aggressive than the videos themselves, creating a false impression of widespread support or general outrage.

The comment section shown in Table I is representative of the pro-Russian narrative profile. This particular video with 839 comments where the most liked comments deny Russian aggression, frame the Czech government as criminals and traitors, and express open support for the Russian Federation – each accumulating hundreds of likes. This pattern illustrates how comment sections amplify the perceived legitimacy of fringe views, making them appear as mainstream consensus to new users encountering the content for the first time.

In contrast, the only dissenting comment "And do You believe the Russians?" received just 18 likes but attracted 44 replies, the highest reply count of any comment shown. This suggests that counter-narrative contributions are not ignored but actively suppressed through collective pushback. The asymmetry between engagement and approval reinforces an echo-chamber dynamic in which dissent becomes visible mainly as a target rather than a legitimate counterpoint.

TABLE I
TOP 5 LIKED COMMENTS FROM A PRO-RUSSIAN TIKTOK VIDEO CIRCULATING IN THE CZECH INFORMATION SPACE

Comment (English translation)	Likes	Replies
"AGGRESSOR is Ukraine and not Russia"	818	15
"The Czech government belongs behind bars as terrorists and traitors ."	868	9
"there is also a documentary by French television which was even broadcast by Czech Television where they show the atrocities of Ukrainians in Donbas... it disappeared even from the archive"	262	12
"Long live the Russian Federation"	80	1
"Whoever has sense knows it, whoever doesn't, supports Zelensky."	773	8

The study also revealed cross-topic contamination. Although each profile was trained on a single theme, content from the other thematic areas began appearing over time – a dynamic particularly favourable for the spread of hybrid threats across multiple narratives simultaneously. The Czech and Slovak information spaces proved closely interconnected, with the same narratives recycled across both environments through translated captions or remixed audio.

Content analysis across all three profiles revealed three recurring production patterns: AI-generated or heavily edited videos, professionally produced clips, and spontaneous single-shot recordings. Despite these stylistic differences, all three shared a common mechanism – narrative recycling. The same claims and visual formats reappeared as re-uploads, remixes, and edits with altered captions or audio tracks. This repetition reinforces the impression of widespread presence and reaches different user segments with each iteration, making the spread of hybrid threat narratives more effective than any single viral video could achieve.

V. CONCLUSION

This paper examined online radicalisation from theoretical, legal, and practical perspectives, confirming a persistent gap between DSA obligations and actual platforms' behaviour. The practical case study showed that TikTok's algorithm can narrow the content feed toward radicalising material within approximately 60 minutes of targeted interaction, regardless of whether the content is political, manipulative, or radicalising in nature. These findings are not an isolated observation – they reflect a broader and growing challenge. Online extremism and radicalisation are increasingly difficult to contain precisely because the mechanisms that drive them are embedded in the normal functioning of platform algorithms, not in exceptional or illegal behaviour. Legal frameworks such as the DSA represent a significant step forward, but were not designed to respond at the speed at which algorithmic systems operate. By the time a risk assessment is completed, an enforcement proceeding is initiated, or a mitigation measure is implemented, thousands of users may already have been exposed to narrowing content streams. This gap between legal response and algorithmic reality remains one of the most pressing limitations of current regulatory approaches to online radicalisation.

REFERENCES

[1] T. Sear and N. Johnson. "Unprecedented Reach and Rich Online Journeys Drive Hate and Extremism Globally," *arXiv preprint*, 2023. Available: https://arxiv.org/abs/2311.08258

[2] C. McCauley and S. Moskalenko. "Mechanisms of Political Radicalization: Pathways Toward Terrorism," *Terrorism and Political Violence*, vol. 20, no. 3, pp. 415–433, 2008.

[3] J. Whittaker, "Online Radicalisation: What We Know," *EU*, 2022. [Online]. Available: https://home-affairs.ec.europa.eu/system/files/2023-11/RAN-online-radicalisation_en.pdf [Accessed: Feb. 28, 2025].

[4] R. Peels. "Cognitive and Behavioral Radicalization: An Explanatory Framework," *Terrorism and Political Violence*, 2025.

[5] J. Ware. "The Third Generation of Online Radicalization," *The George Washington University*. 2023. Available: https://extremism.gwu.edu/sites/g/files/zaxdzs5746/files/2023-06/third-generation-final.pdf [Accessed: Feb. 28, 2025].

[6] D. Hartmann, S. M. Wang, L. Pohlmann, and B. Berendt. "A Systematic Review of Echo Chamber Research: Comparative Analysis of Conceptualizations, Operationalizations, and Varying Outcomes," *Journal of Computational Social Science*, vol. 8, no. 2, 2025.

[7] E. Rodilosso. "Filter Bubbles and the Unfeeling: How AI for Social Media Can Foster Extremism and Polarization," *Philosophy & Technology*, vol. 37, no. 2, 2024.

[8] K. Lerman, D. Feldman, Z. He, and A. Rao. "Affective Polarization and Dynamics of Information Spread in Online Networks," *npj Complexity*, vol. 1, no. 8, 2024.

[9] R. Zhou. "Understanding the Impact of TikTok's Recommendation Algorithm on User Engagement," *International Journal of Computer Science and Information Technology*, vol. 3, no. 2, pp. 201–208, 2024.

[10] C. Chen, J. Kang, P. Sajjadi, and S. S. Sundar. "Preventing Users from Going Down Rabbit Holes of Extreme Video Content: A Study of the Role Played by Different Modes of Autoplay," *International Journal of Human-Computer Studies*, vol. 190, pp. 1–14, 2024.

[11] M. Mosnar, A. Skurla, B. Pecher, et al. "Revisiting Algorithmic Audits of TikTok: Poor Reproducibility and Short-term Validity of Findings," in *Proc. of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2025, pp. 3357–3366.

[12] I. Böhm and S. Lolagar. "Open source intelligence: Introduction, legal, and ethical considerations," *International Cybersecurity Law Review*, vol. 2, no. 2, 2021.

[13] M. Lakomy. "Open-source intelligence and research on online terrorist communication," *Media, War & Conflict*, vol. 17, no. 1, pp. 23–40, 2024.

[14] A. Oleart and L. Bouza García. "From Self to Co-regulation in the EU's Approach to Disinformation: The Framing Power of Big Tech Business Lobbies in the Lead to the Digital Services Act," *International Review of Public Policy*, vol. 7, no. 2, 2025.

[15] S. Mündges and K. Park. "But Did They Really? Platforms' Compliance with the Code of Practice on Disinformation in Review," *Internet Policy Review*, vol. 13, no. 3, 2024.

[16] OECD. "Disinformation and Russia's War of Aggression Against Ukraine: Threats and Governance Responses," *OECD Policy Responses on the Impacts of the War in Ukraine*, pp. 1–33, Nov. 2022.

[17] A. Turillazzi, M. Taddeo, L. Floridi, and F. Casolari. "The Digital Services Act: An Analysis of Its Ethical, Legal, and Social Implications," *Law, Innovation and Technology*, vol. 15, no. 1, pp. 83–106, 2023.

[18] European Commission. "Digital Services Act: Commission Designates First Set of Very Large Online Platforms and Search Engines," *EU*, 2024. Available: https://ec.europa.eu/commission/presscorner/detail/en/ip_23_2413 [Accessed: Feb. 28, 2025].

[19] European Commission. "Questions and Answers on the Digital Services Act," *EU*, 2024. Available: https://ec.europa.eu/commission/presscorner/detail/en/qanda_20_2348 [Accessed: Feb. 28, 2025].

[20] J. van de Kerkhof. "The DSA's Tower of Babel: On Digital Services Coordinators and Freedom of Expression," *European Journal of Risk Regulation*, pp. 1–26, 2025.

[21] European Commission. "Commission Opens Formal Proceedings Against TikTok under the Digital Services Act," *EU*, 2024. Available: https://ec.europa.eu/commission/presscorner/detail/en/ip_24_926 [Accessed: Feb. 28, 2025].

[22] S. H. Taylor and K. St. Cyr Brisini. "Parenting the TikTok Algorithm: An Algorithm Awareness as Process Approach to Online Risks and Opportunities," *Computers in Human Behavior*, vol. 150, pp. 1–9, 2024.



Proč čekat na budoucnost, když ji můžeš vytvářet? V brněnském R&D centru společnosti Honeywell pracujeme na technologiích, které pomáhají určovat směr celých odvětví! Zastav se na našem stánku a dozvíš se víc.

www.honeywell.com

careers.honeywell.com

THE FUTURE IS WHAT WE MAKE IT | Honeywell

PŘÍLEŽITOSTI PRO STUDENTY A ABSOLVENTY

V Honeywellu věříme, že budoucnost je taková, jakou ji společnými silami vytvoříme. V našem R&D centru v Brně se zaměřujeme na oblasti Aerospace a Industrial Automation.

KOHO HLEDÁME?

- Studenty a absolventy všech oborů FEKT VUT
- Zájem o dlouhodobou spolupráci a ochota dále rozvíjet znalosti z oboru
- Stáže: možnost pracovat 16–20 hodin týdně (flexibilní směny)
- Komunikativní znalost angličtiny (na úrovni B2 a vyšší)

JAKÉ POZICE U NÁS MŮŽETE NAJÍT

- Electrical/Electronics Engineer
- Hardware Engineer
- PLC Controls Engineer
- Systems Engineer
- ASIC/FPGA V&V Engineer
- Engineering Support

careers.honeywell.com
honeywell.com
Adela.Budjacova@Honeywell.com

Centralized Dispatch of Photovoltaic Plants, AVEVA Implementation and Operating Data Analysis

Jan Kalina
Department of Electrical Power Engineering
Faculty of Electrical Engineering and Communication
Brno University of Technology
Brno, Czech Republic
247387@vut.cz

Michal Ptacek
Department of Electrical Power Engineering
Faculty of Electrical Engineering and Communication
Brno University of Technology
Brno, Czech Republic
ptacekm@vut.cz

***Abstract*—This paper presents the concept and practical implementation of a technical dispatch center for renewable energy sources operated by ČEZ, a.s. in Málkov and its links to selected utility-scale photovoltaic power plants. It outlines the broader context of photovoltaic expansion in the Czech Republic, describes the Málkov dispatch center, and characterizes the AVEVA environment used for supervision, diagnostics, alarm handling, and remote control. Three photovoltaic plants integrated into the dispatch center are compared in terms of their technological systems and dispatch integration. The analytical part is based on supervisory and local plant data available at the time of paper preparation. The evaluation indicates a strong relationship between electricity generation and cumulative irradiance at the monitored photovoltaic sites and shows that panel temperature affects normalized plant power output. The alarm evaluation further identifies three practically relevant groups of recurring AVEVA alarms, providing a basis for improved alarm prioritization and more effective dispatch supervision.**

Keywords—renewable energy sources, technical dispatch center, photovoltaic power plant, AVEVA, SCADA

I. INTRODUCTION

The Czech power system is undergoing a gradual transformation associated with the increasing share of renewable electricity sources. According to [1] and [10], installed photovoltaic capacity and annual photovoltaic generation have grown significantly in recent years, while conventional fossil-based generation has been declining. This development is linked to both national and European decarbonization goals and to the expected long-term increase in electricity demand.

The growing role of photovoltaic generation also brings new technical and dispatch-related requirements. As discussed in [2]–[4], large photovoltaic power plants are characterized by variable and only partly predictable output, which increases the importance of reliable monitoring, voltage support, coordinated control, and clear interfaces between plant-level control systems and higher-level supervision. In distribution networks, these requirements are reflected in the need for dependable data transfer, alarm handling, and the implementation of active-power and reactive-power control functions.

One practical response to these requirements is the establishment of a dedicated technical dispatch center for renewable sources in Málkov, operated by ČEZ, a.s. As stated in [5], this workplace concentrates the supervision and control of selected renewable plants into one location and creates a unified environment for dispatching, diagnostics, and incident evaluation.

This paper therefore focuses on the concept and practical role of centralized technical dispatching for photovoltaic power plants. It describes the Málkov dispatch center and the AVEVA environment used there, compares selected photovoltaic plants connected to this dispatch structure, and presents the analytical part of the work based on production data, meteorological quantities, and recurring alarms evaluated from the dispatch perspective.

II. MÁLKOV TECHNICAL DISPATCH CENTER

The ČEZ OZE Málkov technical dispatch center was commissioned in May 2025 as a centralized workplace for the supervision and control of a renewable electricity portfolio operated by ČEZ, a.s. The facility is located in northern Bohemia and was created through the reconstruction of a former dispatch and substation site, as described in [5], which reduced construction demands while making use of existing energy infrastructure.

From the dispatch perspective, the Málkov center represents an important integration point for renewable sources supervision. As stated in [5], it enables continuous monitoring of connected plants, supports standardized testing of technological systems, facilitates incident evaluation, and creates conditions for coordinated real-time intervention when required. Fig. 1, based on [5], shows the control room with multi-monitor operator workplaces and a large projection wall displaying aggregated portfolio views and infrastructure maps, which helps dispatchers maintain a clear overview of the supervised portfolio.



Fig. 1. Málkov technical dispatch center.

The dispatch center receives measured values, status signals, and alarm information from local plant control and monitoring systems and processes them in a higher-level supervisory environment. According to [6], this arrangement was designed to make selected information easily available for dispatch supervision and control and also enables selected remote commands related mainly to active-power limitation, voltage support, reactive-power control, and the handling of abnormal or emergency states in cooperation with distribution system operators.

III. AVEVA AS THE OPERATOR ENVIRONMENT

In the Málkov dispatch center, the AVEVA system serves as the main operator environment for supervision, diagnostics, and control of the renewable portfolio. As described in [6], its structure combines portfolio-level overview screens with detailed plant views, allowing the dispatcher to move from an aggregated portfolio perspective to the diagnosis of specific technological components within a single integrated environment.

The main plant screen is organized into several functional parts, including the upper navigation area, the alarm strip, the central visualization field, and the right-hand information panel. The right panel contains identity and performance indicators together with the TSRF status semaphore, where T denotes technology, S inverters, R regulation, and F auxiliary functions. The colored semaphore fields provide an immediate indication of whether the monitored area is in a normal, deviating, or fault condition. Clicking the semaphore opens a short recent history of signal changes, as illustrated in Fig. 2 based on [6].



Fig. 2. Example of an AVEVA main plant screen for FVE Vrskmaň.

The individual AVEVA modules are organized according to dispatch and diagnostic logic. As explained in [6], the Intro screen links current production with meteorological conditions, while the Passport screen combines a table of plant parameters with an interactive map of strings and inverters for easier fault localization. Single-line diagrams and inverter screens provide more detailed technological insight, meter screens connect measured and billing quantities, the Meteo screen shows weather-related trends, and the UPS screen monitors backup-power autonomy. The control screen is used for dispatch commands and selected control actions.

In addition, reference [6] shows that AVEVA includes aggregated maps, alarm lists, and tools for time-based evaluation of plant behaviour. In the context of this paper, the system is important not only as the main interface used by dispatchers, but also as the environment in which production-related data and alarm information are interpreted from the dispatch perspective.

IV. SELECTED PHOTOVOLTAIC PLANTS

The paper further focuses on three photovoltaic plants integrated into the Málkov dispatch center: FVE Křižany, FVE Okrouhlička, and FVE Vrskmaň. These sites were selected because they are located in different parts of the Czech Republic while at the same time representing technologically comparable utility-scale photovoltaic installations. Their overall concept is similar, but each site also exhibits specific technical differences, which makes them suitable for comparing production data, alarm occurrence, and dispatch-relevant behaviour.

The compared plants are characterized not only by their basic parameters as described in [7]–[9], but also by the structure of their technological systems. This includes the overall system concept and topology, inverter arrangement, low-voltage and medium-voltage parts, transformer stations, protection systems, meteorological measurement, and local control and monitoring systems.

All three installations are based on string-inverter architecture and are designed for unattended operation with remote supervision. Table I based on [7]–[9] summarizes the basic technical parameters of the selected plants. As stated in [7]–[9], their technological systems provide automated control, online monitoring, and transfer of measured and status data to higher-level dispatching environments. The local systems are connected to AVEVA in Málkov, where dispatchers can supervise plant behaviour, evaluate alarms, and perform selected control actions.

TABLE I. TECHNICAL COMPARISON OF SELECTED PLANTS

Plant	Technical Parameters			
	DC power	***Panels***	***Inverters***	***Blocks***
Křižany	4.2636 MWp	7,752 x 550Wp	14x 250 kW	2
Okrouhlička	1.1962 MWp	3,480 x 565 Wp	6 x 250 kW	1
Vrskmaň	4.78668 MWp	8,472 x 565 Wp	15 x 250 kW	2

Although the plants differ in installed power, block arrangement, connection voltage, and several technical details, they share a number of features that are important from the dispatch perspective. These include meteorological measurement for interpreting production behaviour and the implementation of key control functions such as active-power limitation and voltage or reactive-power regulation in accordance with distributor requirements, as described in [7]–[9]. This combination of similarity and partial difference makes the selected plants suitable for comparative evaluation.

V. SELECTED RESULTS OF OPERATING-DATA AND ALARM EVALUATION

The analytical part of the paper is based on data exported from local RELIANCE control systems of the selected photovoltaic plants and on alarm information evaluated in the AVEVA environment. Local plant data were used as the primary source for production and meteorological evaluation because they provide direct access to historical measurements suitable for detailed comparison of the selected sites. The complete production-oriented evaluation is defined over a one-year period from April 2025 to March 2026. However, for the purposes of this paper, the monthly comparison is limited to validated data available at the time of paper preparation, and the incomplete March 2026 data are therefore not included in the presented graph. The alarm part is based on the evaluation of the twelve most frequent alarms identified during the first months of dispatch-center operation.

A. Monthly production in relation to cumulative irradiance

One of the main analytical findings is the close relation between monthly electricity generation and cumulative irradiance at all three monitored sites. The evaluated data show that the seasonal course of generation generally follows the seasonal course of available solar potential and that cumulative irradiance is one of the principal factors determining monthly production behaviour. For the paper, this dependence can be illustrated by Fig. 3 in a form of a common scatter plot of monthly generation versus cumulative irradiance for all three plants, with incomplete March 2026 data omitted. This figure along with coefficient of determination directly shows that higher cumulative irradiance is associated with higher monthly generation at each site, while differences among the three trends reflect local meteorological and plant-specific influences.

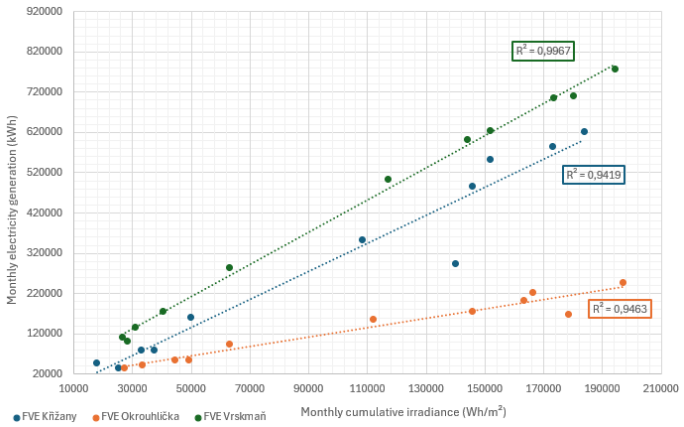

Fig. 3. Monthly electricity generation vs. cumulative monthly irradiance for the selected photovoltaic plants.

In addition to the qualitative relationship, plant performance can also be assessed through the performance ratio, which compares achieved energy yield with available irradiance. The evaluated results indicate that the selected plants differ not only in meteorological conditions but also in the effectiveness with which the available solar potential is converted into electricity.

B. Influence of panel temperature on normalized power

A second selected result concerns the influence of panel temperature on normalized power output. Normalized power output was used in order to reduce the direct influence of varying irradiance and to assess more clearly how power output changes with panel temperature. The evaluated data were based on 15-minute averages from a representative day with stable irradiance conditions, which made it possible to compare the selected sites under a similar meteorological conditions. The evaluated results indicate that increasing panel temperature is associated with a decrease in normalized power. This trend is most visible at FVE Vrskmaň, where the dependence is clearly developed over a broader temperature interval as shown by Fig.4. A similar trend was also observed at FVE Křižany.

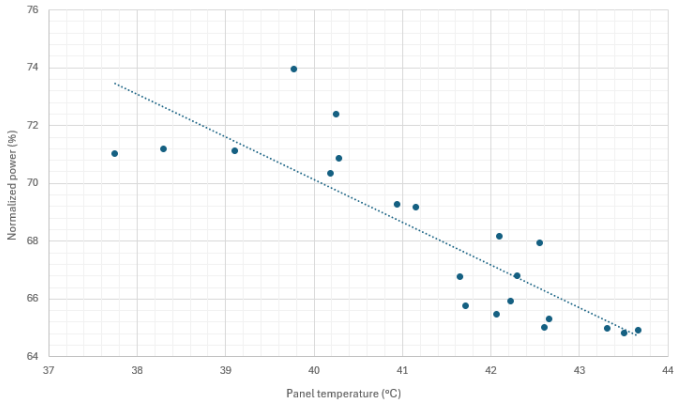

Fig. 4. Normalized power output vs. panel temperature for FVE Vrskmaň on 13 June 2025

At FVE Okrouhlička, the same dependence is less pronounced as shown by Fig.5, mainly because a narrower panel-temperature range was reached during the evaluated period. The generally lower temperatures may be attributed to this location's higher elevation. For this reason, the paper presents these two sites in separate figures.

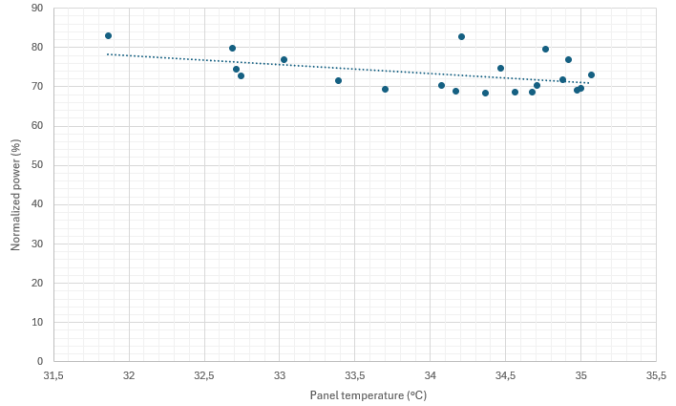

Fig. 5. Normalized power output vs. panel temperature for FVE Okrouhlička on 13 June 2025

C. Interpretation of the most frequent AVEVA alarms

The alarm evaluation naturally complements the production-oriented part of the analysis. The twelve most frequent alarms indicate three principal groups of states. The first group consists of technology-related alarms, such as ZF01, ZF02, ZU31, ZU32, and XA26, which have the highest technical relevance because they may indicate actual deviations or faults in part of the plant technology. The second group includes communication and supervisory alarms, especially YB32, ZZ01, ZQ22, and ZQ22 FB, which show that a significant part of the alarm picture is related to the stability of data transfer between inverters, PLCs, SCADA, and TELEOZE. The third group consists of alarms associated with transition or accompanying states, such as XU12 to XU14, XQ41, XW23, ZQ02, and XD02, whose occurrence is often linked to start-up, communication recovery, planned intervention, or control-mode changes and therefore does not usually represent a fault requiring immediate action. From the dispatch perspective, the key conclusion is that the most frequent alarm is not necessarily the most important one. The practical value of the alarm analysis therefore lies mainly in distinguishing between alarms that may indicate a real technological problem and alarms that are primarily informative or related to transition states.

D. Practical implications for dispatch supervision

Taken together, the selected analytical results show that dispatch-center data are useful not only for immediate supervision, but also for the technical interpretation of longer-term plant behaviour. The relation between monthly generation and cumulative irradiance helps distinguish meteorological influence from site-specific effects, the temperature-normalized comparison provides an additional view of plant performance under comparable irradiance conditions, and the alarm analysis supports clearer prioritization of recurring signals in the AVEVA environment. In this way, the analytical part contributes to more effective dispatch supervision and to a better interpretation of deviations that may appear in day-to-day plant monitoring.

VI. Conclusion

The paper shows that the increasing role of photovoltaic generation is not only a matter of installed capacity, but also of the quality of monitoring, diagnostics, and coordinated control. In this context, the Málkov technical dispatch center represents more than a standard supervision workplace, since it integrates portfolio-level monitoring, plant-level diagnostics, alarm handling, and selected remote interventions within one framework supported by the AVEVA environment.

By combining the description of the dispatch center with a comparison of three selected photovoltaic plants, the paper links the system-level view of technical dispatching with the specific technological characteristics of individual installations. The comparison of geographically different but technologically comparable plants also creates suitable conditions for the mutual evaluation of production behaviour, supervisory signals, and recurring alarms.

The selected analytical results indicate a close relation between monthly electricity generation and cumulative irradiance and show that panel temperature can contribute to differences in normalized plant power output even under comparable irradiance conditions. The alarm evaluation further shows that the most frequent AVEVA alarms can be divided into technology-related alarms, communication and supervisory alarms, and transition-state alarms, which differ significantly in their technical relevance for dispatch supervision.

The practical contribution of the paper therefore lies not only in the description of the Málkov dispatch concept and the AVEVA environment but also in showing how production data, meteorological quantities, and alarm information can be interpreted together to support clearer alarm prioritization, more effective supervision, and better understanding of deviations observed in photovoltaic plant behavior.

Acknowledgment

The author would like to express sincere gratitude to Petr Zelenka for providing materials and expert advice and thank the staff at the Málkov site for their cooperation, practical insights, and valuable information related to the operation of the technical dispatch center.

References

[1] L. Lepič, J. Malý, E. Majling, A. Lazarević, J. Šalavec, et al., „Elektroenergetika v ČR v roce 2024“ [Power Industry in the Czech Republic in 2024], OEnergetice.cz & Ušetřeno.cz, 2025. [Online]. Available: https://oenergetice.cz/public/static/pdf/elektroenergetika-v-cr-v-roce-2024.pdf. Accessed: Oct. 5, 2025.

[2] D. Harazim, „Vybrané problémy a dopady na řízení přenosové soustavy v období dekarbonizace energetiky“ [Selected Issues and Impacts on Transmission System Control in the Period of Energy Decarbonization], M.S. thesis, Czech Technical Univ. in Prague, Prague, Czech Republic, 2024.

[3] J. Dvořáček, „Autonomní a dispečerské řízení distribuovaných zdrojů v distribuční síti VN“ [Autonomous and Dispatch Control of Distributed Resources in MV Distribution Networks], M.S. thesis, Brno Univ. of Technology, Brno, Czech Republic, 2021.

[4] M. Dědek, „Připojování menších obnovitelných zdrojů energie do distribuční sítě“ [Integration of Small Renewable Energy Sources into the Distribution Grid], M.S. thesis, University of West Bohemia, Pilsen, Czech Republic, 2023.

[5] ČEZ, a.s., „Obnovitelné zdroje ČEZ budou řízeny z jednoho místa“ [ČEZ Renewable Sources Will Be Controlled from One Place], press release, 2025. [Online]. Available: https://www.cez.cz/next/cez/cs/pro-media/tiskove-zpravy/obnovitelne-zdroje-cez-budou-rizeny-z-jednoho-mista-220369. Accessed: Oct. 11, 2025.

[6] Š. Černý, „ČEZ Dispečink OZE Málkov – Návod k obsluze“ [ČEZ Renewable Dispatch Center Málkov – Operating Manual], Martia a.s., internal document, vol. 1, 2025.

[7] M. Kratochvíl, M. Kovařík, and L. Krejčí, „Místní provozní předpis FVE Křižany“ [Local Operating Procedure of PV Plant Křižany], internal document, 2025.

[8] M. Kratochvíl, P. Horký, and B. Ducháčková, „Místní provozní předpis FVE Okrouhlička“ [Local Operating Procedure of PV Plant Okrouhlička], internal document, 2025.

[9] J. Beran and M. Švarc, „Dodavatelský provozní předpis FVE Vrskmaň“ [Supplier Operating Procedure of PV Plant Vrskmaň], internal document, 2024.

[10] „Elektroenergetika v ČR v roce 2025” [Power Industry in the Czech Republic in 2025], OEnergetice.cz, 2026. [Online]. Available: https://oenergetice.cz/energostat/report/rocni/2025.Accessed: Feb. 27, 2026

Feed-Forward Curves Calculation for a Control Valve and Circulation Pump Performance Improvement During Ancillary Services Provision

Tomas Rojik
Faculty of Mechanical Engineering
Jan Evangelista Purkyne University in Usti nad Labem
Usti nad Labem, Czech Republic
mojeadresa@icloud.com

Jan Koudelka
Faculty of Mechanical Engineering
Jan Evangelista Purkyne University in Usti nad Labem
Usti nad Labem, Czech Republic
Brno University of Technology
Faculty of Electrical Engineering and Communication
Brno, Czech Republic
jan.koudelka@ujep.cz

***Abstract*—The paper gives the overview of feed-forward curves calculation in a system with a plate heat exchanger. It specifies the necessary inputs concerning heat source, heat exchanger, control valve, circulation pump and hydraulic circuit. Subsequently, it summarizes the calculation procedure in four steps. For three-way mixing valves, the procedure is extended to include the pressure balance of parallel flow paths (heating network and bypass). The calculated curves are compared with the actual settings used in the PLC of a real electric boiler unit that provides ancillary services (mFRR). The comparison shows a significant difference caused mainly by the control strategy and uncertain hydraulic parameters, which confirms the need for on-site calibration.**

Keywords—feed-forward curves, control valve positioning, ancillary services

I. Introduction

Feed-forward curves are a practical tool for improving the speed of a control loop response in heating systems with valves and pumps. A standard PID controller only starts to act after it detects a difference between the desired and actual temperature. A feed-forward curve, on the other hand, sets the valve and the pump to roughly the right position before any temperature error appears.

In practice, this means the control system immediately moves the valve opening and pump speed to values that match the new output level when the required heat output changes suddenly. This is the case, e.g., when an electric boiler responds to a grid balancing signal. The PID controller is then used only for correcting the small remaining difference.

The use of feed-forward curves has a serious impact on the transient phenomena in the heat system, as the transition is much faster and smoother. The improved performance characteristics have a significant impact on providing system response, especially in case of ancillary services provision. Due to better performance of the system, ancillary services providers may better comply with transmission system operator's requirements, and improve the frequency response and stability of the system.

The paper structure is as follows. Section II summarizes the required data for calculating the curves. Section III describes the calculation procedure. Section IV gives an overview on implementation and Section V gives an example of calculation and implementation for a specific unit with discussion.

II. Required Data for Calculation

Due to the complexity of the system, data about many of its parts need to be known for feed-forward curves calculation. This includes the heat source, the heat exchanger, the circulation pump, and the hydraulic circuit. The required data are specified hereinafter, and they are gathered from the equipment manufacturers and operators. The general layout of such a system is shown in Fig. 1.

A. Heat Source

The basic requested input is the range of thermal output of the heat source. Not only the information about maximum and minimum output, but also the control granularity (number of power steps), and how fast the output can change must be known.

Additional requested data about the heat source concerns the primary circuit: rated flow rate, and temperature difference across the source at rated power.

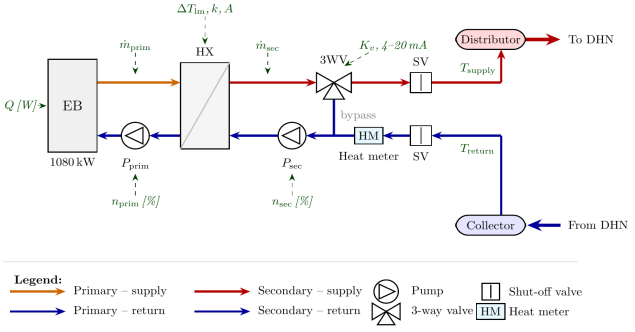

Fig. 1. Schematic diagram of one electric boiler circuit showing the key components and the calculated variables.

Heat Exchanger

The plate heat exchanger is a key component, because it determines the relationship between flow rates, temperatures, and transferred power. Input data obtained mostly from the manufacturer's documentation are:

- rated thermal output and the corresponding temperature differences on both sides,
- rated flow rates on the primary and secondary sides,
- pressure drop through the exchanger at rated flow (on both sides),
- heat transfer characteristic — how the overall heat transfer coefficient depends on flow rate [1].

The pressure drop of a plate heat exchanger grows roughly with the square of the flow rate. This relationship is important for the hydraulic calculation that follows.

B. Control Valve

For the three-way control valve, the following data are required:

- rated flow coefficient (flow at pressure difference of 1 bar for fully opened valve),
- flow characteristic (linear / equal-percentage / logarithmic),
- nominal size.
- stroke (maximum travel of the plug),
- positioning time (time for full stroke from 0 to100 %),
- control signal (typically 4–20 mA or 0–10 V).

The valve actuator is controlled by a 4–20 mA signal from the PLC, where 4 mA corresponds to 0 % stroke (fully open to the heating network, no bypass flow) and 20 mA corresponds to 100 % stroke (fully open to the bypass, all flow recirculates through the heat exchanger). The feed-forward curve therefore maps the required thermal output directly to this control signal.

The flow characteristic of a valve describes how the relative flow depends on the relative stroke (opening). With a linear characteristic, the relationship is straightforward: at 50 % opening, 50 % of the rated flow passes through the valve. With an equal-percentage characteristic, the relationship is exponential — the same change in stroke always causes the same percentage change in flow. The choice depends on the process being controlled.

An important parameter is the valve authority, defined as the ratio of the pressure drop across the fully open valve to the total pressure drop of the circuit. The authority affects the actual (installed) characteristic of the valve. At low authority, a linear characteristic becomes distorted towards the open end, and an equal-percentage characteristic moves closer to linear.

C. Circulation Pump

Requested data for a pump with variable speed control (frequency converter) are:

- Q–H characteristic (head vs. flow curve) at rated speed,
- speed control range (typically 30–100 % of rated speed),
- pump power input at the rated operating point,
- pump efficiency as a function of flow and speed.

To recalculate the characteristic for different speeds, the affinity laws are used [2]: flow is proportional to speed ($Q \sim n$), head is proportional to the square of speed ($H \sim n^2$), and power is proportional to the cube of speed ($P \sim n^3$). This allows to derive characteristics for any speed from a single Q–H curve.

D. Hydraulic Circuit

The last group of input data describes the hydraulic circuit:

- total pressure drop of piping, fittings, and the heat exchanger (excluding the valve) at rated flow,
- static pressure of the system (if relevant),
- required outlet temperature to the heating network and the return temperature.

These data are used to build the circuit characteristic (pressure drop vs. flow), which intersects the pump characteristic at the operating point.

III. Calculation Procedure

The calculation of feed-forward curves consists of several steps that build on each other. For each value of required output (from minimum to maximum), a complete thermal-hydraulic calculation is performed.

A. Required Flow Rate Calculation

From the required thermal output and the known temperatures, the required mass flow rate through the secondary circuit $\dot{m}$ (kg/s) is calculated from (1), where Q (W) is the required thermal output, c_p is the specific heat capacity of water (J/(kg·K)), and ΔT (K) is the temperature difference between the supply and return water.

$$Q = \dot{m} c_p \Delta T \tag{1}$$

The supply temperature is given by the setpoint (the desired outlet temperature to the heating network). The return temperature depends on the heat demand and operating conditions — for the feed-forward calculation, a constant or tabulated return temperature is typically used.

B. Heat Exchanger Calculation

Once the required output and the secondary flow rate are determined, the corresponding flow rate and temperature difference on the primary side of the exchanger are to be calculated. At partial load (below rated output), the log-mean temperature difference and the overall heat transfer coefficient of the exchanger both change. Therefore, a set of equations has to be solved iteratively:

- heat transfer equation for the exchanger (2), where A (m^2) is the area, k (W/(m·K)) is heat transfer coefficient, ΔT_{lm} (K) is the logarithmic mean temperature difference, and Q (W) is the heat transfer rate.

$$Q = k A \Delta T_{lm} \tag{2}$$

- heat balance on the primary side and the secondary side,
- the dependence of the heat transfer coefficient k on flow rates.

The result is the flow rate on both sides of the exchanger and all four temperatures (inlet/outlet, primary/secondary) for each power level.

C. Hydraulic Calculation — Valve Opening

The required flow coefficient of the valve K_v is calculated from the known secondary flow rate and the circuit parameters, according to (3), where Q_v (m³/h) is the volume flow rate, and Δp_v (bar) is the pressure drop across the valve, i.e. the difference between the available pump pressure and the pressure drop of the rest of the circuit.

$$K_v = \frac{Q_v}{\sqrt{\Delta p_v}} \tag{3}$$

From the known flow coefficient K_v and the valve's flow characteristic (linear or equal-percentage), one can derive the required stroke, i.e. the percentage of valve opening. For a linear characteristic the relationship is simply (4), where h is the actual stroke, H is the maximum stroke, and K_{vs} is the rated flow coefficient.

$$\frac{h}{H} = \frac{K_v}{K_{vs}} \tag{4}$$

For an equal-percentage characteristic, the relationship is logarithmic. In both cases, we must use the installed (distorted) characteristic that takes into account the valve authority in the given circuit.

In the case of three-way mixing valves, the hydraulic calculation becomes more complex [3]. The valve divides the flow into two parallel paths: one through the heating network and one through the bypass (short circuit). Each port of the valve has its own flow coefficient — for a linear three-way valve, $K_{v,A} = K_{vs}(1 - x)$ for the network port and $K_{v,B} = K_{vs} \cdot x$ for the bypass port, where x is the relative valve position. The real flow distribution between the two paths depends not only on the valve position but also on the pressure balance of the parallel branches. The bypass path typically has significantly lower hydraulic resistance than the heating network. Due to this, the pump tends to draw more flow through the bypass, even if the bypass port is only slightly open. This effect significantly reduces the effective control range of the valve. Therefore, the hydraulic network equations must be solved together with the thermal calculation.

D. Pump Speed Determination

The pump speed is chosen so that the pump delivers the required flow against the total pressure drop of the circuit (including the valve, the exchanger, and the piping). In circuits with three-way valves, the pump speed and the valve position are mutually coupled — a change in one affects the flow distribution in the parallel branches and therefore also the operating point of the other. For this reason, steps 3 and 4 must be solved simultaneously, not sequentially. The procedure is:

- For each flow rate, calculate the total circuit pressure drop (sum of piping, exchanger, and valve losses).
- Using the pump's Q–H curve and the affinity laws, find at what speed the pump delivers the required flow against the given pressure drop.
- Express the speed as a percentage of rated speed (0–100 %), which corresponds to the control signal for the frequency converter.

At low speeds, one must respect the minimum allowed pump speed (typically around 30 % of rated speed) and check that the operating point lies within the allowed region of the Q–H diagram.

IV. Feed-Forward Curves Building and Implementation

The calculation from section III is repeated for a set of required output values from minimum to maximum (for example in steps of 5–10 % of rated output). For each point, two values are obtained: valve opening (%) and pump speed (%). These points are plotted on a graph and fitted with a suitable curve.

These curves are implemented in the control system (PLC) either as a point table with linear interpolation, or as an approximation function (polynomial, piecewise linear). The choice depends on the capabilities of the particular PLC platform.

The calculated feed-forward curves must always be verified when the system is put into operation. Real conditions may differ from the calculation due to inaccurate input data, changes in return temperature during operation, older piping with higher roughness, or manufacturing tolerances of the pump and valve.

Tuning means that at various power levels, the actual outlet temperature is measured and compared with the setpoint. If the deviation is systematic, the corresponding points of the feed-forward curve are adjusted. The goal is to minimise the error that the PID controller has to compensate.

V. Application of Feed-Forward Curves for a Specific Unit

The previously described procedure of feed-forward curves calculation and implementation was applied to a real unit with electric boilers. The heat source consists of two ELTERM EP 1080 G2 electric boilers, each rated at 1080 kW, giving a total installed capacity of 2160 kW. Each boiler operates in 30 discrete power steps of 36 kW. The boilers are connected to a district heating system (CZT) through plate heat exchangers (Alfa Laval CB210-140H-F), with circulation pumps on the primary side, and three-way mixing valves (LDM RV224) and variable-speed circulation pumps (Grundfos) on the secondary side. The entire system is controlled by a Domat Control System PLC, which receives power setpoint commands via Modbus TCP from the energy trader's dispatch centre as part of the manual Frequency Restoration Reserve (mFRR) ancillary service [4]. Scheme of the system from SCADA system is shown in Fig. 2. The unit is included in ancillary services provisions, i.e. it responds to control signals from transmission system operator – response is shown in Fig. 3. As can be seen from the measured data in Fig. 3, the mFRR activation commands arrive at irregular intervals that can be very short — often only a few minutes apart.

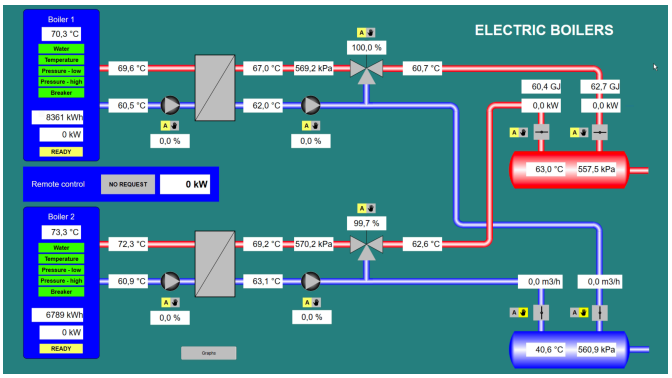

Fig. 2. Process scheme of the electric boiler unit (SCADA screenshot).

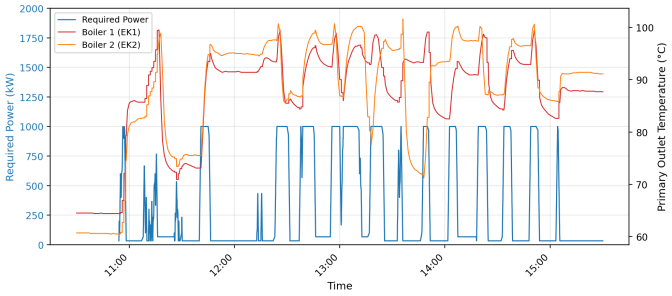

Fig. 3. Measured data from the electric boiler unit during mFRR activation.

When the time between consecutive activations is this brief, the boiler restarts with the water in the circuit still heated from the previous run. Without feed-forward curves, the PID controller starts from its default position and must gradually find the correct valve opening and pump speed. During this transient period, the water temperature inside the boiler can rise above 107 °C, triggering the boiler's internal high-temperature protection and causing an unplanned shutdown. These repeated thermal trips not only interrupt the ancillary service provision but also increase mechanical stress on the equipment. With feed-forward curves, the valve and pump are immediately set to positions that correspond to the current power level, so the controller only needs to compensate for the small difference between the pre-calculated and actual operating point. This significantly reduces the settling time, prevents dangerous temperature overshoots, and improves the stability of the outlet temperature to the district heating network. A formal quantitative comparison of the system response with and without feed-forward curves is planned as future work.

Fig. 4 shows a comparison between the calculated feed-forward curves and the values that are currently set in the PLC of the real plant. There is a clear difference between them. The calculation shows that the three-way valve becomes ineffective above approximately 400 kW (it is already fully open to the network), and the pump speed increases from 30 % to 100 %. However, in the real system the pump runs at a high constant speed (80–100 %) and the valve is used for flow control over the whole power range (from 70 % down to 5 %). The main reason for this difference is the control approach: the real system intentionally keeps the pump at high speed in order to create sufficient pressure drop across the oversized valve (K_{vs} = 160 m³/h for a maximum flow of approximately 46 m³/h).

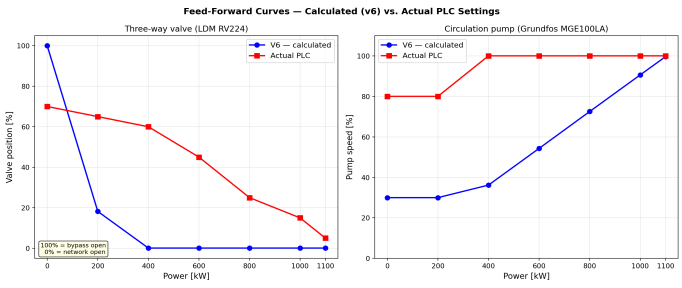

Fig. 4. Comparison of calculated and actual feed-forward curves for the three-way valve (left) and circulation pump (right).

This increases the valve authority and allows the valve to control the flow split between the network and the bypass. Other factors also contribute to the difference. The pressure drop of the CZT network and the pump Q–H had to be estimated, which adds uncertainty to the calculation. This comparison shows that calculated feed-forward curves should be treated as a starting point. Curves need to be checked and adjusted using real operating data. It also shows that the choice of control strategy — whether to save pump energy or to increase valve authority — has a significant effect on the shape.

VI. Conclusion

Feed-forward curves are an effective tool for improving the dynamic response of heating system control loops. Their calculation requires a thorough understanding of every component in the hydraulic circuit — from the heat source through the heat exchanger to the control valve and the pump.

The key steps are: determining the required flow rate from the heat balance equation, performing a thermal calculation of the heat exchanger, and a hydraulic calculation that determines valve position and pump speed by solving the pressure balance of the two parallel paths (heating network and bypass). In circuits with three-way valves and low valve authority, this is important because bypass resistance significantly influences flow distribution. The resulting curves should be verified by measurements during commissioning and fine-tuned to real operating conditions.

Well-designed feed-forward curves significantly reduce overshoot and settling time during sudden changes in output. This is especially important in systems that operate in grid balancing services, where the required output changes suddenly based on an external command.

References

[1] S. Kakaç, H. Liu, and A. Pramuanjaroenkij, *Heat Exchangers: Selection, Rating, and Thermal Design*, 3rd ed. Boca Raton, FL, USA: CRC Press, 2012.

[2] J. B. Rishel, T. H. Durkin, and B. L. Kincaid, *HVAC Pump Handbook*, 2nd ed. New York, NY, USA: McGraw-Hill, 2006.

[3] IEC 60534-2-1:2011, *Industrial-process control valves — Part 2-1: Flow capacity — Sizing equations for fluid flow under installed conditions*, International Electrotechnical Commission, 2011.

[4] ENTSO-E, "mFRR Implementation Framework for the Exchange of Balancing Energy from Frequency Restoration Reserves with Manual Activation," 2020.

Finite Element Analysis of Magnetic Fields in a High-Voltage Synchronous Machine

1st Dominik Heral
Dep. of Power Electrical and Electronic Engineering
Brno University of Technology
Brno, Czechia
Dominik.Heral@vut.cz

2st Martin Světlík
Dep. of Power Electrical and Electronic Engineering
Brno University of technology
Brno, Czechia
Martin.Svetlik@vut.cz

***Abstract*—This paper deals with the numerical modeling methodology of a salient-pole synchronous generator using Ansys Maxwell software. The study describes the development of a two-dimensional geometric model, emphasizing the application of rotational periodicity to optimize computational efficiency. The core of the research focuses on defining the discretization methodology for the computational domain and configuring transient analysis within a nonlinear magnetic environment. The objective of this paper is to present a procedure for investigating the influence of pole shoe geometry on magnetic field distribution and monitoring local magnetic circuit saturation, particularly in critical regions of the stator and rotor.**

***Index Terms*—Synchronous generator, Salient poles, Ansys Maxwell, Finite Element Method, Magnetic flux density.**

I. Introduction

In modern power engineering, synchronous generators play an indispensable role in electricity generation and the stabilization of the power grid. Especially in the field of hydropower, salient-pole machines are utilized due to their ability to operate efficiently at low speeds with a high number of pole pairs. However, increasing demands for efficiency, reliability, and power quality necessitate an ever-deeper understanding of internal electromagnetic phenomena. This paper focuses on the simulation of a generator for specific high-voltage levels, where the purity and quality of the output voltage represent essential parameters [1], [2], [3].

The electromagnetic design of these machines is highly complex due to the pronounced geometric asymmetry of the rotor magnetic circuit. Unlike non-salient pole machines, the salient-pole construction requires accounting for differing magnetic reluctances in the direct (d) and quadrature (q) axes. In addition to these geometric specifics, the nonlinear nature of ferromagnetic materials leads to local saturation of the magnetic circuit. This nonlinearity makes it considerably difficult to accurately determine machine parameters using classical analytical calculations [2].

While traditional analytical methods, based on simplified equivalent circuits, provide a good overview of the machine's global parameters, they fail to accurately capture local phenomena. These include, in particular, the saturation of stator teeth due to leakage fluxes and field distortion caused by stator slotting. For these purposes, the presented work utilizes numerical modeling via the Finite Element Method (FEM) within the Ansys Maxwell environment. This method allows for the solution of Maxwell's equations with high spatial resolution and provides a detailed insight into the magnetic field distribution, which is an essential basis for the design optimization of modern synchronous machines [4].

Authors gratefully acknowledge financial support from the Ministry of Education, Youth and Sports under institutional support and BUT specific research programme (project No. FEKT-S-26-9122)

II. Methodology

This chapter describes the theoretical foundation and computational procedures utilized in the simulation of the synchronous machine. It focuses on the physical principles of electromagnetic field solutions, the specifics of time-varying phenomena, and methods that allow for enhanced accuracy while maintaining efficient computational time [5], [6], [7].

- Physical Foundation and Numerical Solution: the design methodology for rotating electrical machines within the Ansys Maxwell environment is based on the numerical solution of Maxwell's equations. Due to the complex geometry and material nonlinearity, the FEM is employed. This method discretizes the model into a finite number of elements, where the governing equations are solved locally to ensure high precision in the resulting magnetic field distribution.
- Transient Analysis: to accurately capture the real-world operation of the generator, the use of a transient solver is essential. Unlike static or magnetostatic simulations, this approach accounts for time-varying phenomena resulting from the mechanical rotation of the rotor relative to the stator. This enables the monitoring of the dynamic evolution of induced voltage, currents, and electromagnetic torque at every time step of the simulation.
- Rotational Periodicity and Boundary Conditions: synchronous machines are characterized by significant geometric and electromagnetic symmetry. To reduce the computational cost and save computing power, periodicity conditions known as Independent and Dependent boundaries are defined. These conditions allow for the

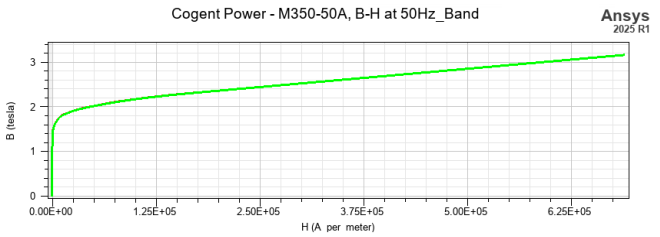

Fig. 1. Magnetization characteristic (B-H curve) of M350-50A electrical steel.

modeling of only a representative sector of the machine. The software then mathematically ensures that the magnetic flux transitions seamlessly across the boundaries of this sector, effectively simulating the behavior of the entire machine.

- Discretization Strategy (Mesh): the density of the computational mesh is critical for the accuracy of the results and must be distributed non-uniformly throughout the model. Fine elements are required primarily in the air-gap region and at the edges of the stator teeth. These areas experience the most significant gradients in magnetic flux and serve as the basis for calculating electromagnetic forces. In the homogeneous parts of the magnetic cores, where the field distribution is more uniform, coarser elements can be utilized. This approach significantly accelerates the overall simulation without compromising accuracy in critical regions.
- Nonlinear Material Properties: the accuracy of the electromagnetic field calculation in a synchronous machine is directly dependent on the accurate representation of the magnetization characteristics of the materials used. Given the high magnetic flux density values in the stator teeth and pole shoes, a constant permeability cannot be assumed. Therefore, nonlinear B-H curves of ferromagnetic steels are implemented into the model, allowing the transient solver to account for local saturation of the magnetic circuit. This approach is crucial for the correct determination of leakage fluxes and, subsequently, for the precise calculation of induced voltage and electromagnetic torque, both of which are directly influenced by iron nonlinearity.

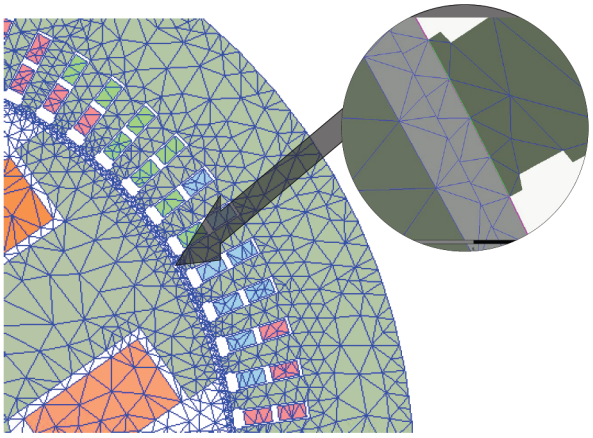

Fig. 2. Mesh discretization detail in the air-gap and pole shoe region.

III. Implementation

This chapter describes the application of the proposed methodology to the specific design of a synchronous generator. The objective of the implementation was to develop a numerical model that meets real-world operational requirements for high-voltage systems and to define simulation parameters such that the results provide an accurate representation of the machine's behavior [6], [7].

A. Design Parameters and Material Definitions

The first step of the implementation involved creating a 2D geometric model that accurately reflects parameters of the proposed machine. The subject of the study is a salient-pole synchronous generator with a power rating of 6–9 kVA, designed to operate at a high-voltage level of 6 kV.

The magnetic circuit geometry is based on a slotted stator core and a rotor with a defined number of pole shoes to ensure optimal electromagnetic coupling. To optimize computational resources, the modeled sector represents one-sixth of the machine, containing 24 stator slots and a single rotor pole shoe. Calculated for the entire machine, this configuration corresponds to a design with 144 slots and 6 poles.

Winding Analysis and Selection to achieve high output voltage quality, a double-layer lap winding was implemented. A crucial design feature in this configuration is the application of a short-pitch winding. The coil is not wound with a full pitch (corresponding to the pole pitch) but is intentionally shortened by one slot. This specific pitch was selected based on analytical requirements to suppress the 5th and 7th harmonic components of the magnetomotive force. In a salient-pole machine, where the magnetic field is naturally distorted by the non-uniform air gap, such short-pitching is essential to mitigate higher harmonic content and ensure a sinusoidal induced voltage waveform, which is critical for this voltage level.

Material Nonlinearity and Specifications to ensure an accurate representation of the machine's behavior under saturated conditions, the selection of ferromagnetic materials was critical, as they directly define the magnetic conductivity of the entire circuit. For the stator core laminations and rotor, M350-50A steel was selected. This material represents an industrial standard offering an optimal balance between cost and specific core losses, reflecting the objective of achieving high generator efficiency while maintaining economic competitiveness.

The full nonlinear B-H characteristic of this material was imported into the Ansys Maxwell simulation environment. This allows the transient solver to precisely calculate local changes in permeability across various sections of the stator teeth and rotor poles based on the instantaneous magnetic saturation. The magnetization characteristic of M350-50A steel in Fig.1 shows a saturation knee in the 1.7 T region.

Full nonlinear data were used for the simulation to faithfully capture the machine's behavior even under extreme saturation above 2.0 T. Copper was chosen for the winding components, including both the stator and excitation windings. The machine shaft, serving primarily a mechanical function, is defined using 1010 steel, which completes the magnetic circuit in the rotor region.

B. Computational Domain Discretization Strategy

The quality and fidelity of the numerical solution using the Finite Element Method are directly dependent on the selected density and structure of the discretization mesh. The computational mesh was designed to reflect the nonlinear magnetic behavior of materials and to ensure computational stability during the dynamic motion of the rotor relative to the stator. A detailed view of the generated mesh for the simulated machine sector is shown in Fig.2

The final model was discretized into a total of 3,196 elements. An initial mesh estimate of approximately 1,900 elements served exclusively for the baseline verification of the geometry and the configuration of boundary conditions. For the final transient analysis, it was essential to significantly refine the mesh, as under-discretization in such a complex salient-pole machine would lead to unacceptable numerical errors and instability in the results, making it impossible to accurately evaluate the time-domain waveforms of the induced voltage and electromagnetic torque.The meshing strategy was primarily targeted at the critical domains of the magnetic circuit where the most rapid changes in magnetic potential occur. Fine discretization was applied within the air-gap region, with primary emphasis placed on achieving the maximum possible mesh symmetry.

IV. Discussion

This chapter interprets the results of the transient analysis with respect to the specific design features of the generator. The discussion focuses on the correlation between the selected discretization strategy, the nonlinear material properties, and the resulting quality of the electromagnetic waveforms [8], [9].

- Magnetic Field and Saturation Analysis: The magnetic flux density distribution shown in Fig. 3 reveals that the most pronounced saturation occurs within the stator teeth and the rotor pole shoes. While the magnetic induction in most parts of the circuit ranges between 1.8 T and 2.0 T, peak local values in critical regions (specifically at the edges of stator teeth and pole shoe necks) reach up to 2.4 T. This saturation level directly corresponds to the nonlinear behavior of the utilized B-H characteristic in Fig. 1. The value of 2.4 T confirms that deep saturation occurs in the exposed parts of the magnetic circuit, significantly exceeding the magnetization knee of the M350-50A material. The performed Fourier transform in Fig. 5 subsequently verifies that despite the implementation of short pitch winding, low order harmonic components (particularly the 3rd and 5th orders) remain visible in the time domain waveform. The presence of these components is triggered precisely by this extreme local iron saturation, which alters the air gap flux distribution so fundamentally that the selected winding configuration cannot fully eliminate these nonlinear distortions.
- Induced Voltage Evaluation: the time-domain waveform of the induced voltage shown in Fig.4 exhibits a periodic character; however, it contains a noticeable ripple caused by stator slotting and the high saturation of the magnetic circuit. Although a double-layer winding with short-pitched coils was implemented to suppress the 5th and 7th harmonic components, the deep saturation exceeding the 2 T threshold introduces additional higher harmonics into the waveform. For applications requiring a higher quality sinusoidal output, a subsequent optimization step would be necessary, focusing on the shaping of the pole shoe or further refinement of the winding pitch.
- Impact of Discretization on Accuracy: the implementation of a highly symmetric and fine mesh within the air-gap region, as shown in Fig.2, proved to be crucial for computational stability. Despite the high level of

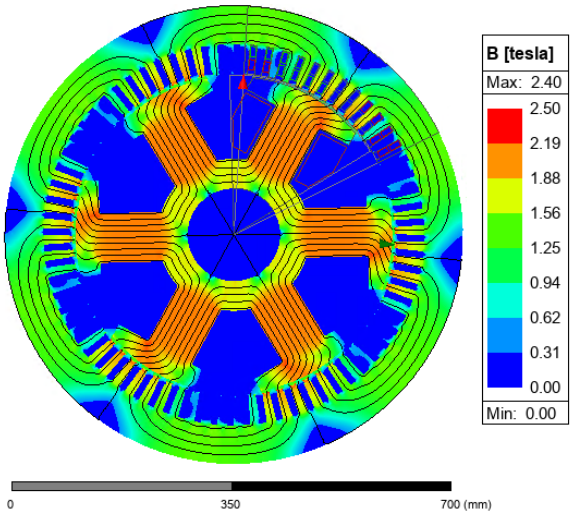

Fig. 3. Magnetic flux density distribution B in the machine cross-section at rated excitation.

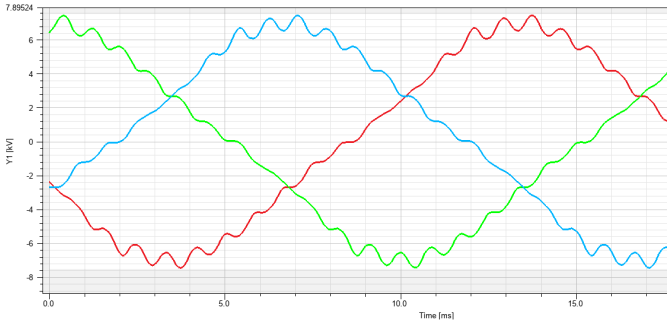

Fig. 4. Time-domain waveform of the three-phase no-load induced voltage showing the effect of stator slotting

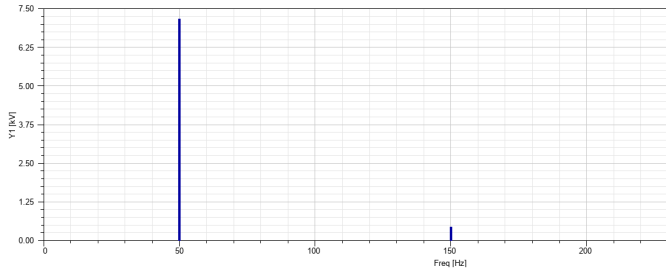

Fig. 5. Harmonic analysis (FFT) of no-load induced voltage at rated excitation.

saturation, the results exhibit no discernible numerical noise that could be mistaken for parasitic torques. The observed oscillations in the waveforms are therefore of purely physical origin, resulting from the interaction between the 144 stator slots and the 6 rotor poles within a nonlinear environment.

V. Conclusion

The presented work successfully demonstrated a methodology for the transient modeling of a high-voltage salient-pole synchronous generator within the Ansys Maxwell environment. The utilization of a two-dimensional model with the application of rotational periodicity (Independent/Dependent boundary conditions) proved to be a highly effective approach. This allowed for detailed discretization of critical machine regions while maintaining acceptable computational costs.

Saturation Analysis and Impact on Field Distribution the magnetic flux density distribution shown in Fig.3 revealed that saturation levels reaching up to 2.4 T occur in the stator tooth necks and the narrowed sections of the rotor pole shoes. This induction level significantly exceeds the knee of the magnetization characteristic of the M350-50A material. Such deep saturation causes an increase in magnetic reluctance at critical points, leading to undesirable distortion of the magnetic flux distribution and an increase in leakage fields within the air-gap region. To eliminate this phenomenon, a suitable solution appears to be the modification of the tooth geometry (e.g., widening the rotor and stator necks) to reduce local magnetic loading.

Induced Voltage Evaluation the time-domain waveform of the induced voltage shown in Fig.4 exhibits a high-frequency ripple, which directly contradicts the requirement for a purely sinusoidal output. Although the implementation of a short-pitched winding suppressed the lower harmonic components (5th and 7th), the influence of stator slotting and iron nonlinearity generates further distortion components. To ensure optimal operation and to achieve an ideal voltage waveform for this 6 kV generator, the next phase of the simulation will involve modifying the shape of the pole shoes and further optimizing the stator winding configuration. These measures are aimed at an even more significant suppression of higher harmonic components and achieving smoother machine operation.

Verification of the Numerical Model the achieved computational stability confirms the validity of the selected methodology. Due to the symmetrical arrangement of the computational mesh within the air-gap region presented in Fig.2, the resulting waveforms are free from parasitic numerical errors. Consequently, all oscillations observed in the simulation are of purely physical origin, determined by the interaction between stator slotting and the nonlinear material properties. As a result, the model is correctly validated for subsequent optimization steps.

The developed 2D transient model of the synchronous generator has successfully validated the simulation methodology for high-voltage machines using nonlinear material characteristics. The work has provided key insights into the behavior of the magnetic circuit under deep saturation (up to 2.4 T) and confirmed that, even with the implementation of a short-pitched winding, further geometric optimization is required to achieve ideal performance parameters.

The primary contribution of this work is the development of a stable and reliable simulation model which, through symmetrical discretization, eliminates numerical errors and enables the accurate study of electromagnetic phenomena within the machine. This model represents an essential foundation for the future design phase, which will focus on reshaping the pole shoes and reconfiguring the stator winding. The objective of these future steps is to achieve a purely sinusoidal voltage waveform and minimize electromagnetic interference, thereby bringing the model closer to the requirements for modern and efficient power equipment.

References

[1] V. Hrabovcová and P. Rafajdus, *Elektrické stroje – Teória a príklady*, 2nd ed. Žilina: EDIS, 2015.

[2] J. Pyrhönen, T. Jokinen, and V. Hrabovcová, *Design of Rotating Electrical Machines*, 2nd ed. Hoboken, NJ: Wiley, 2014, ISBN 9780470695166.

[3] Č. Ondrůšek, *Elektrické stroje. Elektronická skripta FEKT VUT v Brně*. Brno: FEKT VUT v Brně, 2021, elektronická skripta.

[4] V. Hrabovcová, P. Rafajdus, M. Franko, and P. Hudák, *Meranie a modelovanie elektrických strojov*, 3rd ed. Žilina: EDIS – vydavateľské centrum Žilinskej univerzity, 2014.

[5] P. C. Sen, *Principles of Electric Machines and Power Electronics*. Hoboken, NJ: Wiley Global Education, 2013, ISBN 9781118804155.

[6] F. Kiani and H. Tahanian, "A New Method for Fast and Accurate Calculation of Saturation Curves for Large Synchronous Generators," in *Proc. 2024 4th Int. Conf. on Electrical Machines and Drives (ICEMD)*, 2024, pp. 1–6, doi: 10.1109/ICEMD64575.2024.10963561. Struktura převodu

[7] Y. Zheng, F. He, and Q. Wang, "Simulation of electromagnetic field based on Simplorer and Maxwell software," in *Proc. 2018 Chinese Control and Decision Conf. (CCDC)*, 2018, pp. 3461–3466, doi: 10.1109/CCDC.2018.8407722.

[8] B. Kedjar, A. Merkhouf, and K. Al-Haddad, "Co-simulation for Finite Element Model Calibration of Synchronous Generators Connected to an Infinite Bus," in *Proc. 2022 Int. Conf. on Electrical Machines (ICEM)*, 2022, pp. 257–263, doi: 10.1109/ICEM51905.2022.9910748.

[9] L. P. Eschenmann and B. Ponick, "Influence of the Pole Shape of Salient-Pole Synchronous Machines used as Traction Motors," in *Proc. 2025 Int. Conf. on Electrical Drives and Power Electronics (EDPE)*, 2025, pp. 1–10, doi: 10.1109/EDPE66853.2025.11224297.

Simple Active Attitude Determination and Control System for CubeSat

1st Marek Ďurďa
Faculty of Electrical Engineering and Communication
Brno University of Technology
Brno, Czech Republic
256466@vut.cz

2nd Šimon Sloboda
Faculty of Electrical Engineering and Communication
Brno University of Technology
Brno, Czech Republic
simon.sloboda@vut.cz

***Abstract*—This paper presents the design and ongoing development of a custom Attitude Determination and Control System (ADCS) for CubeSats. The system utilizes a magnetometer, gyroscope, and custom sun sensors for attitude determination, fused via a Multiplicative Extended Kalman Filter (MEKF) using singularity free quaternions. Reaction wheels and magnetorquers, managed by a Proportional Derivative (PD) controller, serve as actuators. The paper highlights the core mathematical models, the custom hardware and software architecture, and the current challenges of this in progress project.**

***Index Terms*—ADCS, CubeSat, nanosatellite, attitude control, reaction wheels, magnetorquers, Kalman filter, MEKF, quaternions, sensor fusion.**

I. Introduction

The Attitude Determination and Control System (ADCS) is an essential subsystem for measuring and actively controlling a satellite's spatial orientation. This paper details the ongoing development of a custom, simplified CubeSat ADCS.

The proposed system determines its orientation using a magnetometer, gyroscope, and custom sun sensors. To ensure precise and stable attitude estimation, sensor data is fused via a Multiplicative Extended Kalman Filter (MEKF) based on quaternion kinematics. Active control is achieved through reaction wheels and magnetorquers. We present the system's custom hardware architecture, proprietary C based software libraries, and actuator control algorithms. Currently, the prototype operates with simplified environmental models, omitting dynamic changes in the geomagnetic or Sun position vectors.

II. Sensors

To determine an absolute 3D orientation in space, the Attitude Determination and Control System (ADCS) requires at least two non collinear reference vectors. In this system, these are represented by the Sun vector and the geomagnetic field vector. While the gyroscope provides fast and highly dynamic measurements of angular velocity, it cannot determine absolute orientation on its own. Furthermore, integrating gyroscope data inherently leads to an accumulating error over time, known as drift or bias. To continuously correct this drift, absolute attitude sensors are necessary. However, these reference sensors are not ideal, they are inherently noisy and highly susceptible to external environmental disturbances, which necessitates the use of robust sensor fusion algorithms.

A. Sun Sensor

The sun sensor system utilizes six independent sensor boards, intended to be installed on each face of the CubeSat. Each board features a BPW34S photodiode connected to a TSV991 operational amplifier. The operational amplifier is configured as a transimpedance amplifier to convert the generated photocurrent into a measurable voltage. The TSV991 was chosen specifically for its extremely low input bias current and rail to rail capabilities.

B. Magnetometer

An IIS2MDCTR is used as magnetometer. Because magnetometers are extremely sensitive devices, they are easily affected by parasitic magnetic fields and electromagnetic interference from surrounding electronics. To eliminate these effects, the PCB layout was strategically designed without a large continuous ground plane under the sensor component, ensuring that the measured geomagnetic field is not shielded or severely distorted.

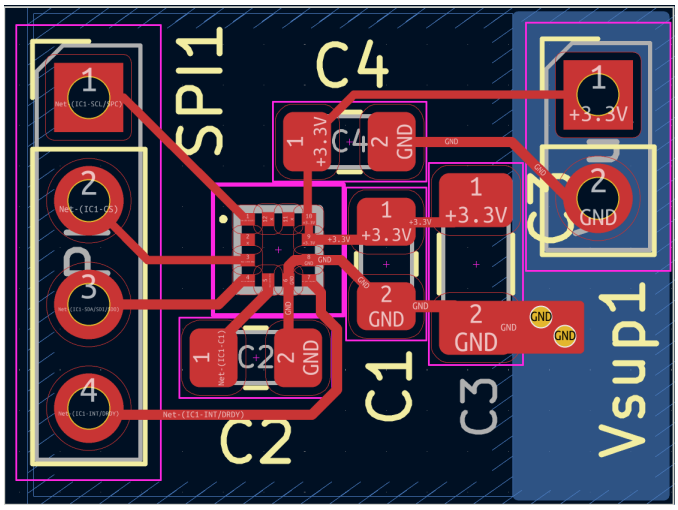

Fig. 1. Magnetometer PCB layout designed without a ground plane to prevent signal shielding.

C. Gyroscope

For angular velocity sensing, a breakout board featuring the MPU9250 IMU is utilized. Although the MPU9250 is an

older generation component, it was selected due to its good availability and its strong performance parameters compared to other modern IMUs, particularly low white noise characteristics.

III. Attitude Determination

Accurate attitude determination is a prerequisite for reliable spacecraft control. While Euler angles provide an intuitive representation of 3D orientation, they suffer from singularities known as gimbal lock when the pitch angle approaches ±90°. To ensure universal stability and reduce computational complexity, this ADCS uses quaternions to represent rotations.

A. Quaternion based orientation

Quaternions are numbers of four elements where first element is the real part and three others are imaginary. We can rotate vector by quaternion with equation below.

$$\boldsymbol{v_{rot}} = \boldsymbol{q} \cdot \boldsymbol{v} \cdot \boldsymbol{q^*} \tag{1}$$

Quaternions are widely used for determining orientation in most modern systems, especially those, rotating more than 90 degrees. For purpose of integration of quaternion in time from angular velocity, equation below is used.

$$\boldsymbol{\Delta q} = cos(\frac{1}{2}||\boldsymbol{\omega}||\Delta t) + \frac{\boldsymbol{\omega}}{||\boldsymbol{\omega}||} \cdot sin(\frac{1}{2}||\boldsymbol{\omega}||\Delta t) \tag{2}$$

B. Sensor Fusion using MEKF

The system relies on a Multiplicative Extended Kalman Filter (MEKF) to fuse the high rate, but drifting, gyroscope data with the noisy, but absolute, reference vectors from the magnetometer and the sun sensors. Unlike a standard Extended Kalman Filter, which is additively updated, the MEKF separates the nominal state (quaternion and gyroscope bias) from the error state. The filter only estimates the error state vector $\boldsymbol{\delta x}$, which consists of the angle error $\boldsymbol{\delta\theta}$ and the bias error $\boldsymbol{\delta b}$:

$$\boldsymbol{\delta x} = \begin{bmatrix} \boldsymbol{\delta\theta} \\ \boldsymbol{\delta b} \end{bmatrix} \tag{3}$$

By operating on the error state, the MEKF maintains linear properties near the origin.

C. Prediction Step

During the prediction phase, the angular velocity $\boldsymbol{\omega_{meas}}$ measured by the gyroscope is first corrected by subtracting the current bias estimate $\boldsymbol{b_k}$. The corrected angular velocity is then used to integrate the nominal quaternion forward in time. Concurrently, the error covariance matrix $\mathbf{P}$ is updated:

$$\mathbf{P_{k+1}^{-}} = \mathbf{F}\mathbf{P_k}\mathbf{F^T} + \mathbf{Q} \tag{4}$$

where Q is the process noise covariance matrix and F is the system transition matrix, defined as:

$$\mathbf{F} = \begin{bmatrix} \mathbf{I_{3\times3}} & -\mathbf{I_{3\times3}} \cdot \Delta t \\ \mathbf{0_{3\times3}} & \mathbf{I_{3\times3}} \end{bmatrix} \tag{5}$$

This matrix essentially models how the orientation error grows proportionally to the integrated bias error over the sampling time $\Delta\mathbf{t}$.

D. Correction Step

In the correction phase, the absolute reference vectors (the geomagnetic field and the Sun direction) are used to correct the accumulated drift. First, the known reference vectors are rotated into the body frame using the predicted quaternion to obtain predicted measurements $\boldsymbol{v_{pred}}$. The measurement residual $\boldsymbol{z}$ is then calculated as the difference between the actual sensor measurements and $\boldsymbol{v_{pred}}$.

The measurement sensitivity matrix $\mathbf{H}$ is constructed using the skew symmetric matrices of the predicted vectors:

$$\mathbf{H} = \begin{bmatrix} [\mathbf{v_{predMAG}}]\times & \mathbf{0_{3\times3}} \\ [\mathbf{v_{predSUN}}]\times & \mathbf{0_{3\times3}} \end{bmatrix} \tag{6}$$

Using standard Kalman filter equations, the Kalman gain K is computed, yielding the new error state $\boldsymbol{\delta x} = \mathbf{K} \cdot \boldsymbol{z}$. In the final step (error injection), the bias is additively corrected, while the nominal quaternion is multiplicatively updated using the error angle $\boldsymbol{\delta\theta}$. After the update, the error state is reset to zero for the next iteration.

E. Simulations

Although the MEKF algorithm is implemented in C for the STM32 MCU, it was also implemented in MATLAB to verify its function.

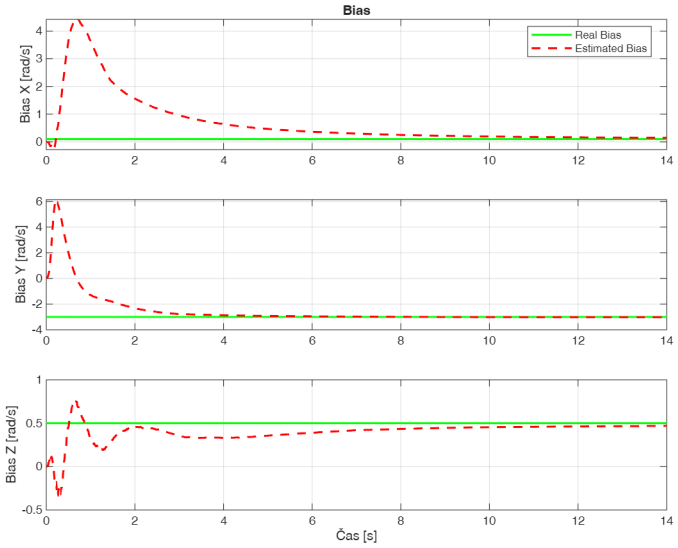

Fig. 2. Bias simulations of implemented MEKF algorithm

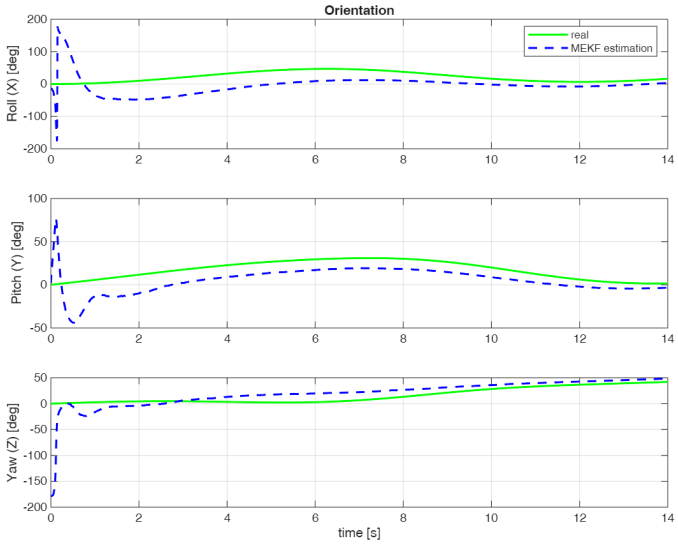

Fig. 3. Orientation simulations of implemented MEKF algorithm

In simulations, we can see synthetic data from the sun, magnetic sensor, and gyroscope with added noise and drift. The bias chosen for this simulation is constant and its real values are bias = [0.1, -3, 0.5], which the filter correctly converges to. Also, we can see that the filter orientation values track the true synthetic orientation. This simulation proves that the filter is working. However, we cannot yet estimate how the filter will perform with real sensors. An important fact is that noise matrices were filled with the exact same noise parameters that were added to the data. In the real world, these values will have to be tuned.

IV. Actuators

The primary purpose of the actuators is to change and maintain the desired spacecraft orientation based on the determined attitude data. In active ADCS systems for CubeSats, reaction wheels and magnetorquers are commonly used. These two types of actuators are combined to compensate for the parasitic properties of reaction wheels, such as the accumulation of residual angular momentum caused by bearing friction and external disturbance torques.

A. Reaction Wheels

The proposed ADCS utilizes three reaction wheels arranged in a standard orthogonal configuration. This topology was selected to simplify both hardware design and software implementation regarding the control moment distribution. The mechanical actuation is provided by 1503 brushless DC (BLDC) motors, which are driven by DRV10975 sensorless motor drivers. The DRV10975 utilizes Field Oriented Control (FOC) to determine the rotor position via the back electromotive force (BEMF) induced in the motor windings.

B. Magnetorquers

Because reaction wheels accumulate momentum over time, magnetorquers are required for momentum dumping and initial detumbling maneuvers. Coreless magnetic coils were chosen for this design. The absence of a ferromagnetic core eliminates remanent magnetic induction, thereby minimizing parasitic interference with the highly sensitive onboard magnetometer. The coils are driven by DRV8837C H-bridges and are controlled using the B-dot algorithm. This algorithm calculates the required magnetic dipole moment μ based on the rate of change of the measured geomagnetic field vector $\boldsymbol{B}$:

$$\boldsymbol{\mu} = k \cdot \frac{d\boldsymbol{B}}{dt} \tag{7}$$

where k is a positive tuning gain. The driver circuitry for both the BLDC motors and the magnetic coils is integrated into a custom designed breakout PCB, shown in Fig. 4.

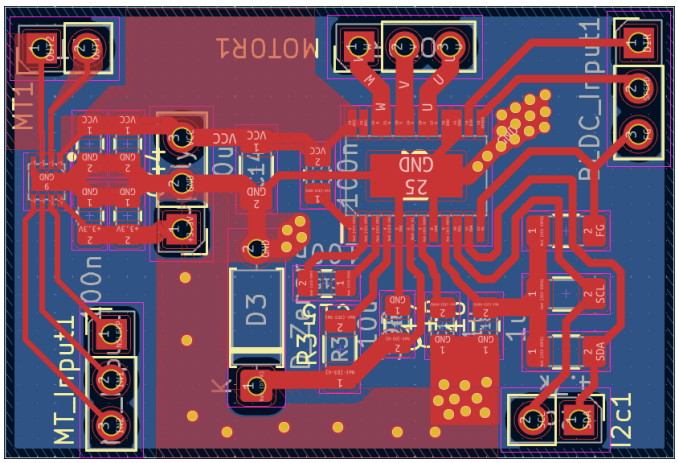

Fig. 4. Actuator driver PCB integrating the BLDC and H bridge controllers.

To construct rectangle coil, following equation was derived:

$$\mu = \frac{U}{\rho \frac{4}{S}} \cdot a \tag{8}$$

Where a is length of side of the coil, S is cross section of wire, ρ is resistivity of wire material and μ is magnetic dipole moment of final coil. There was also derived equations to calculate time of rotation of satellite in small discrete angle steps, which were then implemented in MATLAB.

$$\phi_{n+1} = \omega_n \cdot t + \frac{\mu \cdot B \cdot sin(\alpha)}{2 \cdot J_{sx}} t^2 \tag{9}$$

where angular velocity is sum of all angular velocities from previous steps. For coil with $0.2025 Am^2$ calculated time of rotation is $37 sec$.

V. Attitude Control

Once the current spatial orientation is continuously estimated by the MEKF, the control algorithms compute the required torques to reach and maintain the target attitude. The control loop operates directly with quaternions to maintain mathematical stability across all rotation angles.

A. Reaction Wheel Control Law

To target a specific orientation, the error quaternion q_E is first calculated from the target quaternion q_T and the currently estimated quaternion q_S. The error quaternion represents the rotation required, to align the satellite's current frame with the desired reference frame:

$$\boldsymbol{q_E = q_T \otimes q_S^{-1}} \tag{10}$$

Using the vector part of the error quaternion (q_{1E}, q_{2E}, q_{3E}) and the scalar part (q_{4E}), along with the angular velocity rates (p, q, r) obtained from the corrected gyroscope data, a Proportional Derivative (PD) control law is implemented for the three orthogonal axes:

$$\begin{aligned} T_{cx} &= 2K_x q_{1E} q_{4E} + K_{xd} p \\ T_{cy} &= 2K_y q_{2E} q_{4E} + K_{yd} q \\ T_{cz} &= 2K_z q_{3E} q_{4E} + K_{zd} r \end{aligned} \tag{11}$$

where K and K_d are the proportional and derivative tuning gains, respectively. These gains are specifically derived based on the moments of inertia of both the satellite and the reaction wheels.

VI. System Architecture

The hardware architecture of the proposed ADCS is highly integrated and optimized for the CubeSat form factor. The core of the computational unit is an ARM Cortex-M4 based STM32F407VET6 microcontroller. This MCU was selected for its hardware Floating Point Unit (FPU), which is crucial for the efficient real time execution of the matrix operations.

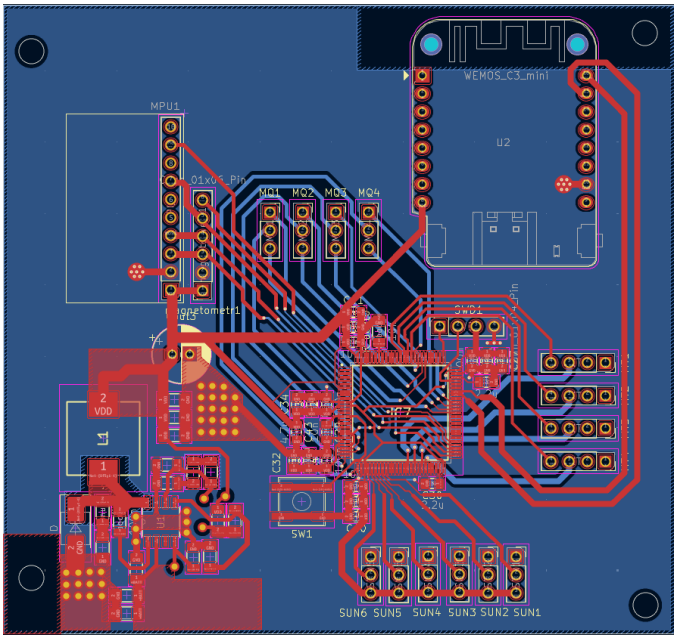

Fig. 5. control ADCS unit with STM32F407VET6 MCU

Power management is handled by an onboard DC/DC buck converter utilizing the TPS54331D integrated circuit, which reliably steps down the primary 7.2V battery voltage to the logic levels required by the sensors and the MCU.

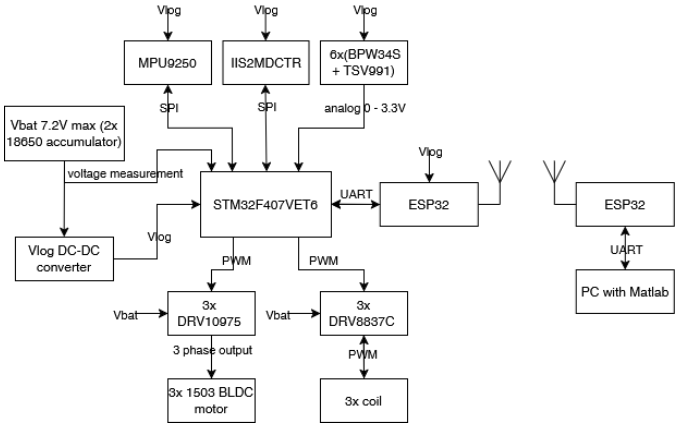

Fig. 6. Complete block diagram of the designed ADCS hardware architecture.

For testing purposes, the system incorporates a wireless communication link using ESP32 modules via the ESP NOW protocol. This interface allows fetching orientation and actuator data, and sending control commands from a PC running MATLAB. Plastic models for the designed PCBs and parts were also developed. These models meet the CubeSat standard requirements and utilize the PC104 standard for PCB size and hole pitches.

VII. Conclusion

This paper outlined the architectural design and theoretical foundations of a custom built ADCS for CubeSats. The core mathematical models, including quaternion kinematics and the MEKF for robust sensor fusion, have been successfully established and implemented in custom software libraries. Furthermore, the hardware architecture encompassing the sensor suite, control MCU, and actuator drivers has been designed. Physical PCBs and 3D printed structural models compliant with the PC 104 standard have already been manufactured.

As the project is currently in the active development phase, immediate future work will focus on the complete physical integration of the hardware components and testing of the control system.

References

[1] M. J. Sidi, *Spacecraft Dynamics and Control: A Practical Engineering Approach*. Cambridge, UK: Cambridge Univ. Press, 1997.

[2] C. Cappelletti, S. Battistini, and B. K. Malphrus, *CubeSat Handbook: From Mission Design to Operations*. London, UK: Academic Press, 2021.

[3] J. Sol, "Quaternion kinematics for the error-state Kalman filter," *arXiv preprint arXiv:1711.02508*, Nov. 2017.

[4] M. Hampsey, "The Multiplicative Extended Kalman Filter," Jul. 2020. [Online]. Available: https://matthewhampsey.github.io/blog/2020/07/18/mekf. [Accessed: Mar. 14, 2026].

[5] STMicroelectronics, "STM32F405xx, STM32F407xx: ARM Cortex-M4 32b MCU+FPU Datasheet," DocID022152, 2016.

[6] The CubeSat Program, "CubeSat Design Specification (1U-12U)," Rev. 14.1, California Polytechnic State University, San Luis Obispo, CA, USA, 2022.

Web-Based 3D Simulator for Industrial Automation Processes

Filip Majchrák
Faculty of Electrical Engineering and Communication
Brno University of Technology
Brno, Czech Republic
256346@vutbr.cz

Jakub Arm
Faculty of Electrical Engineering and Communication
Brno University of Technology
Brno, Czech Republic
arm@vut.cz

***Abstract*—This paper presents the design and implementation of a web-based 3D simulator for industrial automation processes. The simulator enables visualization of technological processes in a virtual environment and allows interaction with control logic implemented in a programmable logic controller (PLC). Communication between the simulator and the control system is realized using the Modbus TCP protocol, while communication tunneling with the web interface is provided through WebSocket communication. The functionality of the developed system is demonstrated on a virtual production line used for sorting cylindrical pins according to their thickness. The presented technology enables the virtual commissioning of PLC program.**

***Index Terms*—Industrial automation, simulation, virtual commissioning, WebSocket, Modbus TCP, Three.js**

I. Introduction

In recent years, concepts related to Industry 4.0 have been increasingly applied in industrial automation. These concepts emphasize the digitalization of manufacturing processes and the integration of physical systems with their digital counterparts. One of the important approaches in this area is the use of virtual models of technological systems, which allow the behavior of the system to be analyzed and tested before it is deployed in real operation. This approach is commonly used in applications such as virtual commissioning and the development of digital twins of manufacturing systems [1], [4].

Virtual models of production systems make it possible to simulate the behavior of technological processes, test control algorithms, and verify different operational scenarios without interfering with the real equipment. Such simulations can reduce commissioning time and improve the overall quality of the control system design [2].

In recent years, web technologies have also become increasingly used for the visualization and simulation of complex systems. Modern web-based graphics libraries enable the creation of interactive three-dimensional visualizations directly in a web browser, making simulation tools more accessible without the need for specialized software [3].

This paper focuses on the design and implementation of a simple 3D simulator of an automation process. The simulator enables visualization of system behavior in a three-dimensional environment and allows interaction between objects, sensors, and control logic to be simulated. It also provides an overview of the input and output states of the control system and allows the simulated process to be monitored in real time.

II. System Architecture

The developed simulator is designed as a modular system that combines a web-based user interface, a 3D visualization environment, and a communication layer responsible for data exchange with a software PLC (softPLC). In the presented implementation, the control logic is executed in the Beckhoff TwinCAT softPLC environment. The main goal of the system is to provide a platform for simulating industrial automation processes and visualizing their behavior in real time.

The architecture of the simulator consists of several main components. The first component is a web-based user interface that runs directly in a standard web browser. The user can observe the simulation, monitor system variables, and visualize changes of input and output signals of the control system.

The server part of the application is implemented using the Node.js runtime environment. The server communicates with the softPLC using the Modbus TCP protocol and processes the exchanged data. The processed data are then transmitted to the web client through a WebSocket connection, enabling data synchronization between the control logic and the simulation environment.

Another important component is the 3D simulation environment that visualizes the technological process. It contains objects representing elements of the production system such as conveyors, sensors, and transported objects. The state of these objects is continuously updated according to the values of the input and output signals received from the control system.

The proposed architecture enables the integration of the virtual simulation environment with the control logic and allows testing of automation scenarios without the need for a physical production system.

III. Used Libraries

Several JavaScript libraries were used during the development of the simulator in order to simplify the implementation of 3D visualization, physics simulation, and system communication.

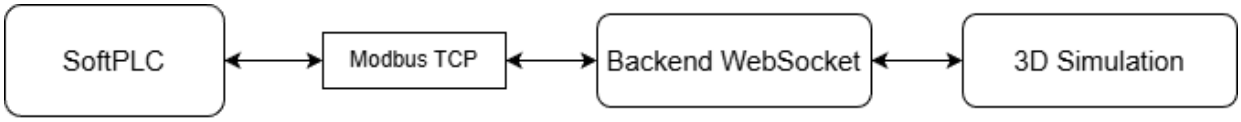

Fig. 1. System architecture of the developed simulator.

The Three.js library is used for rendering the 3D scene and managing objects, camera, and lighting in the virtual environment [5]. Physijs provides a physics layer for Three.js and enables simulation of object movement and collisions [6].

For data communication between the client and the server, the WebSocket library ws is used [10]. Communication with the softPLC is implemented using the jsmodbus library, which provides support for the Modbus TCP protocol [8]. The server part of the application is implemented using the Express framework [9].

IV. 3D Simulation Environment

The 3D simulation environment is responsible for visualizing the technological process and updating the state of the virtual scene in real time. The visualization is implemented using the Three.js library, which enables rendering of interactive three-dimensional scenes in a web browser.

The simulation runs in a continuous update loop in which the scene is rendered and the state of simulated objects is updated. During each cycle, the simulator processes the current values of input and output signals and synchronizes the virtual environment with the control logic.

The virtual scene contains several functional elements including a conveyor, transported objects, and virtual sensors. The conveyor moves objects along a predefined path, while the sensors detect the presence of objects within specific detection areas and generate corresponding input signals for the control system.

To simulate realistic object movement and interactions, the environment uses the Physijs physics engine, which provides rigid-body physics for Three.js scenes [6]. This enables simulation of object movement, collisions, and simple physical interactions within the production process.

V. Control Logic and Communication

The behavior of the simulated technological process is controlled by a software PLC (softPLC), which processes input signals and generates corresponding output signals. In the proposed system, the softPLC used for testing is implemented in the Beckhoff TwinCAT environment.

Communication between the simulation environment and the softPLC is implemented using the Modbus TCP protocol, which is commonly used in industrial automation systems [7]. The Modbus TCP protocol organizes data into several types of registers, including coils, discrete inputs, input registers, and holding registers.

In the developed simulator, selected Modbus addresses are used to exchange signals between the control logic and the simulation environment. Each Modbus address corresponds to a specific variable in the simulator, such as a sensor signal, actuator state, or process parameter.

To simplify system configuration, the mapping between Modbus registers and simulation variables is ... defined in a separate JavaScript Object Notation (JSON) configuration file, a lightweight format for structured data. This allows users to modify signal assignments without changing the simulator source code. When a simulation scene is loaded, the server application reads the JSON configuration and the WebSocket communication layer uses it to map Modbus registers to the corresponding simulation variables.

The communication between the simulator and the PLC forms a closed control loop. Virtual sensors in the simulation act as PLC inputs, while the outputs generated by the PLC control the behavior of actuators in the virtual environment. The simulator periodically reads PLC outputs and updates the scene, while sensor states are written back to the corresponding Modbus input registers.

The communication is performed on localhost while OS and communication impact with delays and uncertainties. The communication period is set to value 500 ms. According to the measurement of 40 communication cycles, the period is 2.49 ms with jitter 48.57 ms using standard Windows OS.

The register mapping used in the developed simulator is summarized in Table I.

TABLE I
Modbus register mapping used in the simulator

Address	Type	Signal
32768	Input	inputs.start
32769	Input	inputs.conv
32770	Input	inputs.conv2
32771	Input	inputs.p1
32772	Input	inputs.p2
32773	Input	inputs.p3
32768	Output	outputs.conv1end
32769	Output	outputs.conv2end
32770	Output	outputs.s1
32771	Output	outputs.s2
32772	Output	outputs.s3
32773	Output	outputs.p1_ex
32774	Output	outputs.p2_ex
32775	Output	outputs.p3_ex
32776	Output	outputs.p1_rec
32777	Output	outputs.p2_rec
32778	Output	outputs.p3_rec
32768	Holding register	outputs.dist1 (scaled ×100)

In the implemented system, Modbus addresses starting from 32768 are used for simulation signals. This range corresponds to the internal Modbus TCP memory area defined in the TwinCAT documentation.

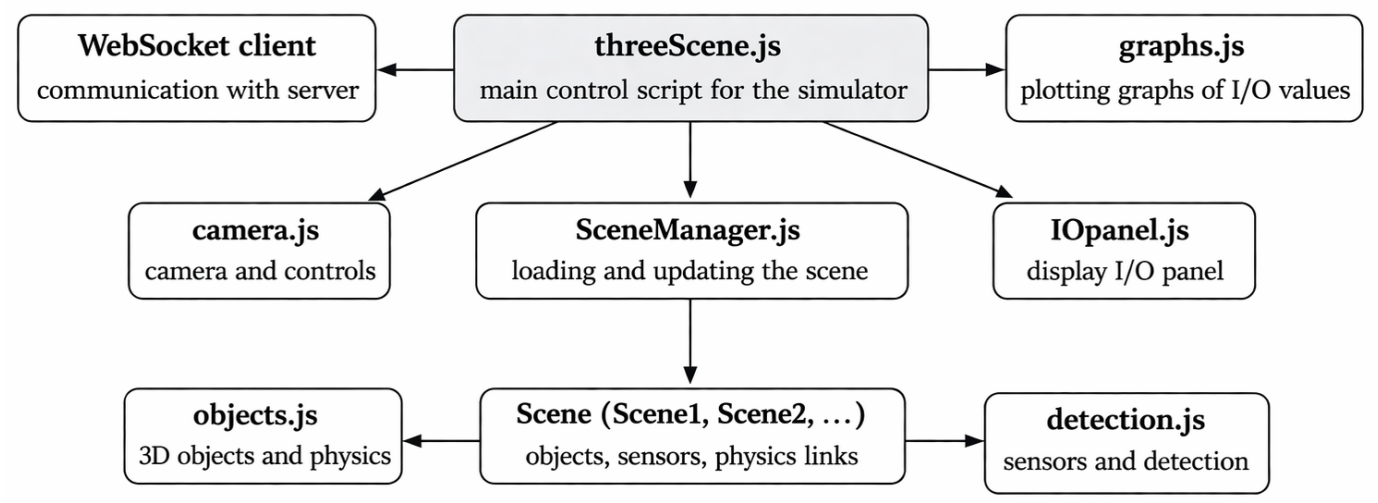

Fig. 2. Software structure of the simulation environment.

VI. Experimental Scenario and Results

To verify the functionality of the developed simulator, a virtual production line for sorting pins according to their thickness was implemented. The simulated system consists of a conveyor transporting cylindrical pins through several sensor stations representing industrial detection points.

The graphical user interface of the simulator is shown in Fig. 3.

The simulator continuously exchanges input and output values with the control logic using the Modbus TCP/IP protocol. These values are processed by the server application and transmitted to the web-based simulation interface through a WebSocket connection, ensuring data synchronization between the control logic and the virtual environment.

In addition to visualizing the technological process, the simulator also monitors several performance parameters including rendering performance expressed in frames per second (FPS), frame rendering time, communication latency between the simulator and the control logic, and communication processing time. The FPS was measured in 400 cycles as value (16,7 +- 0,87).The results depend on the performance of the host PC. In this case, the measurements were carried out on a test virtual machine.

The developed application was tested in a demonstration scenario focused on sorting cylindrical pins according to their thickness. The results confirmed that the simulator is able to visualize the production process in real time and correctly reflect changes of the control signals generated by the softPLC. The measured values of rendering performance and communication delay indicate that the proposed solution is suitable for educational purposes and for verification of simple automation tasks.

Acknowledgment

The completion of this paper was made possible by the grant No. FEKT-S-26-8988 - ”Advanced Methods in Cybernetics, Robotics, and Artificial Intelligence” financially supported by the Internal science fund of Brno University of Technology.

VII. Conclusion

This paper presented the design and implementation of a web-based 3D simulator for industrial automation processes. The simulator enables visualization of a technological process in a virtual environment and allows interaction with control logic implemented in a softPLC.

Communication between the simulation environment and the control system is implemented using the Modbus TCP protocol, while frequent data transfer to the web interface is ensured through a WebSocket connection. The communication period is 2,49 ms with 48,57 jitter.

A virtual production line for sorting pins according to their thickness was implemented as an experimental scenario to verify the functionality of the system. The results demonstrate that the simulator is capable of visualizing the technological process in real time while maintaining stable communication with the control logic according to the virtual commisssioning architecture.

The developed simulator is intended primarily as a tool for teaching PLC programming and basic concepts of industrial automation. It allows students to observe the interaction between control logic and a simulated technological process without the need for physical hardware.

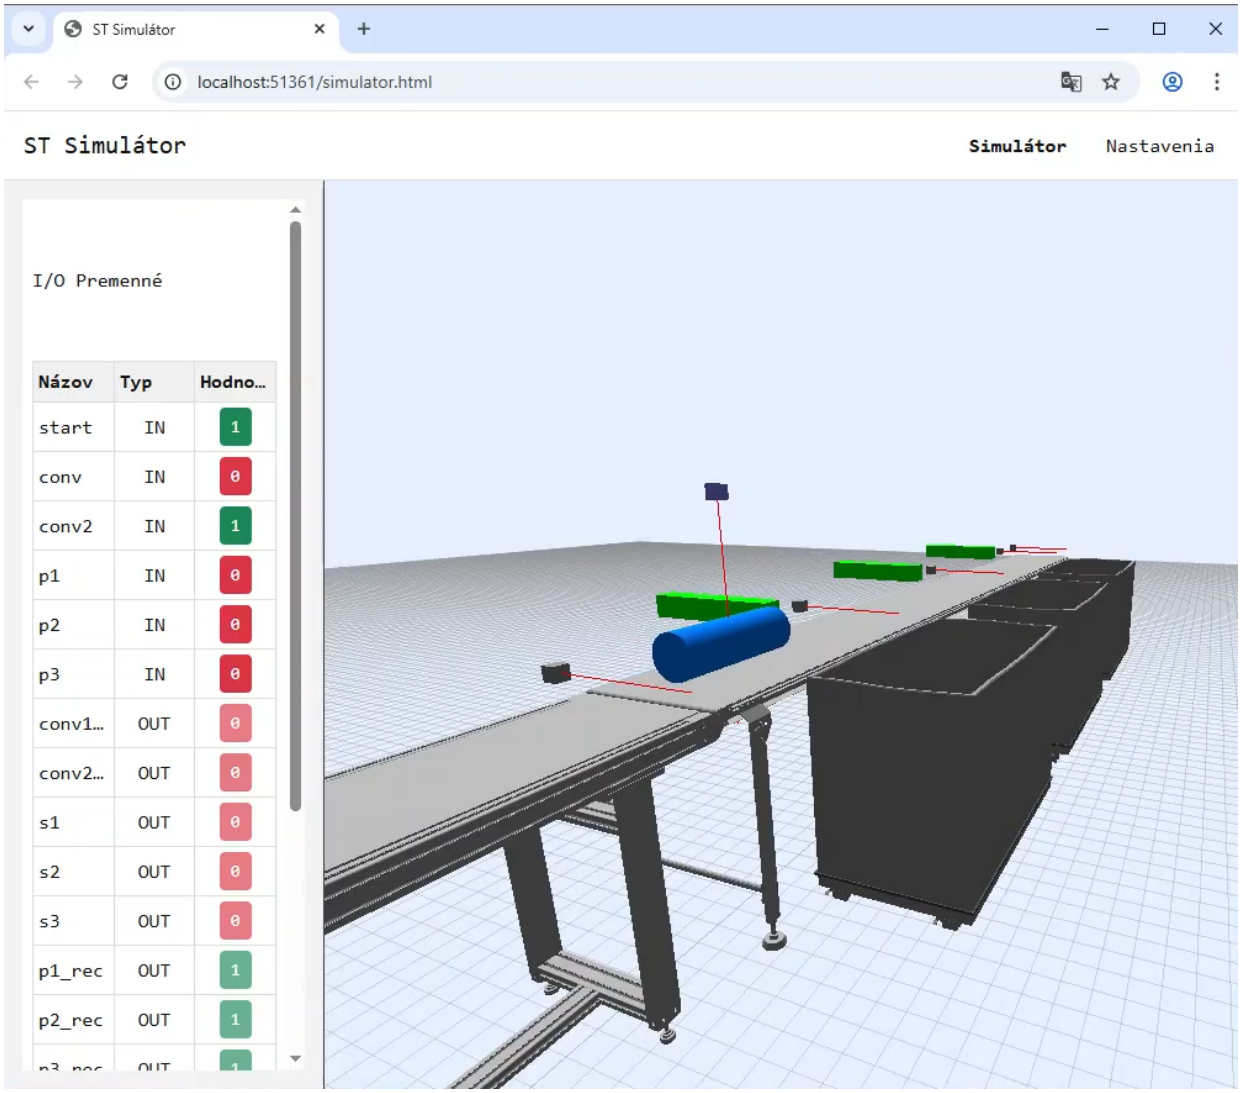

Fig. 3. Main window of the simulator with the 3D scene and I/O panel.

References

[1] STRIFFLER, N.; et al. *Concepts and Trends of Virtual Commissioning*. Computers in Industry. Elsevier, 2023. Dostupné z: https://www.sciencedirect.com/science/article/pii/S0278612523002145. [cit. 2026-03-12].

[2] SCHAMP, M.; HOEDT, S.; CLAEYS, A.; AGHEZZAF, E. H.; COTTYN, J. *Impact of a Virtual Twin on Commissioning Time and Quality*. IFAC PapersOnLine, Volume 51, Issue 11, 2018, s. 1047–1052. Elsevier. Dostupné z: https://www.sciencedirect.com/science/article/pii/S2405896318315970. [cit. 2026-03-12].

[3] QIN, R.; FENG, B.; XU, Z.; ZHOU, Y.; LIU, L.; LI, Y. *Web-based 3D Visualization Framework for Time-varying and Large-volume Oceanic Forecasting Data Using Open-source Technologies*. Environmental Modelling & Software. Elsevier, 2021. Dostupné z: https://www.sciencedirect.com/science/article/pii/S1364815220309658. [cit. 2026-03-12].

[4] SHARMA, A.; MISHRA, M.; GUPTA, S.; MALLIK, A.; SINGH, R. *Digital Twins: State of the Art Theory and Practice*. Sustainable Production and Consumption. Elsevier, 2022. Dostupné z: https://www.sciencedirect.com/science/article/pii/S2452414X22000516. [cit. 2026-03-12].

[5] DIRKSEN, J. a kol. *Three.js — JavaScript 3D Library*. Online. 2024. Dostupné z: https://threejs.org/. [cit. 2026-03-12].

[6] CHANDLERPRALL. *Physijs — Physics plugin for Three.js*. Online. 2023. Dostupné z: https://github.com/chandlerprall/Physijs. [cit. 2026-03-12].

[7] BECKHOFF Automation. *TF6250 TwinCAT 3 Modbus TCP – Product Documentation*. Online. 2024. Dostupné z: https://download.beckhoff.com/download/Document/automation/twincat3/TF6250_TC3_Modbus_TCP_EN.pdf. [cit. 2026-03-12].

[8] CLOUD AUTOMATION. *node-modbus — Modbus TCP/RTU library for Node.js*. Online. 2023. Dostupné z: https://github.com/Cloud-Automation/node-modbus. [cit. 2026-03-12].

[9] HOLOWAYCHUK, T. a kol. *Express — Fast, unopinionated, minimalist web framework for Node.js*. Online. 2024. Dostupné z: https://expressjs.com/. [cit. 2026-03-12].

[10] STANGVIK, E. O. a kol. *ws — Simple to use, blazing fast and thoroughly tested WebSocket client and server for Node.js*. Online. 2024. Dostupné z: https://github.com/websockets/ws. [cit. 2026-03-12].

Prediction of contamination level in a water treatment plant using neural networks

1st Vojtěch Kropáč
Dept. of Control and Instrumentation
Brno University of Technology
Brno, Czech Republic
247406@vutbr.cz

2nd Dominik Říčánek
Dept. of Control and Instrumentation
Brno University of Technology
Brno, Czech Republic
203334@vutbr.cz

***Abstract*—This paper focuses on the development of data-driven soft-sensor models using Multilayer Perceptron and Recurrent Neural Network architectures for the prediction of ammoniacal nitrogen in small wastewater treatment plants operating with Sequential Batch Reactor technology. Utilizing real industrial data, a high-dimensional feature space was optimized through a multi-stage selection process, including XGBoost importance and drop-column analysis. Models were evaluated using metrics such as Mean Squared Error, Mean Absolute Error and the coefficient of Determination. The results demonstrate that the proposed models are capable of reliably predicting the selected variable without the direct use of the corresponding sensor, which represents a cost-effective alternative that may lead to reduced costs associated with the acquisition and maintenance of sensor equipment.**

Keywords—Artificial Neural Network, Wastewater Treatment Plant, Regression, Parameter Prediction, MLP, RNN, Soft-sensor

I. Introduction

Ensuring access to clean water resources is one of the most critical global challenges today. Driven by rapid industrialization, urbanization, and continuous technological advancement, the pollution of natural water bodies has accelerated significantly. To mitigate this issue, effective and financially sustainable wastewater treatment processes are essential. In this context, wastewater treatment plants (WWTPs) represent a critical infrastructure for safeguarding water resources.

The complex technological processes within WWTPs require the continuous monitoring of various physicochemical and biological parameters, such as dissolved oxygen (DO) concentration, pH, oxidation-reduction potential (ORP), as well as the levels of total phosphorus (P) and nitrogenous compounds, specifically ammoniacal nitrogen ($N-NH_4^+$) and nitrates ($N-NO_3^-$)[1]. However, the installation and maintenance of specialized physical sensors to measure these specific substances are both financially and technically demanding. For typical small and medium-sized WWTPs, the acquisition and reliable operation of these systems represent a significant economic burden, which substantially increases the costs associated with the optimal control of biological processes.

To overcome these economic and technical limitations, this paper proposes the development of data-driven soft-sensors utilizing Artificial Neural Networks (ANNs), namely a Multilayer Perceptron (MLP) and a Recurrent Neural Network (RNN) [2],[3]. By leveraging real-time data from standard and cost-effective instrumentation, including flow meters, temperature probes, pH sensors, and aeration controls, the proposed models can accurately estimate the concentrations of hard-to-measure variables, specifically ammoniacal nitrogen ($N-NH_4^+$).

II. Methodology and Data Processing

The prediction models were designed and evaluated using real operational data from the Modřice WWTP. The technological layout of the facility consists of two main components, namely an accumulation tank, which collects and buffers the influent wastewater, and a Sequencing Batch Reactor (SBR) tank, which performs the actual biological treatment and nutrient removal via the activated sludge process. The purification process within the SBR operates in a highly dynamic, cyclical manner. A typical treatment cycle progresses through a sequence of distinct operational phases. These include filling, anaerobic and anoxic mixing, aerobic aeration, stripping, settling, and excess sludge wasting, culminating in the discharge of the treated effluent [4].

A. Data Collection and Exploratory Analysis

The measured dataset consists of signals from several technological sensors with a one-minute sampling period. In total, approximately 24,400 samples were recorded in CSV format over a 22-day period with several temporal gaps, resulting in 18 days of active measurements. Monitored parameters and their corresponding sensor designations are summarized in Table I. Notably, the Oxidation-Reduction Potential (ORP) is monitored by two independent sensors, namely QIC66 and QIC67. This setup is utilized to ensure measurement redundancy and to account for inhomogeneity within the SBR tank.

Prior to modeling, the raw data underwent a preprocessing phase. First, missing values and obvious measurement errors

that could destabilize the neural network training process were detected and removed.

Subsequently, an exploratory data analysis was performed. The properties of individual signals were examined using dynamic statistics and visualizations of time series. The distribution of the measured values for each sensor was evaluated using histograms combined with Q-Q (quantile-quantile) plots. This analysis revealed that the data distribution for most parameters does not correspond to a normal distribution.

TABLE I. OVERVIEW OF MONITORED TECHNOLOGICAL PARAMETERS

Sensor Tag	Parameter	Unit
FIC63	Water Flow Rate	m³/h
QIC64	Dissolved Oxygen Concentration	mg/l
TI65	Temperature	°C
QIC66	Oxidation-Reduction Potential	mV
QIC67	Oxidation-Reduction Potential	mV
QIC68	pH	[-]
FIC69	Air Flow Rate	m³/h
QIC70	Ammonia Nitrogen Concentration ($c_{N-NH_4^+}$)	mg/L
QIC71	Nitrate Nitrogen Concentration ($c_{N-NO_3^-}$)	mg/L
LI72	Water Level	m

For example, the QIC64 sensor, measuring DO, and the QIC70 sensor, measuring $N-NH_4^+$, exhibited strong right-skewness. To suppress this asymmetry, which could cause issues during ANN training, a logarithmic transformation was applied to these sensor outputs.

Furthermore, it was necessary to modify the *SBR_St* signal, which acts as an indicator of the current treatment phase in the SBR tank. Since it is a categorical variable, the One-Hot Encoding method was utilized. This approach successfully prevented any incorrect mathematical interpretation of these states during the neural network training.

B. Feature Engineering and Selection Pipeline

The biological treatment in an SBR is a highly dynamic, time-dependent process. To capture this dynamic behavior, a broad vector of synthetic features was generated, encompassing rolling statistical moments, dynamic rate-of-change indicators, phase-specific transition states, complex physicochemical ratios, and cyclical time encodings. This approach yielded over 1200 potential features.

To reduce dimensionality and eliminate multicollinearity, which could destabilize the gradient-based training of the neural network, the following feature selection methods were sequentially applied:

1) ***XGboost GroupKFold Selection:***

An Extreme Gradient Boosting (XGBoost) algorithm was deployed for the preliminary evaluation of feature importance [5]. The relevance of individual variables was assessed using the Gain metric, which quantifies the average reduction in Mean Squared Error (MSE) during decision tree node splitting. To ensure model robustness a GroupKFold cross-validation was implemented.

2) ***Statistical Filtration:***

The evaluated features were ranked according to their average Gain across all cross-validation folds. To filter out unstable noise and variables with highly fluctuating predictive contributions, a stability criterion was applied to retain only features where the mean Gain exceeded its standard deviation, meaning $Gain_{mean} > Gain_{Std}$. Based on this rule, the top 300 most robust features were selected.

3) ***Correlation Analysis:***

To eliminate redundant information, the remaining features were subjected to a pairwise correlation analysis. Groups of highly correlated features carrying duplicate variance were reduced to a single, most representative element. This step significantly mitigated multicollinearity, which is highly undesirable for the stability of ANNs.

4) ***Drop-Column Importance:***

Following the design of the base ANN architecture, a final feature selection was conducted using the drop-column importance method. This approach determines the true predictive value of each remaining variable by evaluating the degradation in the network's performance when that specific feature is entirely removed from the dataset, yielding the final optimal subset of inputs.

III. MODEL ARCHITECTURE AND TRAINING

The final optimized feature vector was subsequently standardized using Z-score standardization. The transformed dataset was then utilized to train two distinct types of artificial neural networks, namely a Multilayer Perceptron (MLP) and a Recurrent Neural Network (RNN), both implemented within the PyTorch framework.

A. Architecture of Models

. Fig. 1 illustrates the proposed MLP and RNN models. The MLP architecture consists of three hidden layers, an input of 28 features, and dropout layers with a rate of $p = 0.25$. The RNN model features a single hidden layer, an input of 26 features, and a dropout layer with a rate of $p = 0.1$.

B. Training

The Huber loss was selected as the cost function for both models. This specific function was chosen over the standard Mean Squared Error (MSE) due to its inherent robustness to outliers. This characteristic is essential for modeling raw wastewater treatment data, where sensor noise and anomalous spikes are a common occurrence.

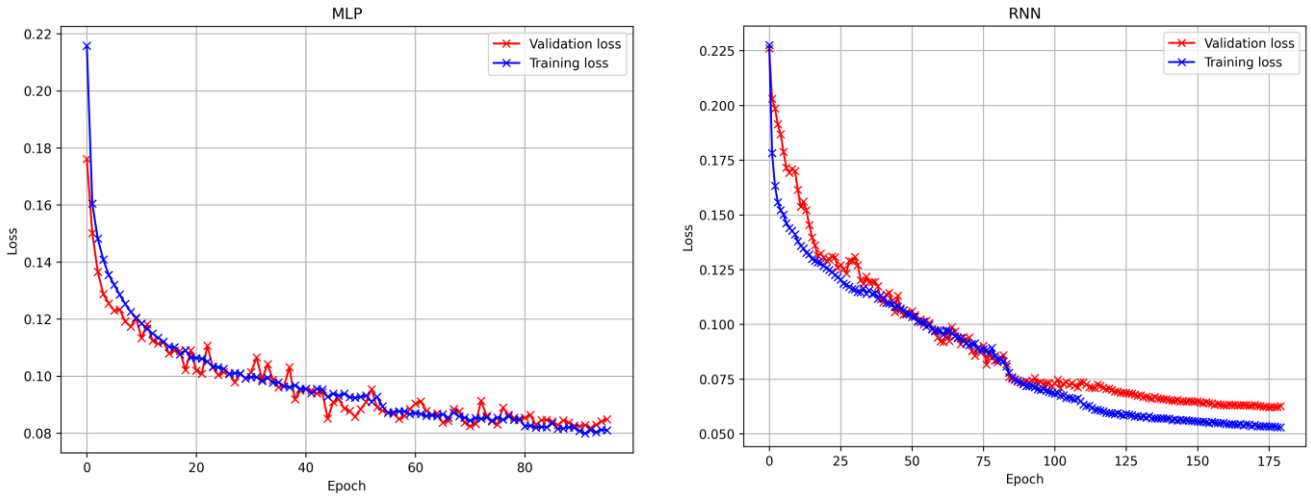

Fig. 2. Training and validation loss curves for proposed MLP (left) and RNN (right) models

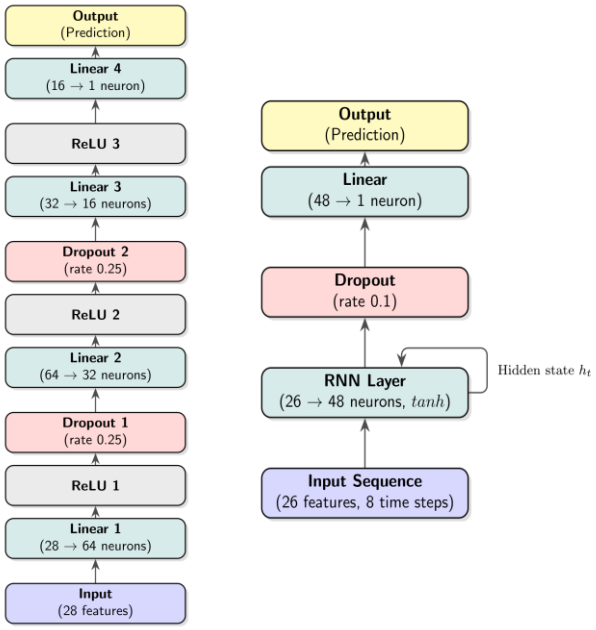

Fig. 1. Architecture of created models: MLP (left), RNN (right)

The Huber loss function is mathematically defined as follows:

$$L_\delta(y,d) = \begin{cases} \frac{1}{2}(y-d)^2, & if \quad |y-d| < \delta \\ \delta \cdot |y-d| - \frac{1}{2}\delta^2, & else \end{cases} \qquad (1)$$

Where y and d represent predicted and desired target values and δ is a hyperparameter that defines the transition point between quadratic and linear error penalization.

Weight optimization for both architectures was performed using algorithms from the Adam family [6]. Specifically, MLP was trained using the AdamW optimizer with an initial learning rate of 0.0015 and an L2 regularization penalty (weight decay) set to 0.001. In contrast**,** the Vanilla RNN utilized the standard Adam optimizer with an initial learning rate of 0.001.

The training progression was continuously monitored using an independent validation dataset. To ensure stable convergence, a dynamic learning rate scheduler ReduceLROnPlateau was employed, scaling down the learning rate whenever the validation loss and the overall training process exhibited stagnation. Furthermore, to strictly prevent overfitting, an Early Stopping mechanism was implemented, halting the training if the validation error failed to improve over a predefined number of epochs [7].

The Early Stopping protocol successfully intercepted the training of both models, concluding the optimization phase at epoch 96 for the MLP and at epoch 180 for the Vanilla RNN, as depicted in Fig. 2.

IV. Results

The predictive performance of both architectures was evaluated using an independent test dataset, with the comparative results summarized in Table II. To quantitatively assess the accuracy of the trained models, standard regression metrics, namely Mean Squared Error (MSE), Mean Absolute Error (MAE) and the Coefficient of Determination (R^2), were employed.

TABLE II. Predictive Performance of the Evaluated ANNs

Model	MSE	MAE	R^2
MLP	1.83	0.82	0.85
RNN	1.69	0.77	0.88

A. Discussion

A detailed analysis of the time-series predictions reveals that the baseline MLP model reliably captures the global macroscopic trend of the biological process. However, as

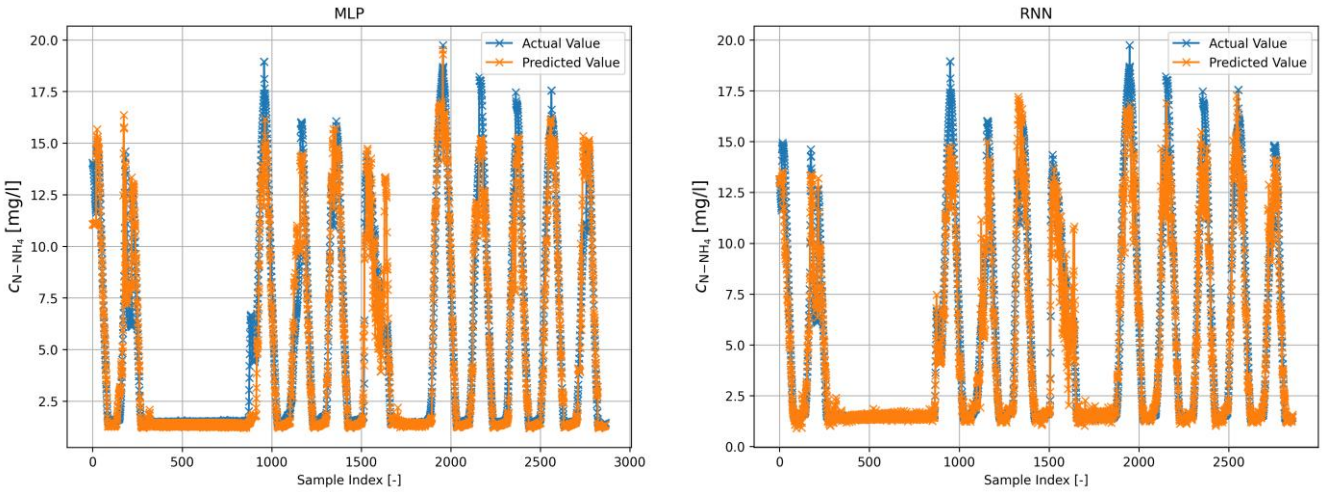

Fig. 3. Comparison of the actual and predicted $c_{NH_4^+}$ for the MLP (left) and RNN (right) models on the testing dataset

illustrated in Fig.3, it often struggles to promptly recognize sudden shifts in the process dynamics, tending to lag behind or partially ignore these rapid transitions. Although it does not always capture the full magnitude of every extreme peak perfectly, it is noticeably more successful at reaching these highest concentration values compared to the RNN.

On the other hand, the RNN demonstrated a visually smoother trajectory and superior alignment with the overall operational dynamics, as depicted in Fig. 3**.** Thanks to its internal memory, the RNN was able to detect and react to sudden trend changes much more effectively than the MLP. It adapted to genuine rapid shifts in the biochemical process while simultaneously smoothing out high-frequency technological noise. The primary trade-off of this inherent temporal smoothing was a dampened response to isolated anomalies. The RNN struggled to reach the highest peak concentrations, underestimating their absolute values more noticeably than the MLP.

These results clearly indicate that while the MLP is more aggressive in tracking isolated extremes, the RNN's capability to capture deeper temporal dependencies provides a significantly more stable and robust overall prediction.

Overall, while both architectures achieve predictive accuracies well within the acceptable standards for industrial time-series forecasting, the RNN proves to be slightly superior in modeling the continuous biochemical process. As summarized in Table II, this advantage is quantitatively reflected by a 7.6% reduction in MSE and a 6.1% reduction in MAE compared to the baseline MLP model.

V. Conclusion

This paper presented the development and evaluation of data-driven soft-sensor models for predicting ammoniacal nitrogen ($N-NH_4^+$) concentrations in a small-scale wastewater treatment plant operating with SBR technology. Based on real operational data, two neural network architectures, MLP and RNN were designed, optimized, and tested.

The results demonstrated that both models were capable of providing reliable predictions. The RNN architecture slightly outperformed the baseline MLP.

In conclusion, the proposed soft-sensors represent a capable and cost-effective alternative to traditional hardware sensors. By reliably estimating crucial technological parameters without the need for direct physical measurement, this approach can significantly reduce the operational expenses associated with sensor maintenance, cleaning, and frequent calibration.

Future research will focus on extending this data-driven approach to predict other critical wastewater parameters, specifically concentration of nitrates and DO, to further reduce the reliance on physical hardware sensors.

Acknowledgment

The completion of this paper was made possible by the grant No. FEKT-S-26-8988 - "Advanced Methods in Cybernetics, Robotics, and Artificial Intelligence" financially supported by the Internal science fund of Brno University of Technology.

References

[1] HLAVÍNEK, Petr; MIČÍN, Jan a PRAX, Petr. *Příručka stokování a čištění*. Brno: NOEL 2000, 2001aa

[2] NIELSEN, M. A. Neural Networks and Deep Learning. Determination Press, 2015

[3] MIENYE,I. D.; SWART,T.G.;OBAIDO,G.RecurrentNeuralNetworks: A Comprehensive Review of Architectures, Variants, and Applications.. In: Information, 2024, vol. 15, no. 9, art. 517.

[4] P. A. Wilderer, R. L. Irvine, and M. C. Goronszy, *Sequencing Batch Reactor Technology*. London, U.K.: IWA Publishing, 2001.

[5] T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, San Francisco, CA, USA, 2016, pp. 785–794

[6] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in *3rd International Conference on Learning Representations (ICLR)*, San Diego, CA, USA, 2015

[7] L. Prechelt, "Early Stopping — But When?," in *Neural Networks: Tricks of the Trade*, G. B. Orr and K.-R. Müller, Eds. Berlin, Heidelberg: Springer, 1998, pp. 53–67

HÝBEJ SVĚTEM S ABB

Zajímá tě, jak se řídí elektrárny a rozvodny?
Jak se chrání a stabilizují elektrické sítě?
Nebo jak se automatizují výrobní provozy a zefektivňuje spotřeba energie v budovách i průmyslu?

V ABB na tom v Česku pracujeme každý den – a **hledáme lidi jako ty.**
Budeš pracovat na reálných projektech – vyvíjet software, testovat zařízení, optimalizovat procesy.

A když se osvědčíš, **můžeš vyrazit na stáž do zahraničí** – Španělsko, Německo, Norsko a další.

Kontaktujte naši náborářku:

Petra Tůmová
petra.tumova@cz.abb.com
+420 733 693 337

Více informací

TOP ZAMĚSTNAVATELÉ 3. místo AUTOMOBILOVÝ & STROJÍRENSKÝ PRŮMYSL 2026

Buďte s námi o krok napřed.



Litemate: An Assistive Tool for Video Image Color Correction

Dominik Chalupka
Brno University of Technology
FEEC, Department of Radio Electronics
Technicka 12, 616 00 Brno, Czech Republic
253225@vutbr.cz

Ladislav Polak
Brno University of Technology
FEEC, Department of Radio Electronics
Technicka 12, 616 00 Brno, Czech Republic
polakl@vut.cz

***Abstract*—*Litemate* is a tool for the post-production software *DaVinci Resolve Studio* that assists in detecting exposure and white balance (WB) shifts and evaluating the necessary corrections. The tool is implemented in the form of a script. The objective of this paper is to investigate and select an appropriate analysis method, design the processing workflow, and develop techniques for comparing and processing the acquired data in order to determine the required corrections. Subsequently, the correction assistance is implemented within the *Resolve* application. The design concept of the tool and the functionality of its individual scripts are presented. Relevant colorimetry techniques and theoretical background are also summarized. The exposure and WB analysis and correction results of the tool prototype are tested and briefly evaluated.**

***Index Terms*—Video post-production, color grading, *DaVinci Resolve Studio*, histogram analysis, exposure, colorimetry**

I. Introduction

One of the key stages of digital image post-processing is color correction and grading. These processes allow the transformation of the color and lighting characteristics of a digital image, enhancing its artistic qualities while also correcting various errors introduced during recording. However, this process can be complex, and if not performed properly, it may result in an over-processed image. Color correction and grading can significantly influence the final visual quality of the image [1], [2], [3]. As the need for video increases, so does the need for processing and accelerating production. Therefore, any automated aid in post-production is welcomed. In the case of long video productions, however, only a limited number of tools are available that dynamically assist with changing color conditions throughout the footage. This limitation motivated the development of the tool presented in this paper. To address this issue, a script plugin for *DaVinci Resolve Studio*[1] [4] was developed. This software is an industry standard for color correction and video editing. The primary goal of the plugin is to accelerate the color correction process [5].

The remainder of this paper is organized as follows. Section II provides a brief overview of colorimetry, color models, and correlated color temperature (CCT). Section III describes the implemented techniques for exposure and white balance (WB) analysis using histograms. Next, the correction methods based on the primary operators (*lift, gamma, gain*) and their implementation are introduced. Section IV and Section V describe the overall functionality of the tool, the approach used for its development, and the results, respectively. The paper concludes with Section VI.

This work was supported by the BUT project no. FEKT-S-26-8967.

[1]https://www.blackmagicdesign.com/products/davinciresolve/training

II. Theory Background of Colorimetry

The International Commission on Illumination adopted a set of color matching functions known as the CIE 1931 color space, or CIE XYZ, which define a chromaticity diagram that describes the relationship between wavelengths of the visible spectrum [6], [7]. This diagram is labeled as the CIE *x,y* chromaticity diagram and serves as the basis for defining color spaces, as illustrated in Fig. 1.

The chromaticity temperature of light describes the color appearance of white light and corresponds to the color emitted by an ideal black-body radiator at a given temperature expressed in Kelvins [6]. This specification is known as the CCT. In the CIE *x,y* diagram, the relationship between chromaticity and temperature is represented by a curve known as the Planck's curve or CCT curve, shown in Fig. 2.

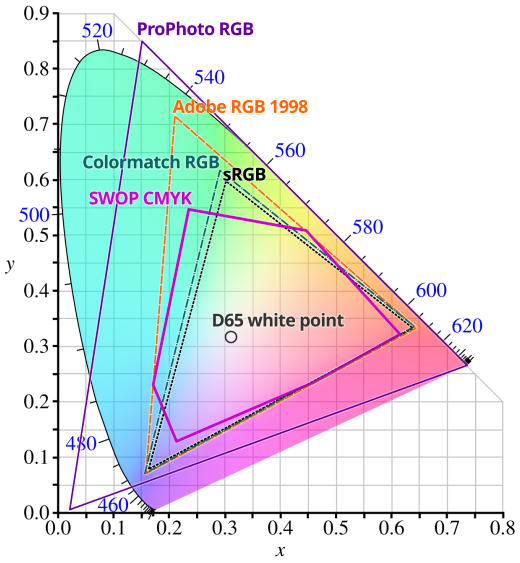

Fig. 1. Chromaticity diagram CIE *x,y* [6]

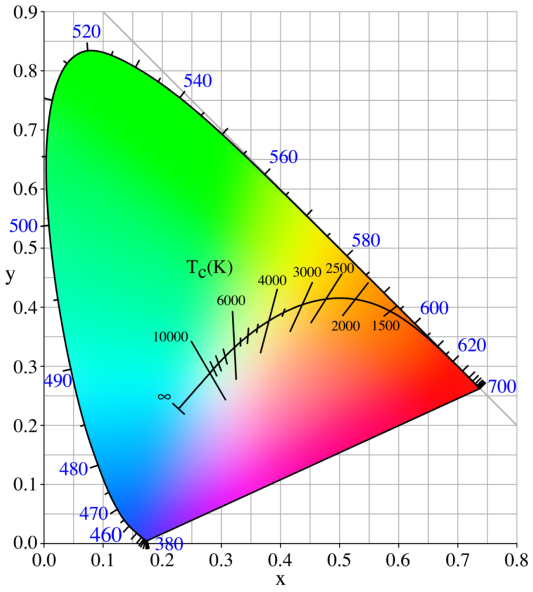

Fig. 2. Color temperature curve in CIE *x,y* diagram [7]

A black body is a theoretical object that absorbs all incident electromagnetic radiation and, depending on its temperature, emits electromagnetic radiation according to Planck's law. At lower temperatures, the emitted light appears warmer, exhibiting more orange tones. As CCT values increase, the light appears cooler with a stronger blue component [8], [9].

III. Analysis and Correction of Image Changes

A. Histogram analysis

A histogram is a type of graph that represents the distribution of pixel intensity values within a digital image. The *x-axis* contains all possible intensity values, while the *y-axis* shows the number of pixels corresponding to each intensity. This quantity may be expressed either as an absolute value or relative to the total number of color subpixels in the frame. The histogram can be described as follows [10]

$$h(g) = \#pix(g) \tag{1}$$

where g denotes the pixel intensity and *#pix* is the number of pixels with that intensity value. The range of intensity values is determined by the bit depth, starting at 0.

Since a digital image consists of RGB components, obtaining a pure luminance distribution for clearer visualization of a frame's lighting characteristics requires computing a luminance value from the RGB components. This can be expressed as [6], [7]

$$L = RL_R + GL_G + BL_B \tag{2}$$

where L denotes the luminance value, R, G, and B represent the values of the individual color components, and L_R, L_G, *and* L_B are the luminance coefficients of these components defined by the used color space as shown in Table I.

TABLE I: Standard luminance coefficients [7]

Format standard	Color space	Use case	Coefficients		
			L_R	L_G	L_B
ITU-R BT.709	Rec.709	TV	0.2126	0.7152	0.0722
ITU-R BT.601	SDTV	TV	0.299	0.587	0.114
ITU-R BT.2020	UHDTV	TV	0.2627	0.6780	0.0593
SMPTE RP431-2	DCI-P3	Film	0.2627	0.6780	0.0593
IEC 61966-2-1	sRGB	Graphic	0.2126	0.7152	0.0722

Note that luminance is labeled L when in relation to color spaces, rather than Y [3]. These coefficients reflect the characteristics and boundaries of the color space to ensure perceptually appropriate luminance levels when displayed. The constants are derived from calculations based on the CIE tristimulus values [6], [7].

B. Operators lift, gamma, gain

One of the most basic methods of luminance manipulation involves the operators *lift*, *gamma*, and *gain*. [2] Each of these operators affects the entire luminance range, not just a portion. The amount of influence is then dictated by the shape of the function [2], [7].

The *lift* operator primarily affects dark regions of the image. Its influence decreases linearly as the luminance value increases. The *gamma* operator mainly affects mid-tone values. Its influence gradually decreases toward both the highlights and the shadows. The *gain* operator acts in the opposite manner to *lift*. It primarily affects the highlights, while the shadows remain barely unchanged. Similar to the *lift* operator, its function is linear.

Although each operator has its own function and equation describing how it modifies the histogram [2], these primary controls can be unified into a single function that simultaneously manipulates all luminance regions. This allows the resulting transformation to be represented as a single curve. The unified primaries function [6], [7] is given by

$$Y_c = ((\frac{Y_a}{2^{n_{bits}}} + lift(1 - \frac{Y_a}{2^{n_{bits}}}))gain)^{\frac{1}{gamma}} * 2^{n_{bits}} \tag{3}$$

where Y_c is the output value, Y_a is the input value, *lift*, *gamma*, and *gain* are the corresponding operator parameters, and n_{bits} denotes the bit depth.

C. White balance correction

Based on the concept of chromaticity temperature [6], [7], white balancing can be interpreted as selecting a white point on the *Planck curve*. In digital video processing, this corresponds to adjusting the RGB component values so that objects that are naturally white appear visually neutral under the given lighting conditions. This adjustment simultaneously determines the overall color balance of the frame, as the color cast affects the entire image [1], [2]. For white balance correction, a method similar to the *Temp* and *Tint* controls in *DaVinci Resolve Studio* is applied [4]. These controls modify the image by applying gain adjustments to the RGB components in a way that approximates the trajectory of the *CCT curve*.

IV. Development of the Tool

The purpose of the developed tool (also called plugin) to automatically detect changes in exposure and white balance within a video, mark the time at which these changes occur, evaluate the values required for correction, and provide the colorist with a timeline in *DaVinci Resolve Studio* containing markers with recommended correction values. The plugin is primarily intended for long-form videos where visual consistency is expected. It focuses on detecting and correcting occasional undesired variations. Typical use cases include talking-head videos or long recordings of theater performances.

The tool operates as follows: within *Resolve*, the user selects a video and initializes a communication script that prepares the timeline and the necessary data for the plugin to perform the analysis. The user is then prompted to specify parameters:

- *Reference frame* defines a point in time (or equivalently a frame number) corresponding to a frame with the desired exposure and white balance values that should be maintained throughout the video.
- *Transition length* specifies the duration of the interval over which the correction is gradually applied. The purpose of this parameter is to smooth the correction and make the transition visually more natural for the viewer.
- *Tolerance setting* defines the acceptable range of variation. If the detected change in exposure or white balance (WB) exceeds this threshold, a correction is triggered. Larger tolerance values allow greater variations before a correction is applied.
- *Correction choice* allows the user to select which type of correction should be applied: exposure, WB, or both.

After these parameters are defined, the analysis can be executed. Once the detection process is complete, the user is presented with the results in the form of graphs showing the detected changes, together with the corresponding image histograms used in the analysis. If the user is satisfied with the results, the detected changes and recommended corrections can be imported into *Resolve* as timeline markers.

Python was selected as the programming language for the development of this tool. The main reason for this choice is its native support in *DaVinci Resolve* [4] as well as its widespread use in image processing applications. In addition, a large number of resources and libraries are available that facilitate development. Visual Studio Code is used as a development environment supporting Python and the installation of various extensions. The environment is also widely used and provides an intuitive and user-friendly interface. GitHub is used for version control and storage of the source files [11].

V. Tests and Results Evaluation

For testing purposes, BRAW recordings captured with the *Blackmagic Pyxis* digital cinema camera were used. The recordings were transcoded from the raw format into the DNxHR HQX 10-bit codec, which is fully supported by the PyAV library used in the implementation. The test video contains professional exposure and sharpness charts as well as a *Calibrite ColorChecker Classic* table.

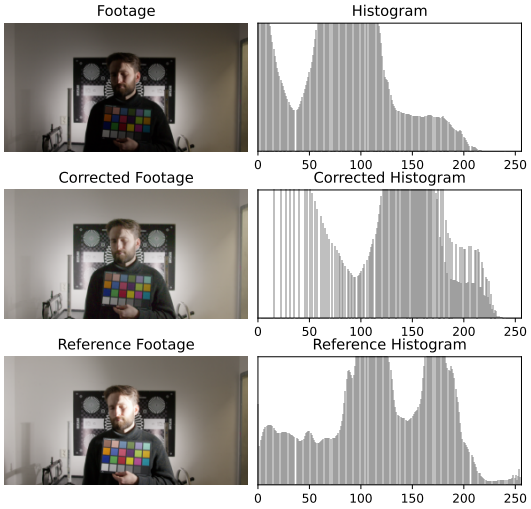

Fig. 3. Exposure analysis and correction using median value

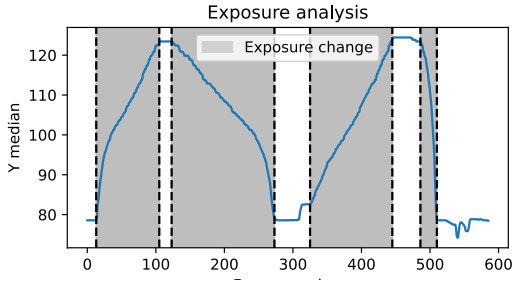

Fig. 4. Exposure analysis in time

A. Results of exposure analysis

Figure 3 presents an example of exposure correction. The left column shows the frames, while the right column shows the corresponding histograms. The first row contains the analyzed frame, the second row the corrected frame, and the third row the reference frame. It is visible that the reference frame is significantly brighter than the analyzed frame. The correction value was obtained from the difference between the mean and median values. After its application, an improvement can be observed both in the image and in the histogram of the corrected frame. Some changes in the appearance of the subject are caused by lighting conditions, and not all such variations can be fully corrected in post-production [1], [9].

The analysis is capable of detecting the onset and stabilization of exposure changes over time and marking them accordingly. It is necessary to fine-tune the number of consecutive frames that must remain within the threshold to consider the exposure stable. After this, the results proved to be reliable. The detected start and end of an exposure change are marked in *DaVinci Resolve* (blue markers indicate the start of a change, red markers indicate stabilization). Near frame 300 and after frame 500, a small fluctuation can be observed (see Fig. 4). Since it did not exceed the defined thresholds, this behavior represents a correct evaluation within the defined tolerance.

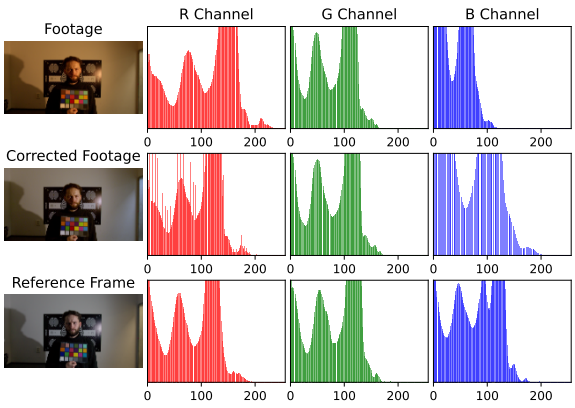

Fig. 5. WB analysis and correction using median value

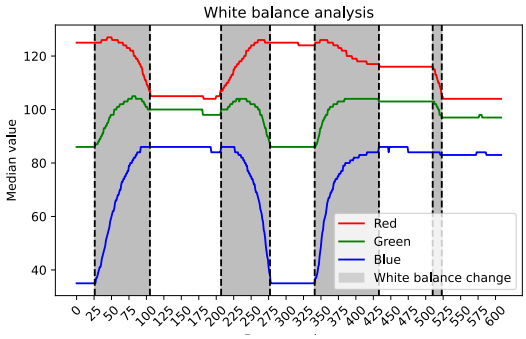

Fig. 6. WB analysis in time

B. Results of WB analysis

To obtain correction values for the gain operators of the individual RGB channels, the ratio of the histogram medians [3] was used M_r/M_a, where M_a denotes the median of the analyzed frame and M_r the median of the reference frame. Adaptation of the chromatic transformation is not implemented in this prototype for the sake of simplicity. Figure 5 shows an example of WB analysis and correction following the same structure as the exposure analysis (first column – frames, second to fourth columns – channel histograms; first row – analyzed frame, second row – corrected frame, third row – reference frame). The analysis was again performed using median and mean values of the histograms.

When evaluating the results, the median value proved to be slightly preferable to the mean value due to its less "aggressive" correction behavior. Overall, the analysis performed as expected and was able to reduce the differences in white balance across the frames. It should be noted that the changes present in the test footage are intentionally extreme, and their visual impact on the subjects could not be fully corrected by white balance adjustment alone [1], [9]. The graph shown in Fig. 6 illustrates the successful detection of changes over time and the corresponding markers, which were again imported into *DaVinci Resolve*.

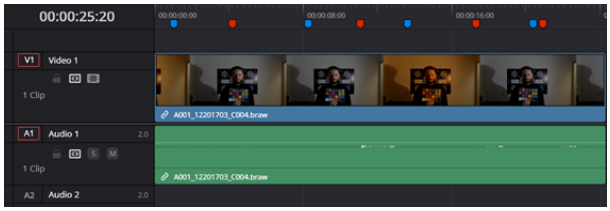

Fig. 7. Marker placement in DaVinci Resolve

VI. Conclusion

The goal of this work was to design a prototype tool that assists a colorist in detecting and correcting unexpected changes in exposure and white balance in video, while demonstrating the functionality of the proposed design and identifying its limitations. The tool operates as a script extension for the post-production software *DaVinci Resolve Studio*. It is implemented in Python 3.6 (for the *DaVinci Resolve API*) and Python 3.12 (for the analysis component). The exposure correction functions, based on primary operators, were tested and normalized to ensure reliable results. The implemented WB correction function showed similarly reliable results. The analysis was evaluated using mean and median histogram intensity values.

The analysis currently operates slowly and requires further optimization. Future work will focus on improving performance through better hardware utilization, particularly GPU acceleration, and a more efficient d ependency b uild process. A limitation of the current implementation is the lack of API support for the color page. As a result, the tool assists the colorist by placing timeline markers with recommended corrections (see Fig. 7) rather than applying them automatically. Despite this limitation, the tool still significantly reduces the time required for an otherwise manual task.

References

[1] B. Brown. *The filmmaker's guide to digital imaging: for cinematographers, digital imaging technicians, and camera assistants*. New York: Focal Press, 2015. ISBN 978-0-415-85411-5.

[2] S. Hullfish. *The art and technique of digital color correction*. Amsterdam: Focal Press/Elsevier, 2012. ISBN 978-0240817156.

[3] E. J. Giorgianni, T. E. Madden. *Digital color management: encoding solutions*. 2nd ed. Wiley-IS&T series in imaging science and technology. Chichester, U.K.: J. Wiley, 2008. ISBN 978-0470512449.

[4] C. Roberts, S. Hall. *The Beginner's Guide to DaVinci Resolve 18*. Blackmagic Design Learning Series, 9.3.2023. ISBN-13 979-8987267110.

[5] J. E. Solem. *Programming computer vision with Python*. Sebastopol: O'Reilly, 2012. ISBN 978-1449316549.

[6] N. Ōta, A. R. Robertson. *Colorimetry: fundamentals and applications*. Wiley-IS&T series in imaging science and technology. Chichester, England: John Wiley, 2006. ISBN 978-0-470-09472-3.

[7] C. A. Poynton. *Digital video and HD: algorithms and interfaces. 2nd ed*. Amsterdam: Morgan Kaufmann, an imprint of Elsevier, c2012. ISBN 978-0123919267.

[8] A. Siess. [Data set]. *Kelvin (color temperature) to RGB conversion table. Zenodo*. [cit. 2026-03-08]. Available from: https://zenodo.org/records/14230959.

[9] H. C. Box. *Set Lighting Technician's Handbook: Film Lighting Equipment, Practice, and Electrical Distribution. 5th ed*., Focal Press, 2014. ISBN 978-0240515236.

[10] R. C. Gonzales, R. E. Woods *Digital Image Processing, 4th global ed*., New York: Pearson, 2018. ISBN 978-1-292-22304-9.

[11] D. Chalupka. *Litemate* [software repository] [cit. 2026-03-08] Available from: https://github.com/dominik-juze-haus/litemate

A Compact RF Power Meter with 40 dB Dynamic Range from DC to 4 GHz

Jakub Šťastný
Department of Control and Instrumentation
Brno University of Technology
Brno, Czech Republic
267152@vutbr.cz

Karel Juřík
Department of Theoretical and Experimental Electrical Engineering
Brno University of Technology
Brno, Czech Republic
karel.jurik@vut.cz

***Abstract* — This paper presents the design of a compact RF power meter providing a 40 dB dynamic range from DC to 4 GHz. The device utilizes an ADL5904 power detector and a digital step attenuator to achieve accurate measurements across a broad range of frequencies. The dynamic range of 40 dB was experimentally verified through measurements of signals from –30 to 10 dBm. Evaluation shows high accuracy for known frequencies, while results for unknown frequencies remain satisfactory for providing power estimates.**

Keywords* — *RF power measurement, ADL5904, power detector, microcontroller, STM32, digital attenuator, RMS power

I. Introduction

Measurement of an electrical field power is essential across a wide range of applications in electrical engineering and physics, from the design of radio frequency equipment to plasma research.

This paper presents the design of a compact, wideband device capable of measuring the power of an electrical field over a broad frequency range. The device targets measurements from 10 MHz to 2.6 GHz, while the chosen components theoretically allow operation from DC to 4 GHz.

The main contribution of this paper is demonstrating how a relatively simple, low-cost design can achieve accurate power measurements. Further suggestions to improve accuracy of the next iteration of the device are also mentioned.

The remainder of the paper is structured as follows: section II briefly describes the hardware and the software design of the device; section III presents and evaluates the measurement results; section IV concludes the paper and suggests improvements for the next iteration of the device.

II. Device Implementation

A. Hardware description

As described in Fig. 1, the device consists of separate components, which communicate with each other. The whole device is controlled by a single microcontroller, which runs real-time operating system (RTOS) to handle concurrency.

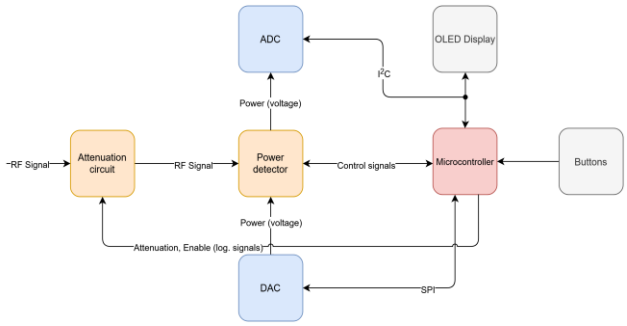

Fig. 1. System diagram of the device

1) Power detector

In the case of this application, power detectors are circuits, which receive high frequency signals as an input and output its power as a dc voltage. Naturally, this implies that when using such a circuit, one must convert the output dc voltage into power units through means of some transfer characteristics.

For our purposes, we have decided to use the ADL5904, manufactured by Analog Devices. “The ADL5904 is a dual-function radio frequency (RF) TruPwr™ detector that operates from dc to 6 GHz. It provides rms power measurement along with a programmable envelope threshold detection function” [1]. Its RF input range is -30dBm to 15dBm. For this range, the transfer characteristic is almost linear, as shown in Fig. 2.

However, it also depends on frequency. For precise measurements, the frequency must be known beforehand or detected from the signal. For simplicity, no automatic frequency detection method was included.

To use the threshold detection functionality for a given power level, a voltage must be applied to the VIN- pin of the detector. This voltage must be determined from a separate characteristic [1]. Although the threshold feature is planned to

be implemented programmatically, the design still includes components to allow potential hardware implementation.

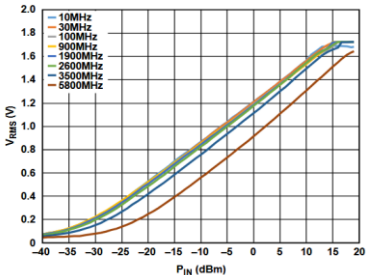

Fig. 2. V_{RMS} vs. Input Level (P_{IN}) for Various Frequencies [1]

2) Attenuation circuit

To increase the dynamic range of the device, the design incorporates the DAT-31A+ digital step attenuator manufactured by Mini-Circuits. It provides adjustable attenuation from 0 to 31 dB in 1dB increments. It is controlled via a 5-bit parallel control interface. Its frequency range is dc to 4 GHz [2]. Therefore, it lowers the frequency limit from 6 GHz to 4 GHz, which still meets our requirements.

3) Microcontroller

To control the device, a NUCLEO-F303K8 development board by STMicroelectronics was selected. It integrates the STM32F303K8 microcontroller and supporting circuitry that simplifies the process of programming and debugging. The microcontroller itself features 64 kB of Flash memory, 12 kB of static RAM, several communication interfaces (CAN, I^2C, and SPI), and many other peripherals [3], making it more than powerful enough for the intended application.

4) Analog-to-Digital and Digital-to-Analog converters

The microcontroller itself contains integrated ADC and DAC peripherals. However, the internal ADC provides a maximum resolution of only 12 bits, with a reference voltage of 3.3 V. To achieve a measurement resolution of at least 0.01 dB, higher precision is required. With a detector slope of approximately 36 mV/dB, as specified in [1], 0.01 dB corresponds to a voltage change of about 0.36 mV. Given an input voltage range of up to 1.5 V, a minimum of 14 bits is required, leaving no margin for noise or other non-ideal effects. Therefore, an external ADC, the MCP3422 manufactured by Microchip Technology, is used. It features two differential input channels with configurable resolution of up to 18 bits and communicates via I^2C [4]. To safely meet the accuracy requirements, the ADC is configured for 16-bit resolution.

To utilize the threshold functionality, the hardware must also be capable of applying a voltage to the VIN- pin of the detector, which can be achieved using a DAC. Although the internal 12-bit DAC would be sufficient, an external DAC is used to allow easy replacement with a higher-resolution device if higher precision becomes desirable. For this purpose, the MCP4822 by Microchip was selected. The device contains two output channels and uses SPI as a communication interface [5].

5) Display and buttons

Measurement results are displayed on a small 128 × 32 OLED module with an SSD1306 driver, which communicates via I^2C. The device is controlled using four membrane buttons.

6) Power supply

The device operates from a 3.3 V power supply, with the worst-case current consumption estimated to be approximately 100 mA. To meet these requirements, the NCP718ASN330T1G fixed low-dropout voltage regulator manufactured by onsemi was selected. It can provide up to 300 mA of output current, with an input voltage range from 2.5 V to 24 V. It also integrates several protection features, such as thermal shutdown and current limiting [6]. A standard DC power connector is included to allow the device to be powered from commonly available consumer power adapters.

B. Schematic and printed circuit board design

A detailed description of the schematic and PCB design would exceed the scope of this paper. Instead, only the crucial design considerations are discussed. The schematic and PCB files are available in [7].

The finished PCB can be seen in Fig. 3. During the design process, manufacturing limitations such as the overall dimensions and the use of a maximum of two layers were imposed.

The bottom layer functions as a ground plane. Its integrity is particularly maintained beneath the RF trace, providing a low-inductance path for high-frequency return currents [8]. The RF path was implemented as a microstrip with a controlled impedance of 50 Ω to ensure proper impedance matching. The PCB thickness was reduced from the typical 1.6 mm to 0.8 mm, which reduces the required width of the microstrip. The resultant PCB has a width of 80 mm and a height of 75 mm.



Fig. 3. Finished PCB

C. Software

1) Device modes

The device operates as a finite-state system, consisting of four operating modes. These are:

- **Measurement mode**, which is used to enable or disable the measurement process. When measurement is enabled, the measured value is displayed in all modes except calibration mode. The measurement can be toggled using the two horizontally placed buttons.

- **Threshold mode,** in which the threshold value can be set in the range from -29 dBm to 14 dBm using the horizontally placed buttons. The value can be incremented or decremented in 0.5 dBm steps. When the measured power level reaches the threshold value, the threshold LED is illuminated.
- **Attenuation mode,** during which the attenuation level can be incremented or decremented using the horizontal buttons. The attenuation range is 0 dB to 31 dB in 1 dB steps.
- **Calibration mode**, which displays the voltage measured by the ADC. This information is used during device calibration.

This behavior is internally modelled as a finite-state machine, with four states. A state change is requested using the "next" and "previous" commands. A diagrammatic representation is shown in Fig. 4.

More specifically, the state machine is implemented as a structure, with transitions to the next or previous state performed using dedicated functions.

The state machine alone is insufficient to capture all the necessary data in the device. Therefore, an additional structure is used to store the device information, primarily including the current attenuation level, the current threshold value, and the currently measured power.

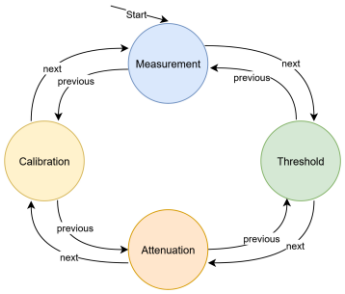

Fig. 4. Representation of the device as a finite-state machine

2) Tasks

The device must process user input, display information, and measure the voltage at the output of the power detector. This behavior is implemented using three concurrent tasks synchronized with a queue:

- **Input processing task** periodically polls for button presses, which are debounced using a simple algorithm. Based on the detected button press, it places a corresponding event into the queue.
- **Measurement task** periodically reads the voltage measured by the ADC. Using calibration data, it converts the measured voltage to dBm and places an event into the queue to refresh the displayed value. This task can be suspended by the **main task**.
- **Main task** remains blocked by the queue until an event is received. It then determines whether to update the state machine, adjust threshold or attenuation values, suspend or resume the measurement task, or refresh the display.

Concurrency in the system is handled by the CMSIS-RTOS2. The CMSIS (Common Microcontroller Software Interface Standard) is a set of application programming interfaces designed to standardize software interfaces across different ARM Cortex-based devices [9]. CMSIS-RTOS2 specifies a generic interface for RTOS kernels [10]. In this application, FreeRTOS is used as the underlying kernel.

3) Calibration

The device can be calibrated by segmenting the transfer characteristic. First, the voltages at the output of the detector are measured over a range of power levels and frequencies, producing discrete points on the transfer characteristic. Each pair of neighboring points defines a segment, which is represented by a linear function.

All points are stored in the microcontroller's data section of the firmware. Upon device startup, the slope and intercept are calculated for each segment and stored in the microcontroller's memory.

III. Measurements

The goal of the measurement was to compare the power measured by the device with the reference power values.

A. Measurement setup

The device's input was connected via an RF coaxial cable to a ROHDE & SCHWARZ SMC100A signal generator, which was configured to produce a harmonic signal at a specific power level and frequency. The generator's output power setting was used as a reference for comparison. This setup is illustrated in Fig. 5.

Nine equally spaced calibration points were measured in the range from −30 dBm to 10 dBm for five different frequencies. Signals at the same frequencies were subsequently used for the measurements. The list of the tested frequencies is shown in TABLE I.



Fig. 5. Photograph of the measurement setup

The device was configured to expect an input signal with a frequency of 760.7 MHz (f_3). For signals with frequencies

different from this value, this setup simulates a scenario in which the signal frequency is unknown beforehand and a transfer characteristic corresponding to an arbitrary calibration frequency is used.

B. Measurement results

TABLE II contains the reference power value, denoted as P_r and the deviation from the reference value, denoted as Δ. The deviation was calculated with (1), where P_m denotes the value measured by the device.

$$\Delta = P_m - P_r \quad (1)$$

Table III presents the maximum error (MAX), root mean square error (RMSE) and mean absolute error (MAE) calculated for each frequency individually as well as for all frequencies combined.

TABLE I. FREQUENCIES USED FOR CALIBRATION AN MEASUREMENT

f_1	f_2	f_3	f_4	f_5
160.7 MHz	460.7 MHz	760.7 MHz	1.061 GHz	1.361 GHz

TABLE II. DEVIATIONS FROM REFERENCE POWER VALUE FOR EACH FREQUENCY

	f_1	f_2	f_3	f_4	f_5
P_r (dBm)	**Δ (dB)**	**Δ (dB)**	**Δ (dB)**	**Δ (dB)**	**Δ (dB)**
-29.54	0.51	0.48	-0.03	-0.23	-0.30
-25.42	0.42	0.16	-0.04	-0.21	-0.27
-18.68	0.42	0.19	0.00	-0.15	-0.20
-14.21	0.42	0.23	0.02	-0.13	-0.15
-7.65	0.35	0.21	0.00	-0.12	-0.09
-5.47	0.35	0.20	0.00	-0.11	-0.06
0.12	0.34	0.13	-0.01	-0.06	0.07
6.61	0.04	-0.30	-0.04	-0.44	-0.29
9.14	0.01	-0.27	-0.08	-0.45	-0.29

TABLE III. MAX, MAE AND RMSE VALUES CALCULATED FOR ALL FREQUENCIES SEPARATELY AND COMBINED

	f_1	f_2	f_3	f_4	f_5	$f_{1\text{-}5}$
MAX (dB)	0.51	0.48	0.08	0.45	0.30	0.51
MAE (dB)	0.32	0.24	0.02	0.21	0.19	0.20
RMSE (dB)	0.36	0.26	0.03	0.25	0.21	0.25

The dependence of the deviation D on the reference power P_r is shown in Fig. 6.

Table III shows that when the selected frequency matches the input, measurement is highly accurate, with a maximum deviation of 0.08 dB and a mean absolute error of 0.02 dB. Fig. 6 further shows that the deviation for f_3 = 760.7 MHz remains nearly constant.

When the signal frequency is unknown, measurement accuracy degrades. The maximum deviation is 0.51 dB, while the mean absolute error is 0.20 dB, indicating the device can still provide a reasonable power estimate.

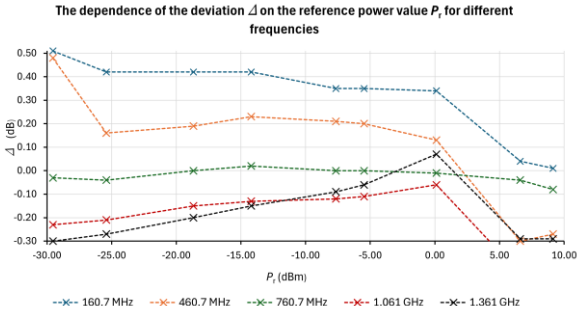

Fig. 6. The dependence of the deviation D on the reference power value P_r for different frequencies

IV. CONCLUSION

A compact device for RF power measurement was designed and built, and its performance was evaluated by comparing measured power values with reference values over five different frequencies and a power range from -10 dBm to 30 dBm. The results show that the device achieves high accuracy when the signal frequency is known and still provides a satisfactory estimate when the frequency is unknown. In future iterations, adding an automatic frequency detection method could further improve measurement accuracy.

V. REFERENCES

[1] *ADL5904 (Rev. B) - ADL5904: Data Sheet*, Rev. B. 2017. Accessed: Mar. 09, 2026. [Online]. Available: https://www.analog.com/media/en/technical-documentation/data-sheets/ADL5904.pdf

[2] *Surface Mount Digital Step Attenuator DAT-31A+ Series*, Rev. C. c2017. Accessed: Mar. 09, 2026. [Online]. Available: https://www.minicircuits.com/pdfs/DAT-31A-PP+.pdf

[3] *Arm ® Cortex® -M4 32b MCU+FPU, up to 64KB Flash, 16KB SRAM, 2 ADCs, 3 DACs, 3 comp., op-amp 2.0 - 3.6: Datasheet - production data*, Rev 8. 2018. [Online]. Available: https://www.st.com/resource/en/datasheet/stm32f303k8.pdf

[4] MCP3422/3/4: 18-Bit, Multi-Channel ΔΣ Analog-to-Digital Converter with I2C™ Interface and On-Board Reference, Revision C. 2009. Accessed: Mar. 09, 2026. [Online]. Available: https://ww1.microchip.com/downloads/aemDocuments/documents/OTH/ProductDocuments/DataSheets/22088c.pdf

[5] *MCP4802/4812/4822: 8/10/12-Bit Dual Voltage Output Digital-to-Analog Converter with Internal VREF and SPI Interface*, Revision B. 2015. Accessed: Mar. 09, 2026. [Online]. Available: https://ww1.microchip.com/downloads/en/DeviceDoc/20002249B.pdf

[6] *NCP718 - Voltage Regulator - Low Dropout, Wide Input Voltage, Low Iq*, Rev. 11. 2023. Accessed: Mar. 09, 2026. [Online]. Available: https://www.onsemi.com/pdf/datasheet/ncp718-d.pdf

[7] J. Stastny, "RF-meter," GitHub repository, https://github.com/stastny1/RF-meter, accessed on: Mar. 15, 2026.

[8] S. Maniktala, *Switching power supplies A to Z*. Burlington: Newnes, 2006.

[9] "CMSIS: Introduction." Accessed: Mar. 13, 2026. [Online]. Available: https://arm-software.github.io/CMSIS_6/latest/General/index.html

[10] "CMSIS-RTOS2: Overview." Accessed: Mar. 13, 2026. [Online]. Available: https://arm-software.github.io/CMSIS_6/latest/RTOS2/index.html

Temperature-Controlled Chamber for Comparing Temperature Sensors

Martin Zatloukal
Department of Radio Electronics
FEEC, Brno University of Technology
Brno, Czech Republic
256608@vutbr.cz

Jan Kufa
Department of Radio Electronics
FEEC, Brno University of Technology
Brno, Czech Republic
kufa@vut.cz

***Abstract*—This paper presents the design and experimental validation of a custom, cost-effective temperature-controlled chamber for the characterization and comparative analysis of ambient temperature sensors. Driven by an ESP32 microcontroller, the system utilizes an asymmetric software PID controller to manage bidirectional thermal regulation via a Peltier module and power resistors within a compact 3D-printed enclosure. To ensure high signal integrity and minimize parasitic thermal inertia, the architecture employs a modular hardware design featuring a specialized sensor PCB without copper pours. Experimental results demonstrate exceptional stability, achieving a steady-state error of only 0.02 °C and a minimal maximum overshoot of 0.07 °C. The proposed platform provides a highly accessible, high-precision alternative to expensive commercial climatic chambers, making it an ideal solution for educational laboratories and rapid prototyping.**

***Index Terms*—temperature-controlled chamber, ambient temperature sensors, ESP32 microcontroller, Peltier module, PID control**

I. Introduction

Accurate ambient temperature measurement is critical across industrial and scientific applications. However, evaluating the dynamic performance of various sensor technologies-such as thermal inertia (τ) and packaging effects [1] requires strictly defined and repeatable temperature transients [2]. Standard educational laboratories typically lack the active thermal regulation necessary for such precise characterization.

To address this gap, this paper presents the design and realization of a compact, cost-effective temperature-controlled chamber. Driven by an ESP32 microcontroller [3], the system employs a Peltier module [4] and power resistors [5] to achieve active, bidirectional thermal regulation within a custom 3D-printed enclosure. By integrating an advanced software PID controller and a specialized low-thermal-mass dedicated printed circuit board (PCB), the proposed platform enables simultaneous, high-precision comparative analysis of analog and digital sensors, offering an accessible alternative to expensive commercial climatic chambers.

II. System Architecture

Based on the experimental evaluation of the initial prototype, a two-part hardware architecture was designed to physically separate the main control unit from the measurement chamber. This separation primarily optimizes the system's thermodynamics: minimizing the active chamber's internal volume significantly accelerates temperature transients, while isolating the power electronics prevents parasitic heat dissipation from affecting the measurements.

Furthermore, this modularity ensures the 3D-printed parts remain compact. To withstand elevated temperatures during active heating, the chamber enclosure is manufactured from polyethylene terephthalate glycol (PETG) filament rather than standard polylactic acid (PLA), owing to its superior thermal resistance. The complete block diagram is illustrated in Fig. 1.

A. Modular Design and Signal Integrity

The control unit serves as the core of the platform, housing the power supply, the ESP32 microcontroller, power switching MOSFET modules [6], and the external 16-bit analog-to-digital converter (ADC) ADS1115 [7]. Because all signals - including analog voltages and digital buses for integrated sensors - must travel between the two units, signal integrity is a critical factor. The interconnect system utilizes foiled/unshielded twisted pair (F/UTP) cables. Specifically, a PCB equipped with two RJ45 connectors is used to break out the signals to screw terminals, from which they are routed to the respective components. To further ensure stability over the external cabling, the circuits incorporate localized RC filtering for the analog voltage dividers.

B. Thermal Chamber Design

The chamber features an active measurement area of 170 × 140 mm. Thermal regulation is driven by a TEC1-12706 Peltier module (maximum cooling power Q_{max} = 60 W, maximum temperature differential ΔT_{max} = 67 °C [4]) and an array of 20 W power resistors for rapid Joule heating [5].

To minimize parasitic thermal inertia, the sensor PCB entirely omits copper pours. Additionally, an integrated switchable fan allows comparing natural and forced convection. During standard operation, its turbulent airflow counteracts spatial temperature gradients induced by localized thermal sources, rapidly homogenizing the internal climate and ensuring uniform thermal distribution across all sensors within seconds.

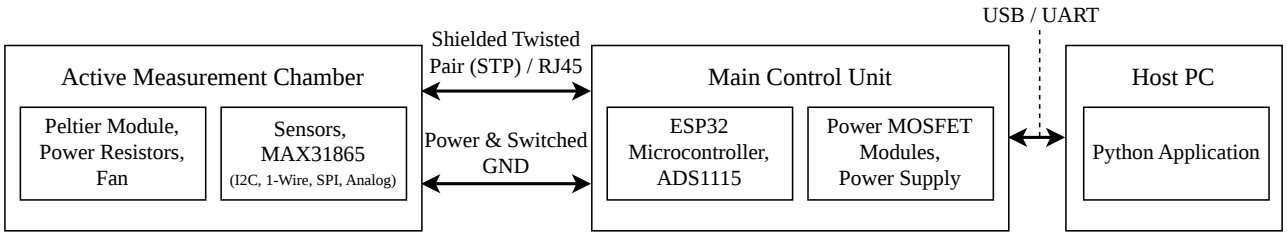

Fig. 1. Block diagram of the proposed two-part hardware architecture.

C. Sensor Array Selection

The measurement chamber is equipped with a custom PCB designed to accommodate a diverse array of sensors.

The core of the measurement array is the TMP117 sensor (I2C). It serves as the absolute reference variable for the software PID controller and establishes the baseline truth for subsequent data analysis. To ensure maximum reliability, the TMP117 is traceable to the National Institute of Standards and Technology (NIST) and features a typical accuracy of ± 0.05 °C within the -20 °C to 50 °C range. Furthermore, its extremely low active current consumption of 3.5 μA minimizes parasitic self-heating [8].

To evaluate the impact of physical packaging on thermal inertia, the system includes two DS18B20 sensors utilizing the 1-Wire protocol [9]. While identical in their silicon components, one is housed in a standard TO-92 plastic package, whereas the other is enclosed in a heavier waterproof stainless-steel probe.

The sensor array further includes an RTD PT1000 platinum resistance thermometer [10]. To ensure highly accurate readout, this sensor is interfaced via an Adafruit 3648 amplifier board (based on the MAX31865 IC), which communicates with the microcontroller over the SPI bus [11].

The purely analog section is designed to be evaluated simultaneously by both an external 16-bit ADC (ADS1115) and the internal ADC of the ESP32 microcontroller. This dual-routing architecture allows students to directly compare measurement precision and noise levels between a dedicated high-resolution ADC and a standard integrated microcontroller peripheral.

This analog section features two distinct voltage dividers. The first consists of a 10 kΩ NTC thermistor [12] paired with a precision 10 kΩ reference resistor [13]. The second is a fixed reference divider composed entirely of two 10 kΩ precision resistors [13]. The primary educational purpose of this fixed divider is to practically demonstrate the negligible temperature coefficient of standard precision resistors compared to the highly temperature-dependent NTC thermistors. Additionally, it serves as a static hardware reference to monitor the baseline noise and stability of the A/D conversion independently of environmental thermal changes.

III. Software and Control System

The platform is built upon a distributed control architecture. A desktop application handles the high-level control loop, user interface, and data logging, while the ESP32 microcontroller serves as the real-time hardware interface. This separation ensures that the complex computational tasks do not interfere with low-level hardware timing.

A. Firmware and Hardware Safety

The ESP32 microcontroller asynchronously polls the sensor array and transmits the data to the host PC via a USB serial connection at a 1 Hz update rate. In return, it receives calculated pulse-width modulation (PWM) duty cycle values to drive the power MOSFETs.

Because the array of resistive heating elements can deliver tens of watts of thermal power, a critical safety mechanism is implemented within the microcontroller firmware. A hardware watchdog continuously monitors the serial communication. If the connection is lost or the host PC fails to send an update within a 3-second window, the ESP32 automatically forces all PWM outputs to zero. This fail-safe mechanism prevents thermal runaway and protects the measurement device from destruction.

B. Temperature Control Algorithm

The core of the thermal regulation is a discrete software-based PID controller. To ensure loop stability, the primary process variable (the TMP117 temperature) is processed through an exponential moving average (EMA) filter. Simultaneously, to mitigate internal ESP32 ADC noise and 50 Hz mains interference for analog sensors, an Alpha-Trimmed Mean filter is applied. This statistical filter captures 101 burst samples over a 30 ms window, discards 50 extreme outliers, and averages the remaining median values.

Because Joule heating via power resistors exhibits a faster thermodynamic response than the thermoelectric cooling of the Peltier module [1], the algorithm employs an asymmetric PID design, dynamically switching between two distinct sets of constants (K_p, K_i, K_d). These parameters were tuned heuristically; proportional gain was initially adjusted to drive the temperature near the target, followed by empirical fine-tuning of the integral and derivative terms to achieve a critically damped response.

Furthermore, standard anti-windup strategies are implemented to prevent overshoot. The integral effort is paused within a ± 0.1 °C deadband, and the accumulated term is clamped to ± 45 %, providing sufficient steady-state power to compensate for environmental losses.

C. Data Logging and User Interface

The desktop application provides a graphical interface for real-time process monitoring. For educational purposes, a background thread continuously logs all incoming sensor telemetry and exports the dataset to a CSV file. This enables students to perform offline post-processing, such as evaluating the sensor step responses and calculating their respective thermal time constants (τ).

IV. Educational Methodology

The platform provides hands-on experience in evaluating sensor characteristics through three specific experimental tasks.

A. Signal Processing and Sensor Linearity

This task focuses on resistive temperature measurement and signal conversion. The system outputs raw voltage values for the NTC thermistor divider, the reference resistor, and an equivalent voltage derived from the RTD's SPI converter. Students manually apply Ohm's law and respective transfer functions to calculate final temperatures.

Furthermore, students evaluate analog-to-digital conversion precision using the fixed reference divider. By toggling the "Oversampling" checkbox in the desktop graphical interface, they activate a digital filter within the ESP32 firmware. This allows them to directly observe the reduction in quantization noise on the internal 12-bit ADC and compare it against the highly stable readout of the external 16-bit ADS1115.

B. Dynamic Response and Thermal Inertia

Students calculate the thermal time constant (τ) by subjecting the chamber to a thermal step response. By comparing two identical DS18B20 sensors - one in a standard TO-92 plastic package and the other in a stainless-steel probe - they quantitatively measure the time delay introduced by the protective packaging's thermal mass.

C. Static Accuracy and Environmental Monitoring

Students evaluate the system in steady-state conditions (e.g., 20 °C, 30 °C, and 40 °C). They determine the absolute deviations of all sensors against the TMP117 reference to distinguish between offset and linearity errors. Additionally, using the BME280 sensor data, students calculate the dew point inside the chamber [14]. This demonstrates the thermodynamic constraints of active thermoelectric cooling and the physical necessity of preventing water condensation on electronic components.

V. Experimental Validation

To verify the proposed hardware architecture and control algorithms, a functional prototype was constructed. The internal realization of the two-part system, following the modular design described in Section II-A, is shown in Fig. 2 and Fig. 3.

The dynamic performance of the system and the software PID controller were experimentally validated through a thermal step response test. The system was commanded to heat the active chamber from a stabilized baseline of 29.6 °C to a target setpoint of 35.0 °C.

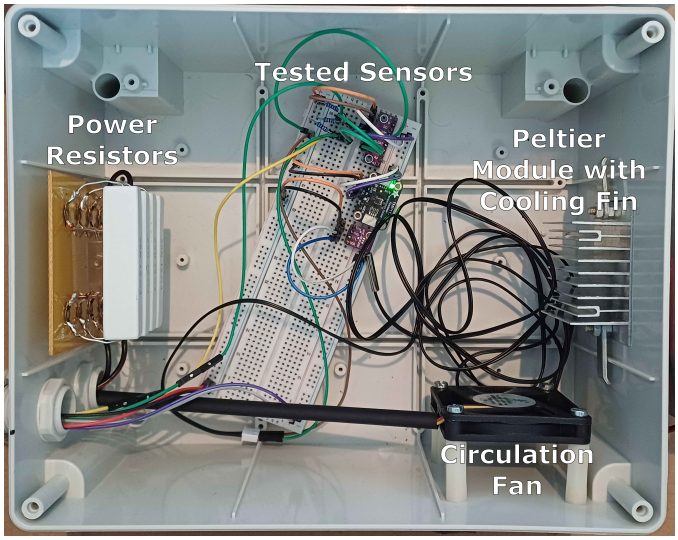

Fig. 2. Internal view of the experimental measurement chamber prototype used for validation.

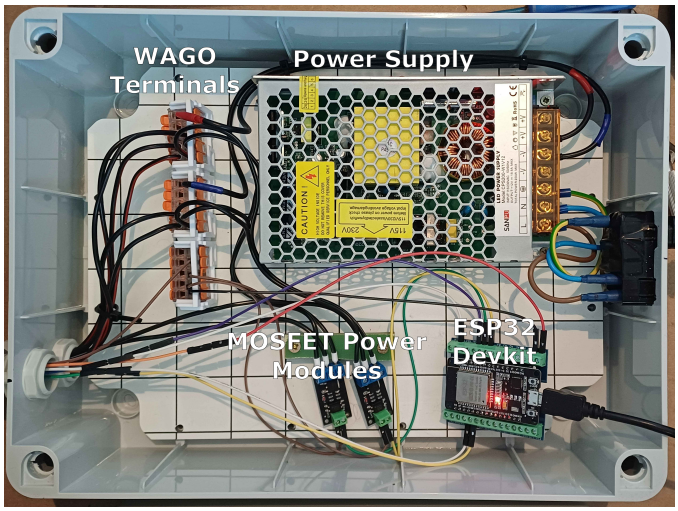

Fig. 3. Internal view of the experimental main control unit prototype used for validation.

Figure 4 shows the graphical user interface of the PC control application, displaying the system's thermal response. The controller reached the target temperature with a minimal maximum overshoot of only 0.07 °C. Because this peak deviation falls strictly within a ±0.1 °C tolerance band, the system's rise time effectively equals its settling time (t_s), which was measured at exactly 640 seconds for this 5.4 °C thermal step.

Once fully settled, the controller maintained the 35.0 °C setpoint with a remarkable steady-state error (e_{ss}) of just 0.02 °C. Furthermore, thanks to the high-resolution digital transmission of the TMP117 reference sensor, the steady-state temperature oscillation (noise floor) was confined to a minimal ±0.02 °C envelope (fluctuating strictly between 34.98 °C and

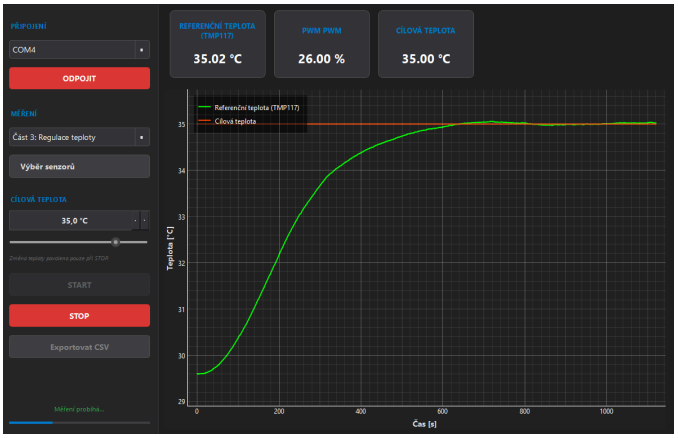

Fig. 4. Real-time thermal step response captured by the PC control software, demonstrating a critically damped PID control with a 640 s settling time and minimal steady-state error.

35.02 °C). These metrics confirm the platform's capability to deliver the highly precise and stable environmental conditions required for rigorous sensor evaluation.

The dual-ADC architecture and the digital oversampling filter were experimentally validated using a fixed voltage reference.

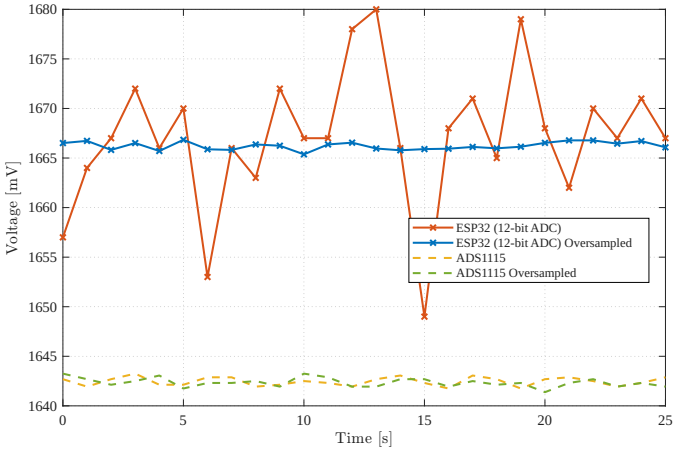

Fig. 5. Experimental evaluation of ADC quantization noise and digital filtering.

Fig. 5 shows that while the raw ESP32 data exhibits significant quantization noise, the software-applied oversampling effectively smooths the signal to match the stability of the dedicated 16-bit ADS1115.

VI. Conclusion

The proposed platform was validated using an experimental prototype. The separation of the control unit and the active chamber isolated the measurement environment from parasitic heat dissipation. In a thermal step response test from 29.6 °C to 35.0 °C, the system achieved a steady-state error (e_{ss}) of 0.02 °C, a maximum overshoot of 0.07 °C, and a settling time of 640 seconds. The noise floor was measured within ±0.02 °C.

Validation of the dual-ADC architecture further confirmed the system's educational and technical value. The implementation of a software-based oversampling algorithm allowed the internal 12-bit ESP32 ADC to match the stability of a dedicated 16-bit ADS1115 reference. These results prove that the platform is capable of high-precision sensor characterization and real-time digital signal processing demonstrations.

Since the ESP32 natively supports Wi-Fi and Bluetooth communication, this represents a direct path for future system enhancement. Future work will focus on finalizing the hardware design and refining the software user interface to prepare the system for deployment in regular laboratory courses.

References

[1] FRADEN, Jacob. Handbook of Modern Sensors [Online]. 5th ed. Springer International Publishing Switzerland: Springer, 2016 [Accessed: 2025-12-21]. ISBN 978-3-319-19302-1. Available: DOI: 10.1007/978-3-319-19303-8.

[2] SINANOVIC, Nermin; PEVEC, Simon; NJEGOVEC, Matej; BUDINSKI, Vedran; DONLAGIC, Denis. High Resolution Short Response Time Fiber-Optic Temperature Sensor. IEEE Sensors Journal [Online]. 2025-2-1, vol. 25, no. 3, pp. 4655-4661 [Accessed: 2025-12-21]. ISSN 1530-437X. Available: DOI: 10.1109/jsen.2024.3517640.

[3] ESP32 Series Datasheet [Online]. 2024. [Accessed: 2025-12-26]. Available: https://documentation.espressif.com/esp32_datasheet_en.pdf.

[4] TEC1-127xx Thermoelectric Cooler [Online]. 2022. [Accessed: 2025-12-26]. Available: http://www.handsontec.com/dataspecs/peltier.pdf.

[5] PASSIVES, SR. Cement Resistor [Online]. [n.d.]. [Accessed: 2025-12-26]. Available: https://www.tme.eu/Document/d2d66de62be8f045452af502fd3b634f/SR-Passives-CR-W.pdf.

[6] International Rectifier. LR7843 Datasheet: HEXFET Power MOSFET [Online]. [n.d.]. [Accessed: 2026-03-14]. Available: https://www.alldatasheet.com/html-pdf/735664/IRF/LR7843/53/1/LR7843.html.

[7] Texas Instruments. ADS1115 Datasheet [Online]. 2024. [Accessed: 2026-03-14]. Available: https://www.ti.com/lit/ds/symlink/ads1115.pdf.

[8] TMP117 Datasheet [Online]. [n.d.]. [Accessed: 2025-12-26]. Available: https://www.ti.com/lit/ds/symlink/tmp117.pdf.

[9] DS18B20 Datasheet [Online]. [n.d.]. [Accessed: 2025-12-26]. Available: https://www.analog.com/media/en/technical_documentation/data-sheets/ds18b20.pdf.

[10] Adafruit Industries. Platinum RTD Sensor - PT1000 - 3 Wire 1 meter long Product ID: 3984 [Online]. 2022. [Accessed: 2026-03-14]. Available: https://www.alldatasheet.com/datasheet-pdf/view/2196518/ADAFRUIT/3984.html.

[11] Adafruit Industries. Universal RTD-to-Digital Converter (MAX31865) - Product ID: 3648 [Online]. [n.d.]. [Accessed: 2026-03-14]. Available: https://www.alldatasheet.net/datasheet-pdf/pdf/2186442/ADAFRUIT/3648.html.

[12] NTC Thermistors: Application Note [Online]. Vishay Intertechnology, 2021 [Accessed: 2025-12-26]. Available: https://www.vishay.com/docs/29053/ntcappnote.pdf.

[13] Metal Film Resistors: General Type, MFR Series [Online]. Revision: 201304. 2013. Available: https://www.laskakit.cz/user/related_files/2304140030_yageo-mfr-25fte52-10k_c172967.

[14] BME280 Datasheet [Online]. Bosch Sensortec GmbH, [n.d.] [Accessed: 2025-12-26]. Available: https://www.bosch-sensortec.com/media/boschsensortec/downloads/datasheets/bst-bme280-ds002.pdf.

Functional Profiling of Bacterial Genomes Using Constrained Ordination and Dimensionality Reduction

A. Fialová, K. Sedlář
Department of Biomedical Engineering
Faculty of Electrical Engineering and Communication
Brno University of Technology
Czech Republic
240499@vut.cz, sedlar@vut.cz

***Abstract*—Modern bacterial taxonomy is transitioning from phenotypic characterization to genome-wide sequence comparisons. While metrics like Average Nucleotide Identity (ANI) provide robust species delineation, they often lack functional context. This study evaluates the use of genome-wide functional profiles, specifically Enzyme Commission (EC) numbers, for bacterial classification. Using a dataset of 100 diverse bacterial genomes, we mapped functional repertoires and analyzed them through linear and non-linear dimensionality reduction techniques. Our results confirm that enzymatic profiles carry a strong taxonomic signal. Constrained ordination and non-linear embedding showed high congruence with established nomenclature, notably highlighting the functional identity between *Escherichia* and *Shigella*.**

***Index Terms*—Bacterial taxonomy, functional annotation, EC numbers, dimensionality reduction, CAP, t-SNE, ANI.**

I. Introduction

In recent decades, microbiology has been transformed by Next-Generation Sequencing, enabling the study of entire genomes rather than individual genes. Historically, bacterial taxonomy relied on phenotypic traits such as morphology, staining, and biochemical activity [1], but these methods often lacked stability and objectivity due to trait variability and inconsistent results within species.

The introduction of 16S rRNA sequencing provided a molecular clock for phylogeny, but its resolution is limited for closely related species [2]. Modern taxonomy now relies on genome-wide comparisons, with Overall Genome Relatedness Indices (OGRIs), such as Average Nucleotide Identity (ANI) and digital DNA-DNA hybridization (dDDH), serving as high-resolution, reproducible standards [3], [4].

While sequence-based metrics robustly classify phylogeny, they do not capture a bacterium's functional landscape. Understanding biological roles requires complementing genomic data with functional annotation, which assigns biological meaning and biochemical context to genes [5].

Standardized systems like the Enzyme Commission (EC) nomenclature map genes to specific catalytic reactions, enabling large-scale reconstruction of an organism's metabolic potential [6] and exploration of the "metabolic space" defining a taxon.

Functional annotation depends on reference databases, whose coverage varies across taxa, so potential biases and incomplete reconstructions must be considered.

This study explores whether bacterial metabolic repertoires, represented by EC numbers, can reliably reflect taxonomy. By converting genomes into high-dimensional functional matrices, we assess how evolutionary history aligns with metabolic phenotype, supporting a functionally grounded microbial classification.

II. Materials and Methods

A. Dataset Selection

A dataset of 100 bacterial genomes from 16 genera was assembled from the NCBI Assembly database. Genome quality was assessed using standard completeness and contamination metrics reported for each assembly. Only genomes with an estimated completeness greater than 95% and contamination below 5% were retained, consistent with commonly applied genome quality thresholds in comparative genomics studies. The genomes were divided into three groups:

- Group A (Closely Related): This group consists of 20 strains belonging to the *Escherichia* and *Shigella* genera. These were selected to provide a high-resolution test case for distinguishing between taxa that are known to be phylogenetically nearly identical but clinically distinct.
- Group B (Moderately Related): Comprising 70 genomes, this tier represents seven major bacterial groups, including *Salmonella enterica, Pseudomonas aeruginosa, Staphylococcus aureus, Bacillus subtilis, Streptococcus pneumoniae, Enterococcus faecalis,* and *Mycobacterium spp.*, with 10 strains per group.
- Group C (Outgroup): A diverse set of 10 genomes from various phyla, such as *Bacteroides, Cyanobacteria*, and *Deinococcus*, was included to provide the broad evolutionary contrast necessary for robust ordination analysis.

B. Functional Annotation Pipeline

Raw genomic sequences were annotated using the KEGG Automatic Annotation Server (KAAS) [7] via the bi-

directional best hit (BBH) approach, in which predicted proteins are matched against reference sequences and retained only when reciprocal best matches occur between query and database entries. Genes were assigned to KEGG Orthology (KO) groups, representing orthologous genes with equivalent functions, using KEGG Release 116.0 (October 2025). KO identifiers were then mapped to Enzyme Commission numbers via the KEGG API, producing a binary matrix of thousands of enzymes that serves as a high-dimensional functional "fingerprint" for each of the 100 genomes.

C. Data Matrices and Dimensionality Reduction

The processed data were arranged into matrices for dimensionality reduction. Although absolute enzyme counts were initially obtained, a binary transformation was primarily used: presence = 1, absence = 0. This reduces the influence of gene copy variation and annotation artifacts, emphasizing the shared functional repertoire among taxa. While simplifying the functional landscape and ignoring dosage or expression differences, this approach provides a stable signal suitable for taxonomic-scale comparisons, less sensitive to strain-level variation.

III. RESULTS AND DISCUSSION

A. Pan-genome Composition and Diversity

Pan-genome analysis of 100 bacterial strains revealed high genomic plasticity. The core genome, enzymes shared by all organisms, was extremely small ($\sim$10 enzymes), which is a direct consequence of the broad taxonomic diversity spanning multiple phyla. Most enzymes belong to the accessory genome, carrying the primary taxonomic signal. These universally conserved enzymes primarily correspond to fundamental cellular processes required for basic metabolism. Genera such as *Bacillus* have a notably larger accessory genome, which may reflect a higher propensity for horizontal gene transfer or adaptation to diverse environments.

B. Exploratory Analysis

Unconstrained ordination via Principal Component Analysis (PCA) and Principal Coordinates Analysis (PCoA) provided the first evidence that functional profiles align with taxonomy.

- PCA: The first two principal components explained 33.18% of the total variance (Fig. 1). The visualization showed a clear separation of major groups, though some overlap occurred in the dense "functional core" where housekeeping enzymes are located.
- PCoA: Binary metrics proved superior to Euclidean distances for analyzing sparse enzyme matrices. Jaccard-based PCoA successfully clustered major genera, while Dice ordination (Fig. 1) revealed core metabolic relationships through a characteristic "arch" distribution, a commonly observed effect in ordination methods applied to ecological presence/absence data, where a dominant gradient structure in the dataset is projected onto curved axes. Both metrics confirmed the close functional proximity of *Escherichia* and *Shigella*.

Both Jaccard and Dice metrics yielded similar genus-level clustering, but Dice was chosen for visualization as it better separated closely related taxa by emphasizing shared presences in sparse enzyme matrices.

C. Constrained Ordination

To test the hypothesis that taxonomy dictates the functional repertoire, we performed constrained ordination using Canonical Correspondence Analysis (CCA) and Canonical Analysis of Principal Coordinates (CAP). Unlike exploratory methods, these techniques constrain the ordination axes to the provided taxonomic labels (genera).

- CCA: The CCA analysis showed that the chosen taxonomic constraints were significant (Fig. 1), explaining 11.14% of the total variance. Based on the permutation test, the model reached a p-value of $p < 0.001$, confirming that the distribution of EC numbers across the dataset is not random but follows genus-specific patterns.
- CAP: CAP was employed to overcome some of the linear limitations of CCA (Fig. 1). Although the first two CAP axes explained only 8.27% of the total variance, this proportion is typical for high-dimensional ecological datasets and still captures a biologically meaningful signal across 100 genomes with thousands of enzymes, with the model reaching statistical significance at $p < 0.001$.

D. Non-linear Dimensionality Reduction

The t-SNE algorithm [8] was applied to capture non-linear relationships that linear methods like PCA might overlook. The analysis (Fig. 2) produced visually distinct clusters, effectively separating even diverse genera, demonstrating the suitability of non-linear embedding for large functional datasets. The perplexity was set to 30, a standard value balancing local and global structure preservation [9].

E. Comparative Analysis of Clustering Metrics

To objectively evaluate which method best represents biological reality, we applied the k-means clustering algorithm to the coordinates generated by both linear (CAP) and non-linear (t-SNE) approaches. The number of clusters (k = 16) was set according to the number of genera represented in the dataset. The quality of the resulting clusters was then assessed using standardized metrics:

- Homogeneity: Assesses if each cluster contains only members of a single taxon.
- Completeness: Measures if all members of a given taxon are assigned to the same cluster.
- V-measure: The harmonic mean of homogeneity and completeness, representing overall taxonomic alignment.

TABLE I: Comparison of clustering metrics.

Metric	CAP	t-SNE
Homogeneity	0.7118	0.7453
Completeness	0.6539	0.6774
V-measure	0.6816	0.7097
Purity	0.7027	0.6892

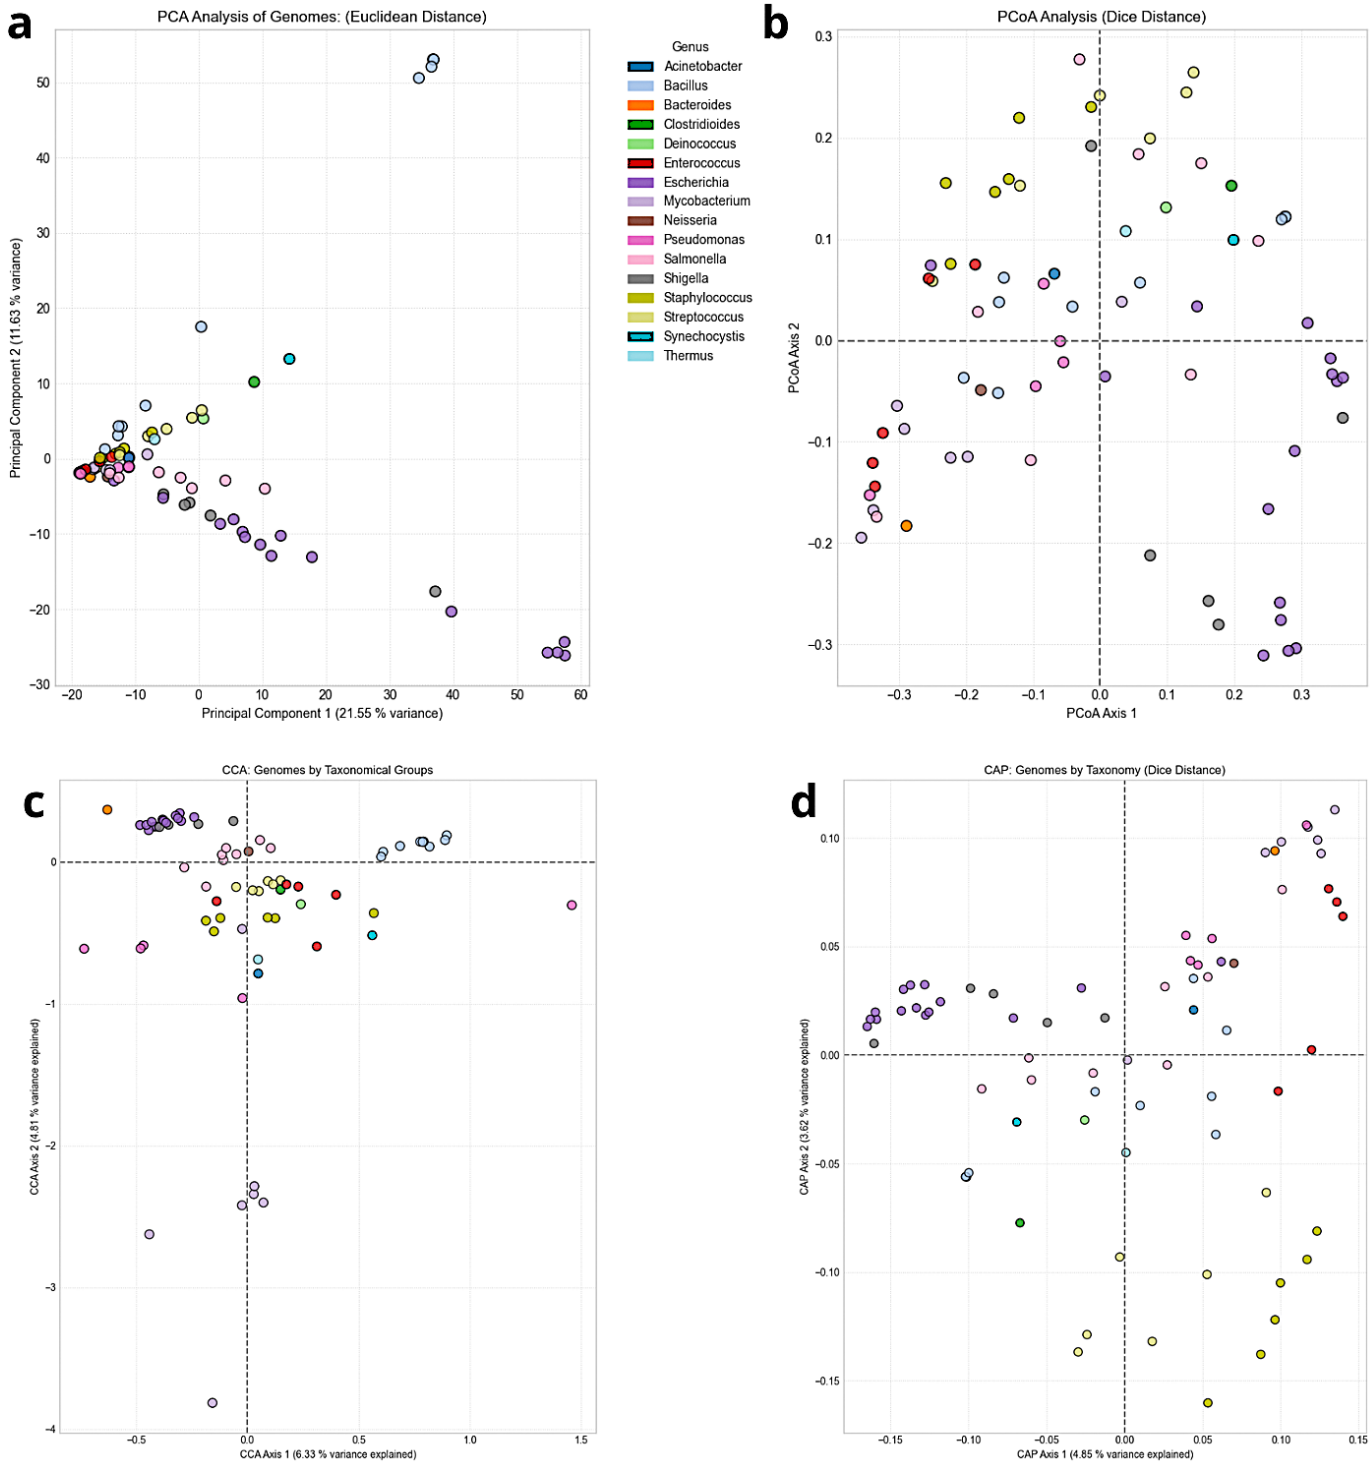

Fig. 1: Results of PCA analysis (a), PCoA analysis using Dice metrics (b), CCA analysis (c) and CAP analysis using Dice metrics (d).

The table I shows that t-SNE outperforms CAP across most clustering metrics, suggesting that it is more effective in capturing bacterial functional diversity based on enzymatic profiles. However, CAP exhibits a slightly higher Purity score, indicating improved consistency within predefined bacterial groups.

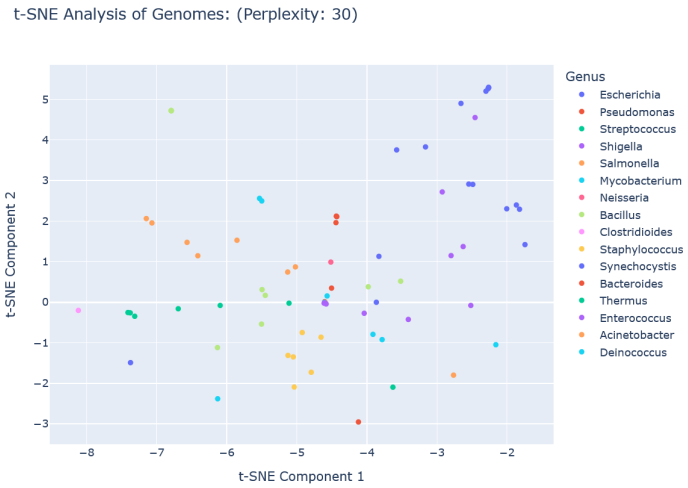

Fig. 2: Results of t-SNE analysis.

F. Limitations

Several limitations of the present study should be acknowledged. First, the dataset comprises 100 genomes, which, while sufficient for methodological exploration, may not fully represent the diversity of bacterial metabolic capabilities. Second, the functional annotations rely on the KEGG database, whose coverage varies across taxonomic groups. As a result, genomes from poorly characterized lineages may have incomplete functional representations. Finally, horizontal gene transfer can influence metabolic repertoires independently of vertical evolutionary relationships. Although the binary transformation of enzyme presence was intended to reduce noise from such events, residual HGT signals may still affect the observed clustering patterns.

IV. CONCLUSIONS

In conclusion, this study has demonstrated that the metabolic potential of bacteria, represented by Enzyme Commission numbers, serves as a robust biological signal for taxonomic classification. The findings confirm that functional profiles are not randomly distributed but are deeply rooted in evolutionary history, as evidenced by the significant grouping of genera in both linear and non-linear ordination spaces. Specifically, the application of non-linear methods like t-SNE proved highly effective for identifying local functional clusters and capturing nuances that linear methods might obscure. Furthermore, constrained ordination techniques, such as CCA and CAP, provided the necessary statistical verification to conclude that genus-level taxonomy is a primary driver of functional diversity, explaining a significant portion of the total variance. The high functional congruence observed within the *Escherichia-Shigella* complex further supports the ongoing taxonomic discussion regarding their potential reclassification as a single evolutionary lineage, given their nearly identical metabolic fingerprints. Future work will focus on integrating these analytical workflows into an automated bioinformatics package. This tool will aim to provide a streamlined transition from genomic assemblies to high-resolution functional and phylogenetic insights, enhancing the speed and accuracy of microbial characterization in the genomic era.

FUNDING

This study was supported by the grant project GACR 25-18143S.

REFERENCES

[1] VANDAMME, P; POT, B; GILLIS, M; DE VOS, P; KERSTERS, K et al. Polyphasic taxonomy, a consensus approach to bacterial systematics. Online. Microbiological Reviews. 1996, vol. 60, no. 2, s. 407-438. ISSN 0146-0749. https://doi.org/10.1128/mr.60.2.407-438.1996. [cit. 2025-11-30].

[2] JANDA, J. Michael a ABBOTT, Sharon L. 16S rRNA Gene Sequencing for Bacterial Identification in the Diagnostic Laboratory: Pluses, Perils, and Pitfalls. Online. Journal of Clinical Microbiology. 2007, vol. 45, no. 9, s. 2761-2764. ISSN 0095-1137. https://doi.org/10.1128/jcm.01228-07. [cit. 2025-11-30].

[3] CHUN, Jongsik; OREN, Aharon; VENTOSA, Antonio; CHRISTENSEN, Henrik; ARAHAL, David Ruiz et al. Proposed minimal standards for the use of genome data for the taxonomy of prokaryotes. Online. International Journal of Systematic and Evolutionary Microbiology. 2018, vol. 68, no. 1, s. 461-466. ISSN 1466-5026. https://doi.org/10.1099/ijsem.0.002516. [cit. 2025-11-30].

[4] RICHTER, Michael a ROSSELLÓ-MÓRA, Ramon. Shifting the genomic gold standard for the prokaryotic species definition. Online. Proceedings of the National Academy of Sciences. 2009, vol. 106, no. 45, s. 19126-19131. ISSN 0027-8424. https://doi.org/10.1073/pnas.0906412106. [cit. 2025-11-30].

[5] EJIGU, Girum Fitihamlak a JUNG, Jaehee. Review on the Computational Genome Annotation of Sequences Obtained by Next-Generation Sequencing. Online. Biology. 2020, vol. 9, no. 9, s. 295. ISSN 2079-7737. https://doi.org/10.3390/biology9090295. [cit. 2025-10-04].

[6] TIPTON, Keith a BOYCE, Sinéad. History of the enzyme nomenclature system. Online. Bioinformatics. 2000, vol. 16, no. 1, s. 34-40. ISSN 1367-4811. https://doi.org/10.1093/bioinformatics/16.1.34. [cit. 2025-10-04].

[7] KANEHISA, Minoru; FURUMICHI, Miho; TANABE, Mao; SATO, Yoko a MORISHIMA, Kanae. KEGG: new perspectives on genomes, pathways, diseases and drugs. Online. Nucleic Acids Research. 2016, vol. 45, no. D1, s. D353-D361. ISSN 0305-1048. https://doi.org/10.1093/nar/gkw1092. [cit. 2026-03-30].

[8] VAN DER MAATEN, Laurens a Geoffrey HINTON. Visualizing data using t-SNE. Journal of Machine Learning Research. Online. 2008, 9, 2579–2605. https://www.jmlr.org/papers/volume9/vandermaaten08a/vandermaaten08a.pdf. [cit. 2026-03-30].

[9] LIU, Honghua; YANG, Jing; YE, Ming; JAMES, Scott C.; TANG, Zhonghua et al. Using t-distributed Stochastic Neighbor Embedding (t-SNE) for cluster analysis and spatial zone delineation of groundwater geochemistry data. Online. Journal of Hydrology. 2021, vol. 597, s. 126146. ISSN 0022-1694. https://doi.org/10.1016/j.jhydrol.2021.126146. [cit. 2026-03-30].

Hardware and Software Augmentation of Bioreactors for Culturing Bacteria

1st Borys Aleksandrenko
Department of Biomedical Engineering
Faculty of Electrical Engineering and Communication
Brno University of Technology
Brno, Czech Republic
240481@vutbr.cz

2nd Martin Súkeník
Institute of Physical and Applied Chemistry
Faculty of Chemistry
Brno University of Technology
Brno, Czech Republic
martin.sukenik@vut.cz

3rd Karel Sedlář
Department of Biomedical Engineering
Faculty of Electrical Engineering and Communication
Brno University of Technology
Brno, Czech Republic
sedlar@vut.cz

***Abstract*—Modern microbial cultivation increasingly relies on semi-automated and robotic platforms equipped with integrated sensing and control to enable reproducible, scalable, and data-rich experiments. Open-source bioreactors represent a particularly attractive class of such systems due to their affordability, transparency, and extensibility. In this work, three representative open-source bioreactor platforms: Chi.Bio, Pioreactor and eVOLVER, are analyzed and compared with respect to their hardware modularity, software architecture, sensing capabilities, scalability, and suitability for advanced cultivation tasks.**

As a practical contribution, this work focuses on extending the sensing capabilities of the Chi.Bio bioreactor, a compact and cost-effective open-source cultivation platform originally limited primarily to optical and thermal measurements. Continuous pH monitoring was integrated using a commonly available laboratory pH electrode, addressing one of critical chemical parameters influencing microbial growth.

The integrated system was evaluated in a two-day *in vitro* cultivation experiment using *Escherichia coli* (CCM 3954). pH was externally modified through stepwise additions of a 30% NaOH solution, and the relationship between pH and optical density, used as a proxy for cell concentration, was assessed. Each alkaline addition resulted in a discrete increase in pH accompanied by a moderate decrease in optical density. The results demonstrate that the developed pH module provides stable and reproducible measurements, integrates seamlessly with the Chi.Bio software ecosystem, and enables straightforward data retrieval for downstream analysis

***Index Terms*—open-source bioreactor, Chi.Bio, Pioreactor, pH sensing, microbial cultivation, automated cultivation, bioprocess monitoring, sensor integration, *in situ* measurement**

I. Introduction

Microbial cultivation is fundamental to biotechnology, synthetic biology, and microbiological research. Controlled cultivation systems enable the study of cellular physiology, metabolic regulation, and population dynamics under defined environmental conditions [1], [2]. Bioreactors provide the infrastructure required to regulate key parameters such as temperature, aeration, nutrient availability, and pH, which is essential for reproducibility and quantitative analysis of biological experiments [3].

Access to advanced bioreactor systems, however, remains limited in many academic laboratories due to high costs and proprietary hardware [4]. In response, several open-source bioreactor platforms have emerged, aiming to provide affordable, modular, and extensible cultivation systems. Platforms such as Chi.Bio and Pioreactor demonstrate that automated microbial cultivation can be achieved using low-cost hardware and open software architectures [5], [6].

Among cultivation parameters, pH plays a critical role in microbial physiology by influencing enzyme activity, membrane transport, and metabolic processes. Despite its importance, pH sensing is not included in many open-source bioreactor designs.

This work therefore focuses on the design, implementation, and evaluation of a pH sensing module integrated into an open-source bioreactor platform.

II. Open-source bioreactors analysis

A comparison of selected open-source bioreactor platforms is summarized in Table I. Among the evaluated systems, Pioreactor provides particularly strong support for rapid sensor integration due to its well-documented API, extensive tutorials, and active development community.

Despite these advantages, the implementation presented in this work was performed using the Chi.Bio platform, which was available within the laboratory infrastructure and enabled immediate experimental testing. Chi.Bio also offers sufficient hardware extensibility and a modular Python-based software architecture, allowing integration of additional sensors through its expansion interface.

The pH parameter was selected for integration due to its critical role in microbial growth and metabolism. Reliable pH monitoring can significantly expand the experimental capabilities of open-source cultivation systems, particularly for studies involving non-model microorganisms or processes sensitive to acid–base conditions.

III. Materials and Methods

A. Components

A practical example of pH sensor integration into the Chi.Bio platform has recently been presented by Denton *et al.* [7], providing a valuable reference for the conceptual design proposed in this work. The pH sensing module is based

The study was supported by the grant project GACR GM25-17459M of The Czech Science Foundation.

TABLE I
COMPARISON OF FUNCTIONALITY, SENSING CAPABILITIES, AND USABILITY OF SELECTED OPEN-SOURCE BIOREACTOR PLATFORMS

Feature	Chi.Bio	Pioreactor	eVOLVER
Primary use case	Single-vessel cultivation	Single-vessel cultivation	High-throughput cultivation
Culture volume	∼30 mL	∼20–40 mL	∼20–40 mL per vial
Heating / temperature control	Built-in heater, closed-loop	Built-in heater, closed-loop	Built-in per-vial heating
Optical density (OD)	Infrared OD	Infrared OD (configurable)	Infrared OD per sleeve
Spectrometry / fluorescence	Integrated spectrometer and LEDs	External LEDs, optional sensing	No
Magnetic stirring	Yes	Yes	Yes
Liquid handling	Integrated pumps	Integrated pumps	Centralized pump array
pH sensing (built-in)	No (modifiable)	No (modifiable)	No (modifiable)
Sensor extensibility	Moderate (I^2C)	High (I^2C, plugin-based)	Moderate (custom)
Software architecture	Python + web UI	Python + modular API	Distributed Python system
Parallelization	Moderate (eight)	Moderate (limit not defined)	High (tens–hundreds)
Approximate cost (USD)	∼1190	∼405 ±20	∼10,000 (estimate)
Documentation quality	Good	Excellent	Moderate
Community support	Moderate	Strong	Active but fragmented
Assembly complexity	Low	Low	High

on the EZO™ pH Circuit manufactured by Atlas Scientific (Fig. 1A), a compact embedded module for electrochemical measurements. The circuit provides a measurement range of 0.001–14.000 with a resolution of 10^{-3} and accuracy of ±0.002. The typical measurement time is approximately 800 ms. The module supports universal pH probes, one-, two- or three-point calibration, and communicates via the I^2C protocol. The operating voltage is 3.3–5 V and measurement data are returned in ASCII format [8].

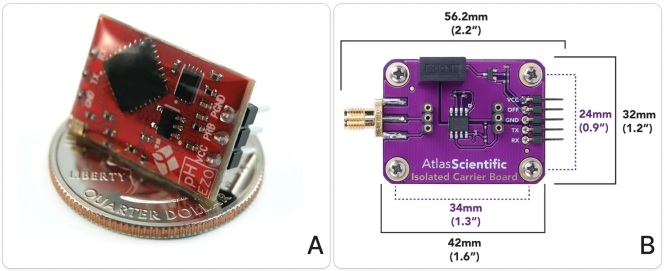

Fig. 1. Overview of the EZO™ pH sensing hardware, adapted from [8]. (A) EZO™ pH circuit module. (B) Electrically Isolated EZO™ Carrier Board with mechanical dimensions and probe interface.

To enable electrical isolation and provide a standardized interface for the pH probe, the EZO™ circuit was mounted on an Electrically Isolated EZO™ Carrier Board (Fig. 1B), equipped with a female SMA connector for probe connection. Consistent with the Chi.Bio hardware architecture, the pH sensing module was connected to the bioreactor via the external 5-pin header port (Fig. 2B).

B. Hardware-Level Modifications and Final Design

The initial configuration proposed by Denton *et al.* [7] combines an Arduino Nano with the Atlas Scientific EZO™ pH circuit mounted on an Electrically Isolated EZO™ Carrier

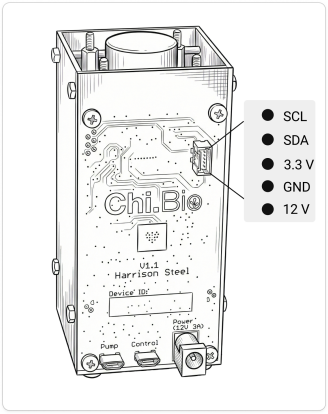

Fig. 2. Close-up view of the Chi.Bio bioreactor expansion port with a 5-pin connector enabling external hardware integration.

Board (Fig. 3A). In this arrangement the microcontroller provides both power and communication with the sensing module. During testing, however, the configuration proved unstable due to the limited current capability of the Arduino Nano 3.3 V regulator. When the EZO™ circuit was connected, reliable USB serial communication could not be maintained.

To ensure sufficient power delivery, the pH circuit was subsequently tested using an STM32 Nucleo-32 development board based on the STM32L432KC microcontroller (Fig. 3B). The higher current capacity of the onboard 3.3 V regulator enabled stable operation of the sensing module and allowed validation of measurement functionality. Communication with the pH circuit was configured using the I^2C protocol required for integration with the Chi.Bio platform. Bus activity and timing were verified using a logic analyzer (Fig. 3C).

For final integration, the external microcontroller was removed to reduce hardware complexity. The pH circuit was

connected directly to the Chi.Bio Mini CT 5POS expansion header and powered from the 12 V supply available on the connector. The 3.3 V pin of the header was not used due to the limited current capability of its onboard voltage regulator. Instead, a compact DC–DC step-down converter (AIMTEC AMSR1-783.3-NZ) was used to generate the required 3.3 V supply for the EZO™ circuit. The resulting wiring scheme enabling direct I²C communication with the bioreactor is shown in Fig. 3D.

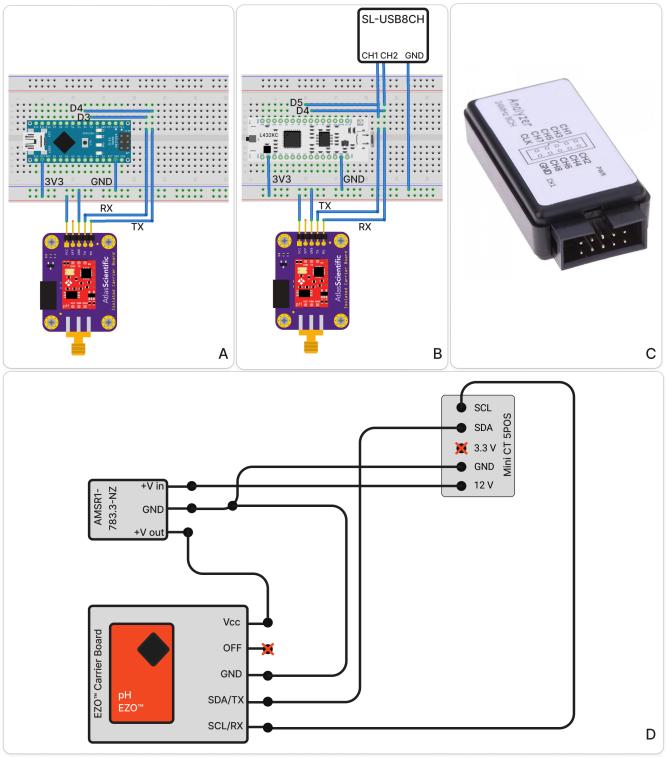

Fig. 3. Wiring schemes of the pH sensing module integration. (A) Initial connection of the Atlas Scientific EZO™ pH circuit to an Arduino Nano. (B) STM32 Nucleo-32 development board used for stable power supply and communication testing. (C) Logic analyzer used for inspection of the I²C communication bus. (D) Final wiring scheme connecting the EZO™ pH circuit to the Chi.Bio Mini CT 5POS expansion header with a DC–DC step-down converter (AIMTEC AMSR1-783.3-NZ).

C. Final Hardware Implementation

The final hardware implementation integrates the Atlas Scientific EZO™ pH circuit and the DC–DC step-down converter into a compact external module connected to the Chi.Bio expansion header. A representative view of the assembled module is shown in Fig. 4.

To protect the electronics during routine laboratory operation, the sensing module was placed in a dedicated enclosure produced using a custom 3D-printed design. A protective case was necessary because cultivation workflows involve frequent handling of liquid reagents, which may expose sensitive electronic components.

The enclosure contains two internal compartments: one housing the EZO™ pH circuit and the other accommodating the DC–DC voltage converter used to generate the 3.3 V supply. The internal layout allows secure routing of soldered connections and provides strain relief for the attached cables.

The enclosure was fabricated from polylactic acid (PLA) using a Prusa MK4S 3D printer with a layer thickness of 0.2 mm and a nozzle temperature of 210 °C.

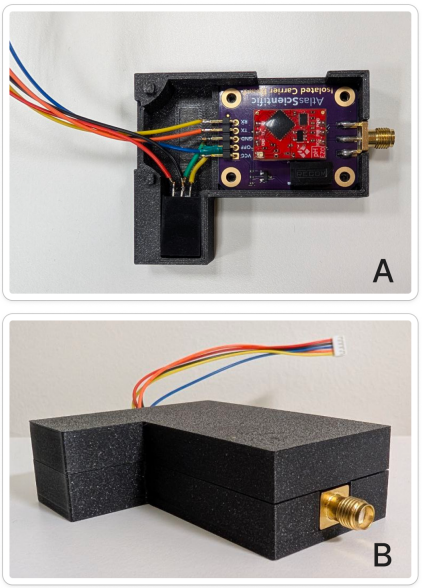

Fig. 4. Final hardware implementation of the pH sensing module. (A) Open enclosure showing internal placement of the EZO™ pH circuit and DC–DC converter. (B) Side view of the assembled module with the enclosure closed.

D. Software-Level Modifications

The Chi.Bio platform uses a BeagleBone Black single-board computer running a Python-based control system with a web interface for experiment management. The software architecture consists of a backend implemented in `app.py`, responsible for hardware communication and data acquisition, and a frontend defined by `index.html` and supporting JavaScript scripts for visualization and user interaction.

Integration of the pH sensing module required modifications to both backend and frontend components. The backend was extended using the *atlas_i2c* library to enable I²C communication with the Atlas Scientific EZO™ pH circuit. Additional variables were introduced into the internal system data structure to store pH measurements, and a dedicated acquisition routine was implemented to periodically read sensor values and include them in the recorded data output.

Frontend scripts were modified to display real-time pH measurements and configuration parameters within the existing web interface. Minor adjustments were also introduced to remove the mandatory dependency on an attached pump module, allowing experiments to run without active liquid handling hardware.

IV. Prototype Evaluation and Discussion

Validation was performed during a two-day *in vitro* cultivation experiment using *Escherichia coli* (CCM 3954). The aim

was to assess the behavior of the integrated pH measurement during microbial growth and controlled pH perturbations.

Cultivation was performed in two Chi.Bio bioreactors operated in parallel to avoid interference between infrared optical density (OD) measurements and the inserted pH electrode. Each reactor contained 20 mL of modified Ashby's medium inoculated with *E. coli*, maintained at 37 °C and continuously mixed using a magnetic stirrer. After biomass accumulation, pH was increased by stepwise additions of 30% NaOH: 10 μL at 21.3 h, 20 μL at 22.6 h, 30 μL at 24.6 h, 80 μL at 25.5 h, and 200 μL at 26.5 h. The experimental setup is shown in Fig. 5.

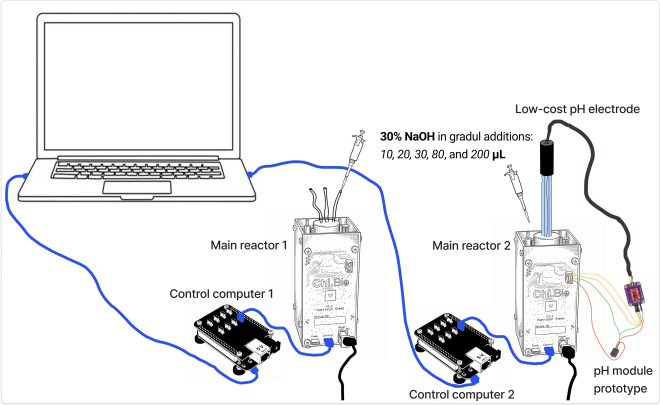

Fig. 5. Schematic representation of the experimental setup used for validation of the integrated pH module during *E. coli* cultivation in the Chi.Bio bioreactor platform.

Recorded datasets were exported and analyzed to evaluate the relationship between OD and pH (Fig. 6). Each NaOH addition produced a discrete increase in pH accompanied by a moderate decrease in OD, likely reflecting short-term physiological stress caused by alkaline shifts.

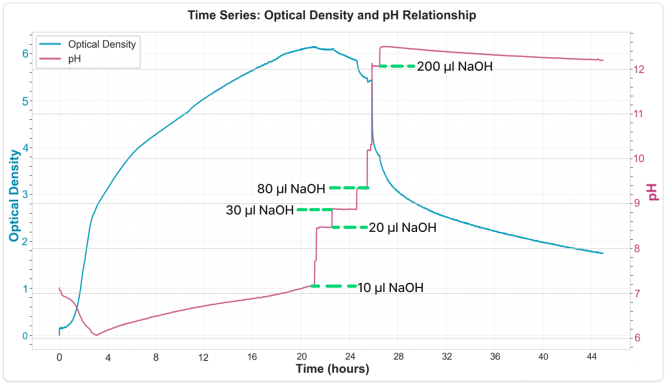

Fig. 6. Optical density and pH recorded during *E. coli* cultivation with stepwise NaOH additions used to induce controlled alkaline shifts.

The experiment confirmed stable pH signal acquisition and successful integration of the sensing module into the Chi.Bio data pipeline. Although the validation was limited by the absence of a reference culture and by the open-vessel pH electrode insertion, which introduces a potential contamination risk, it represents a sufficient validation of the proposed prototype. The prototype therefore provides a suitable platform for future improvements, including sterile probe integration and more rigorous validation experiments.

V. Conclusion

This work presented the design and integration of a pH sensing module for the open-source Chi.Bio bioreactor platform. The proposed solution extends the sensing capabilities of the system by enabling continuous pH monitoring using a commonly available laboratory electrode connected through an Atlas Scientific EZO™ pH circuit. Both hardware and software modifications were implemented to achieve seamless integration with the existing Chi.Bio architecture. Experimental validation using *Escherichia coli* cultivation confirmed stable pH signal acquisition and successful incorporation of pH data into the standard measurement pipeline. The developed module provides a low-cost and accessible approach for extending open-source cultivation platforms. Future work will focus on improved sterile probe integration and more comprehensive validation experiments.

References

[1] SEGERITZ, Charis-P. and Ludovic VALLIER. *Cell Culture: Growing Cells as Model Systems In Vitro*. In: Morteza JALALI, Francesca Y. L. SALDANHA and Mehdi JALALI, eds. *Basic Science Methods for Clinical Researchers*. Academic Press, 2017, pp. 151–172. ISBN 978-0-12-803077-6. Available from: https://doi.org/10.1016/B978-0-12-803077-6.00009-6.

[2] ROUF, Akhmad, Varsha KANOJIA, H. R. MANJUNATH NAIK, Bazilla NASEER and Tahiya QADRI. *An overview of microbial cell culture*. *Journal of Pharmacognosy and Phytochemistry*, 2017, vol. 6, pp. 1923–1928. Available from: https://api.semanticscholar.org/CorpusID:55929069.

[3] WANG, Si-Jing and Jian-Jian ZHONG. *Chapter 6. Bioreactor Engineering*. In: *Bioprocessing for Value-Added Products from Renewable Resources*, vol. 1, 2007, pp. 131–161. ISBN 978-0-444-52114-9. Available from: https://doi.org/10.1016/B978-044452114-9/50007-4.

[4] PALLADINO, F., P. R. F. MARCELINO, A. E. SCHLOGL, Á. H. M. JOSÉ, R. d. C. L. B. RODRIGUES, D. L. FABRINO, I. J. B. SANTOS and C. A. ROSA. *Bioreactors: Applications and Innovations for a Sustainable and Healthy Future—A Critical Review*. *Applied Sciences*, 2024, vol. 14, article 9346. Available from: https://doi.org/10.3390/app14209346.

[5] Pioreactor. *Pioreactor: Open-source bioreactor platform*. Online. Available from: https://pioreactor.com/. [cit. 2026-03].

[6] STEEL, H., R. HABGOOD, C. L. KELLY and A. PAPACHRISTODOULOU. *In situ characterisation and manipulation of biological systems with Chi.Bio*. *PLOS Biology*, 2020, vol. 18, no. 7, article e3000794. Available from: https://doi.org/10.1371/journal.pbio.3000794.

[7] DENTON, M. C. R., N. P. MURPHY, B. NORTON-BAKER, M. LUA, H. STEEL and G. T. BECKHAM. *Integration of pH Control into Chi.Bio Reactors and Demonstration with Small-Scale Enzymatic Poly(ethylene terephthalate) Hydrolysis*. *Biochemistry*, 2024, vol. 63, no. 13, pp. 1599–1607. Available from: https://doi.org/10.1021/acs.biochem.4c00149.

[8] ATLAS SCIENTIFIC. *EZO-pH™ Embedded pH Circuit: ISO 10523 Compliant*. Version 6.1, revised February 2024. Online. Available from: https://atlas-scientific.com/embedded-solutions/ezo-ph-circuit/.

Interactive rehabilitation device for hand and wrist

Dominika Kubová
Department of Biomedical Engineering
FEEC, Brno University of Technology
Brno, Czech Republic
xkubov09@vutbr.cz

Martin Králík
Department of Biomedical Engineering
FEEC, Brno University of Technology
Brno, Czech Republic
kralikmartin@vut.cz

Martin Mézl
Department of Biomedical Engineering
FEEC, Brno University of Technology
Brno, Czech Republic
mezl@vut.cz

***Abstract*—This paper presents an interactive rehabilitation device for hand and wrist therapy. The system provides objective, repeatable measurements during repetitive motor training after trauma, orthopedic procedures, or neurological impairments. Grip force is measured using force-sensitive resistor (FSR) sensors, while wrist motion is tracked by an inertial measurement unit (IMU), and signal acquisition and control are implemented on an ESP32-class microcontroller platform. The work describes the complete sensing-to-application pipeline, including embedded processing, wireless communication, and real-time visualization for therapeutic tasks. A practical implementation with dedicated electronics and a handheld form factor is presented and pilot-tested. The proposed system demonstrates the feasibility of combining force and motion sensing in a compact rehabilitation tool suitable for the objective monitoring of therapy exercises.**

***Index Terms*—hand rehabilitation, grip, dynamometry, grip strength, biofeedback, force-sensitive resistor (FSR), inertial measurement unit (IMU)**

I. Introduction

Hand and wrist rehabilitation is essential after trauma, orthopedic surgery, and neurological disorders. Impairments in grip strength [1] and wrist mobility can significantly limit a patient's ability to perform daily activities. Effective rehabilitation, therefore, requires repetitive motor training focused on restoring strength, coordination, and range of motion. In clinical practice, therapists often use simple tools such as therapy balls or grip trainers to support these exercises. However, conventional tools provide little objective feedback, while many commercial smart systems are costly or difficult to adapt to specific therapeutic workflows. This lack of visible progress often leads to decreased patient motivation and poor adherence to the rehabilitation plan. Integrating biofeedback through interactive game-based applications has shown potential to transform monotonous exercises into engaging tasks, thereby improving long-term therapeutic outcomes [2].

This work develops a sensing concept that combines FSR-based grip force measurement with wrist-orientation tracking using an inertial measurement unit (IMU) in a practical device [3]. The presented system includes custom PCB design, mechanical realization of a physical prototype, and unified firmware–application communication on an ESP32 platform.

The aim of this work is to develop and validate a practical sensing system that enables objective monitoring of rehabilitation exercises. This paper describes the hardware architecture, sensor integration, and the data communication pipeline, and presents a prototype implementation designed to validate whether the integrated architecture can reliably capture both force and motion signals. For ergonomic comfort and physiologically natural hand positioning, the device uses a conical form factor that measures whole-hand grip while placing the sensors inside the device body.

II. System Design and Requirements

A. Functional Requirements

The following requirements were defined for the practical system:

- grip-force sensing in at least two contact points,
- three-axis wrist orientation estimation,
- wireless data transmission to a host application,
- low weight for repeated training,
- universal ergonomic shape (one size),
- battery-powered operation with USB charging.

In addition to these primary requirements, the platform was designed with modularity in mind. This enables firmware updates, sensor replacement, and the addition of new game scenarios without any redesign of the electronics.

From a clinical perspective, time-dependent signal behavior is as important as peak force values [4]. Therefore, the software analyzes complete exercise profiles, including ramp-up time, holding stability, and relaxation dynamics, providing therapists with a comprehensive view of the patient's neuromuscular control.

B. Complete System Block Diagram

The system architecture is organized into three functional layers: sensing, communication, and application-level feedback. The complete block diagram, illustrating the data flow from the patient to the final application, is shown in Fig. 1.

The sensing layer consists of dual FSR sensors and a six-axis IMU. Analog signals from the force sensors are sampled using the integrated Analog-to-Digital Converter (ADC) of the transmitter ESP32, while kinematic data from the IMU are acquired in parallel via a digital interface.

Wireless transmission is implemented using the ESP-NOW protocol, a wireless communication stack that operates in the 2.4 GHz band [5]. This architecture utilizes a dedicated ESP32 receiver acting as a hardware bridge (dongle) between the rehabilitation device and the PC. This design choice provides

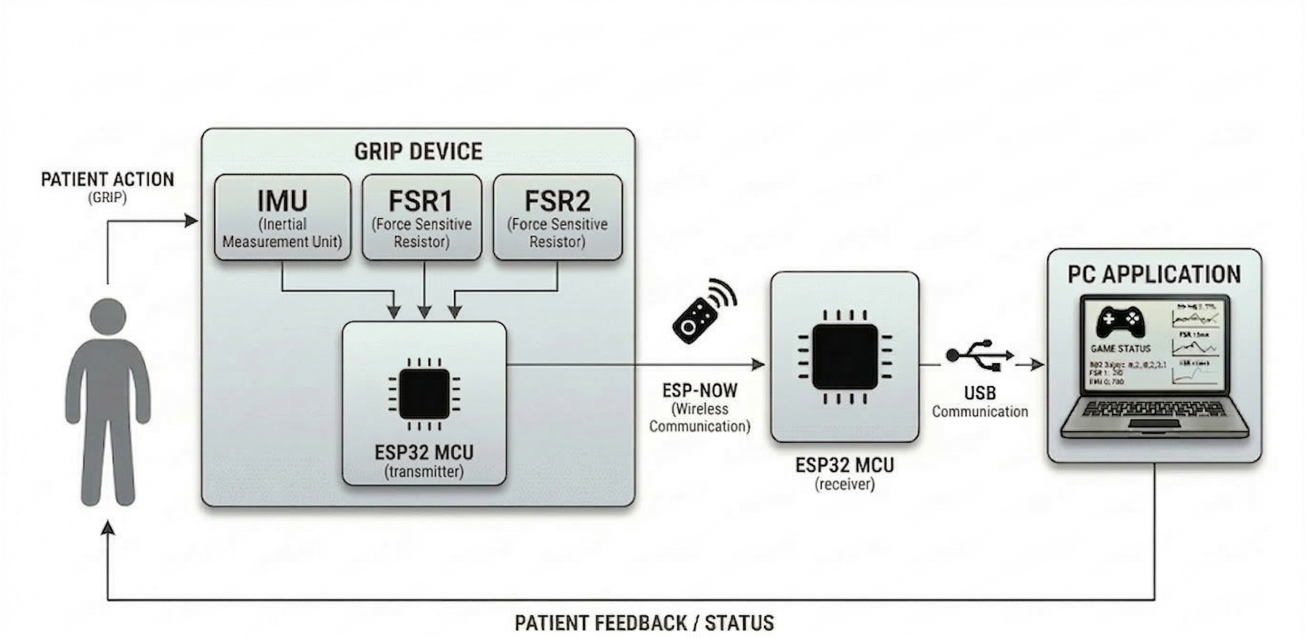

Fig. 1. Block diagram of the proposed rehabilitation system.

several advantages over standard Wi-Fi or Bluetooth Low Energy (BLE) solutions:

- Elimination of pairing complexity: The transmitter and receiver are hard-coded with unique MAC addresses, ensuring an immediate "plug-and-play" experience without the need for manual device discovery or complex pairing procedures.
- Robustness in clinical environments: By bypassing the host computer's native wireless stack and enterprise network restrictions (such as firewalls or SSID authentication), the system ensures a stable, low-latency connection.
- Streamlined data handling: The receiver forwards incoming packets via a robust UART-to-USB serial interface. This allows the host application to treat the incoming data as a standard serial stream, simplifying software development and increasing system reliability across different operating systems.

The received data are then visualized in real-time, providing the patient with immediate biofeedback while simultaneously being stored for long-term session logging and progress evaluation.

III. Practical Implementation

A. PCB Design and Electronics Integration

A dedicated PCB was designed to reduce wiring complexity, improve robustness, and enable reproducible assembly. The implemented PCB is shown in Fig. 2.

The ESP32-C3 Super Mini microcontroller module was selected due to its compact form factor, enabling integration into

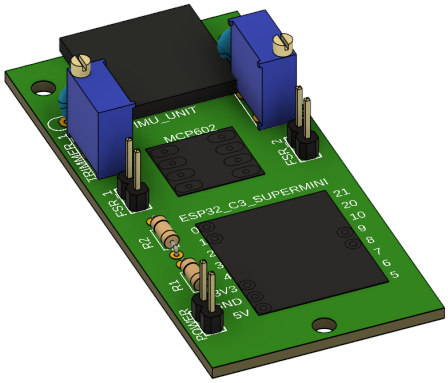

Fig. 2. Designed and assembled PCB of the rehabilitation device.

the limited internal space of the portable rehabilitation device [6]. Instead of connecting the force sensors directly to the microcontroller inputs, signal conditioning is implemented using the MCP602 operational amplifier. This specific component was selected based on the FSR manufacturer recommendations for sensor integration. The primary advantage of the MCP602 in this application is its support for single-supply operation, which eliminates the need for a symmetrical power supply and is ideal for portable, battery-powered devices. Each FSR (Tekscan FlexiForce A301) is connected to the inverting input of the operational amplifier, which is referenced to ground [7].

The non-inverting input is supplied with a stable 1 V reference voltage generated by a simple resistive voltage divider. To provide adjustable gain and sensitivity, a multi-turn cermet trimmer is placed in the feedback path. This configuration allows for precise calibration of the system response, with the trimmer currently set to 80 kΩ in accordance with the sensor manufacturer recommendations to optimize the dynamic range [8].

This approach partially compensates for the inherent non-linearity of FSR sensors and provides a more stable voltage response across the useful measurement range. A dedicated measurement series was performed to derive calibration curves, which were then used to convert measured voltage values to equivalent mass values in grams. Based on the obtained curves, the variability between individual FSR sensors was found to be negligible. The resulting analog voltage is sampled by the analog-to-digital converter (ADC) of the ESP32-C3 Super Mini, where the signal is digitized and transmitted as raw measurement data for further processing in the host application.

Wrist motion sensing is implemented using a BMI270 inertial measurement unit that communicates with the ESP32-C3 Super Mini via the I^2C interface. This six-axis MEMS sensor was selected for its ultra-low current consumption and integrated noise reduction features, offering 16-bit resolution for both the accelerometer and gyroscope to detect subtle motor changes during rehabilitation [9]. From a spatial arrangement perspective, the IMU was positioned on the PCB to reside as close as possible to the geometric center of the conical device body.

Power management is implemented using a single-cell 18650 Li-Ion battery with a nominal voltage of 3.7 V and a capacity of 3000 mAh. Charging and voltage stabilization are handled by an IP5306 SoC, which utilizes the CC/CV (Constant Current/Constant Voltage) charging method with a maximum current of 2.1 A [10]. This integrated circuit features a built-in boost converter that steps up the battery voltage to a stable 5 V rail to power the microcontroller and peripheral sensors [10]. Furthermore, the architecture incorporates comprehensive protection against overcharging, deep discharge (below 3.0 V), and short circuits.

From a manufacturing perspective, the board was designed using commonly available components to keep prototyping costs low and simplify replication in educational and clinical research environments.

B. Device Model and Mechanical Realization

A conical form factor was selected following consultations with an orthopedic specialist to ensure optimal ergonomics and a physiologically natural hand fit during repeated grip exercises. The geometry, with base diameters of 6 cm and 3.5 cm and a total length of 15 cm, was designed as a universal size suitable for patients with both small and large hands. The realized physical model is shown in Fig. 3.

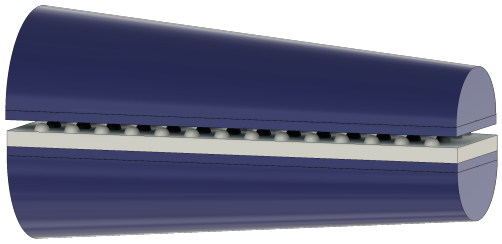

Fig. 3. Physical model of the rehabilitation device.

The enclosure is fabricated using 3D printing in PLA and is divided into two rigid halves that house the internal electronics. The internal architecture is organized such that one half contains the main PCB and sensing units, while the second half houses the battery and power management circuitry, allowing for balanced weight distribution. Between these halves, a flexible intermediate layer is integrated and 3D-printed from Prusament Resin Flex80 to provide a realistic compressive feel during grip training, giving the patient the tactile sensation of actually squeezing a compliant object. This flexible layer is represented by the white section in the model and utilizes a structured surface to enhance tactile feedback. The internal geometry is engineered so that the compression of this layer acts directly upon the active areas of the FSR sensors for accurate force detection.

IV. Conclusion

This paper presents a practical implementation of an interactive rehabilitation system for hand and wrist therapy. The proposed solution integrates a custom-designed PCB, embedded sensing electronics, and a mechanically optimized handheld device into a compact platform capable of capturing both grip force and wrist motion during therapeutic exercises. The system combines force-sensitive resistors for grip-force measurement with inertial sensing for wrist orientation tracking, enabling simultaneous monitoring of multiple parameters relevant to rehabilitation training.

The developed hardware architecture emphasizes compact dimensions, modularity, and reliable signal acquisition during repetitive hand movements. By integrating the sensing electronics directly into a dedicated PCB and using a compact ESP32-C3 microcontroller platform, the device achieves a portable form factor suitable for practical rehabilitation scenarios. The implemented signal conditioning using an operational amplifier improves the stability of FSR measurements and supports consistent acquisition of force profiles during exercise.

The next development step focuses on the software layer of the system. A desktop application is being developed to provide graphical visualization of measured signals and to support interactive rehabilitation tasks. In addition to real-time feedback, the application will include a game-based training

environment designed to improve long-term patient motivation and encourage consistent rehabilitation practice.

In future work, the device will be tested on a larger group of participants to verify its performance under practical rehabilitation conditions. This evaluation will focus on measurement consistency, usability, and user acceptance in both clinical and home-based settings. Based on these results, the system can be further extended with additional motion-analysis functions, richer exercise tracking, and adaptive training scenarios derived from measured performance data.

ACKNOWLEDGMENT

I would like to express my sincere gratitude to Assoc. Prof. Tomáš Tomáš, M.D., Ph.D., Head of the First Department of Orthopedic Surgery, St. Anne's University Hospital Brno (FNUSA), for expert consultations on the ergonomic design of the device. Special thanks also go to Ing. Vratislav Harabiš, Ph.D., for his assistance with the 3D printing and fabrication of the device enclosure.

REFERENCES

[1] A. S. Lu, S. J. Swartz, and J. G. Keefe, "Development and validation of an isometric handgrip strength measurement system using a thin-film force sensor," *BMC Musculoskeletal Disorders*, vol. 16, p. 156, Jun. 2015, doi: 10.1186/s12891-015-0612-4.

[2] M. Gandolfi, N. Valè, E. Dimitrova, and et al., "Effectiveness of a hand-bike circuit training on upper limb motor function and quality of life in persons with chronic spinal cord injury: A randomized controlled trial," *Brain Sciences*, vol. 11, no. 1, p. 45, Jan. 2021, doi: 10.3390/brainsci11010045.

[3] C. S. Nam, J. W. Kim, and S. H. Park, "A wearable hand rehabilitation system using force-sensing resistors and inertial measurement units," *Sensors*, vol. 23, no. 22, p. 9047, Nov. 2023, doi: 10.3390/s23229047.

[4] T. Silva-Santos, R. S. Guerra, R. Valdiviesso, and T. F. Amaral, "Hand grip force–time curve indicators evaluated by dynamometer: A systematic review," *Nutrients*, vol. 16, no. 12, p. 1951, Jun. 2024, doi: 10.3390/nu16121951.

[5] Espressif Systems, "ESP-NOW user guide," 2024. [Online]. Available: https://docs.espressif.com/projects/esp-idf/en/latest/esp32/api-reference/network/esp_now.html. [Accessed: Mar. 14, 2026].

[6] Espressif Systems, "ESP32-C3 Super Mini datasheet," [Online]. Available: https://dl.artronshop.co.th/ESP32-C3%20SuperMini%20datasheet.pdf. [Accessed: Mar. 12, 2026].

[7] Tekscan, "FlexiForce A301 datasheet, Rev. J," 2024. [Online]. Available: https://www.mouser.com/catalog/specsheets/Tekscan_FLX_A301_2_0Datasheet_RevJ.pdf. [Accessed: Mar. 12, 2026].

[8] Tekscan, "Best practices in electrical integration of the FlexiForce sensor," 2024. [Online]. Available: https://lp.tekscan.com/hubfs/FLX/FLX-Best-Practice-Electrical-Integration_2.0.pdf. [Accessed: Mar. 12, 2026].

[9] Bosch Sensortec, "BMI270 smart low-power 6-axis motion sensor," Rev. 1.5, 2023. [Online]. Available: https://www.laskakit.cz/user/related_files/bst-bmi270-ds000.pdf. [Accessed: Mar. 14, 2026].

[10] Injoinic Technology Co., Ltd., "IP5306," Shanghai, China, 2016. [Online]. Available: https://www.laskakit.cz/user/related_files/ip5306.pdf. [Accessed: Mar. 14, 2026].

Nastartuj svou kariéru

e·on



Studentské programy

- Letní energetická akademie – poslední týden v červenci
- Letní brigády – od června do září
- Stipendijní program
- Diplomové/bakalářské práce
- Praxe a brigády

Absolventské programy

- Trainee program
- Junior technik

E.ON Kariéra

eon.cz/kariera

Design of a Pulsed Laser Diode Driver for Optical Fiber Sensor Measurement

Matej Grega
Department of Microelectronics, Faculty of Electrical Engineering and Communication
Brno University of Technology
Brno, Czech republic
240839@vut.cz

***Abstract*—Fiber optic sensor systems often require pulsed excitation of a laser source. This paper presents the design of a laser diode (LD) driver capable of generating current pulses with arbitrary shape, amplitudes up to 2 A, and durations up to 10 µs. The proposed driver achieves a settling time of 110 ns to within 2 % without current overshoot and communicates with the master system through a digital interface. The design emphasizes low-noise operation and high stability and repeatability of the generated pulses. To ensure safe LD operation, the module integrates a modified overcurrent protection (OCP) mechanism that evaluates the thermal stress of the LD during the pulse and disconnects the LD under excessive conditions. The implemented driver was experimentally characterized, and the results confirm its suitability for integration into a fiber optic sensor evaluation unit.**

***Index Terms*—laser diode driver, current source, pulsed current generation, set-point kick suppression**

I. Introduction

Fiber-optic measurement systems are widely utilized in numerous applications, including structural health monitoring [1], gas detection in explosion- or fire-hazardous environments [2], monitoring seismic activity during underground gas storage and extraction [3], and surveillance of protected areas, power lines, and product pipelines. These systems typically employ a laser diode as a monochromatic source of optical radiation due to its compact size and ease of integration into the optical evaluation unit. A laser diode is a semiconductor device characterized by a steep volt–ampere characteristic, with its emitted optical power proportional to the drive current. This behavior necessitates driving the diode with a current source to maintain a stable operating point.

This paper presents the design of a laser diode (LD) current driver module for use in an optical fiber sensor evaluation unit, which requires driving the LD with short pulses ranging from 1 µs to 10 µs and amplitudes of up to 2 A. The pulse shape is arbitrary, necessitating digital control of the instantaneous current. The system requires a minimum current amplitude resolution of 1 mA, a time resolution of at least 100 ns, and no current overshoot when a new value is set. For reliable operation in measurement systems, the current source must maintain stable performance under such pulsed conditions across varying temperatures, with a targeted settling time of no more than 200 ns. This settling time corresponds to the bandwidth of the current source, and further reduction is deemed unnecessary given the specific application for which the module is designed.

Development of the presented laser diode driver module was financially supported by the NETWORK GROUP, s.r.o.

II. Design of The Current Source

Conceptually, the current source can be implemented using several topologies. A simple solution with the minimum number of components in the signal path was sought to maximize the achievable dynamic and noise performance. For this reason, a current sink topology with output current sensing via a ground-referenced shunt resistor was selected. During the design of the current sink control and its feedback loop, particular attention was paid to dynamic performance, especially to the response of the output current to a step change in the control signal, which must occur without overshoot or transient spikes.

A. Controller Topology and Implementation

The concept of a controller based on an operational amplifier (op-amp) is illustrated in Fig. 1 a). As described in [4], step changes in the control signal in such a topology result in spikes in the regulated current referenced to as set-point kick. This effect arises from the simultaneous change in voltage at the op-amp output and at its inverting input when the control signal changes abruptly. Consequently, the integrating capacitor experiences only minimal charging or discharging during this transition and therefore does not effectively suppress rapid changes in the output signal.

Undesirable set-point kick can be eliminated by employing a controller with a summing op-amp configuration, as shown in Fig. 1 b). In this arrangement, both the control and feedback signals are applied to the same summing node, ensuring that the response to a change in the control signal is identical to the response to a change in the feedback signal. During operation, both inputs of the op-amp remain at ground potential. Consequently, any change in the op-amp output voltage is associated with charging or discharging of the integrating capacitor, which effectively suppresses transient spikes.

The topology mitigating the effect of set-point kick is also employed in the proposed LD driver module. The schematic of the driver topology principle is shown in Fig. 2. A drawback of the summing op-amp topology is the inversion of the input

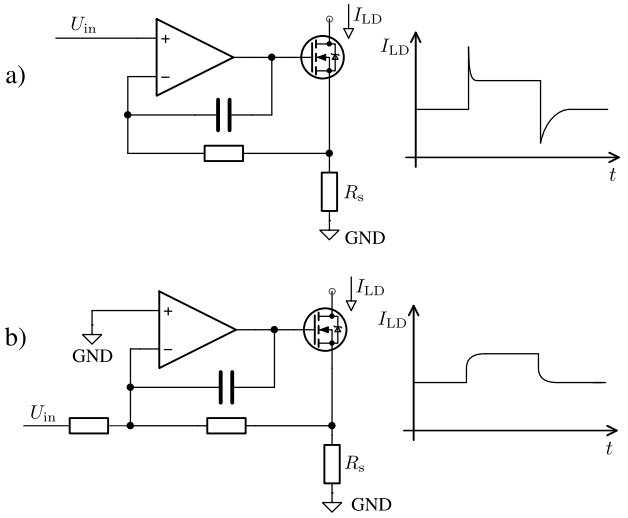

Fig. 1. a) Current source topology utilizing non-inverting op-amp PI controller. b) Current source topology mitigating the effect of set-point kick.

signals; therefore, proper operation requires an inverted control signal. To address this, the reference potential of the input signal $V_{\mathrm{I_set}}$ is shifted from ground to a reference voltage of 2 V. With this modification, a zero input voltage corresponds to the maximum output current, which decreases with increasing input voltage and reaches 0 A at $V_{\mathrm{I_set}} = 4$ V. Additional level shifting is implemented in the digital-to-analog converter (DAC) stage to ensure zero LD current at a zero DAC code.

B. Dynamic Control Loop

The output current is sensed using a shunt resistor R_{s} referenced to ground. A common approach is to employ a low-value shunt resistor with minimal voltage drop and amplify the resulting voltage using a differential amplifier. However, this solution is not suitable in the present case because the current must be regulated with a rise time of approximately 70 ns. Achieving this would require an amplifier with a bandwidth on the order of tens of megahertz at gains ranging from tens to hundreds. Such a design would in turn require an operational amplifier with a gain–bandwidth product in the hundreds of megahertz to several gigahertz. Moreover,

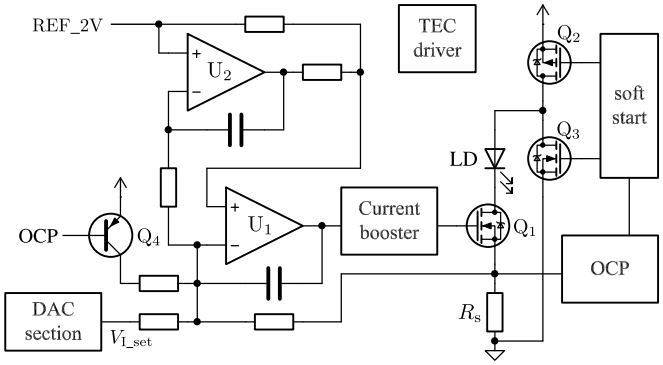

Fig. 2. Schematic diagram of the LD driver topology

an amplifier operating with such high gain and bandwidth would introduce significant noise in the feedback loop, which would be directly transferred to the output current. For these reasons, the feedback signal representing the sensed current is obtained directly from the shunt resistor R_{s} without additional amplification. Consequently, the shunt resistor must provide a sufficiently large voltage drop while simultaneously defining the amplitude of the control signal. A shunt resistor value of 1 ohm was selected, corresponding to a control signal $V_{\mathrm{I_set}}$ range of 2–4 V.

The core of the current regulation loop is the operational amplifier U_1 shown in Fig. 2. Its function is to respond to rapid changes in the set current while introducing minimal noise. Several candidate devices were evaluated using LTspice simulations and experimentally verified with optimized feedback components to identify the most suitable configuration. The selected operational amplifier features a gain–bandwidth product exceeding 150 MHz, a slew rate greater than 200 V/µs, and an input voltage noise density below 6 nV/$\sqrt{\mathrm{Hz}}$. With this configuration, the driver achieves a rise time (10–90 %) of 40 ns and a fall time of 70 ns when generating a 1 A current pulse.

The LD current is regulated by MOSFET Q_1, which must be carefully selected for pulsed-current operation. Particular attention was given to transient behavior, as the MOSFET gate must be charged and discharged rapidly to preserve the dynamic performance of the control loop. In pulsed operation, achieving the required switching speed typically requires gate currents on the order of hundreds of milliamperes. The selected device therefore features a low input capacitance of 195 pF and a total gate charge of 4.5 nC to minimize this requirement; however, the required drive current still exceeds the output capability of typical operational amplifiers. Consequently, a discrete class-AB follower stage was employed as a current booster.

C. Current Booster Stage

The class-AB follower was designed to ensure that it does not introduce significant additional noise into the current control loop. This required careful selection of the transistors and their operating point. The noise performance of bipolar junction transistors varies among devices, particularly due to base resistance r_{B}. Paper [5] provides a relation for determining the optimal collector current for minimum noise

$$I_{\mathrm{C_opt}} = \frac{V_{\mathrm{T}}}{R_{\mathrm{s}} + r_{\mathrm{B}} + R_{\mathrm{E}}} \cdot \frac{\beta}{\sqrt{1+\beta}}, \tag{1}$$

which was used to determine the operating point of the selected complementary transistor pair with base resistances of 40 Ω and 17 Ω. The source output resistance driving the base R_{s} and the external emitter resistance R_{E} were determined through LTspice simulations.

D. Static Control Loop

In addition to dynamic performance, the stability of the optical source under static conditions is critical, as it directly

affects measurement repeatability. Consequently, all components involved in optical signal generation need to maintain long-term stability and exhibit minimal temperature dependence. The most critical passive component of the presented module is the shunt resistor determining the LD current magnitude, since its temperature coefficient directly influences the absolute current value. To minimize this effect, a resistor with a temperature coefficient of resistance (TCR) of 2 ppm was selected.

The control and feedback signals are combined through a resistive network, where maintaining a stable resistance ratio is essential for accurate operation. For this reason, resistors with a TCR of 10 ppm were used in this network and placed in close proximity on the PCB to ensure good thermal matching.

A significant source of error is the input offset voltage of operational amplifier U_1, particularly its temperature dependence (offset drift). The requirement for a high-speed op-amp in the feedback loop limits the achievable offset and drift performance; for the selected device, these parameters may exceed ±10 mV and ±6 µV/°C, respectively. To mitigate this limitation, a composite configuration was implemented in which the offset of the fast op-amp U_1 is compensated by a precision operational amplifier U_2 with a negligible offset of 0.25 µV and an offset drift of approximately 5 nV/°C.

Operational amplifier U_2 compares the voltage at the inverting input of U_1 with the reference potential REF_2V, which should ideally be identical. Any detected deviation is used to adjust the voltage at the non-inverting input of U_1, thereby compensating for the offset. This arrangement introduces a secondary feedback loop within the current source. To maintain stability, this loop must operate significantly slower than the primary current-regulation loop. At the same time, complete offset compensation of U_1 is required; therefore, U_2 is configured as an integrator. Because the compensation signal has a small amplitude, the output of U_2 is applied through a resistive divider and additionally filtered by a capacitor (not shown in the Fig. 2) to minimize noise.

E. Over Current Protection

Laser diodes operated in pulsed regimes may experience instantaneous currents exceeding their nominal continuous ratings; therefore, an effective protection mechanism is required to prevent device damage. The primary failure mechanism is thermal breakdown of the junction; consequently, the protection scheme must account for the thermal processes occurring within the laser diode. The junction temperature increases with the energy delivered during the current pulse and decreases as heat is dissipated to the surrounding environment.

To address this behavior, a modified OCP scheme based on monitoring the delivered energy was developed. The protection system is designed to satisfy the following requirements:

- evaluation of the delivered energy, approximating the LD junction temperature during a pulse,
- integration of the energy delivered by the drive current,
- subtraction of energy corresponding to cooling when the LD is not significantly heated,
- latching of the OCP in the fault state once the energy limit is exceeded,
- edge-triggered electronic reset to ensure proper operation independent of the reset signal level.

The OCP monitors the LD current as the voltage across a shunt resistor R_s, which is converted into a proportional current used to charge an integrating capacitor. The capacitor voltage is then evaluated to determine whether the protection threshold has been reached. Although the LD continuously dissipates thermal energy to the environment, a decrease in junction temperature occurs only at sufficiently low drive currents. Therefore, when the LD current falls below a threshold of 15 mA, the integrating capacitor is slowly discharged to represent the cooling process.

1) Soft Start and Power Switching: To ensure stable LD operation during current pulses, a bank of ceramic capacitors with a total capacitance of 100 µF is placed close to the LD connected to its supply line. While this capacitance stabilizes the supply, it also stores sufficient energy to potentially damage the LD. Therefore, the designed LD power switch rapidly discharges the capacitor bank when the OCP is activated. A dead time of approximately 0.2–0.3 µs is implemented to prevent cross-conduction between transistors Q2 and Q3.

A soft-start mechanism is used to gradually charge the capacitor bank when the LD supply is enabled. The capacitors are charged with a current of 120 mA, reaching at least 99 % of the nominal voltage within 47 ms.

F. Digital-to-Analog section

The control signal for the LD current source is generated by a digital-to-analog converter (DAC), enabling arbitrary waveform generation. The DAC selection was driven by several requirements: sufficient resolution to control the output current with steps not exceeding 1 mA, an update rate comparable to the current-source rise time ($\geq$10 MHz), and support for an external reference voltage. The use of an external reference ensures that all reference voltages within the LD driver module are derived from a common source and therefore exhibit identical temperature dependence, providing higher stability than a DAC with an internal reference. Based on these requirements, a 12-bit DAC (LSB = 0.5 mA) with a 20 MHz update rate was selected.

The DAC output current is converted to a voltage and level-shifted to meet the requirements of the current source, generating the control signal $V_{\mathrm{I_set}}$ corresponding to the desired output current.

G. Temperature stabilization of Laser Diode

The presented module is designed as a LD driver, integrating all necessary functions to ensure stable operation under varying conditions, including temperature stabilization of the LD. A common approach for temperature control is the use of a thermoelectric cooler (TEC) in the form of a Peltier module, coupled with the LD and thermally insulated from the ambient environment.

The temperature stabilization circuit typically comprises a full H-bridge for both heating and cooling via the TEC, a temperature-sensing element for the LD, and an analog or digital controller. In the proposed measurement unit, a TEC driver is implemented as a compact PCB module (System-on-Module, SoM) with dimensions of $21 \times 12.4 \times 3.1$ mm^3. The module supports TEC drive currents up to 1.5 A and incorporates a digital PID controller configurable via a UART interface.

III. Measurement and Testing

The static and dynamic parameters of the designed LD driver module were measured to verify its suitability for the fiber optic sensor evaluation unit. Measurements were performed in two operating modes: static and dynamic.

In the static mode, the temperature dependence of the set LD current was evaluated. The measurement was conducted with the LD current set to 10 mA, as higher currents than 15 mA are limited by the OCP. Within the temperature range of 15–25 °C, the temperature coefficient of the set current was measured to be 0.15 %/°C.

Dynamic measurements focused on the response of the regulated LD current to changes in the digitally programmed current amplitude. The parameters were characterized using a rectangular LD current pulse with an amplitude of 2 A and a duration of 10 µs. Fig. 3 shows the LD current and the control voltage $V_{\mathrm{I_set}}$ driving the current source during the pulse. The measured performance parameters of the LD driver are summarized in Table I.

The soft-start and OCP functions were further verified using a 1 A current pulse with a duration exceeding the OCP limit. Correct activation of the protection mechanism and shutdown of the LD power supply were observed. Following OCP triggering, the LD supply voltage decreased from the nominal 12 V to below 1 V within 1.2 µs. This interval includes a response time of 240 ns and an additional 250 ns dead time between disconnecting the LD from the supply and actively shorting the supply node to ground.

IV. Conclusion

This paper presented the design of a LD driver module intended for a fiber optic sensor evaluation unit that employs pulsed current excitation of the LD. The proposed module is capable of generating arbitrary current pulses with amplitudes up to 2 A and durations of up to 10 µs. To protect the LD from thermal damage, the module integrates a modified OCP mechanism. Particular attention was devoted to ensuring the low noise, stability and repeatability of the generated pulses during the design process.

TABLE I
Measured parameters of LD driver module

parameter	value
current consumption 5 V (excluding TEC)	85 mA
current consumption –5 V	9 mA
LD current rise time (10–90 %)	50 ns
LD current fall time (90–10 %)	48 ns
LD current pulse settling time (2 %)	110 ns

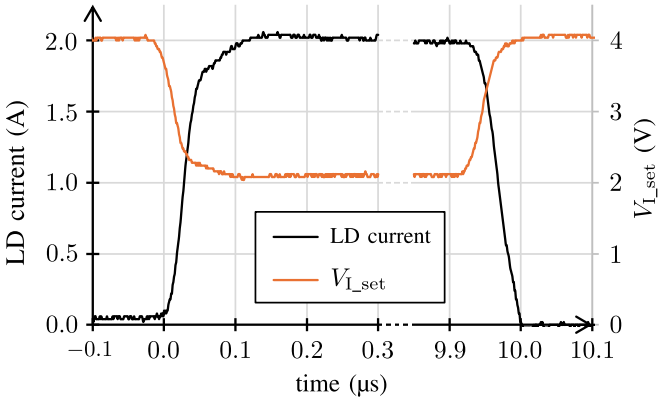

Fig. 3. Rising and falling edge of a 2 A 10 µs LD current pulse

The LD driver module was constructed, experimentally characterized, and tested under operating conditions representative of the intended measurement system, confirming the suitability of the proposed design. The measured dynamic performance, including a settling time of 110 ns to within 2 % during a 2 A LD current pulse with a rise time of 50 ns without an overshoot, demonstrates the effectiveness of the selected topology for pulsed operation.

The measured temperature coefficient of the LD current was 0.15 %/°C, which is acceptable for the target application where temperature variations are limited; however, further reduction remains a potential area for future improvement. Additional performance gains may also be achieved by optimizing the DAC section, which was identified as the primary factor limiting the overall settling time of the LD current.

V. Acknowledgement

I would like to express my sincere gratitude to my supervisor doc. Ing. František Urban, CSc. and Ing. Radek Helán, Ph.D. for their professional advice and guidance throughout this project.

References

[1] A. N. D. Alhussein, M. R. T. M. Qaid, T. Agliullin, B. Valeev, O. Morozov, and A. Sakhabutdinov, “Fiber Bragg Grating *Sensors*: Design, Applications, and Comparison with Other Sensing Technologies”, Sensors, vol. 25, no. 7, Apr. 2025, doi: 10.3390/s25072289.

[2] T. He, W. Wang, B. -G. He, and J. Chen, “Review on Optical Fiber Sensors for Hazardous-Gas Monitoring in Mines and Tunnels”, *IEEE Transactions on Instrumentation and Measurement*, vol. 72, 2023, doi: 10.1109/TIM.2023.3273691.

[3] M. Johnson Singh et al., “Applications of fibre Bragg grating sensors for monitoring geotechnical structures: A comprehensive review”, *Measurement*, vol. 218, 2023, doi: 10.1016/j.measurement.2023.113171.

[4] M. Kopka, *Elektronická zátěž s digitálním řízením* [Electronic Load with Digital Control], Bachelor’s thesis, Brno Univ. Technol., Faculty of Electrical Engineering and Communication, Dept. Microelectronics, Brno, Czech Republic, 2024. [Online]. Available: https://www.vut.cz/studenti/zav-prace/detail/160210

[5] W. Marshall Leach, “Fundamentals of low-noise analog circuit design”, *Proceedings of the IEEE*, vol. 82, no. 10, Jan. 1994, doi: 10.1109/5.326411.

Augmented reality for assisted testing of radar sensors

1st Miloslav Kužela
Department of Radio Electronics
FEEC, Brno University of Technology
Brno, Czech Republic
240648@vut.cz

2nd Tomáš Frýza
Department of Radio Electronics
FEEC, Brno University of Technology
Brno, Czech Republic
fryza@vut.cz

***Abstract*—This paper describes the possibility of utilizing an *Augmented Reality* (AR) headset for assisted data collection for training datasets of radar sensors used in automotive applications, that are used for detecting movement underneath a rear bumper to automatically open the trunk of a car. It delves into a proposal of an application developed for a Meta Quest 3 headset, using a Unity game engine. The application is designed to simplify the data collection process by providing guidance and visualization for an untrained user. This would also allow for a more efficient way of collecting data, as it would free up human resources due to the current requirement of a trained operator to be present during the data collection process. The main focus of this paper is in the detection of a bumper using the headset's cameras and future plans for the finalized application.**

***Index Terms*—Automotive, Augmented Reality, Radar Sensors, Data Collection, Virtual Reality, Machine Vision**

I. Introduction

An accurate and comprehensive dataset is crucial for training machine learning models. The constant requirement for larger quantity of data makes any simplification or automatization of the data collection process highly desirable. The company Alps Alpine [1], which the project in this paper is for, manufactures radar modules which are designed to be installed in a car's rear bumper to allow for a hands free opening of the trunk. Such a radar outputs large quantities of data, which is ideal to use with machine learning algorithms, but the collection of data, with sufficient quality, can be, at times, difficult. Such a process contains a lot of different scenarios with specific movements and locations. This in turn requires a trained operator to be present when trying to collect data from untrained users. Due to the recent popularity of *Augmented Reality* (AR) and the appearance of AR capable headsets an idea of using such a device to help with the data collection process came up. This could help simplify the data collection process and improve its time effectiveness.

II. Existing solution and Motivation

The training is currently organized and run by a minimum of two personnel. One trained operator is responsible for guiding an untrained user that is providing the data while operating the data collection software as well as managing a spreadsheet with file categorization. This process is time consuming as the number of different scenarios, which includes different kicking movements and locations which, depending on the car manufacturer, can become quite large. This also requires a visual guide on the ground indicating the precise positions relative to the sensor, which is usually done by pieces of tape as seen in Fig. 1. As a results the length of such a session can last longer than an hour. The collection process sometimes needs to be done even for a car model that has yet to be announced or manufactured. The only data available are *Three Dimensional* (3D) models of the car's bumper, or just the specific mount for the radar. This then results to only a small part of a bumper, which makes the data collection difficult as the person that provides the movement data cannot see the entirety of the bumper and in reality kicks only into the free air.

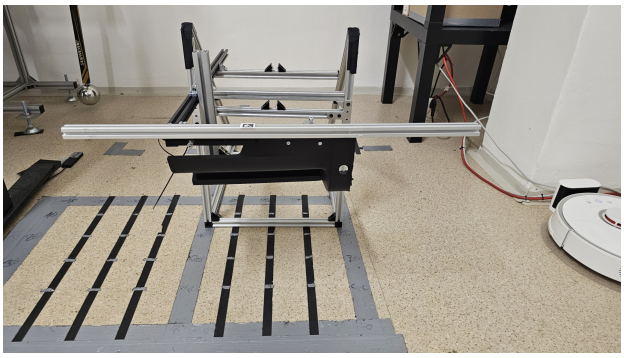

Fig. 1. Testing setup with a small part of a bumper.

III. Proposed solution

A. Hardware

The Meta Quest 3 is a standalone *Virtual reality* (VR) and AR headset manufactured by Meta [2]. It is equipped with a *System On a Chip* (SoC) which allows it to run applications without the need of another computer [3]. To enable artificial reality the headset has 2 cameras on the front which are accessible to applications, allowing them to overlay virtual objects on the real world. There are multiple development platforms that can be used, with Meta providing the necessary *Software Developments Kits* (SDKs) and documentation for each of them.

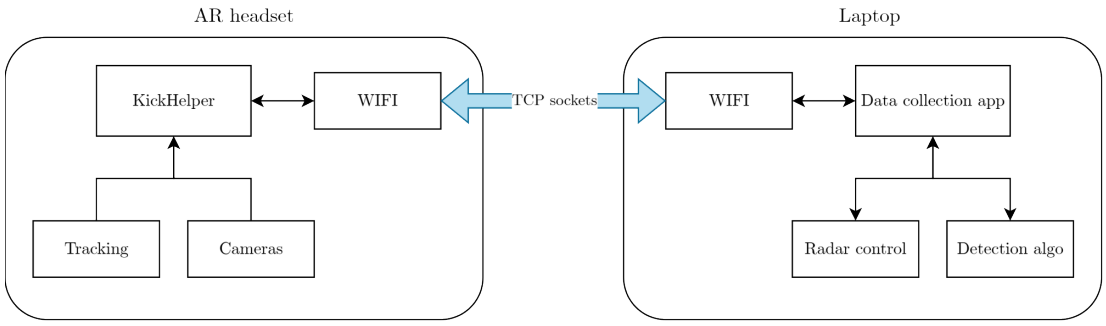

Fig. 2. Application system diagram.

B. Software

The application is developed using a Unity game engine. This engine uses C# language as its main scripting language and allows for a simple scene organization and development using its graphical interface [4]. The application is supposed to detect a predefined marker and load a 3D model of a bumper and place it in the correct position and orientation, while also providing a visual guidance and instruction for the user to follow. This required a *graphical user interface* (GUI) that allows the user to interact with the application and set basic variables. It is planned that the application would also communicate, with a data collection computer that is connected to the radar sensor, using a WiFi connection. The communication would be used to send short messages to and from the headset, such as feedback from the movement detection algorithm. A diagram of the proposed system is in Fig. 2. Such a system is expected to be using a custom TCP protocol for communication, but the details of the protocol are still being worked on.

IV. Detection problem

A. Consideration of Machine vision

Due to the direct access to the front cameras of the headset, the use of machine vision was considered as a possible solution to the problem of detecting a bumper in a scene. The first consideration was to use an object detection algorithm, this proved as an ineficient and resource intensive solution that would not achieve the goal.

B. Use of markers

Using a marker that can be placed in the environment proved to be more suitable for this application. As the markers can be easily printed and placed and then later detected using a simple algorithm that calculates the position relative to the camera.

C. AruCo Markers

An AruCo marker, seen in Fig. 3, is a visual pattern that can be detected by a computer vision algorithm. It consists of a black square with a unique identifier thanks to a white pattern inside the square. These markers can be easily generated online or by using a compatible machine vision library, such as OpenCV. This solution has been thought of after noticing the work of Takashi Yoshinaga [5]. The advantage of multiple identification patterns allows for assigning different markers to different testing scenarios, such as loading a different 3D model according to the marker's ID.

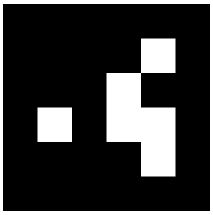

Fig. 3. A 4x4 AruCo marker with an ID of 123.

D. Perspective-n-Point pose estimation

Perspective-n-Point (PnP) is one of the methods used for calculating the position and orientation of a camera in a 3D space, based on a set of points in the 3D world. It is commonly used in computer vision and robotics applications, where it is necessary to determine the position and orientation of a camera or other sensor relative to a known set of points in the environment [6]. The algorithm is a mathematical equation consisting of a set of 3D points, their corresponding 2D projections in the camera image, and a camera matrix, that is also known as the intrinsic matrix. Containing information about the camera, like its focal lengths and the optical center. Lastly distortion coefficients can also be used to account for lense distortions. This equations can then be used to solve for the pose of the 3D objects if we know the camera parameters and it's pose. Example of a simple PnP can be seen in Fig 5.

V. Averaging

When calculating the position and rotation of the marker, the algorithm is affected by noise and inaccuracies. Such inaccuracies can be cause by the camera's resolution, the distance of the marker from the camera, lighting conditions and the constant movement of the user. To reduce the effect of such inaccuracies, the algorithm is not using the position it calculates from the first image frame where the marker is located, but instead it acquires multiple frames and averages

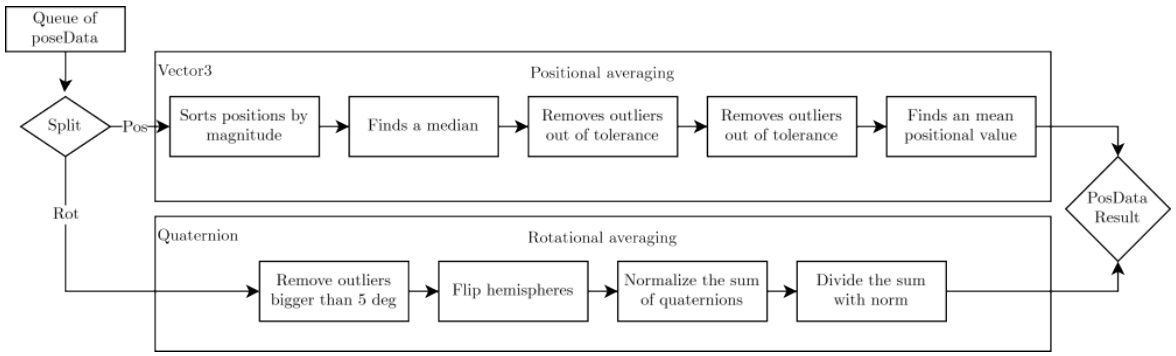

Fig. 4. A diagram describing the averaging process of the positional data.

the positional and rotational data. This in turn results into a more stable and accurate calculated position. The entire process of averaging is described in a diagram see in Fig. 4.

A. Rotation

As the rotation is Unity is represented as a quaternion, the averaging process is not as straight forward as with the position, which is represented as a vector. In the case of this paper the averaging is done by normalizing the sum of all quaternions, which proved to be a simple and an effective solution for such an application. The process is of normalization is described in in in the following equation (1).

$$\begin{aligned} norm &= \sqrt{\sum {q_x}^2 + \sum {q_y}^2 + \sum {q_z}^2 + \sum {q_w}^2} \\ q_{avg} &= \left(\frac{\sum q_x}{norm}, \frac{\sum q_y}{norm}, \frac{\sum q_z}{norm}, \frac{\sum q_w}{norm} \right) \end{aligned} \tag{1}$$

Where q_n are the individual components of the quaternions that are being averaged, q_{avg} is the resulting average quaternion and $norm$ is the calculated normalization factor.

VI. Testing of accuracy

A test was done to determine the accuracy of the proposed detection and averaging algorithm. They were created by making a simple testing pattern. The pattern consisted of 4

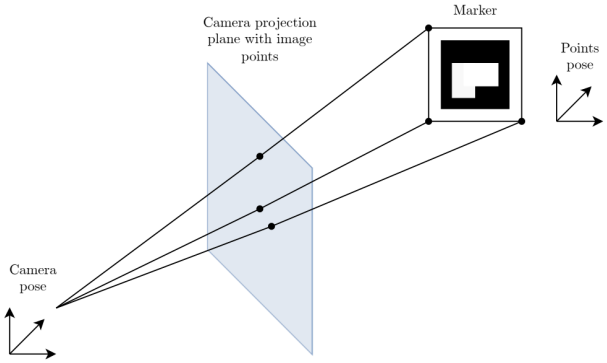

Fig. 5. A simple representation of a PnP problem.

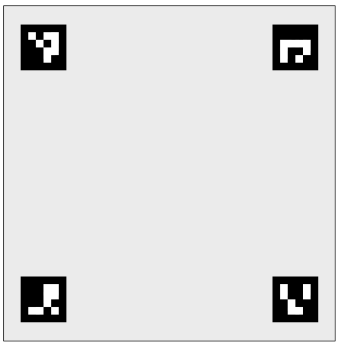

Fig. 6. A testing pattern with 4 markers spaced 10 cm apart.

markers spaced exactly 10 cm apart positioned in a square. Each marker had a different ID. This pattern can be seen in 6. The headset was positioned 100cm away from the center of the pattern and the position of each marker was calculated in two different ways. One where the 50th frame was used and an another one with 100 positions averaged.

A. Cameras

As previously mentioned the headset features two camera streams which can be accessed from an application. Each of these cameras have a max resolution of 1280x1280 with a *Peak Pixel Density* (PPD) of 25, which represents the number of pixels displayed per degree in the center of the user's vision [7].

B. Testing results

Relative positions of the markers were calculated and compared to the real positions on the testing pattern. The results are in a table VI-B. The rotations were expected to be the same across all of the markers. The results show that the proposed averaging process helps with positional error and slightly improves the rotational error as well. Compared to a simulated modeled scenario [8], where the error is expected to be around 10mm the results are deemed acceptable.

VII. Application prototype

The application is still in a development stage, but a basic prototype has been created. It is able to detect a marker and

TABLE I
COMPARISON OF POSITIONAL AND ROTATIONAL ERROR BETWEEN METHODS

Method	Positional error [mm]	Rotational error [°]
50th sample	15.8	4.91
Averaging	11.1	3.55

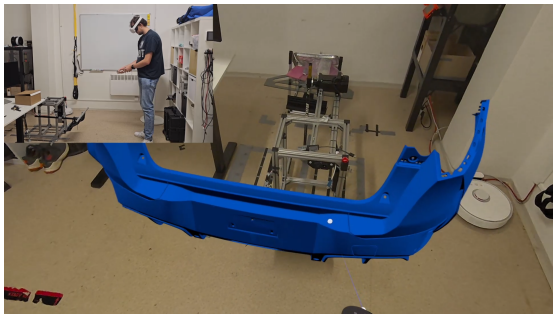

Fig. 7. Point of view from the headset with an overlaid third person view.

load a 3D model of a bumper in real time and place it in a correct position and orientation. An example image can be seen in Fig. 7. The user is guided using a 3D intractable GUI and a floating text that provides basic instructions and messages. This *User interface* (UI) can be seen in Fig. 8. The UI is dynamic, changing text and elements depending on the state as well as featuring a config selection window that loads and displays available configuration files.

A. Detection process

The detection process is encapsulated in a single coroutine that runs in the background when the user starts the detection. If this happens the camera feed is opened and the algorithm start to look for a marker, as soon as a marker is detected the software waits until it collects a sufficient amount of positional data, this data is then processed and averaged using an algorithm described in a previous chapter. The final position and orientation is then used to place a 3D model that is loaded from the devices storage after applying offsets defined in a configuration file.

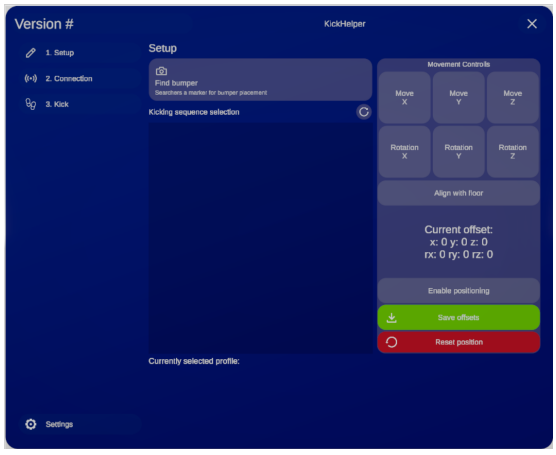

Fig. 8. User interface of the application on a position setup page.

B. Creating offsets

Due to the fact that the marker cannot always be placed at the 3D models place of origin, a system of managing positional and rotational offsets was implemented. This allows the user to place the marker in a convenient location, while still being able to automatically place the 3D model in the correct position and orientation. The offsets are created by using a reposition mode, which allows the user to freely move the 3D model using the headset's controllers on hand tracking. Once correctly positioned the off sets are saved in a *YAML Ain't Markup Language* (YAML) file, which is then later used by the application on subsequent runs.

VIII. CONCLUSIONS

Use of an AR enabled headset has a potential to be used with assisting data collection for large datasets for machine learning applications. The proposed solution of using AruCo markers and machine learning algorithms for detection and averaging proved to be a simple and so far effective solution for helping with the data collection process and user immersion. The application is still in a development stage and it is planned to add more features, such as a guidance system and a communication system with a data collection computer. It is also planned to do further testing with untrained users to determine the effectiveness of such a use case.

ACKNOWLEDGMENT

The Author would like to acknowledge the contribution and valuable time provided by the supervisor, doc. Ing. Tomáš Frýza Ph.D., as well as to thanks my work colleagues and the authors of the sources used in this paper.

REFERENCES

[1] Alps Alpine electronic components and in-vehicle information systems. Online. Available from: https://www.alpsalpine.com/e/. [Accessed-13-March-2026].

[2] Meta Quest 3: Next-Gen Virtual Reality Headset. Online. Available from: https://www.meta.com/quest/quest-3/Discover. [Accessed-13-March-2026].

[3] Meta Quest3: Full Specification. VRcompare. Online. Available from: https://vr-compare.com/headset/metaquest3The. [Accessed-16-December-2025].

[4] John K. Haas. A History of the Unity Game Engine. SciSpace - Paper. Online. 6 March 2014. Available from: https://scispace.com/papers/a-history-of-the-unity-game-engine-5flr86kgabUnity [Accessed-15-March-2026].

[5] YOSHINAGA, Takashi. QuestArUcoMarkerTracking. Online. C#. 10 December 2025. Available from: https://github.com/TakashiYoshinaga/QuestArUcoMarkerTrackingThis. [Accessed-16-December-2025].

[6] SHRESTHA, Rashik. Perspective-n-Point(PnP). Medium. Online. 29 January 2024. Available from: https://medium.com/@rashik.shrestha/perspective-n-point-pnp-f2c7dd4ef1ed. [Accessed-22-December-2025].

[7] Meta Quest 3: Full Specification. VRcompare. Online. Available from: https://vr-compare.com/headset/metaquest3. [Accessed-16-December-2025].

[8] POROYKOV, Anton, KALUGIN, Pavel, SHITOV, Sergey and LAPITSKAYA, Irina. Modeling ArUco Markers Images for Accuracy Analysis of Their 3D Pose Estimation. Proceedings of the 30th International Conference on Computer Graphics and Machine Vision (GraphiCon 2020). Part 2. 17 December 2020. [Accessed-4-April-2026].

Classification of Normal and Anomalous Air Using BME688 Sensor Data

1st Yonatan Agafari
Faculty of Electrical Engineering and Communication
Brno University of Technology
Brno, Czech Republic
260623@vut.cz

2nd Ing. Petr Kriz
Faculty of Electrical Engineering and Communication
Brno University of Technology
Brno, Czech Republic
petr.kriz4@vutbr.cz

***Abstract*—This paper deals with a comprehensive data analysis pipeline for BME688 gas sensor measurements, focusing on gas resistance as the core signal. Data are loaded from aggregated CSV files. For each of the 8 sensors, the feature set combines heater-step gas-resistance values (step1-step10) with summary statistics (mean, std, min, max, median, range). Missing values are imputed using feature-wise medians, features are standardized, and dimensionality reduction is applied through feature selection (SelectKBest) before classification. For visualization, t-SNE is generated at perplexity of 5, with air hidden and Air represented as normal to make class structure easier to interpret.**

For supervised detection, a binary Random Forest model is trained to separate normal and anomaly classes. In the current balanced setup (60 normal and 60 anomaly samples), the model achieves perfect separation (TN=60, FP=0, TP=60, FN=0), corresponding to 100% accuracy, precision, and recall on this evaluation set. These results are strong, while future validation on larger and more diverse data (across sessions, environments, and devices) is important to confirm robustness.

***Index Terms*—bme688, air classification, smell data processing, Random Forest**

I. Introduction

Gas sensors play a crucial role in environmental monitoring, industrial safety, and medical diagnostics. The BME688 sensor array, with its multi-sensor configuration and programmable heater profiles, offers rich data for substance identification and anomaly detection.

Several recent studies have focused on BME688 data processing and applications. Panzardi et al. [1] assessed BME688 performance under controlled outdoor-like conditions, showing that sensors exhibit consistent normalized responses but require individual baseline calibration, and that humidity and operating temperature (e.g., 360–400 °C) strongly affect sensitivity—findings that support the need for calibration and normalization in pipelines like ours. In a food-safety context, machine learning on BME688 gas sensor matrices has been used for meat species recognition (chicken, sheep, cattle) from volatile organic and sulfur compounds, with classifiers such as Gaussian Naive Bayes and ensemble methods achieving high accuracy [2]. Similarly, the BME688 has been applied with machine learning to detect olive oil adulteration and for odor identification in case studies, demonstrating the sensor's suitability for classification and quality control when combined with appropriate preprocessing and learning methods [3], [4]. Building on this body of work, we address anomaly detection on BME688 data with a block-based preprocessing and calibration pipeline and with dimensionality reduction for visualization and outlier detection.

This work presents a systematic approach to analyzing BME688 measurements, with particular emphasis on gas resistance as the key parameter for substance characterization.

The analysis workflow follows a structured pipeline that addresses several critical aspects of sensor data processing such as described in Methodology section.

II. Methodology

The BME688 sensor array comprises eight sensors, each measuring gas resistance in ten steps in the heater profile. For each sensor-step combination, the data are collected as a time series of gas-resistance values (in ohms). The dataset contains measurements of five substances regarded as anomalous: acetone, thinner, Savo, Softasept, and vinegar, as well as measurements of normal air. Hence, the classification task is defined as a binary classification problem with two classes: normal and anomalous.

Label in data	Generic description
Aceton	Acetone (organic solvent, VOC)
Redidlo	Organic solvent–based paint thinner
Savo	Chlorine-based disinfectant
Softasept	Alcohol-based skin disinfectant
Vinegar	Acetic acid solution

The class labels Redidlo, Acetone, Savo, and Softasept are product names used in Czech labeling/context and are kept in their original form throughout this document.

The following subsections describe the key data processing methods used for this binary classification task.

A. Data acquisition

We ran all measurements in a simple cube-shaped chamber about 50 cm on each side (roughly 125 litres of air inside), with the BME688 board sitting inside the chamber for every run. For each substance, we put about 8–10 mL of liquid on the plate inside the chamber and let the sensors log for 10 minutes while the heater stepped through its profile. When that was done, we opened things up and let fresh room air move

through for about 5 minutes, so the previous vapour could clear out. Then we closed it again and recorded another 10 minutes of ordinary room air that stretch is what we used as “normal air” for calibration and baseline statistics (see Section II.B). If you logged temperature and humidity during the sessions, it’s worth adding a line here with the typical range.

BME688 gas sensor arrays are exposed to volatile substances (acetone, vinegar, disinfectants) and to normal air under controlled conditions. Each array provides up to eight sensors and a heater profile with 10 steps, so every scanning cycle yields 80 gas resistance values (8 sensors × 10 heater steps). The chamber is closed with a rigid lid and

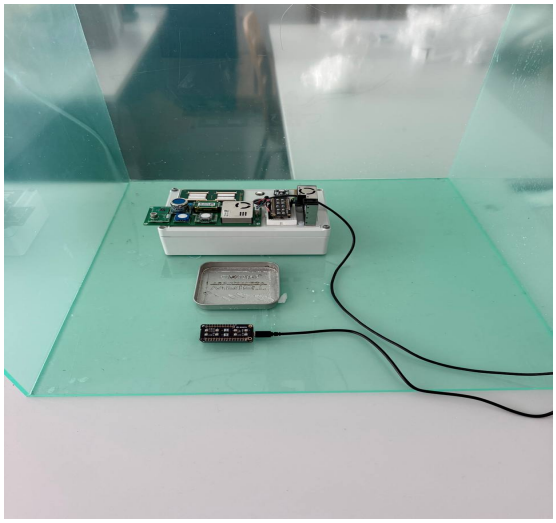

Fig. 1. Data Collection Chamber.

a gasket compressed by clips, so the internal volume stays separate from the room between venting steps. There are two sensor boards are visible inside enclosure: the first supports database-side logging, and the second represents the standard application placement; we use only the recordings from the second board for calibration, features, and classification. Data are stored as raw JSON logs (.bmerawdata) then it converted to CSV format for the next analysis. Normal air is recorded as a baseline for calibration, and the software tries several root directories to find valid normal-air recordings.

B. Data calibration process

Before each measurement the calibration data of ”Normal air” was taken. These measurements are used to establish baseline statistics for each sensor-step combination. For sensor $s \in \{1, 2, \ldots, 8\}$ and heat step $h \in \{1, 2, \ldots, 10\}$, we compute:

$$\mu_{s,h} = \frac{1}{N} \sum_{i=1}^{N} GR_{s,h,i} \tag{1}$$

$$\sigma_{s,h} = \sqrt{\frac{1}{N-1} \sum_{i=1}^{N} (GR_{s,h,i} - \mu_{s,h})^2} \tag{2}$$

$$\text{median}_{s,h} = \text{median}(\{GR_{s,h,i}\}_{i=1}^{N}) \tag{3}$$

where $GR_{s,h,i}$ represents the i-th gas resistance measurement for sensor s at heat step h, and N is the total number of calibration measurements.

This process yields 80 calibration statistics (8 sensors × 10 steps), which serve as the foundation for subsequent normalization.

Gas resistance values are normalized using z-score normalization based on the calibration statistics:

$$GR_{\text{norm}} = \frac{GR - \mu_{s,h}}{\sigma_{s,h}} \tag{4}$$

where GR is the raw gas resistance value, and $\mu_{s,h}$ and $\sigma_{s,h}$ are the mean and standard deviation computed from Normal Air calibration data for the corresponding sensor-step combination.

This normalization approach removes sensor-specific biases and environmental variations, allowing for meaningful comparison across different measurement conditions.

C. Data visualization

Figure 2 gives a quick, intuitive snapshot of the data set by showing both how many samples were collected and how the main sensor features behave for each class (air, Acetone, Redidlo, Vinegar, Softasept, Savo).

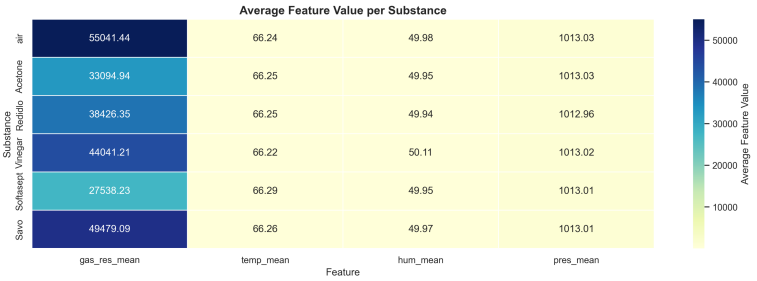

Fig. 2. Summary Statistics Overview.

Air has the most samples (100), while each chemical class has 50. This is deliberate: in gas-sensing tasks, it is common to collect more background (air) data so the model learns the difference between normal/no-odor conditions and odor-present conditions more reliably.

The heatmap on 2 shows class-wise averages of four key features:

- gas_res_mean (average gas resistance),
- temp_mean (average temperature),
- hum_mean (average relative humidity),

- pres_mean (average pressure).

A few patterns stand out. The strongest class-to-class differences appear in gas_res_mean, which ranges from about 27,538 (Softasept) to 55,041 (air). This suggests gas resistance carries the most useful signal for distinguishing substances. The higher value for air also makes physical sense, since exposure to volatile compounds often reduces sensor resistance. In contrast, temperature (66°C), humidity (50%), and pressure (1013 hPa) remain very similar across classes, indicating stable measurement conditions and limited discriminative power from these environmental features alone.

The dimensionality reduction techniques are applied to visualize and analyze the high-dimensional gas resistance data to see a potential for class separation.

The method we have used is called t-Distributed Stochastic Neighbor Embedding (t-SNE) which preserves local neighborhood structure in the embedding by minimizing the Kullback-Leibler divergence between probability distributions in high and low dimensions:

$$KL(P||Q) = \sum_i \sum_j p_{ij} \log \frac{p_{ij}}{q_{ij}} \tag{5}$$

where p_{ij} and q_{ij} are probabilities in high and low dimensions, respectively.

Figure 3 the t-SNE plot with perplexity = 5, the five chemical classes (Acetone, Redidlo, Savo, Softasept, and Vinegar) appear as compact, clearly separated groups. This means the sensor features carry strong class-specific patterns, so these substances are easy to distinguish in the reduced 2D space. By contrast, the normal (air) class is spread out rather than

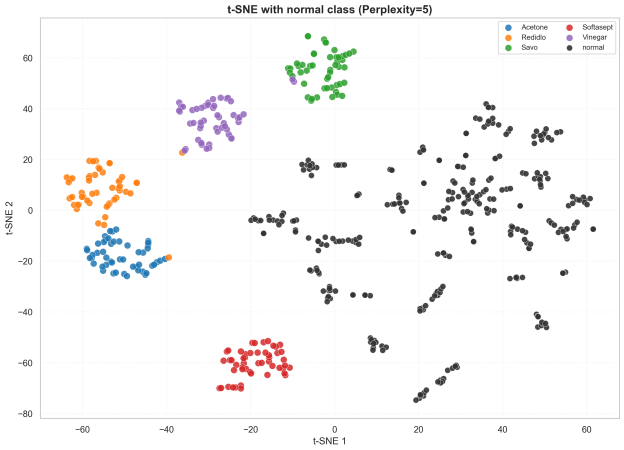

Fig. 3. Dimensionality reduction visualization showing substance separation in 2D space using t-SNE.

forming one tight cluster. That wider spread suggests baseline air measurements vary more within the class, while chemical samples are more consistently distinct.

D. ML data preparation

Raw BME688-derived features (8 sensors × 10 heater steps x statistics) were loaded from CSV file sources. The anomaly class was formed from non-air substances (Acetone, Redidlo, Savo, Softasept, and Vinegar) in the aggregated dataset, while the normal class was taken from Air samples (renamed as normal in analysis). For each of the 8 sensors, gas-resistance features included both heater-step values (step1 to step10) and summary statistics (mean, std, min, max, median, range), yielding a gas-resistance-only feature space. Missing values were imputed using feature-wise medians. Features were standardized, and univariate feature reduction (SelectKBest with ANOVA F-score) was applied before classification. For the anomaly-detection experiment, a balanced subset of 60 normal and 60 anomaly samples was used.

E. Data classification

We implemented a binary Random Forest classifier to separate normal from anomaly. After median imputation and standardization, feature reduction was performed with SelectKBest (up to 120 top features, capped by available dimensions). The Random Forest used n_estimators=600, class_weight="balanced_subsample", and random_state=42. Data were split with stratification (75% train, 25% test) for model fitting, while preprocessing components (scaler and selector) were fit on training data and then applied to other data to avoid leakage. Final reporting for the balanced 60/60 setting includes a confusion matrix and metrics generated over the full balanced set (120 samples), together with saved outputs in final_visualization.

III. Results

The Random Forest classifier correctly classified all 60 normal test samples (TN=60, FP=0) and all 60 anomaly samples (TP=60, FN=0). This yields 100% anomaly recall, 100% anomaly precision, and 100% overall accuracy on the balanced 60-normal / 60-anomaly evaluation set (120 samples total).

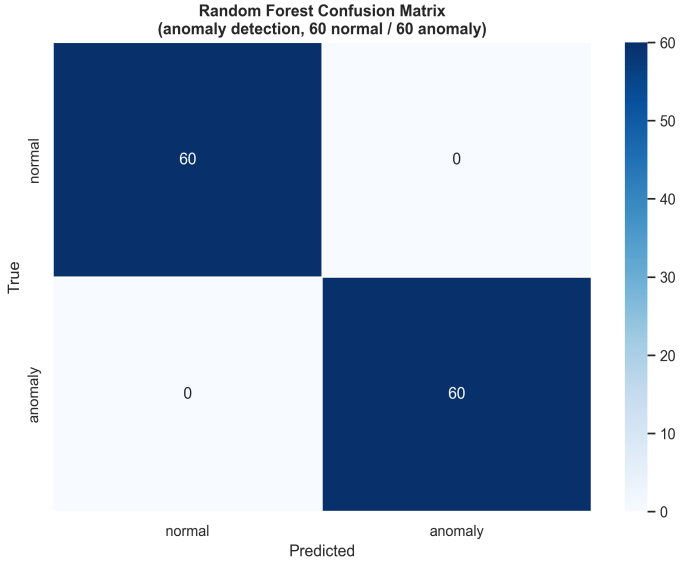

Fig. 4. Confusion Matrix.

IV. Conclusion

No normal sample was incorrectly flagged as anomalous, and no anomaly sample was missed. Therefore, in this 60/60 setup, the model achieved perfect separation between normal and anomaly classes. This indicates that the weighted performance metrics are also all 1.00 (precision, recall, and F1-score).

This work builds a practical anomaly-detection pipeline for BME688 data. Instead of using all sensor signals, we focus on gas resistance features from all 8 sensors across the 10 heater steps, and also include per-sensor summary statistics (mean, std, min, max, median, range). After handling missing values, scaling features, and reducing dimensionality, we visualize the data with t-SNE (with air hidden and Air New treated as normal), and train a binary Random Forest model for normal-vs-anomaly detection. In the current balanced experiment (60 normal + 60 anomaly samples), the Random Forest model classifies every sample correctly: TN=60, FP=0, TP=60, FN=0, which corresponds to 100% accuracy, precision, and recall on this evaluation set.

These are very strong results, but they should be interpreted carefully. The evaluation set is still relatively small, and we need broader testing to confirm robustness across different recording sessions, environments, and devices. A strong next step is to collect more diverse normal/anomaly data and validate the model under cross-session and cross-device conditions, while also comparing it with alternative anomaly scoring methods in reduced spaces (for example PCA/UMAP distance-based scoring).

References

[1] E. Panzardi, A. Fort, V. Vignoli, M. Mugnaini, and A. Flammini, "Assessing BME688 sensor performance under controlled outdoor-like environmental conditions," Sensors, vol. 25, no. 23, art. no. 7102, Nov. 2025, doi: 10.3390/s25237102.

[2] "A new approach for machine learning-based recognition of meat species using a BME688 gas sensors matrix," Electron. J., Chiang Mai Univ. [Online]. Available: https://epg.science.cmu.ac.th/ejournal/journal-detail.php?id=11918

[3] Ismail Hakki Parlak et al., "Machine learning–based detection of olive oil adulteration using BME688 gas sensor matrix", Food Analytical Methods, vol. 18, pp. 1454–1464, 2025. DOI: 10.1007/s12161-025-02803-0.

[4] J. Pereira, A. Mota, P. Couto, A. Valente, and C. Serôdio, "Potential use of BME development kit and machine learning methods for odor identification: A case study," Appl. Sci., vol. 15, no. 17, Art. no. 9687, 2025. DOI: 10.3390/app15179687.

[5] Bosch Sensortec. BME688 Data Sheet. Available: https://www.bosch-sensortec.com

Adaptive Digital Control Techniques for Wide-Range Efficiency Optimization in a Buck Converter

Martin Kopka
Department of Microelectronics, Faculty of Electrical Engineering and Communication
Brno University of Technology
Brno, Czech Republic
240852@vutbr.cz

***Abstract*—This paper presents digital control algorithms for efficiency optimization in a buck converter operating over a wide output-voltage range. The proposed methods adapt to the changing behavior of the power stage to maintain high efficiency across varying operating conditions. Sensorless Diode Emulation and Adaptive Valley Switching algorithms are introduced to reduce power loss during discontinuous conduction mode (DCM) operation, while the Adaptive Frequency Scaling algorithm adjusts the switching frequency according to the output voltage and load current. As a result, the converter operates close to its maximum efficiency over a wide operating range.**

***Index Terms*—buck converter, switching power supply, digital control, efficiency optimization, diode emulation, valley switching**

I. Introduction

The design of high-power consumer power supplies often requires efficient utilization of available space through high-density PCB layouts in order to achieve practical dimensions. At the same time, these power supplies must maintain very high efficiency to mitigate thermal constraints associated with components with inherently limited surface area.

Digital control provides several advantages in this context. One of the main benefits is the ability of the controller to acquire detailed information about the real-time operating state of the power stage. This information can then be processed by advanced control algorithms to improve overall system performance. Because these algorithms are implemented in software, their complexity does not increase the required PCB area. In addition, a large number of parameters can be configured and optimized for a specific power-stage design.

Analog control can be particularly limiting in wide-range converters, where the controller must ensure stable operation under worst-case conditions, often at the expense of performance in other parts of the operating range. In contrast, digital controllers can adapt to changes in power-stage behavior across a wide range of operating conditions [1].

The algorithms presented in this paper were developed for an ultra-high-power-density buck converter intended for USB Power Delivery 3.2 applications supporting the full fixed-supply output range of 5–48 V [2].

The author acknowledges financial support from **onsemi**

II. System Overview

For context, the high-level block diagram of the proposed buck converter is shown in Fig. 1. The GaN-based power stage is designed to achieve high power density, while the adaptive digital control architecture aims to reduce power losses such that the converter can operate without active cooling or additional heat sinks. The digital controller is implemented using an STM32G474 microcontroller.

Output-voltage regulation is achieved using an Average Current Mode (ACM) control scheme. This control structure consists of an inner loop that regulates the average inductor current and an outer loop that generates the current reference required to maintain a constant output voltage.

This paper focuses on three proposed digital algorithms, which modify selected switching parameters during operation to improve efficiency across the operating range.

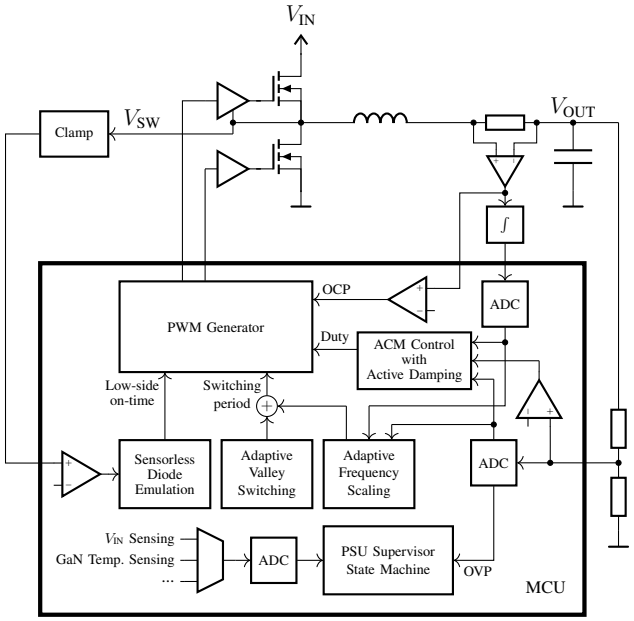

Fig. 1. High-level block diagram of the proposed digital control architecture.

III. Sensorless Diode Emulation

To eliminate the power loss associated with the forward voltage drop of the low-side diode in the buck converter power stage, the diode is typically replaced by a MOSFET. However, this complicates the controller as the low-side switch now requires its own driving signal. The simplest implementation drives the low-side switch in the opposite phase of the high-side switch with the addition of sufficient deadtime for cross-conduction prevention. This operating mode, referred to as forced continuous conduction mode (FCCM), never allows the converter to operate in the discontinuous conduction mode (DCM) as the low-side switch remains on even after the inductor discharges to zero. In this condition, the current reverses through the switch, discharging the output capacitor and recuperating that energy to the input through the high-side switch during its subsequent on-time T_{ON}. This is disadvantageous for efficiency because the negative inductor current does not contribute to output power and only increases current ripple, leading to higher conduction losses.

Most implementations solve this by monitoring the inductor current in real time and turning off the low-side switch at the zero crossing. This approach, referred to as diode emulation, controls the low-side switch in a way that substitutes a conventional diode. However, due to space and power loss constraints and the need for high bandwidth and gain in the current-sensing circuitry, the use of direct inductor current sensing became impractical.

A. Principle of Sensorless Operation

This paper proposes a sensorless diode emulation method that determines the optimal turn-off instant of the low-side switch by measuring the freewheeling interval derived from the switching-node voltage waveform. Its operating principle is based on observing the switching-node voltage V_{SW} using a zero-steady-state-loss clamping circuit and a comparator. The digital controller then processes this information and maintains the optimal on-time of the low-side switch T_{OFF} through a feedback loop. The principle can be explained with reference to Fig. 2, which shows three characteristic switching-node waveforms.

The upper waveform corresponds to an early turn-off of the low-side switch. The inductor continues discharging through the body diode, or through the third-quadrant conduction path in the case of GaN devices, resulting in a negative switching-node voltage and delayed onset of idle-phase ringing.

The third waveform shows the opposite case, where the low-side switch is turned off too late. The inductor current has already crossed zero and become negative, causing current to continue through the body diode of the high-side switch. This is reflected in the switching-node waveform by a pronounced voltage rise above the input voltage V_{IN}. The high-side body diode conducts until the inductor current again reaches zero, after which the idle phase begins and ringing appears.

The middle waveform represents the ideal case, where the low-side MOSFET is turned off exactly at the zero-crossing of the inductor current and the idle phase begins immediately. The three cases differ in the time interval between the turn-off of the low-side switch and the moment when the switching-node voltage rises above a defined threshold. This interval is referred to as the freewheeling time.

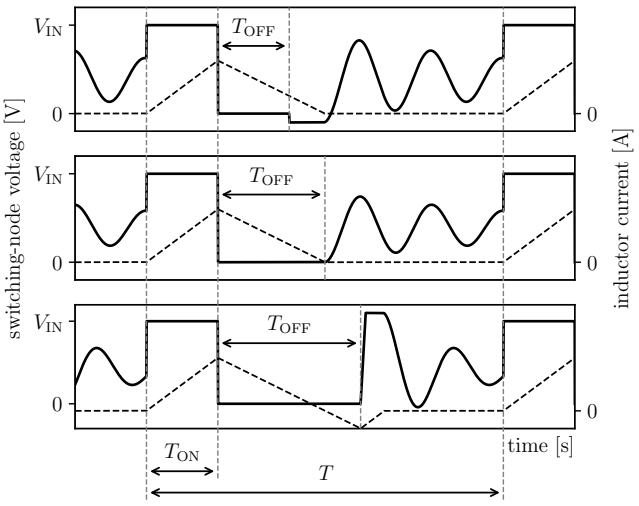

Fig. 2. Waveforms showing the principle of Sensorless Diode Emulation.

B. Feedback Acquisition and Control Loop

Figure 3 shows the circuitry for capturing the freewheeling time, which serves as feedback for the control loop. Here, a MOSFET-based clamping circuit limits the comparator input voltage. The design of this clamp eliminates steady-state power loss associated with conventional resistor-diode clamps.

An SR latch generates a capture window for a capture register. The latch is set at the instant of the low-side MOSFET turn-off. After the switching-node voltage crosses the comparator threshold, the latch is reset and the capture register stores the current digital value of the PWM generator counter. This approach ensures that only the first rising edge after the low-side switch turn-off is captured. Since the low-side switch is controlled by the PWM generator counter, the freewheeling time in counter ticks can be computed by subtracting the captured value from the low-side turn-off compare value. An integral-type controller is then applied to the freewheeling time error and T_{OFF} is adjusted.

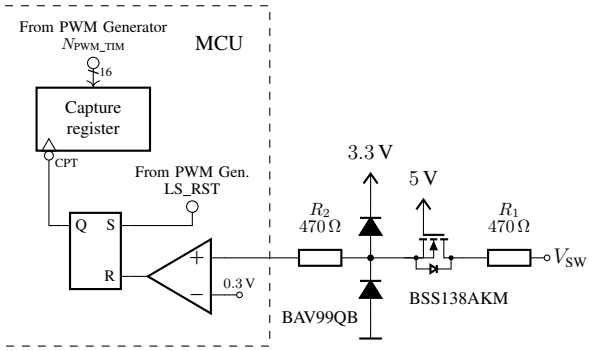

Fig. 3. Sensorless Diode Emulation feedback acquisition schematic.

C. CCM Detection and DCM Test Sequence

When the converter transitions to CCM, the inductor current no longer reaches zero and no ringing is captured. Consequently, the first comparator rising edge detected by the sensing circuitry corresponds to the turn-on of the high-side switch in the following switching period. If not addressed, the algorithm would regulate this delay to its setpoint, effectively introducing additional deadtime before the next magnetization phase and increasing power losses. To prevent this behavior, the algorithm checks whether the captured value corresponds to the subsequent switching cycle. If so, the algorithm is paused and the converter enters FCCM operation.

When the algorithm pauses and transitions to Forced CCM, the low-side switch remains on for the entire remainder of the switching period. Since the switch keeps the switching-node voltage at zero, no ringing occurs and feedback is therefore lost. Even if the output current later decreases below the CCM/DCM boundary, the algorithm would remain stuck in Forced CCM. To prevent this condition, the algorithm performs a DCM Test every 1024 switching cycles during Forced CCM operation. By shortening the low-side switch on-time for one switching cycle, the possibility of observing switching-node ringing is temporarily restored. Data from this single switching cycle are evaluated using the same logic as the CCM detection and the result determines whether the algorithm operation is resumed.

IV. Adaptive Valley Switching

During DCM operation, the inductor current falls to zero before the next switching cycle begins. After the inductor current reaches zero, the switching node exhibits a resonance determined primarily by the output capacitances of the power transistors and the inductor inductance. This resonance produces several voltage maxima (peaks) on the switching-node waveform before the next switching period begins. Initiating the next switching cycle at one of these peaks, as shown in Fig. 4, can significantly reduce switching losses because the drain–source voltage of the high-side switch is lower, resulting in reduced turn-on energy. This condition is referred to as quasi-zero-voltage switching (QZVS) [3].

The switching loss of the high-side switch $P_{\text{SW,HS}}$ is given by

$$P_{\text{SW,HS}} = \frac{1}{2} \cdot (V_{\text{IN}} - V_{\text{SW}}) \cdot I_{\text{OUT}} \cdot (t_R + t_F) \cdot f_{\text{SW}}, \qquad (1)$$

where V_{SW} denotes the switching-node voltage at the instant of switching, I_{OUT} is the output current, f_{SW} is the switching frequency, and t_R and t_F denote the rise and fall times of the switch, respectively.

Valley switching also reduces the output capacitance loss. Under these conditions, the output capacitance C_{OSS} of the high-side switch is partially discharged by the inductor ringing current prior to turn-on. A lower effective drain–source voltage corresponds to less energy stored in C_{OSS}, which would otherwise be dissipated during switching:

$$P_{\text{COSS}} = \frac{1}{2} \cdot C_{\text{OSS}} \cdot (V_{\text{IN}} - V_{\text{SW}})^2 \cdot f_{\text{SW}}. \qquad (2)$$

The dynamic losses of the low-side switch are ignored since it operates under zero-voltage switching (ZVS) conditions.

The proposed algorithm operates without the need for additional hardware and uses a preprogrammed value of the ringing period T_{RING}. This is justified by the relatively flat shape of the ringing peaks, meaning that errors caused by component variation result in negligible efficiency degradation. However, T_{RING} varies with operating conditions, primarily due to the voltage dependence of C_{OSS}. For this reason, the ringing-period setpoint is configurable at runtime, and updated following each output-voltage change.

Since the ringing is fully contained within the idle interval T_{IDLE}, and the waveform always starts at a minimum, alignment of a selected ringing peak with the beginning of the subsequent switching period is achieved when

$$T_{\text{IDLE}} = T_{\text{RING}} \cdot \left(n - \frac{1}{2} \right), \qquad (3)$$

where n denotes the index of the targeted peak. An integral-type controller performs fine adjustment of the switching period T to indirectly control T_{IDLE} and maintain this alignment.

As the output current decreases, both T_{ON} and T_{OFF} shorten. To maintain a constant T_{IDLE}, the switching period must therefore decrease. This would significantly increase the switching frequency under light-load conditions, which is undesirable from an efficiency standpoint.

The nominal switching frequency is defined by a higher-level system (Adaptive Frequency Scaling) as the optimal operating frequency for the given operating point. To maintain this nominal value, the algorithm exploits the periodic nature of the ringing waveform and re-locks to a later peak by incrementing its control-loop setpoint by T_{RING}. Selecting a later peak maintains quasi-ZVS operation of the high-side switch without requiring a large deviation of the switching frequency from its nominal value.

The inverse operation is performed when the output current increases, thereby re-locking to an earlier peak. The algorithm selects the target peak index such that the resulting alignment yields a switching frequency closest to the nominal value.

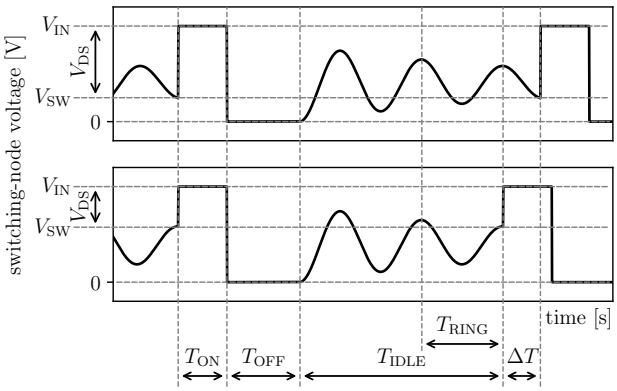

Fig. 4. Adaptive Valley Switching algorithm locking onto a voltage peak.

V. Adaptive Frequency Scaling

The Adaptive Frequency Scaling algorithm exploits the fact that individual power-loss mechanisms dominate at different operating points. High switching frequencies are dominated by switching losses, whereas low switching frequencies increase inductor ripple current and therefore RMS-current-related conduction losses. Consequently, the switching frequency that yields maximum efficiency varies across the operating range.

The system monitors the output voltage and current and adjusts the switching frequency between 40–400 kHz based on a three-dimensional efficiency map implemented as a lookup table obtained from experimental characterization of the power stage. This allows the converter to operate close to the optimal switching-frequency over the entire operating range.

Since the ACM control-loop compensators are implemented as Infinite Impulse Response (IIR) filters whose transfer functions depend on the sampling rate, their parameters must be updated when the switching frequency changes. To avoid computational delays and control-loop disturbances, the adaptive control architecture employs two instances of each compensator. New IIR coefficients are computed into the inactive instance and its internal state is precharged to match the steady-state of the active filter, allowing glitchless switching between configurations. Each output-voltage option includes dedicated CCM and DCM compensator settings to maintain optimal transient performance.

At light-load conditions, the low switching frequency limits achievable control-loop bandwidth. To improve the response to sudden load increases, a Transient Response Booster detects rapid output-voltage drops indicating an upcoming rise in inductor current demand. The system then immediately switches to the maximum switching frequency and Forced CCM operation until the transient is resolved. To avoid delay caused by compensator recomputation, a third compensator instance remains precomputed and ready for immediate activation.

VI. Experimental Results

Although the proposed converter is still under development, preliminary measurements illustrate the contribution of the proposed algorithms. Fig. 5 shows the measured efficiency as a function of output current up to 5 A at an input voltage of 50 V. Two representative output voltages, 20 V and 36 V, were selected from the full operating range.

The baseline case corresponds to forced CCM operation with a fixed switching frequency of 400 kHz across the entire range, selected to accommodate the worst-case output voltage of 28 V. Next, the switching frequency was optimized for the full-load efficiency of each output voltage. Enabling Sensorless Diode Emulation and Adaptive Valley Switching significantly improved efficiency in DCM operation. Finally, Adaptive Frequency Scaling further reduced switching losses by lowering the switching frequency at light load.

The measured full-load efficiency reaches 97.8 % and 98.8 %. At 10 % load, the efficiency improves from 89.2 % to 96.4 % and from 94.7 % to 98.3 %, demonstrating the effectiveness of the proposed efficiency optimization methods.

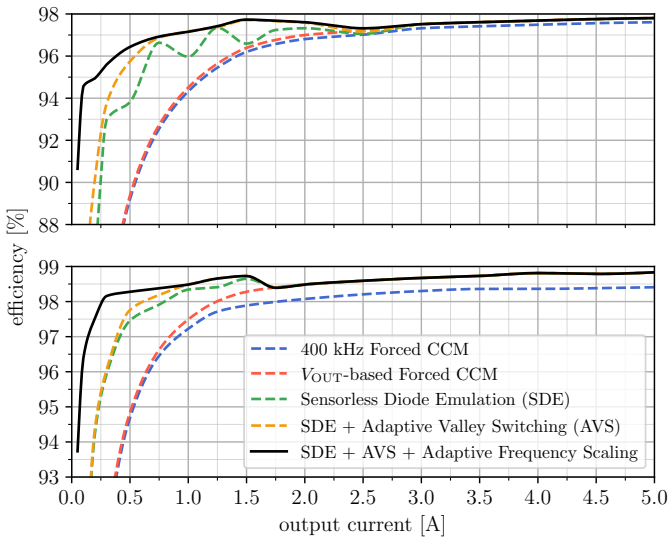

Fig. 5. Power-stage efficiency at the output voltage of 20 V (upper plot) and 36 V (lower plot) and incremental contributions of proposed algorithms.

VII. Conclusion

This paper proposed three digital efficiency optimization methods for a wide-range buck converter.

The Sensorless Diode Emulation algorithm enables DCM operation of a synchronous buck converter without requiring direct inductor current sensing. By relying only on switching-node voltage observation, the method eliminates the need for high-bandwidth current-sensing circuitry.

The Adaptive Valley Switching algorithm maintains near-ZVS conditions of the high-side switch, thereby reducing switching losses. In addition, the algorithm dynamically relocks to different ringing peaks to prevent drift toward unfavorable switching frequencies under varying load conditions.

The Adaptive Frequency Scaling algorithm adjusts the switching frequency according to the output voltage and load current, allowing the converter to operate close to its optimal efficiency point across a wide operating range.

Experimental results confirm the effectiveness of the proposed methods. A more detailed description of the complete converter design will be available in the author's master's thesis titled "Ultra-High Power Density 240W Digital Buck Converter with Adaptive Efficiency Optimization".

Acknowledgment

I would like to thank my supervisor Ing. Roman Prokop, Ph.D. for his guidance during the preparation of this paper.

References

[1] P. Mattavelli and S. Buso, *Digital Control in Power Electronics*. San Rafael, CA, USA: Morgan & Claypool Publishers, 2006, p. 3.

[2] USB Implementers Forum, *Universal Serial Bus Power Delivery Specification*, Rev. 3.2, Ver. 1.1, Beaverton, OR, USA, 2024. [Online]. Available: https://usb.org/document-library/usb-power-delivery [Accessed: Mar. 11, 2026].

[3] F. C. Lee, "High-Frequency Quasi-Resonant Converter Technologies," *Proceedings of the IEEE*, vol. 76, no. 4, pp. 377–390, Apr. 1988, DOI: 10.1109/5.4424.

Design of MIBICO Payload Controller for the CIMER Mission

Tomáš Paulysko
Department of radioelectronics
FEEC, Brno University of technology
Brno, Czech Republic
238731@vutbr.cz

Jiří Veverka
Department of radioelectronics
FEEC, Brno University of technology
Brno, Czech Republic
xvever12@vut.cz

***Abstract*—This paper focuses on the control, monitoring, and maintenance of environmental conditions for biological payloads in the space environment. The article describes methods for monitoring and optimizing the MIBICO (Micro Bio Container) payload developed for the student satellite mission CIMER (Cyanobacteria In Microgravity Environment Research) conducted under the guidance of the European Space Agency (ESA). The main objective of the experiment is to verify the feasibility of rehydrating cyanobacteria in orbit and to monitor their biological activity in low Earth orbit. The work focuses on the design of a modular and energy-efficient electronic system that enables largely autonomous control of the experiment, acquisition of sensor data, regulation of environmental conditions inside the experimental capsules, and communication with the main satellite system.**

***Index Terms*—CubeSat, CIMER, MIBICO, payload controller, microcontroller, biological experiment in space, power budget, space, LEO**

I. Introduction

In recent years, there has been a significant growth in the development of small satellites known as CubeSats. Due to their standardized design and relatively low cost, they enable scientific and technological experiments to be conducted by smaller research teams and universities. One promising area of their application is biological experiments focused on studying the behavior of microorganisms under microgravity and space radiation conditions. This work focuses on the design of a control module for the MIBICO (Micro Bio Container) experimental payload within the student CubeSat mission CIMER (Cyanobacteria In Microgravity Environment Research) based on a 3U platform. The main objective of the experiment is to verify the feasibility of rehydrating dehydrated cyanobacteria in orbit and to monitor their biological activity. The proposed modular and energy-efficient system enables semi-autonomous control of the experiment, acquisition of image and other data from sensors, regulation of environmental conditions inside the experiment, and communication with the main satellite system while respecting the typical constraints of CubeSat platforms.

II. Motivation and General System Overview

Mission goal is to investigate whether dehydrated cyanobacteria can be transported to orbit in an inactive state and later reactivated to resume biological activity. This approach significantly reduces storage and maintenance requirements during launch and transport, while improving the stability and safety of biological samples. Understanding these processes could contribute to the development of future autonomous biotechnological systems for long-term space missions, including biological oxygen production or life-support systems.

A. MIBICO

The primary payload of the mission, MIBICO (Micro Bio Container), contains four experimental cuvettes with cyanobacteria cultures placed on a porous sponge substrate. Before launch, the cultures are dehydrated under controlled laboratory conditions and hermetically sealed inside the experiment container. During the mission, the samples are sequentially rehydrated and their biological activity is monitored using environmental sensors and imaging systems while maintaining controlled temperature and illumination conditions. The experiment focuses on observing the recovery of cyanobacteria after rehydration and studying their behavior in microgravity, providing data on their growth, metabolic activity, and response to the space environment.

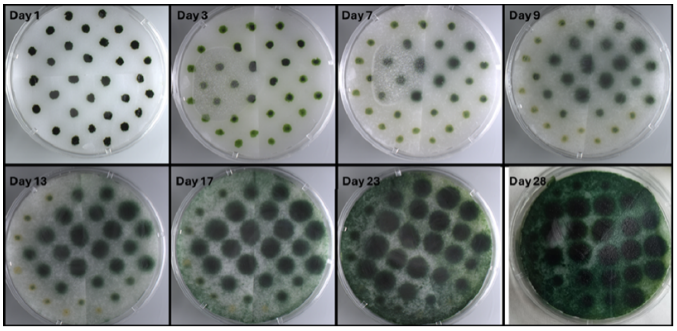

Fig. 1. Growth of cyanobacterial surface area during cultivation [1].

B. General system design overview

The system consists of three main modules: CAU (Camera Unit), PU (Power Unit), and RU (Rehydration Unit). The CAU module serves as the main control unit of the system and is primarily responsible for camera control and data acquisition.

The PU module contains power electronics for efficient current regulation of the experiment illumination. Due to

spatial constraints, the PU module also includes connectors for the RU module, sensors located outside the module, and an interface connector for connection to the main satellite system.

The CAU and PU modules are electrically interconnected using a board-to-board. This connector was selected due to its small size and its use in commercial CubeSat components. The modules are mechanically connected near the connector using an aluminum frame and M3 screws.

The RU module is used for direct monitoring of the experiment, measurement of environmental conditions inside the cuvette, and temperature control of the samples. This module is connected via cabling.

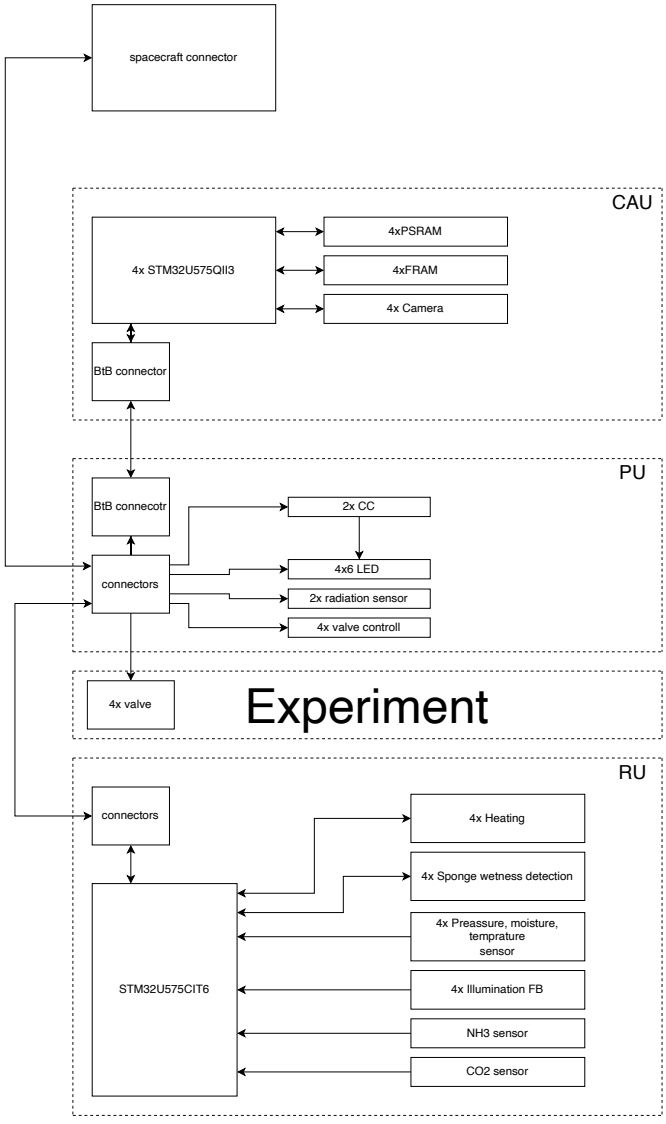

Fig. 2. System block diagram

III. Camera Unit

After several design iterations of the CAU module, which initially considered the use of multiplexers and a single high-performance microcontroller with larger memory, a different architecture was selected. Primarily due to redundancy requirements, and secondarily due to a more feasible PCB design, a solution using a dedicated microcontroller for each of the four cameras was chosen.

A. Cameras

Cameras OV5640 with a resolution of 5 MP and dimensions of 8.5 × 8.5 mm from OmniVision were selected for the CAU module. The lens used with the sensor is a 160-degree wide-angle lens. This lens allows the entire surface of the sponge containing the experiment to be captured sharply at a distance of approximately 30 mm between the sensor and the sample.

B. Microcontrollers and Memory

As the microcontroller (MCU) for camera control and data acquisition from sensors located on the upper bus layer, the low-power microcontroller STM32U575 from STMicroelectronics was selected [3]. The microcontroller allows connection of the camera via the DCMI (Digital Camera Interface). Camera control, memory interfaces, and the connection of additional sensors require a relatively high number of I/O pins.

At the same time, the design must satisfy the strict space constraints of the CubeSat platform. Therefore, a UFBGA package with 132 pins was selected, as it meets both the pin-count and size requirements.

Each MCU is equipped with two types of external memory. The first is 16MBit FRAM [4], which is used for storing configuration data and operational logs. The second is 64Mbit PSRAM [5], which serves as a temporary buffer for image data from the camera and other sensors.

C. Communication

For communication between microcontrollers within the CAU module, communication with RU and rest of the satelite, CAN bus is primarily used. As redundant or optional communication interfaces RS-485 and I2C are also implemented. All of the afore mentioned interfaces utilize the CSP (CubeSat Space Protocol) which provides a decentralized subsystem command and control from anywhere in the spacecraft communication network.

For communication with memeories SPI and QSPI interfaces are used. For sensors I2C and SPI interfaces are used.

IV. Power Unit

The main function of the Power Unit module is voltage regulation for the experiment and current regulation for the experiment illumination. Due to space constraints, two BG51-OEM sensors for the detection of nuclear, beta, and gamma radiation are also placed on the module. The sensors are positioned approximately at the level of the experimental sponges, allowing any unexpected failure of the experiment to be potentially associated with detected radiation events.

Each of the four sections of the experiment is illuminated by six LEDs with a high CRI to provide a light spectrum as close as possible to natural sunlight. To achieve high efficiency, the illumination is driven by a switching power supply operating

as a constant current source using a feedback resistor. Two independent current regulation channels are implemented in the module. Each channel controls half of the LEDs for all four experimental sections.

During normal operation, the LEDs are powered at approximately 25% of the configured maximum output power. This design allows one of the two power channels to fully compensate for the other in case of failure.

V. Rehydration Unit

The RU module is responsible for monitoring environmental conditions inside the experimental cuvettes and collecting data from the cyanobacteria samples. It is designed as a local control and measurement unit placed close to the experiment, this also allows to utilize contact measurmenst. A dedicated MCU also from STM32U575 series controls all circuits on the RU board and gathers data from the connected sensors. Communication with the CAU and the rest of satelite is primarily performed via a CAN bus, while RS-485 and I²C interfaces are implemented as redundant communication channels to increase system reliability.

The module also includes circuits for experiment heating, hydration detection of the biological sample, and feedback control of illumination intensity. Environmental conditions inside the experimental capsules are monitored using BME280 sensors measuring temperature, humidity, and pressure. Sensors for NH3 and CO2 are also implemented to monitor gas concentrations.

A. Experiment monitoring

For monitoring the experiment, contact sensors for detecting sample rehydration and feedback circuits for regulating illumination intensity are implemented. The rehydration detection sensors operate based on a voltage divider driven by an alternating voltage without a DC component. Due to the absence of the DC component, the electrodes do not oxidize and the biological sample is not contaminated.

The feedback circuit for illumination control uses photodiodes placed near the samples to measure the light intensity.

B. Temperature control of the experiment

Six resistors are implemented for temperature control of each of the four samples. The resistors are divided into two branches, each containing three resistors.

Each branch is controlled using N-channel MOSFET transistors, which allow regulation of the power delivered to the resistors using a PWM (Pulse Width Modulation) signal generated by the MCU. For feedback, temperature sensors are placed close to the resistors to monitor the temperature.

The temperature is regulated using a PID controller. The resistors are driven at relatively low switching frequencies to minimize EMI.

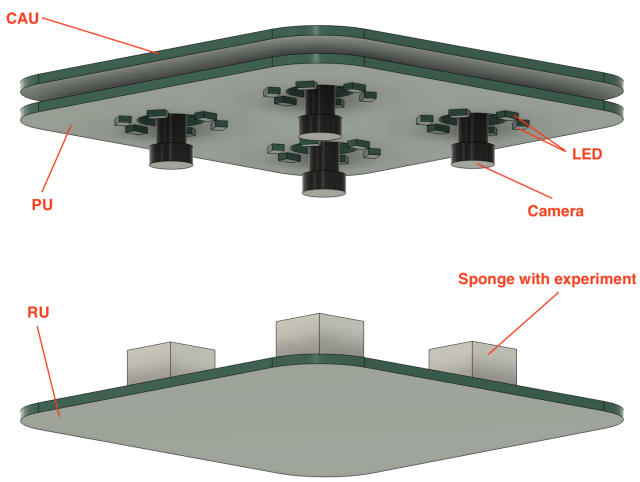

Fig. 3. Component placement

Conclusion

This paper presented the design of a control and monitoring system for the MIBICO biological payload developed for the student CubeSat mission CIMER. The proposed system enables semi-autonomous control of the experiment, monitoring of environmental conditions, acquisition of image and sensor data, and communication with the main satellite system while respecting CubeSat constraints. The architecture is divided into three main modules: the Camera Unit (CAU), Power Unit (PU), and Rehydration Unit (RU). This modular and decentralized approach improves system scalability, redundancy, and fault tolerance. The presented design demonstrates the feasibility of implementing a compact and energy-efficient control system suitable for biological experiments in the space environment and provides a basis for future hardware prototyping and integration into the CIMER mission.

Acknowledgment

I would like to thank to my supervisor Ing. Jiří Veverka, for his professional guidance, valuable consultations, patience, and helpful suggestions during the preparation of this work. I would also like to thank the student organization YSpace for providing the materials, technical information, and resources necessary for the completion of this work.

References

[1] YSPACE. *MIBICO Payload ARPT*. Internal Technical Report. Brno: YSpace, Brno University of Technology in cooperation with Mendelova univerzita in Brno, 2025. Not publicly available.

[2] BOROVIČKA, Matěj. *THERMAL MANAGEMENT FOR MISSION CIMER*. MASTER'S THESIS. Brno: Vysoké učení technické v Brně, 2024.

[3] STMicroelectronics. *STM32U575xx Arm Cortex-M33 32-bit MCU Datasheet*. Geneva: STMicroelectronics, 2024.

[4] Infineon Technologies AG. *CY15B116QI, 16Mb EXCELON™ (F-RAM) Datasheet*. Munich: Infineon Technologies AG, 2022.

[5] AP Memory Technology. *APS6404L-3SQR QSPI PSRAM Datasheet*. Taiwan: AP Memory Technology Corp., 2020.

Spacecraft electrical passivation switch providing a physical gap on the power line

Jan Bažil
Department of Radioelectronics
FEEC, Brno University of Technology
Brno, Czech Republic
jan.bazil@vut.cz

Jiří Veverka
Department of Radioelectronics
FEEC, Brno University of Technology
Brno, Czech Republic
xvever12@vut.cz

***Abstract*—This work focuses on the design of an electrical electrical passivation system for a spacecraft Electrical Power System (EPS) with a focus on telecommunication satellites in the geostationary orbit. The goal is to propose an architectural solution along with a complementary electrical switch providing a physical gap in the spacecraft's power bus. Driving requirements, compliant with relevant legal frameworks and competitive with the state of the art solutions, were identified and a supporting environmental analysis was performed. Further the current state of the art and different mechanisms of electrical passivation are discussed and subsequently a compliant approach is chosen based on a trade-off analysis. Based on the previous choice, an architectural design of the EPS is proposed and a detailed design of a relevant switch based on a serial-parallel system of mechanical relays is presented. A supporting reliability assessment was performed to ensure compliance with strict reliability requirements common in space technologies.**

***Index Terms*—Spacecraft EPS, electrical passivation, space debris**

I. Introduction

With the ongoing commercialization of space, a larger and larger number of spacecraft are being deployed in Earth's orbit, and consequently the amount of orbital space debris is rising rapidly. The majority of deployed spacecraft perform their missions without any relevant incident and, after their end of life, become relatively inert. However, even a single fragmentation event can become a significant contributor to the total amount of space debris. While individual collisions or tests of anti-satellite weapons (ASATs) create a huge number of small fragments, these events are very rare. The majority of fragmentation events are small but numerous, as can be seen in fig. 1, and are commonly related to the explosion of an on-board source of energy [1] [3].

One of these sources of energy is the spacecraft's batteries, commonly utilizing lithium-ion chemistry. These can experience thermal runaway when subjected to an overvoltage, and high temperatures or a puncture at a high State of Charge (SoC). Thermal runaway can produce a number of high-velocity fragments due to battery cell ruptures that can themselves contribute to space debris or consequently damage and fragment the spacecraft itself [2]. To avoid this event, a process of electrical passivation is performed. This process aims to deplete the battery and avoid its continued

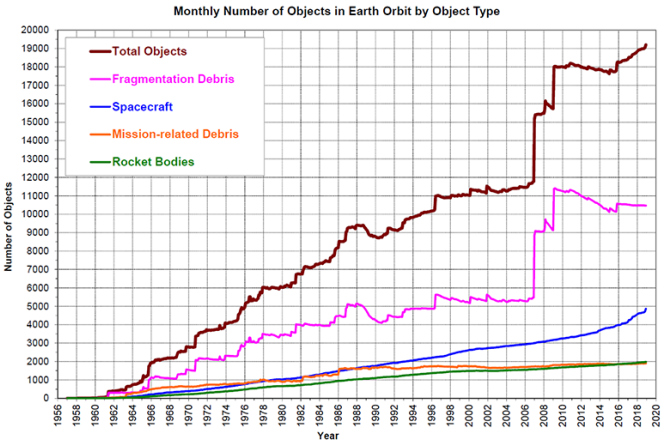

Fig. 1. Number of space debris objects current to 2020 as reported by NASA ODPO

charging by isolating it from the Solar Array (SA) [6]. As the passivation mechanism provides a unique mission-ending function, it needs to be highly reliable, resistant to single failures, and resistant to the space environment [3].

II. State of the Art

A. Universal Passivation Mechanisms

Electrical passivation is commonly implemented through three mechanisms—wire cutters, electromechanical relays, and solid-state transistors [4]. Each of these offers distinct advantages and disadvantages and is suited to different space environments.

Wire and cable cutters in their different forms possess extensive heritage in aerospace engineering. They are reliable and inherently resistant to radiation; however, their behaviour during the cutting process is unpredictable and can result in undesirable short circuits. Furthermore, they are generally employed at the beginning of missions and, as such, their long-term degradation is not well studied [4]. Wire cutters also perform a non-reversible process. This can be an issue in the case of an initialization caused by an internal failure that could otherwise be reversed by the On-Board Computer

(OBC).

Electromechanical relays are an old and proven technology with extensive space heritage. They provide a reliable, radiation-resistant, and reversible solution that can be deployed for high-power applications common in geostationary satellites and are often specifically used in power electronics. However, they are susceptible to arcing and several other issues related to the vacuum environment of space [8].

Lastly, solid-state transistors are common electrical components used in the vast majority of electronics on the ground as well as in space. However, even the most radiation-resistant transistors, such as MOSFETs, can usually withstand only about 30 krad(Si) of Total Irradiation Dose (TID) and are highly susceptible to Single Event Effects (SEEs) [5]. As such, when used, they are implemented with a number of redundancies to mitigate potential radiation-induced failures, or dedicated radiation-hardened transistors are used, which however come at a significant price.

B. Equipment Specific Passivation Mechanisms

Further, the SA and the battery offer separate mechanisms specific to their design. Most satellites employ mechanisms for power extraction optimization such as MPPT or S3R. These can be used to cut the SA power via shunting [4]. A unique opportunity in the SA is that the solar cells, or their sections, can be shorted without any adverse effects. This offers an opportunity for reversible passivation of the SA that would otherwise require cutting the power bus or some other means of physical separation [4].

In batteries, a cell bypass can be utilized. These devices allow for intentional shorting of a cell or a whole brick in case of a cell failure. While some are implemented via electrical fuses and as such cannot be controlled from the OBC, a commandable mechanical solution is also common [7]. However, cell bypasses are only employed in parallel-serial battery configurations and therefore do not offer a universal solution to the outlined problem. They also inherently do not offer an opportunity to discharge the battery before its passivation.

III. Driving Requirements and Environmental Analysis

A requirements specification outlining the final performance characteristics of the solution was composed. Several requirements identified as driving the mechanism selection and switch design are:

1) The passivation switch shall create a physical gap on the power line.
2) The design of the passivation switch shall ensure an activation reliability of at least 0.99.
3) Once activated, the passivation switch shall keep the spacecraft passivated with a reliability of at least 0.95.
4) The passivation device shall be compatible with an operational voltage of up to 200 V.
5) The passivation device shall be compatible with an operational current of up to 250 A.
6) The passivation switch shall be able to sustain a passivated state under a temperature envelope of -150 to +120 °C.
7) The passivation switch shall be compatible with a mission duration of 18 years under the conditions of a GEO orbit.
8) The passivation switch should be able to maintain a passivated state for at least 25 years under the conditions of a GEO orbit.

In addition, a supporting radiation analysis was performed to characterize the expected radiation environment.

A. Radiation Environment Analysis

A radiation analysis executed in SPENVIS determined the required radiation hardness of the chosen solution. Two simulations were performed. The first evaluated the TID accumulated during the spacecraft's 18-year mission, the results of which can be found in figure 2. The second analyzed the TID received during a period of 25 years after the end of the mission, with the results displayed in figure 3. These two cases need to be distinguished as some potential radiation-induced failure modes may be acceptable after the spacecraft has been passivated.

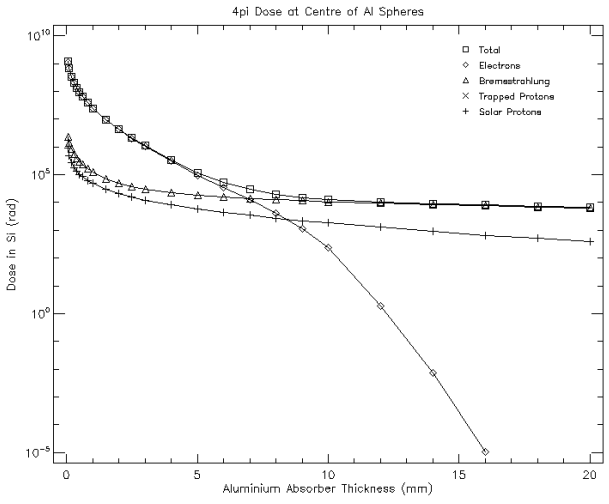

Fig. 2. TID accumulated during the spacecraft's lifetime

As can be expected the TID, disregarding the 11 year solar cycle, grows linearly with time. Assuming an equivalent shielding of 2.5 mm Al, we can say that the device will experience about 98 krad(Si) in the first 18 years and about 152 krad(Si) in the subsequent 25 years.

IV. Architectural Design

A trade-off analysis weighing several factors, such as implementation complexity, added mass and architectural

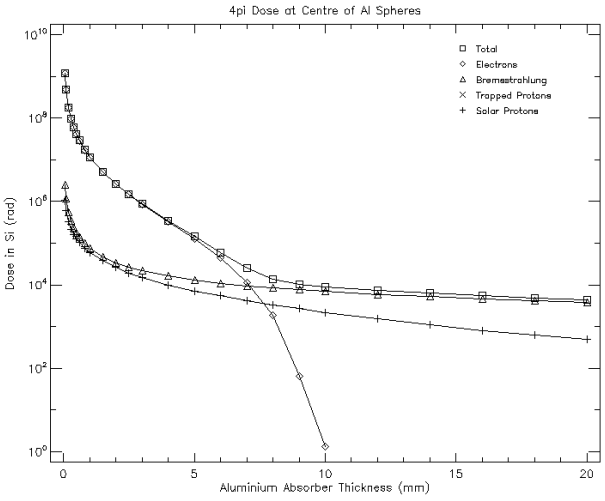

Fig. 3. TID accumulated during 25 years after the end of mission

universality, was performed to compare three architectural approaches - a dedicated passivation unit, passivation implementation in the power conditioning and distribution unit (PCDU) and passivation implemented through individual switches.

Although a dedicated unit was deemed the most architecturally universal solution and an implementation in the PCDU the most mass efficient, the overall best solution was determined to be an implementation via individual switches on the battery and the solar array. An architectural block diagram of a centralized EPS employing this solution can be found in fig. 4 below. As can be seen, the SA passivation switch is placed behind the solar array drive mechanism (SADM) as that allows unimpeded access the power lines of individual SA sections. On the other hand the battery passivation switch is to be implemented along the battery management system (BMS). This placement allows the shortest path to the battery proper and also easy access to the spacecraft's structure.

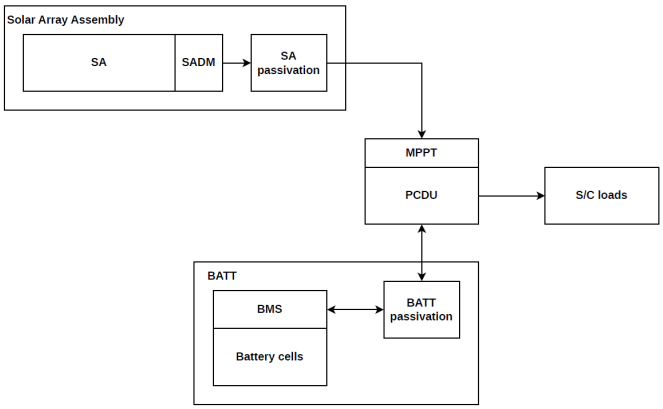

Fig. 4. Passivation implementation in a centralized EPS architecture

V. Preliminary Battery Switch Design

A. Passivation Mechanism

Another trade-off analysis was performed in order to select the optimal mechanism from the selection outlined above in section II-A. The weighing factors included universality (denoting the possibility of implementation in both the battery and the SA), reversibility, radiation tolerance, mechanical endurance, long-term stability and reliability. As already explained, each mechanism brings different strengths and weaknesses, however, based on the analysis, it was deemed that the overall best performance can be expected from mechanical relays. While these tend to degrade over time, their degradation is seen mostly in their cycle life, contact degradation and potential long-term leaking, all of which are not of concern in this device.

B. Mechanism Implementation

To implement the battery passivation switch a serial-parallel system of 4 electromagnetic switches was chosen and can be seen below in fig. 5. This is the minimal configuration compliant with the design requirements. A series of 2 switches would violate common requirements on single-points-of-failure, while a parallel configuration of 2 switches would not reach the desired reliability. Under normal operation, the switches connect the spacecraft battery with the main power bus. When passivation is initiated, the relays switch and disconnect power supply to the spacecraft. Instead connecting the battery to the ground through a block of high ohmic resistors with the goal of slowly discharging the battery.

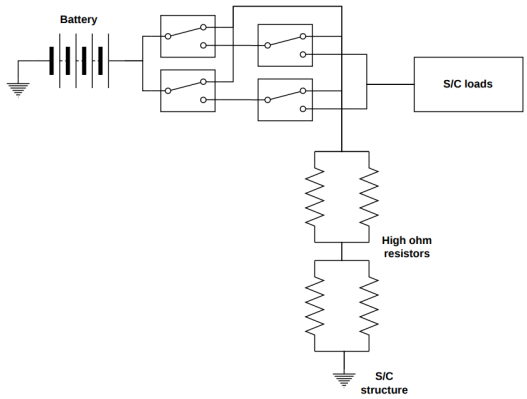

Fig. 5. Battery passivation switch schematic

The specific configuration of two-by-two relays was chosen to provide reliability not only for the passivation initialization and sustainment itself, but also so that the design would be robust against failures threatening to cut the supply before the end of mission.

A similar solution can be implemented for solar array passivation as well. To adjust the design the relays would need to connect positive and negative terminals of the SA sections in order to short them. While architecturally the switch would

remain almost unchanged, the selected components would need to be reviewed as in this case they would need to be compliant with an unregulated power bus since the switch would be implemented before the solar power is conditioned in the PCDU.

To evaluate the feasibility of this solution, a preliminary selection of a specific latching relay had to have been performed. As relays for such high-power applications and extreme environment are often custom made, Tyco Electronic's KILOVAC EV200 for terrestrial applications was chosen as a baseline. This relay does not only complies with the outlined voltage and current requirements, but also is built to withstand high vibration environments what will be experienced during launch, however further testing of its radiation resistance would be necessary.

C. Preliminary Reliability Assessment

A preliminary reliability assessment was performed in support of the design. As reliability data of the preselected component are not publicly available, an estimate according to MIL-HDBK-217F was performed. With the expected environment, performance and part quality taken into account, the predicted failure rate was estimated as $1.72 \cdot 10^{-6}$ per flight hour. In addition, a significant effect on the final reliability will have the design of the control scheme as well as of the battery interface. As both of these are currently unknown, they were omitted from the analysis.

Two main failure modes of interest were studied - a failure to perform the passivation and a failure to sustain it. Both of these are guided by the same reliability block diagram (RBD) presented in figure 6 below, with the only differentiation being the time duration during which the failure can occur. As can be seen in the diagram, the resistor block does not play an active role in either of the failures and only the relays affect the result.

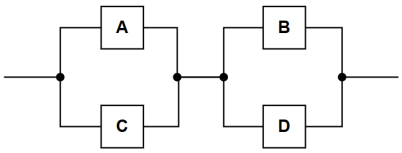

Fig. 6. Reliability block diagram of the battery passivation switch

The RBD was analysed through a decomposition method and the following term describing the system's overall failure rate was found:

$$\lambda_S = (1-(1-\lambda_A)\cdot(1-\lambda_C))\cdot(1-(1-\lambda_B)\cdot(1-\lambda_D)) \quad (1)$$

From this we can find that the failure rate of the system per flight hour is $1.18 \cdot 10^{-11}$. A failure probability for the whole 18 year mission and the subsequent 25 year duration was calculated and an overall reliability was determined. It was found that the failure probability for the initialization of the passivation is $1.86 \cdot 10^{-6}$ corresponding to a reliability of 0.999998. The failure probability of maintenance of a passivated state was determined to be $2.58 \cdot 10^{-6}$ corresponding to a reliability of 0.999997. As can be seen, both values are significantly above the required reliabilities listed in section III. It however needs to be noted that, as discussed above, any simpler configuration would not be compliant.

VI. Conclusion

This paper investigated the motivation behind and the principles of electrical passivation of spacecraft. Different technical approaches to the problem were discussed, including specifics of implementing the switch into the spacecraft's solar array and battery. Driving requirements of the design were defined, and a complementary radiation environment analysis in SPENVIS was performed to determine the expected total irradiation dose. Next, the results of a trade-off analysis were used to determine the optimal architectural approach, which was then defined. Another trade-off analysis served to select a suitable passivation mechanism, and a preliminary battery passivation switch design was created. The reliability of the selected solution was then evaluated using a reliability block diagram to confirm the design's alignment with previously presented requirements. The design was confirmed to be aligned with the requirements and can further be elaborated into a detailed design. Further, at the moment the specific failure mode distribution is not known as it is not given by MIL-HDBK-217F. As such RAC EPRD database will further need to be consulted. Lastly, the chosen part is significantly oversized for the intended use, however it is one of only a few commercially available in this power range.

References

[1] L. Hall, The History of Space Debris, Space Traffic Management Conference, November 2014

[2] I. Eymen. (2025). Mitigation of space-debris from ruptured Li-ion batteries. 10.13009/EUCASS2025-600.

[3] F. Bausier, F. Tonicello, T. Soares. (2014). Passivation of Spacecraft Electrical Power System at End of Mission

[4] F. Bausier, M. Nestoridi, B. Samaniego López, J. Simola, A. Wolahan, J. Austin, T. Soares. (2017). Spacecraft Electrical Passivation: From Study to Reality. E3S Web of Conferences. 16. 13002. 10.1051/e3sconf/20171613002.

[5] W. Ley, K. Wittmanss, W. Hallmann. Handbook of space technology. Chichester: Wiley, 2009. ISBN 978-0-470-69739-9.

[6] S. Hull, W. Schonberg. (2022). Spacecraft passivation – An overview of requirements, principles, and practices as applied to spacecraft pressure vessels. Journal of Space Safety Engineering. 9. 10.1016/j.jsse.2022.08.004.

[7] D. Prevot,A. F. Castric, E. Pasquier, Batteries and Bypass Switch: Saft Heritage and Standard Search, 8th European Space Power Conference, 2008, vol. 661, Art. no. 107.

[8] A. Teverovsky. (1999). Relay Failures Specific to Space Applications. 10.31399/asm.cp.istfa1999p0285.

ŠKODA Kariéra



Stáže a závěrečné práce

Studuješ vysokou školu a zároveň chceš získat pracovní zkušenosti? Zkus odbornou stáž ve ŠKODA AUTO. Pod vedením našich odborníků se zapojíš do inovativních projektů – třeba těch v oblasti elektromobility. Získané znalosti pak můžeš zúročit i ve své závěrečné práci.

Doktorandský program

Aplikuj výsledky svého výzkumu v reálném prostředí, využívej nejnovější technologie a dej své disertační práci nový rozměr. Poskytneme ti flexibilní pracovní dobu a možnost spolupracovat s profesionály ve svém oboru.

Trainee program

Dokončil jsi studia a přemýšlíš kam dál? Trainee program pro absolventy ti pomůže najít obor, který tě bude opravdu bavit a naplňovat. V průběhu jednoho roku budeš na plný úvazek pracovat v mezinárodním týmu, poznáš různá oddělení a vycestuješ na zahraniční stáž.

Volné pozice

Vyber si ze široké nabídky volných pozic, staň se součástí nejprogresivnější české firmy a tvoř s námi budoucnost automobilového průmyslu.

Škoda Auto Kariéra @WeAreSKODA Škoda Auto a.s. Škoda Auto - Career

Comparative Study on the Electrochemical Performance of Lead Foil in Different Electrolytes for Sodium-Ion Batteries

Jan Jilčík
Department of Electrical and electronic Technology
FEEC Brno University of Technology
Brno, Czech Republic
240631@vutbr.cz

***Abstract* – This paper investigates lead as a promising anode material for sodium-ion batteries. The main objective was to evaluate the influence of electrolyte composition on the overall electrochemical stability and kinetics of a system utilizing a 50 µm thick lead foil anode. To mitigate the catalytic decomposition of the electrolyte on the lead electrode surface, two distinct electrolytes were compared. The first system employed carbonate-based solvents with a fluoroethylene carbonate (FEC) additive, while the second used an ether-based solvent, dimethoxyethane. The results indicate that the FEC additive promotes the formation of a more stable Solid Electrolyte Interphase (SEI) layer, resulting in higher overall electrochemical capacity and stability of the cell.**

***Keywords*—anode, lead, lead foil, sodium-ion batteries, electrolyte, FEC, DME**

I. Introduction

Sodium-ion batteries (SIBs) currently represent one of the most promising alternatives to the dominant lithium-ion technology, primarily due to the high natural abundance of sodium, its low cost, and environmental sustainability [1], [2]. Given the growing demand for large-scale stationary energy storage, the development of efficient and low-cost anode materials for NIBs is a key technological challenge [3].

Lead emerges as an attractive candidate for SIB anodes due to its ability to form alloys with sodium, offering a high theoretical capacity that reaches 485 mAh/g upon final saturation and the formation of the $Na_{15}Pb_4$ phase [4]. This value significantly exceeds the capacity of commonly used hard carbon anodes. However, the main obstacle to the commercialization of lead anodes is the extreme volume changes, reaching up to 400 %, that occur during the sodium alloying and dealloying processes These expansions lead to the mechanical degradation of the active material, loss of electrical contact, and a subsequent rapid capacity fade [4], [5].

Another critical issue is the high reactivity of the lead surface, which, combined with a large specific surface area, leads to massive catalytic electrolyte decomposition. This process not only reduces the Coulombic efficiency in the initial cycles but also results in the formation of an unstable and highly resistive passivation layer on the surface [6]. To address these shortcomings, advanced techniques are utilized in practice, such as controlled electrochemical activation [7] the application of specific electrolyte additives like fluoroethylene carbonate (FEC), or the use of ether-based solvents such as dimethoxyethane (DME), which facilitate the formation of a more stable interphase [6], [8]. This work focuses on the experimental comparison of these approaches using a lead foil.

II. Preparation of the Anode Material

The electrodes were prepared using a lead sheet conventionally employed in the manufacturing of lead-acid batteries. Initially, the sheet was subjected to a combined mechanical and chemical cleaning process to remove grease and surface impurities. The surface was scrubbed using fiberglass cloth soaked in ethanol. Following the cleaning procedure, the sheet was rolled down to a thickness of 50 µm. Electrodes were subsequently punched from the resulting foil using a 16 mm punch. To eliminate any residual impurities, the punched electrodes were immersed in ethanol for 48 hours, with the solvent being replaced after the first 24 hours. Finally, the electrodes were dried in a vacuum oven at 60 °C for 48 hours.

III. Material Characterization

The 50 µm thick lead foil was examined using scanning electron microscopy (SEM). The acquired SEM image (Fig. 1) confirmed a continuous and homogeneous surface devoid of any significant macroscopic defects, revealing only rolling marks associated with the material's fabrication. Subsequent energy-dispersive X-ray spectroscopy (EDS) analysis demonstrated the high purity of the sample. The elemental composition was determined to be 96.8 wt. % lead (Pb), only 0.2 wt. % oxygen (O), and 3.0 wt. % carbon (C). The detected carbon content is likely a standard artifact originating from the

carbon tab used for sample mounting and typical atmospheric contamination.

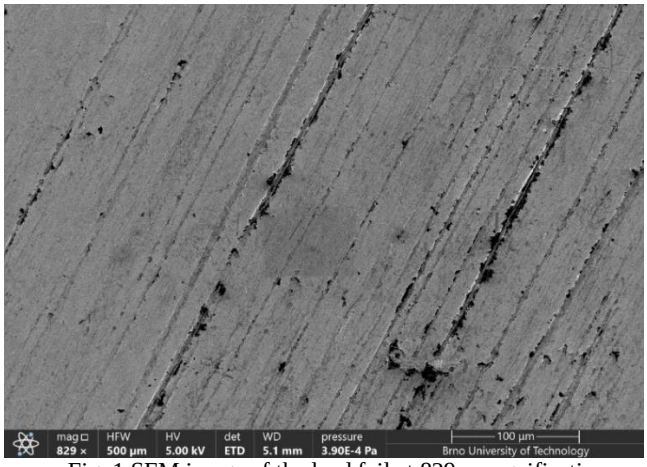

Fig. 1 SEM image of the lead foil at 829x magnification

IV. Electrolytes

Two types of electrolytes were prepared for the electrochemical testing. In both cases, sodium hexafluorophosphate ($NaPF_6$) at a concentration of 1 M was used as the conducting salt.

For the initial testing of the 50 µm thick lead foil, a solvent mixture of ethylene carbonate, propylene carbonate and diethyl carbonate (EC:PC:DEC) in a 1:1:1 volume ratio was employed, supplemented with 5 wt. % of fluoroethylene carbonate (FEC). This specific additive was incorporated into the system primarily to facilitate the controlled formation of a stable SEI layer on the lead electrode surface.

For the subsequent series of measurements with the lead foil, an alternative ether-based solvent was selected. In this case, the sodium hexafluorophosphate ($NaPF_6$) was dissolved in pure dimethoxyethane (DME).

V. Cell assembly

EL-CELL test cells were employed to assemble the half-cells comprising the lead anode and sodium counter electrode. These cells feature a spring-loaded mechanism that ensures constant pressure on the electrodes, while their sealing effectively prevents any air ingress. The cell assembly was carried out in a glovebox under an inert argon atmosphere. Metallic sodium was utilized as the counter electrode. To eliminate the potential formation of micro-short circuits, two separators were incorporated into each cell. The electrolyte volume was kept identical in both configurations, totalling 200 µl (100 µl + 100 µl); specifically, each separator was wetted with 100 µl of the respective electrolyte.

For a comprehensive evaluation, two distinct cell configurations were prepared:

1. Lead foil electrode – 1 M $NaPF_6$ in EC:PC:DEC + 5 wt. % FEC electrolyte, 2 separators
2. Lead foil electrode – 1 M $NaPF_6$ in DME electrolyte, 2 separators

VI. Electrochemical measurements

All electrochemical measurements of the assembled cells were conducted at room temperature using cyclic voltammetry (CV). The measurement parameters were identical for both cell configurations. Cycling was performed within a potential window of 0.05 V to 1.3 V. The scan rate was reduced in a stepwise manner: the initial four cycles were recorded at a scan rate of 5 mV/s, the subsequent four cycles at 1 mV/s, and the final four cycles at 0.1 mV/s.

VII. Results and discussion

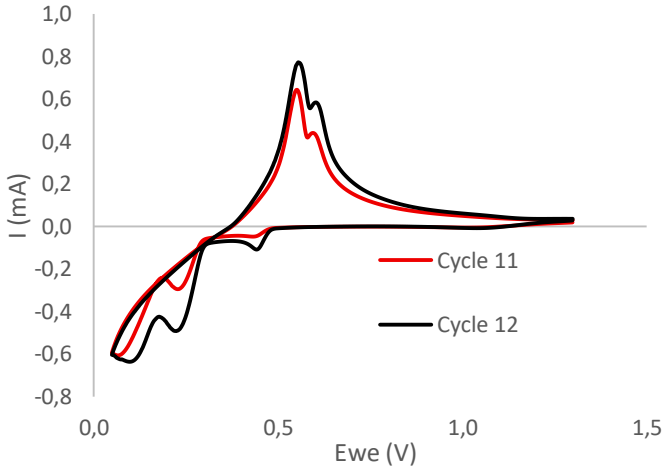

Fig. 2 Cyclic voltammograms of the lead foil electrode in EC:PC:DEC + FEC electrolyte

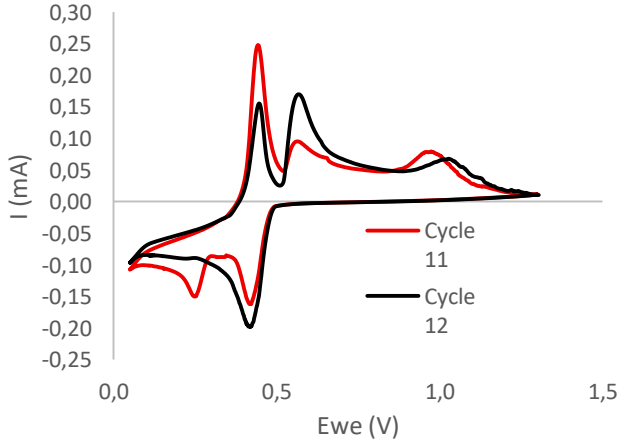

Fig. 3 Cyclic voltammograms of the lead foil electrode in DME electrolyte

To directly assess the influence of the electrolyte composition on the overall performance of the system, two cells were tested under identical conditions. The mass of both electrodes was identical, and each cell utilized two separators along with the exact same volume of electrolyte. The first cell contained the $NaPF_6$ in EC:PC:DEC + 5 wt. % FEC electrolyte, while the second cell employed the $NaPF_6$ in DME electrolyte.

A comparison of the cyclic voltammograms (Fig. 2 and 3) reveals that the system with the EC:PC:DEC + FEC electrolyte exhibits an approximately threefold higher current response compared to the cell with the DME solvent. This indicates that the cell with the EC:PC:DEC + FEC electrolyte achieves higher electrochemical capacity and activity. During the CV cycles in the FEC-containing system, the individual peaks corresponding to specific electrochemical reactions do not change their characteristics. A gradual increase in peak height is observed, indicating a stable environment for the electrochemical processes and a progressive increase in the active surface area of the electrodes due to pulverization. We hypothesize that the FEC additive actively participates in the formation of a stable SEI layer on the lead electrode surface, which effectively mitigates the catalytic decomposition of the electrolyte while demonstrating improved permeability for sodium ions [6].

The positions of the cathodic and anodic peaks in the carbonate-based electrolyte with FEC coincide with the peak positions of the DME system (Fig. 2 and 3). However, these two systems differ in peak shape and current response. Another distinction is that the system with FEC exhibits three distinct cathodic peaks, in contrast to the DME system, which displays two cathodic peaks followed by only a slight current increase at the lowest potential. Conversely, the system with FEC does not exhibit an anodic peak in the region of approximately 1.0 V. Nevertheless, it can be assumed that a third anodic peak (in the region of approximately 1.0 V) will fully develop during subsequent cyclic voltammetry cycles, as the electrochemical activation and peak evolution in this system have been observed to proceed gradually.

Based on a comparison with previously published literature regarding lead anodes, it can be assumed that the observed peaks in the cyclic voltammograms correspond to the progressive formation of intermetallic sodium-lead phases (Na_xPb). The precise position and shape of the measured peaks differ slightly from the results reported in the literature for lead powder anodes. This discrepancy is likely attributed to the morphology of the lead foil. The cathodic peak in the 0.40 – 0.50 V region most likely corresponds to the initial $NaPb_3$ phase, the subsequent peak in the 0.25–0.30 V region corresponds to the NaPb phase, and the peak at the lowest potential (< 0.15 V) presumably encompasses merged transitions into the most highly sodium-saturated phases, namely the Na_9Pb_4 and $Na_{15}Pb_4$ phases. During the anodic scan, the lead foil exhibits a significant shift in the onset of oxidation processes towards higher potentials. While the literature reports the first anodic peaks appearing as early as the 0.2 V region, the current response for the tested foil electrodes only begins in the 0.45 V region. These anodic peaks likely correspond to the progressive reverse desodiation of the lead electrode [8], [9].

In contrast to the first cell, the cell utilizing the DME solvent exhibits distinct cathodic and anodic peaks during the 11th cycle. However, in the immediately subsequent cycle, the cathodic peak in the 0.25 V region disappears, and the intensity of the anodic peaks changes significantly. This loss of activity may indicate the instability of the electrochemical reactions or the onset of mechanical degradation of the electrode and the subsequent loss of electrical contact of the active material due to volume expansion.

VIII. Conclusion

In this work, the influence of electrolyte composition on the electrochemical stability and overall capacity of a 50 µm lead foil in a sodium-ion battery (SIB) system was investigated. The objective was to identify an optimal environment that would prevent premature material degradation caused by both surface passivation and volume changes. Two electrolytes containing the same sodium salt, but different solvents were compared under identical conditions. The first electrolyte utilized a carbonate-based solvent consisting of EC:PC:DEC supplemented with 5 wt. % FEC, while the second employed an ether-based DME solvent.

The results of the electrochemical measurements, specifically cyclic voltammetry, indicate that the electrolyte composition has a profound impact on the performance of the lead foil. Although the system with DME exhibited complete reversible activity, it failed to sustain this activity over the long term, presumably due to volume changes at the electrode surface. In contrast, the system with FEC demonstrated superior electrochemical behaviour. We hypothesize that the FEC actively participates in the formation of a highly stable SEI layer. Consequently, the FEC-containing system achieves a threefold higher current response for the same electrode mass and dimensions. Based on these findings, the application and further investigation of a lead foil in combination with a carbonate-based electrolyte containing the FEC additive emerges as a highly promising direction for the development of stable, high-capacity anode materials for sodium-ion batteries (SIBs).

Acknowledgment

The author would like to thank their supervisor for expert guidance and consultations during the interpretation of electrochemical data.

Declaration of AI tools

During the preparation of this work, author used Gemini Pro 3.1 tool for literature review, text enhancement and readability

improvement. After using this tool, author reviewed and edited the content as needed.

REFERENCES

[1] J. Jilčík, "Výzkum elektrodových materiálů pro sodno-iontové baterie", Accessed: Mar. 01, 2026. [Online]. Available: https://www.vut.cz/studenti/zav-prace/detail/159651

[2] A. Eftekhari and D. -W. Kim, "Sodium-ion batteries: New opportunities beyond energy storage by lithium", *Journal of Power Sources*, vol. 395, 2018, doi: 10.1016/j.jpowsour.2018.05.089.

[3] J. F. Peters, A. Peña Cruz, and M. Weil, "Exploring the Economic Potential of Sodium-Ion Batteries", *Batteries*, vol. 5, no. 1, 2019, doi: 10.3390/batteries5010010.

[4] J. Zhang *et al.*, "Identifying Problematic Phase Transformations in Pb Foil Anodes for Sodium-Ion Batteries", *Journal of The Electrochemical Society*, vol. 171, no. 9, Sep. 2024, doi: 10.1149/1945-7111/ad76e0.

[5] A. Ponrouch, E. Marchante, M. Courty, J. -M. Tarascon, and M. R. Palacín, "In search of an optimized electrolyte for Na-ion batteries", *Energy & Environmental Science*, vol. 5, no. 9, 2012, doi: 10.1039/c2ee22258b.

[6] L. D. Ellis, B. N. Wilkes, T. D. Hatchard, and M. N. Obrovac, "In Situ XRD Study of Silicon, Lead and Bismuth Negative Electrodes in Nonaqueous Sodium Cells", *Journal of The Electrochemical Society*, vol. 161, no. 3, Dec. 2013, doi: 10.1149/2.080403jes.

[7] J. Libich, "Záporná elektroda pro lithno-iontové akumulátory", Dizertační práce, 2015. Accessed: Dec. 29, 2025. [Online]. Available: https://dspace.vut.cz/items/3c75ef3d-6221-48f1-b76d-1e607bd7c7a9

[8] A. Darwiche, R. Dugas, B. Fraisse, and L. Moncondiut, "Reinstating lead for high-loaded efficient negative electrode for rechargeable sodium-ion battery", J*ournal of Power Sources*, Art. no. vol. 304, pp. 1-8, 2015, doi: 10.1016/j.jpowsour.2015.10.087.

[9] C. Kim *et al.*, "A high rate and long-cycle-life anode based on micrometer-sized Pb powder for sodium-ion batteries", *Journal of Alloys and Compounds*, vol. 886, 2021, doi: 10.1016/j.jallcom.2021.161240.

Apparent Power Time Aggregation Algorithms

Martin Poul
Department of Electrical Power Engineering
Brno University of Technology
Brno, Czech Republic
240795@vut.cz

***Abstract*—The increasing share of dynamically switched loads in modern distribution networks causes fluctuations in voltage, current, and with them associated power. The subject of this study is to identify which apparent power time aggregation algorithm provides the most accurate representation of energy transfer under conditions of dynamically switched resistive load. Using a simulation model in the PSCAD environment and a custom software analyzer, three aggregation algorithms are evaluated against a power loss equivalence criterion. The results demonstrate differences in accuracy among the algorithms, highlighting which method best reflects the used power loss criterion.**

***Index Terms*—apparent power, time aggregation, aggregation algorithms, power quality, power loss**

I. Introduction

The current transformation of power engineering towards decentralized and smart grids brings entirely new challenges in the field of energy flow measurement and evaluation. Currently, due to the integration of nonlinear and dynamically switched loads (heat pumps, demand-side management systems, electric vehicle charging stations), voltage and current fluctuations occur at end consumers due to dynamic changes in the consumed power. For the purposes of this paper, the phenomena in which these fluctuations appear will be referred to as dynamic events.

The primary task of aggregation algorithms is to process values measured in basic intervals (e.g., 200 ms for power quality analyzers or 1 s for electricity meters [1]) and provide a single equivalent representative value, such as apparent power, for a longer aggregation time interval (e.g., 1 minute or 15 minutes) [2].

The behavior of the algorithms in an environment with voltage and current fluctuations must be analyzed in detail. Therefore, this paper aims to test the accuracy of aggregation algorithms for the yield of apparent power in a scenario with a switched ideal resistive load.

The replacement of these algorithm methods in modern digital meters (e.g., smart meters) requires only firmware updates, offering a highly cost-effective approach to improving measurement accuracy without the need for hardware replacement [3].

II. Tested Aggregation Algorithms

Several aggregation algorithms can be used to determine the aggregated value of apparent power over a time interval. Let us assume that the aggregation window consists of N discrete samples, where each k-th sample carries information about the RMS value of voltage U_k, current I_k, and apparent power $S_k = U_k \cdot I_k$ in the given basic measurement interval window.

These three algorithms (S_1, S_2, S_3) were selected for comparison because they logically represent the fundamental mathematical approaches available for time aggregation: quadratic averaging of RMS voltage and current values (based on standard [5]), quadratic averaging of apparent power, and simple arithmetic averaging of apparent power. It is expected that the algorithms will yield different results during dynamic events.

A. Quadratic Averaging of RMS Values (Method S_1)

The first algorithm, defined in publication [4] and standard [5] does not aggregate the apparent powers themselves, but calculates the aggregated RMS voltage U_{agg} and the aggregated RMS current I_{agg} independently of each other. The resulting aggregated apparent power is their product.

$$S_1 = U_{\mathrm{agg}} \cdot I_{\mathrm{agg}} = \sqrt{\frac{1}{N}\sum_{k=1}^{N} U_k^2} \cdot \sqrt{\frac{1}{N}\sum_{k=1}^{N} I_k^2} \tag{1}$$

B. Quadratic Averaging (Method S_2)

The second considered algorithm, also based on [4], operates on the basis of quadratic averaging (root mean square) of the calculated discrete samples of apparent power over the basic measurement window.

$$S_2 = \sqrt{\frac{1}{N}\sum_{k=1}^{N} S_k^2} = \sqrt{\frac{1}{N}\sum_{k=1}^{N} (U_k \cdot I_k)^2} \tag{2}$$

C. Arithmetic Averaging (Method S_3)

The third analyzed aggregation algorithm is based on the simple averaging of apparent power values over time. Mathematically, it is a classical arithmetic mean, the output of which is the mean value of apparent power over time.

$$S_3 = \frac{1}{N}\sum_{k=1}^{N} S_k = \frac{1}{N}\sum_{k=1}^{N} (U_k \cdot I_k) \tag{3}$$

This algorithm is selected for comparison because it is one of the possible and commonly used aggregation approaches.

III. Methodology

To verify the aggregation algorithms, a single-phase low-voltage network model was developed in the PSCAD environment, on which a series of simulations focusing on the effects of dynamic and constant loads were performed. The results of the simulations were instantaneous voltage and current values in time, which were further evaluated by a software analyzer.

A. Model Parameters

The simulation network was modeled using standardized impedances from the PSCAD library, as illustrated in the simplified single-line diagrams in Figure 1.

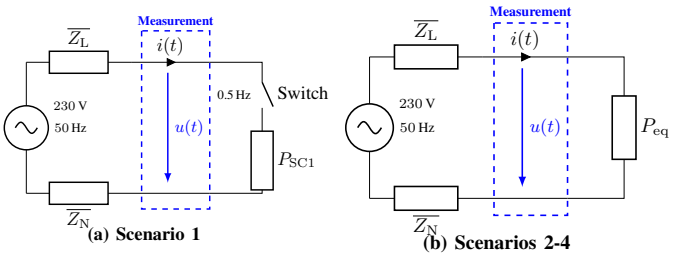

Fig. 1. Simplified single-line diagrams of the simulation model for a) simulation of a dynamic event, b) simulation of an event with constant equivalent load.

The phase line impedance $\overline{Z_L}$ was set to $(0.24 + j0.15)\ \Omega$ and the neutral line impedance $\overline{Z_N}$ to $(0.16 + j0.10)\ \Omega$. The inductance values were set for a nominal frequency of 50 Hz. The simulation had a solution step of 10 μs, hence the waveform sampling frequency was 100 kHz for greater evaluation accuracy. The duration of the simulations was always set to 21 s. However, due to transient phenomena occurring at the start of the simulation, the first second of the record was not used, and a time interval of length $T_{\text{sim}} = 20$ s was evaluated. The ideal resistive load was used for both dynamically switched and constant loads.

The basic premise for comparing the individual aggregation algorithms is based on one of the physical definitions of apparent power S. This definition states that apparent power is the value of constant active power that, when transmitted through a given line, would cause the exact same power loss ΔP as the examined dynamic, asymmetrical, or nonlinear state. The correctness of any algorithm for calculating S can therefore be easily verified by comparing the power loss on the line ΔP [1], [6]. While alternative theoretical definitions of apparent power exist, this power loss equivalence criterion was chosen as the standard of correctness in this paper. Since the evaluation relies on purely simulated data, hardware measurement uncertainties and transducer errors are not considered in this study. This allows for an isolated assessment of the mathematical accuracy of the compared algorithms.

B. Definition of Test Scenarios

The simulation model was divided into four scenarios.

- **Scenario 1 (Dynamic event):** A single-phase load consuming an active power of $P_{\text{SC1}} = 3680$ W ($\cos\varphi = 1$). The load is switched with a 50 % duty cycle (frequency 0.5 Hz, i.e., 1 s on, 1 s off). The voltage and current fluctuations during load switching are shown in Figure 2. The changes in voltage and current are inversely related.

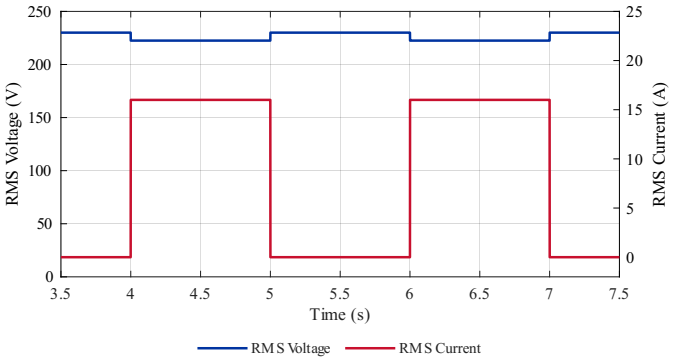

Fig. 2. Voltage and current changes during dynamic load switching in Scenario 1

As the current begins to flow, the voltage in the circuit drops. The aim of this scenario is to measure the power loss on the line (ΔP_{SC1}) and obtain the aggregated values of apparent power S_1, S_2, and S_3 according to the aggregation algorithms presented in Section II.

- **Scenarios 2–4 (Testing of algorithms S_1, S_2, and S_3):** A static constant load is connected to the network, consuming an active power P_{eq} corresponding to the apparent power value determined by the respective aggregation algorithm (1), (2), or (3) (i.e., $P_{\text{eq}} = S_{\text{x}}$). The line losses during this constant consumption are compared with the reference losses from Scenario 1.

C. Software Analyzer and Data Processing Pipeline

The instantaneous voltage and current values in time from the simulation were exported in .csv format and evaluated using a software analyzer developed in Python. The complete data processing pipeline is illustrated in the block diagram in Fig. 3.

The analyzer uses a basic measurement interval of $t_k = 200$ ms, although this can be freely adjusted. The set aggregation interval corresponds exactly to the evaluated length of the waveform, which in this case is $T_{\text{sim}} = 20$ s. First, the computing core handles the loading and validation of the input data. Next, the synchronization and resampling section of analyzer detects zero crossings on the filtered fundamental harmonic voltage waveform. To prevent spectral leakage, analyzer resamples the instantaneous values to a new sampling frequency.

From these resampled data, the components within the basic measurement interval are then calculated using defined computational algorithms. These calculated components include the RMS values of voltage U_k, current I_k, and apparent power S_k used in this paper. An extension of the analyzer is the subsequent implementation of the time aggregation algorithms (1), (2), and (3).

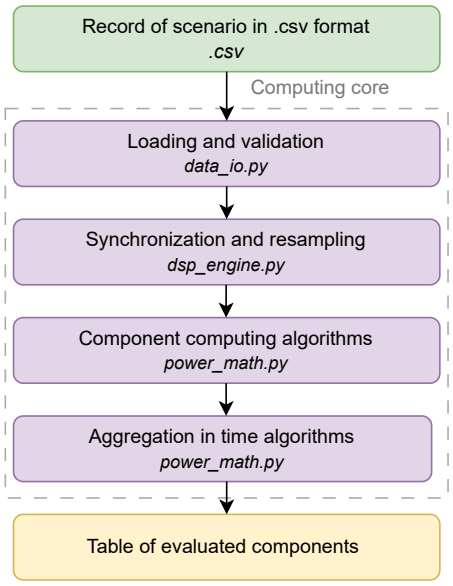

Fig. 3. Software analyzer data processing pipeline.

Furthermore, the formulas used for calculating the power loss in the basic measurement window ΔP_k and the total power loss in the aggregation interval ΔP are applied.

$$\Delta P = \frac{1}{T_{\text{sim}}} \sum_{k=1}^{N} (\Delta P_k \cdot t_k)$$
$$\Delta P = \frac{1}{T_{\text{sim}}} \sum_{k=1}^{N} \left((R_L + R_N) \cdot I_k^2 \cdot t_k \right) \tag{4}$$

This software analyzer was developed entirely by the author of this paper. While the definitions of the evaluated components are based on [4] and [6], the implementation of the time aggregation algorithms (1), (2), and (3), along with the automated evaluation pipeline, are original contributions.

D. Evaluation Process and Post-Processing

In Scenario 1, the aggregated apparent power results S_1, S_2, and S_3 were determined. These were then set in individual Scenarios 2 to 4 for a constant load consuming active power. When evaluating these scenarios, the resulting value of the power loss ΔP was observed. The algorithm whose resulting ΔP is closest to the value determined in Scenario 1, i.e., $\Delta P_{SC1} = 48.47$ W, is considered the correct aggregation formula.

To compare the differences between the individual methods, the absolute and relative deviations from the power loss value ΔP_{SC1}, which is considered the reference, were calculated.

$$\Delta_{\Delta P_{\text{SCx}}} = \Delta P_{\text{SCx}} - \Delta P_{\text{SC1}} \tag{5}$$

$$\delta_{\Delta P_{\text{SCx}}} = \frac{\Delta P_{\text{SCx}} - \Delta P_{\text{SC1}}}{\Delta P_{\text{SC1}}} \cdot 100\,\% \tag{6}$$

where x designates the given Scenario 2 to 4.

IV. Algorithm Analysis Results

After analyzing the instantaneous voltage and current waveforms using the software analyzer, the results were subjected to a more detailed examination. A summary of these results is presented in Table I. The indices SC1 to SC4 designate the given scenario for which the value was calculated. In Scenario 1, all algorithms are evaluated simultaneously, whereas in Scenarios 2–4, a constant equivalent load based on the respective algorithm's result is used.

TABLE I
Comparison of Aggregation Results and Power Losses

Evaluated results	**Scenario 1** P_{SC1}	**Scenario 2** $P_{\text{eq}} = S_1$	**Scenario 3** $P_{\text{eq}} = S_2$	**Scenario 4** $P_{\text{eq}} = S_3$
S_1 (VA)	2508.0	2508.0	2474.5	1818.5
S_2 (VA)	2474.5	2508.0	2474.5	1818.5
S_3 (VA)	1818.5	2508.0	2474.5	1818.5
ΔP_{SC} (W)	48.47	49.05	47.72	25.51

From the results of Scenario 1, i.e., the dynamic event, it was found that the individual presented aggregation algorithms give different results. While aggregation algorithms S_1 and S_2 provided similar results, algorithm S_3 calculated a much lower value. When compared with Scenarios 2 to 4, where no voltage fluctuations occurred over time, it is evident that if the voltage or current magnitude does not change during the aggregation interval, there are no differences among the aggregated results. This comparison of results in per-unit (p.u.) values was plotted in the graph in Figure 4. The base value is $S_{\text{base}} = S_1 = 2508$ VA.

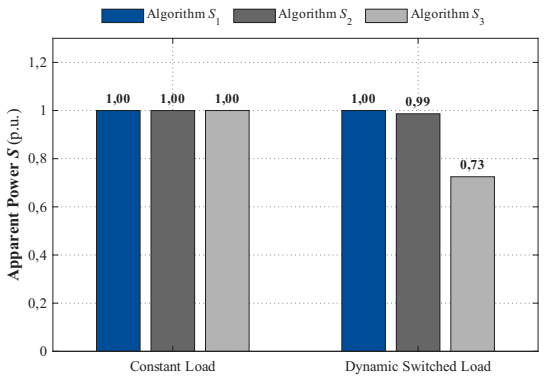

Fig. 4. Comparison of aggregation algorithm results in per-unit (p.u.) values

The comparison and calculation of the relative error showed that the value closest to the power loss $\Delta P_{\text{SC1}} = 48.47$ W is $\Delta P_{\text{SC2}} = 49.05$ W. Thus, according to the definition, the aggregation algorithm S_1 performs the aggregation most accurately, with a deviation from Scenario 1 of $\delta_{\Delta P_{\text{SC2}}} = 1.2$ %. The resulting value of the aggregated apparent power S_2 was, based on the power loss comparison $\Delta P_{\text{SC3}} = 47.72$ W, also close to the value ΔP_{SC1} from Scenario 1. However, the

relative deviation was larger than that of algorithm S_1, namely $\delta_{\Delta P_{\mathrm{SC3}}} = -1.5$ %. The largest deviation was exhibited by the value ΔP_{SC4} when using algorithm S_3. Here, the deviation from the power loss in Scenario 1 was $\delta_{\Delta P_{\mathrm{SC4}}} = -47.4$ %. These differences are summarized in Table II and Figure 5.

TABLE II
POWER LOSS DIFFERENCES IN THE DEFINED SCENARIOS

Evaluated results	**Scenario 1** P_{SC1}	**Scenario 2** $P_{\mathrm{eq}} = S_1$	**Scenario 3** $P_{\mathrm{eq}} = S_2$	**Scenario 4** $P_{\mathrm{eq}} = S_3$
ΔP_{SC} (W)	48.47	49.05	47.72	25.51
$\Delta_{\Delta P_{\mathrm{SC}}}$ (W)	—	0.58	-0.75	-22.96
$\delta_{\Delta P_{\mathrm{SC}}}$ (%)	—	1.2	-1.5	-47.4

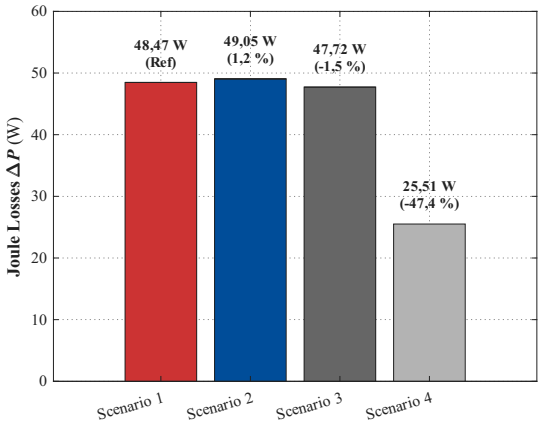

Fig. 5. Comparison of power losses

The differences between the algorithms can be analytically proven. The divergence between S_1 and S_2 is governed by the covariance of the squared voltage and current: $S_2^2 = S_1^2 + Cov(U^2, I^2)$. Since the current increase causes a voltage drop on the line impedance, these variables are negatively correlated ($Cov < 0$), mathematically dictating that $S_2 < S_1$ [7]. Furthermore, the relationship between quadratic method S_2 and the arithmetic mean S_3 is driven by the variance of the apparent power. According to the inequality of arithmetic and quadratic means, the arithmetic average of fluctuating values is inherently lower than their quadratic average [8].

V. CONCLUSION

This paper dealt with the issue of apparent power aggregation in a modern type of network with dynamic consumption. Using a simulation in the PSCAD program and a software analyzer, the behavior of three different aggregation algorithms for calculating the apparent power S was compared. For the individual aggregation algorithms, it was found which of them most accurately calculates the resulting aggregated apparent power over the aggregation interval.

The simulation results proved that in a network with voltage and current fluctuations over time, i.e., in the presence of dynamic events, not all presented aggregation methods provide the same results. While aggregation algorithms S_1 and S_2 provide very similar results, the aggregation algorithm S_3 has a significantly different result. These divergent results are shown in Figure 4.

This was followed by a comparison of which of the aggregation algorithms provides the most accurate result. From the perspective of equivalent power loss, the aggregation algorithm using the quadratic averaging of RMS values S_1 appears to be the most accurate. The algorithm of quadratic averaging of apparent powers S_2 can be considered less accurate. The most significant deviation was given by the aggregation algorithm based on simple arithmetic averaging S_3. These deviations were plotted in Figure 5 and Table II.

However, it is important to acknowledge that the results presented in this study rely on a specific dynamic scenario, which is a purely resistive load switching at a fixed frequency of 0.5 Hz. While this idealized scenario clearly demonstrates the mathematical differences between the algorithms, real-world applications often involve nonlinear current sources and variable reactive power. The conclusions regarding the accuracy of algorithm S_1 are tied to this specific scenario. Future research should include additional scenarios with reactive and highly distorted loads to fully generalize these findings.

From the perspective of the future development of the Smart Grid concept, it is therefore necessary to apply a uniform aggregation algorithm across all measuring devices, otherwise, different results will be reported. Incorrect data from measuring instruments can lead to erroneous results in subsequent calculations that utilize the currently measured aggregated values of apparent power as an input component for further analyzes.

REFERENCES

[1] ČSN EN 61000-4-30 ed. 3:2015, *Electromagnetic compatibility (EMC) – Part 4-30: Testing and measurement techniques – Power quality measurement methods*, Czech Office for Standards, Metrology and Testing, 2015.

[2] M. Calin, A. M. Dumitrescu, M. Asprou, E. Kyriakides, and M. Albu, *Measurement data aggregation for active distribution networks*, in 2013 4th IEEE International Workshop on Applied Measurements for Power Systems (AMPS), Aachen, Germany, 2013, pp. 144-149.

[3] A. J. Berrisford, *A Smarter Meter: IEEE-1459 Power Definitions in an Off-the-Shelf Smart Meter*, in 2015 IEEE International Instrumentation and Measurement Technology Conference (I2MTC), Pisa, Italy, 2015, pp. 1-6.

[4] J. Klusáček, *Definition and Measurement of Electricity Components in Active Distribution Systems*, Ph.D. dissertation, Brno University of Technology, 2025.

[5] ČSN EN IEC 61557-12 ed. 2:2022, *Elektrická bezpečnost v nízkonapěťových rozvodných sítích se střídavým napětím do 1 000 V a se stejnosměrným napětím do 1 500 V - Zařízení ke zkoušení, měření nebo sledování činnosti prostředků ochrany - Část 12: Zařízení pro měření a sledování výkonu (PMD)*, Úřad pro technickou normalizaci, metrologii a státní zkušebnictví, 2022.

[6] IEEE Standard 1459-2010, *IEEE Standard Definitions for the Measurement of Electric Power Quantities Under Sinusoidal, Nonsinusoidal, Balanced, or Unbalanced Conditions*, IEEE Power and Energy Society, 2010.

[7] A. Papoulis and S. U. Pillai, *Probability, Random Variables, and Stochastic Processes*, 4th ed. New York, NY, USA: McGraw-Hill, 2002.

[8] G. H. Hardy, J. E. Littlewood, and G. Pólya, *Inequalities*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 1952.

Compact Redesign and Safety Upgrade of a Bidirectional Power Supply for Flexible Automotive Testing

1st Martin Kyseľ
Faculty of Electrical Engineering and Communication
Brno University of Technology
Brno, Czech Republic
240767@vut.cz

2nd Kamil Jaššo
Faculty of Electrical Engineering and Communication
Brno University of Technology
Faculty of Military Technology
University of Defense
Brno, Czech Republic
kamil.jasso@vut.cz, kamil.jasso@unob.cz

***Abstract*—This paper presents the comprehensive redesign of a 80V bidirectional AC/DC power supply utilized for automotive component testing. The primary objective was to overcome the dimensional limitations of the original setup, housed in an oversized 2000 mm cabinet, which restricted its portability between laboratory spaces. The proposed solution involved migrating the technology into a compact Rittal enclosure (1200 mm height), thereby enabling full mobility through standard facility doors without compromising power capability. Furthermore, the modernization addressed critical safety deficiencies by replacing outdated relay logic with a modular Safety PLC (Mosaic) and integrating a specialized insulation monitoring device (isoHV425) compliant with unearthed IT network standards. The result is a compact, highly mobile, and safety-compliant testing unit prepared for CE certification under Machinery Directive 2006/42/EC, specifically tailored for the dynamic needs of electromobility research.**

***Index Terms*—bidirectional power supply, hardware conversion, safety PLC, insulation monitoring, automotive testing, machinery directive, IT network.**

I. Introduction

The rapid electrification of the automotive sector has fundamentally altered the requirements for laboratory testing infrastructure. Unlike internal combustion engine components, electric vehicle (EV) powertrains—specifically batteries, inverters, and motors—require test equipment capable of bidirectional energy flow. These systems must not only source energy but also act as electronic loads to simulate regenerative braking or battery charging cycles, often returning energy to the grid to maximize efficiency [1].

Despite the advanced power electronics within modern bidirectional sources, such as the EA-PSB 10000 4U, their physical implementation in laboratory environments often lacks the necessary flexibility. High-power units are frequently installed in large-scale industrial cabinets (exceeding 2000 mm in height) which effectively renders them stationary. These dimensions prevent passage through standard laboratory doors, limiting the equipment's utility to a single location [2]. Furthermore, legacy installations often rely on traditional hard-wired safety relays and basic insulation monitoring devices (IMDs) that are ill-suited for the dynamic conditions of modern high-voltage IT networks used in EV testing.

This rigidity contradicts the need for agile testing environments where equipment must be frequently reconfigured or relocated. Moreover, modifying such high-voltage machinery to improve mobility falls under the "significant modification" clause of the Machinery Directive 2006/42/EC, necessitating a complete re-evaluation of safety risks and conformity to ensure operator protection [3].

This paper presents a comprehensive redesign and conversion of a stationary bidirectional AC/DC power supply into a compact, mobile laboratory unit. The objective was to overcome dimensional limitations while significantly upgrading the safety architecture. The proposed solution integrates a modular Safety PLC (Mosaic) and a specialized industrial insulation monitoring system to handle the specific risks of unearthed DC networks. The result is a fully mobile testing station compliant with current safety standards, capable of flexible deployment across various testing facilities.

II. System Re-Engineering

The transformation of the stationary power supply into a mobile unit required a complete overhaul of both the mechanical structure and the electrical safety architecture. The redesign followed the requirements for a "significant modification" under the Machinery Directive 2006/42/EC [3].

A. Mechanical Redesign for Mobility

The primary constraint of the original setup, shown in Fig. 1a, was its enclosure height. With dimensions of $600 \times 2000 \times 1000$ mm, the cabinet exceeded standard laboratory door dimensions, effectively rendering the unit stationary despite being mounted on castors. To address this, a comprehensive market analysis was conducted, leading to the selection of the Rittal VX IT data/server cabinet (model VX 5304116).

The new enclosure, depicted in Fig. 1b, features dimensions of $800 \times 1200 \times 1000$ mm. While the width increased slightly, the critical height reduction to 1200 mm ensures the unit

can easily pass through standard facility doors, fulfilling the mobility requirement. The selection was further justified by the cabinet's compatibility with standard 19-inch profile rails, allowing for the direct transfer of existing power components without the need for custom mechanical adapters.



(a) Old Enclosure (Original Setup) (b) New Enclosure (VX 5304116)

Fig. 1. Comparison of the mechanical design. The height reduction from the original cabinet (a) to the new Rittal VX rack (b) enables mobility through standard doors.

To achieve this compact form factor without compromising power capacity, non-essential components were identified and removed. Specifically, the integrated industrial PC and monitor were eliminated. This decision aligns with the modern "remote operation" philosophy, where test engineers control the unit from a separate control room via Ethernet or CAN bus interfaces to minimize exposure to high-voltage hazards.

B. Safety Architecture Upgrade

The legacy safety architecture relied heavily on traditional hardwired relays, which presented significant limitations regarding flexibility and diagnostic capabilities. In such a configuration, any modification to the safety logic – whether for integrating new test parameters or adapting to different laboratory environments – required complex physical rewiring of the control cabinet. To overcome these constraints and ensure compliance with modern safety standards for mobile machinery, the central control logic was migrated to a modular *Mosaic M1 Safety PLC*.

This transition shifted the safety definition from hardware to software. The logic is now configured via the *Mosaic Safety Designer* environment (see Fig. 2), allowing for the rapid creation of complex safety zones. For instance, the system can be reprogrammed to accept signals from different types of external emergency stops or laser scanners without mechanical intervention. Furthermore, unlike the binary status of relays, the PLC provides real-time monitoring of input/output states, facilitating faster troubleshooting and system verification.

To fully utilize this programmable logic within the mobile unit, specific hardware expansion was necessary. The system was equipped with the *MOR4S8* output module, which manages the redundant disconnection of the main power path via series-connected contactors (-QA1.1, -QA1.2) and controls the active discharge circuit (-KM5). Simultaneously, the *MI8O2* input expansion module was integrated to interface with the laboratory infrastructure, processing signals from external door interlocks and remote safety loops, ensuring that the power supply cannot be energized if the testing room is unsecured [5].

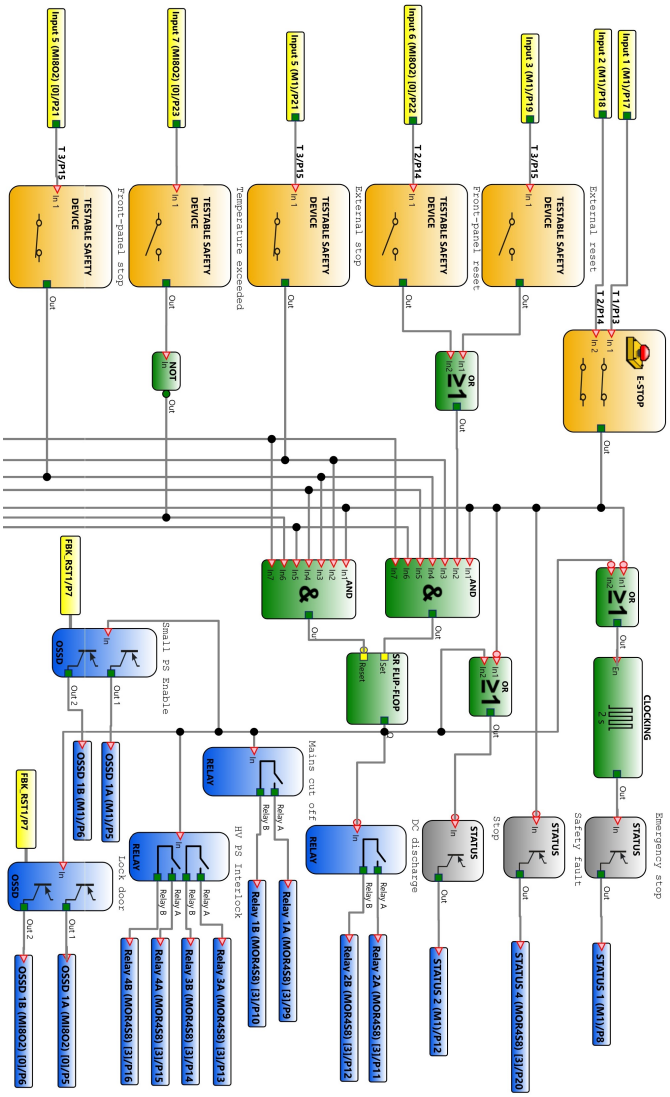

Fig. 2. Configuration of safety logic in Mosaic Safety Designer. The graphical interface allows for the definition of input/output relations and safety functions without physical rewiring.

C. Safety Logic Implementation

Unlike conventional PLC programming languages, the safety logic for the Mosaic PLC was developed using Functional Block Diagrams (FBD), complying with the IEC 61131-3 standard. The programming process involved configuring logical operators (AND, OR) and functional blocks to manage inputs from the Mosaic M1 master unit and the MI8O2 expansion module.

The core logic relies on two primary functional blocks:

- **Reset Logic (OR):** This block integrates signals from both the front-panel reset button and the external remote reset. This ensures redundancy, allowing the safety circuit to be reset from either location only after a fault is cleared.
- **Enable Logic (AND):** A critical safety chain that requires logical '1' from all safety sensors (E-Stop not pressed, doors closed, IMD healthy) to enable the power supply.

A dedicated *SR Flip-Flop* block acts as a state memory. It latches the safe state upon a successful reset and immediately deactivates the outputs if the *Enable Logic* condition fails (e.g., if a door is opened). The output logic then drives two types of signals: OSSD (Output Signal Switching Device) for static semiconductor outputs controlling auxiliary systems, and RELAY outputs (via the MOR4S8 module) for galvanically isolated switching of the main power contactors (-QA1) and the discharge circuit.

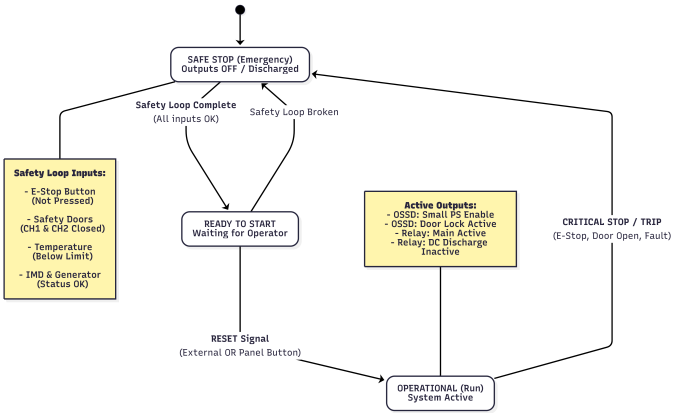

Fig. 3. State transition diagram of the safety PLC logic. The system transitions from Safe Stop to Operation only after a valid reset signal and fulfillment of all safety loop inputs.

D. Advanced Insulation Monitoring

For bidirectional power supplies operating in an IT network (unearthed) configuration, reliable insulation monitoring is critical. The original device (isoPV425) was designed for photovoltaic panels and lacked the necessary coupling for dynamic DC testing. To address this, the Bender isoHV425 was integrated together with the AGH422 coupling device. This system continuously monitors the insulation resistance (R_{iso}) between the active DC conductors and the protective earth (PE). It distinguishes between two fault levels:

1) **Warning (1st Fault):** R_{iso} drops below a warning threshold. The system alerts the operator but maintains power to prevent test interruption.
2) **Alarm (2nd Fault):** A critical fault or short circuit is detected. The Safety PLC immediately disconnects the main contactors and activates the discharge circuit [6].

E. Grid Protection and Automatic Isolation Unit

To enable safe regeneration of energy back to the grid (sink mode), the grid protection architecture required a fundamental upgrade. The original protection unit (EA-ENS2) was deemed unsuitable for the modernized safety concept. As analyzed in the design phase, the ENS2 lacked feedback signaling regarding the disconnection status and was not user-programmable, preventing integration with the new safety logic [2].

Consequently, the obsolete unit was replaced by the VMD460-NA automatic isolation unit. This device was selected specifically for its compatibility with the Mosaic Safety PLC, enabling a bidirectional feedback loop. The VMD460-NA continuously monitors grid parameters, including voltage levels, frequency stability, and phase angle. Upon detecting grid anomalies, it triggers immediate isolation via redundant series-connected contactors (-QA1.1, -QA1.2), ensuring compliance with international standards for power generation systems [7].

III. Results and Verification

The redesigned unit underwent a series of functional and safety tests to verify compliance with the design requirements.

A. Functional Verification

The mechanical integration into the 1200 mm rack was successful, ensuring full portability between laboratories. All wiring was executed according to new schematics created in PCSCHEMATIC (see Fig. 4), with verified segregation of power and signal cables to minimize electromagnetic interference.

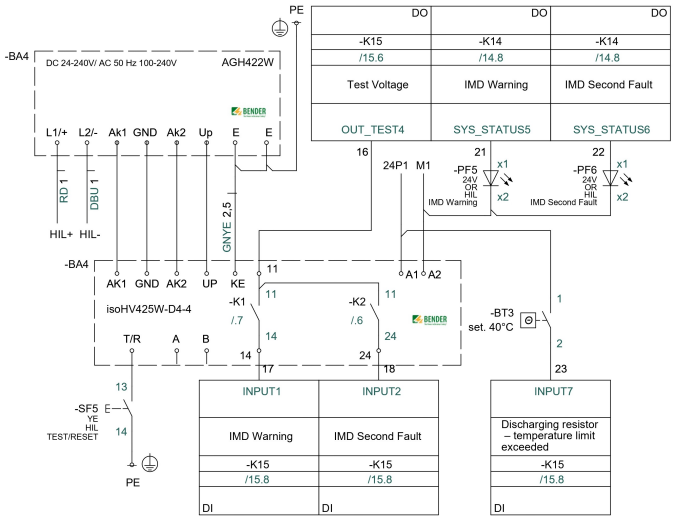

Fig. 4. Simplified schematic of the Insulation Monitoring Device integration created in PCSCHEMATIC. The diagram details the connection of the coupling device (AGH422) and the isoHV425 monitor to the safety PLC.

B. Safety System Validation

Safety functions were validated through fault simulation:

- **Emergency Stop:** Activation of the E-STOP button resulted in the immediate opening of AC contactors (-QA1.1, -QA1.2).
- **Discharge Logic:** Simultaneously with the shutdown, the Safety PLC activated the discharge contactor (-KM5),

connecting a $27\,\Omega$ / 300 W resistor to the DC output. Measurements confirmed that the residual voltage dropped to safe levels (< 60 V) within the required timeframe.
- **Insulation Fault:** Simulated earth faults were correctly identified by the isoHV425, triggering the respective warning or shutdown sequences in the PLC.

C. Compliance Preparation

A complete Technical File was compiled, including the Operating Manual, Maintenance Manual, and Risk Assessment. This documentation serves as the basis for the Declaration of Conformity, confirming that the modified machine meets the essential health and safety requirements of the EU legislation.

An integral part of the documentation is the maintenance schedule defined in the Maintenance Manual (see Table I). This schedule mandates strict periodic inspections of safety-critical components to ensure long-term reliability.

TABLE I
MAINTENANCE SCHEDULE FOR SAFETY DEVICES

Activity Description [months]	**1**	**12**	**24**
Check for damage to covers, closing mechanisms, buttons, movable cables	X		
Check integrity of protective circuit and connection of PE conductors		X	
Check current settings of protective and safety elements		X	
Check condition of electrical equipment, replacement of worn devices			X

IV. CONCLUSION

The comprehensive redesign of the bidirectional AC/DC power supply presented in this paper successfully transformed a stationary, dimensionally restrictive installation into a highly flexible and safe mobile laboratory unit. By migrating the technology into a compact Rittal VX IT enclosure, the critical issue of portability was resolved, allowing the equipment to be easily deployed across various testing facilities without physical barriers.

The integration of advanced safety systems represents a significant upgrade over the original relay-based architecture. The implementation of the Mosaic Safety PLC provides a programmable and modular safety logic capable of handling complex scenarios, such as external emergency stops and door interlocks, which are essential for modern testing rooms. Furthermore, replacing the photovoltaic-specific insulation monitor with the industrial-grade Bender isoHV425 and AGH422 coupling device ensures reliable fault detection in dynamic unearthed DC networks, significantly enhancing operator protection.

The modification was classified as ”significant” under the Machinery Directive 2006/42/EC, necessitating a full conformity assessment. The creation of a complete Technical File, including risk assessments and new operating manuals, confirms that the modernized unit meets stringent EU safety and health requirements. The result is a future-proof, compliant, and versatile testing station ready for the demanding requirements of electromobility component validation.

REFERENCES

[1] EA Elektro-Automatik, ”EA-PSB 10000 4U: Programmable Bidirectional DC Power Supplies - Operating Guide,” 2022.
[2] EA Elektro-Automatik, ”Operating Guide: ENS2,” 2022.
[3] European Parliament and Council, ”Directive 2006/42/EC on machinery,” Official Journal of the European Union, 2006.
[4] International Organization for Standardization, ”ISO 6469-3:2021 Electrically propelled road vehicles – Safety specifications – Part 3: Electrical safety,” 2021.
[5] ReeR Safety, ”Mosaic Safety Integrated Controller - User Manual,” Torino, Italy, 2024.
[6] Bender GmbH & Co. KG, ”ISOMETER® isoHV425 with AGH422 - Data Sheet,” Grünberg, Germany, 2025.
[7] Bender GmbH & Co. KG, ”VMD460-NA Network and system protection - User Manual,” 2021.

Modular High-Power Electronic Load

Jan Gadas
Department of Radio Electronics
FEEC, Brno University of Technology
Brno, Czech Republic
240612@vut.cz

Aleš Povalač
Department of Radio Electronics
FEEC, Brno University of Technology
Brno, Czech Republic
povalac@vut.cz

***Abstract*—This work presents the design and implementation of a low-power digitally controlled electronic load with a maximum power of 100 W. The prototype includes custom analog measurement circuitry, digital signal processing, operating mode implementation, and a custom user interface. Its functionality was verified by simulations and laboratory measurements. The resulting design serves as the basis for the ongoing development of a 5 kW high-power modular electronic load within the follow-up master's thesis.**

***Index Terms*—Electronic load, Thermal stability, SOA, MOSFET, STM32, HMI, feedback loop**

I. Introduction

The growing use of electrical energy sources and storage systems increases the demand for reliable testing methods. Electronic loads enable controlled simulation of power consumption and are therefore essential for evaluating such devices. Their design is challenging, as it involves both digital control and precise analog measurement, as well as proper dimensioning of the power stage and analysis of the thermal stability of the selected semiconductor devices.

II. Semiconductor Thermal Stability

Most electronic loads use a semiconductor device as the power-dissipative element. However, selecting a suitable transistor for this purpose is challenging, because in a DC electronic load it operates in the linear region, where the device is subjected to significantly higher stress than in conventional switching applications. Among the available options, namely BJTs, IGBTs, and MOSFETs, the latter are generally preferred due to their resistive channel structure. Even so, their nonlinear characteristics and unfavorable thermal behavior considerably limit the range of devices suitable for this application.

A. Charge Carrier Concentration vs. Resistivity

The intrinsic carrier concentration increases with temperature as more electrons gain sufficient energy to overcome the bandgap. This reduces the substrate Fermi potential and consequently the MOSFET threshold voltage [1]. In MOSFETs intended primarily for switching applications, the temperature-driven increase in carrier concentration n_i may dominate over the stabilizing effect of the increasing $R_{\mathrm{DS(on)}}$, which can cause thermal imbalance, as shown in Figure 1.

Research described in this paper was supported by the Internal Grant Agency of the Brno University of Technology under project no. FEKT-S-26-8967.

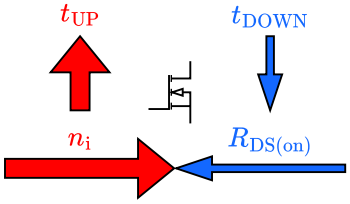

Fig. 1. MOSFET stability balance

B. Thermal Runaway and Hot Spot Formation

A power MOSFET consists of many parallel elementary cells sharing common gate, drain, and source electrodes. Local variations in geometry and doping concentration may disturb uniform current sharing between these cells. Modern process optimization reduces cell pitch and lowers $R_{\mathrm{DS(on)}}$, but also increases sensitivity to such nonuniformities [2]. Cells with lower V_{th} and higher local transconductance tend to carry more current, which increases local heating and junction temperature T_{J}. The resulting positive thermal feedback may eventually cause fatal local hot spots, as illustrated in Figure 2.

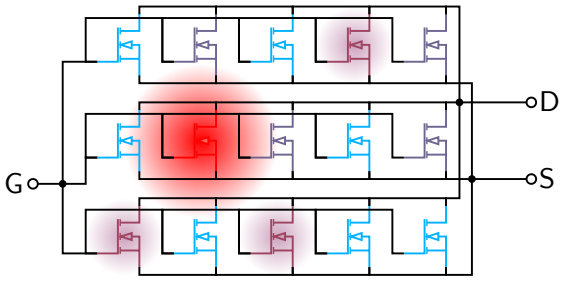

Fig. 2. Visualization of Hot Spot in the MOSFET structure

C. Thermal System Design

Thermal design and cooling are essential parts of electronic load design. The key limiting parameter is the junction temperature T_{J}, which must not exceed its maximum rated value. Thermal design is based on thermal resistances describing sections of the heat path. Expressed in $\mathrm{K/W}$, thermal resistance represents the temperature rise per unit of dissipated power. Figure 3 shows the thermal path from the TO-247 transistor through the thermal interface material, insulating pad, and heatsink, which dissipates heat to the ambient environment.

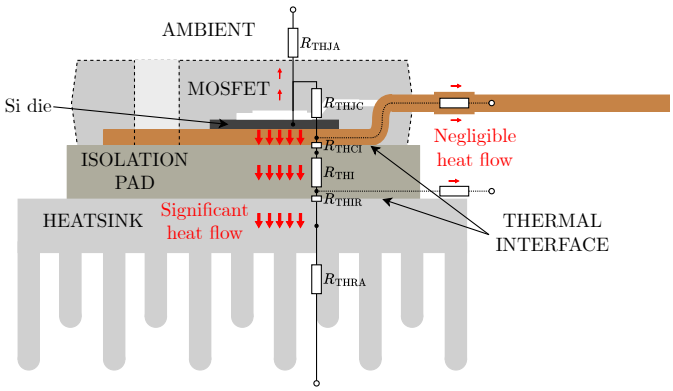

Fig. 3. Thermal model and illustration of heat flow paths

The junction temperature can be calculated from the thermal model in Figure 3 and the dissipated power using equation (1).

$$T_J = T_A + P \cdot (R_{THJC} + R_{THCI} + R_{THI} + R_{THIR} + R_{THRA}) \quad [K] \tag{1}$$

D. SOA Graph and Manufacturer Specifications

The SOA graph defines the allowable combinations of voltage and current for a given pulse duration. As shown in Figure 4, its boundaries are given by the $R_{DS(on)}$ limit at low voltage, the maximum current limit, the power dissipation limit, the thermal instability region, and the breakdown voltage limit. A major concern is that manufacturers have often omitted the thermally unstable region from the SOA graph. Even in current datasheets, DC SOA is commonly specified only for fixed case-temperature conditions. Therefore, stable linear-mode DC operation cannot be guaranteed from the published SOA curve alone and requires derating and thermal-stability evaluation, particularly near the Spirito-effect region [3].

The IXYS Linear L2 series MOSFETs are intended for stable linear operation. Improved thermal stability and suppression of hot-spot formation are achieved by integrating a shunt resistor into the source region of each MOSFET cell [4]. This increases manufacturing complexity and cost, but ensures better stability in DC linear-mode operation. Therefore, these MOSFETs represent a suitable choice for the electronic load.

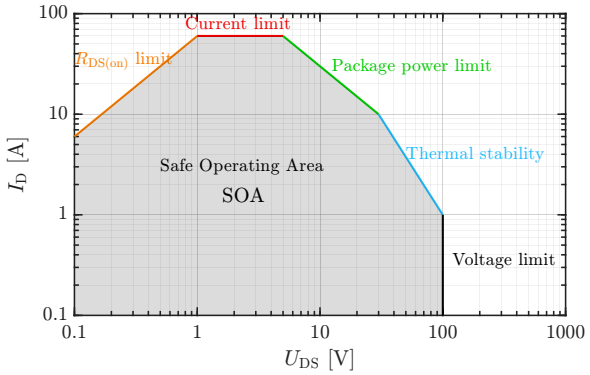

Fig. 4. Description of SOA Graph Boundaries

III. Low-Power Electronic Load Design

As a preliminary step toward the development of a 5 kW high-power modular electronic load, a low-power 100 W prototype was designed and implemented. Its purpose is to verify the proposed topology and validate both the analog and digital parts of the system. The prototype supports the standard operating modes CC, CV, CP, and CR, together with predefined STOP limits and accumulated capacity and energy measurement. Additional functions, such as battery testing and MPPT search, will be implemented in the future.

TABLE I
Low-Power Prototype Electronic Load Parameters

Parameter	Value	Unit
Input voltage	0–500	V
Resolution	10	mV
Min. voltage at max. current	2.5	V
Input current	0–10	A
Resolution	1	mA
Input power	0–100	W
Resolution	10	mW
Input resistance	0.5–10000	Ω
Resolution	0.1	Ω
Maximum parasitic inductance	10	μH

A. MASTER and SLAVE Division

The system architecture is based on a division into MASTER and SLAVE sections, as shown in Figure 5. The SLAVE section is responsible for analog control and measurement acquisition, whereas the MASTER section handles communication and the user-oriented HMI/GUI layer. In the future modular electronic load, the MASTER will form the main bus control unit to which SLAVE modules will be connected.

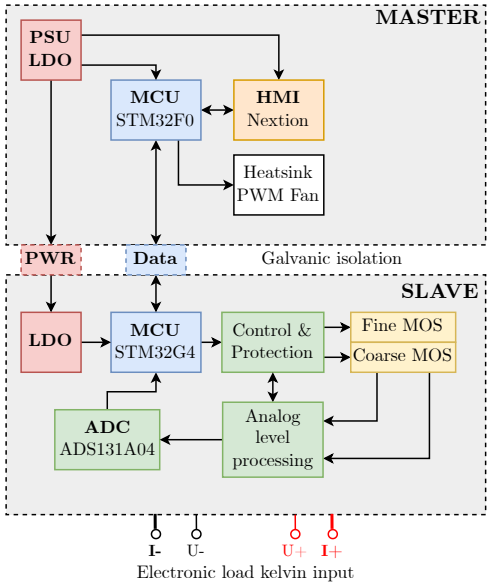

Fig. 5. Low power electronic load prototype block diagram

B. Galvanic Isolation

The negative terminal of the electronic load power supply can be tied to PE (protective earth). If the DUT (Device Under Test) is also tied to PE at its negative terminal, a closed return path may be created through the PE conductor, which can result in a serious fault condition and equipment damage. To avoid this, the communication and power supply links between the MASTER and SLAVE sections are galvanically isolated, as shown in Figure 5.

C. Digital Control

The device control architecture is based on STM32 microcontrollers. The MASTER section uses an STM32F0 and is responsible for bidirectional communication with the Nextion HMI running the custom GUI, as well as PWM fan control. The SLAVE section is built around an STM32G4, selected for high-speed peripherals such as SPI for the ADC. It also evaluates protection conditions and device status and controls the DAC for the analog feedback loops.

D. Coarse and Fine Current Control Topology

To improve current-setting resolution and accuracy, the proposed design employs separate coarse and fine adjustment channels, as shown in Figure 6. The coarse channel covers 0 to 10 A with 100 mA resolution, while the fine channel covers 0 to 100 mA with 1 mA resolution. The SLAVE MCU automatically splits the requested constant-current setpoint into coarse and fine components for the corresponding channels.

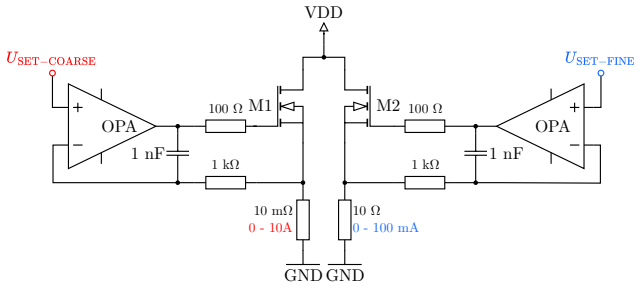

Fig. 6. Constant current feedback loops

E. Analog to Digital Converter

Because the internal MCU ADC does not provide sufficient resolution, the design employs the ADS131A04, a four-channel simultaneous-sampling SPI ADC operating at 32 kSPS per channel [5]. Its differential inputs are driven by fully differential amplifiers, while the remaining signal-conditioning stages use precision amplifiers with sub-µV offset [6], enabling highly accurate voltage and current measurements.

F. Circuit Protection

The proposed electronic load incorporates both digital and analog protection mechanisms. The analog protection responds faster and operates independently of the MCU state. Its main function is to supervise the minimum operating voltage (0.5 V) before enabling the op-amp output driving the MOSFETs, thus preventing load activation when no DUT is connected.

G. Touchscreen-Based Human–Machine Interface

A 5-inch Nextion touchscreen display was selected as the Human–Machine Interface (HMI) [7]. It presents both the measured and user-defined parameters, including cumulative values. In addition, the interface provides control elements such as a keypad, an ON/OFF button, and dedicated status fields for warnings and error messages.

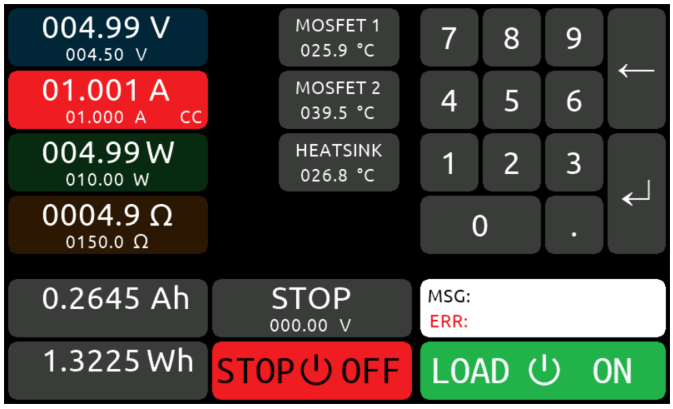

Fig. 7. Custom GUI for Nextion HMI

IV. Low-Power Electronic Load Construction

All analog circuits subject to stability constraints were verified in OrCAD PSpice. The PCB layout was designed in KiCad 9.0, followed by the mechanical design in SolidWorks.

A. PCB Design

The PCB is a four-layer design with separated functional blocks. The left side contains the power supply and MASTER control section, while the right side is dedicated to the SLAVE section, which includes the measuring and control analog circuits, as well as the four-wire input of the electronic load. Insulation slots were added in high-voltage areas to improve electrical isolation. The corresponding electrical schematic is provided in the GitHub repository [1].

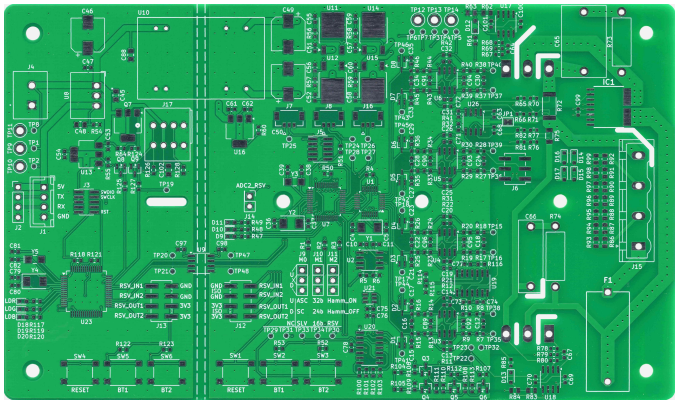

Fig. 8. Top-Side Scan of the Low-Power Electronic Load PCB

[1] https://github.com/Gady159/Gadas-M2E-Schematic

B. Mechanical Construction

The low-power prototype was implemented in a 3D-printed mechanical enclosure designed in SolidWorks. The structure provides support for the PCB assembly and the heatsink, which is mechanically mounted to the board. It also integrates the Nextion display, the power switch, and all input connectors.

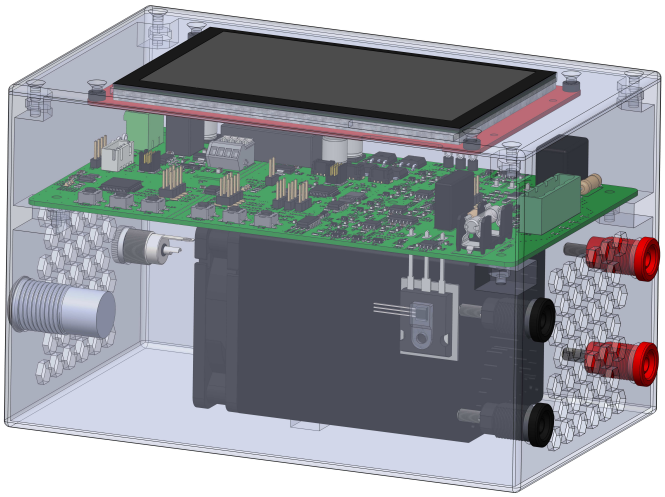

Fig. 9. Electronic Load Construction Preview

C. Measurement and Validation Results

The prototype was experimentally evaluated in both static and dynamic tests focused on accuracy and control performance. In addition, gain and offset calibration was performed, resulting in voltage measurement accuracy better than the required specification. Protection mechanisms were also verified, together with the response to sudden changes in DUT parameters and load settings. Practical validation was further carried out through the complete discharge of a 21700-type Li-ion cell in combined CC–CV operation.

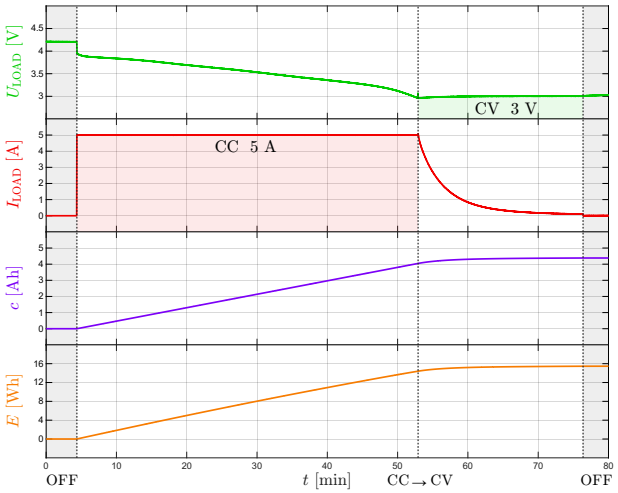

Fig. 10. Practical test discharging 21700 Li-ion battery

V. Conclusion

The project addressed the fundamental principles of electronic loads, with particular emphasis on the thermal stability of semiconductor devices in linear operation. Based on the acquired knowledge and the evaluation of suitable discrete devices, a low-power 100 W prototype was designed and implemented. The system includes a four-layer PCB with separate MASTER and SLAVE sections and a custom graphical user interface for a touchscreen display. During the bring-up, the prototype was calibrated and experimentally verified through static and dynamic tests, response measurements under varying operating conditions, and practical Li-ion cell discharge experiments. The resulting prototype represents the main outcome of this work and serves as the basis for the ongoing development of a modular high-power electronic load with a target power of 5 kW. As shown in Figure 11, the illustration presents not only the water-cooling system composed of cooling and distribution blocks but also the assembly of seven printed circuit boards forming the core of the modular 5 kW electronic load.

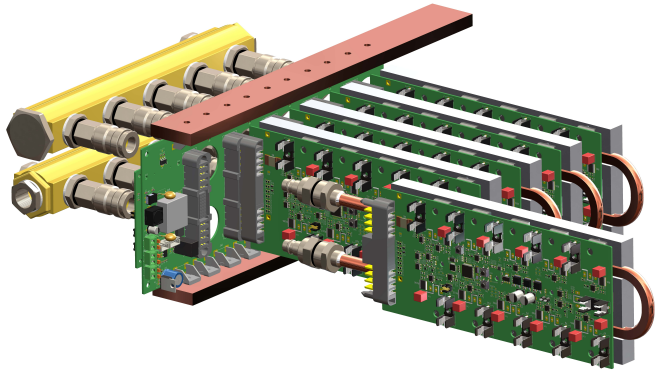

Fig. 11. Electro-Mechanical Assembly of the Modular 5kW Electronic Load

References

[1] STMICROELECTRONICS. *AN2386 Application note: How to achieve the threshold voltage thermal coefficient of the MOSFET acting on design parameters*. Online. 2006. Available at: https://www.st.com/resource/en/application_note/an2386-how-to-achieve-the-threshold-voltage-thermal-coefficient-of-the-mosfet-acting-on-design-parameters-stmicroelectronics.pdf.

[2] NEXPERIA. *IAN50006 Application Note: Power MOSFETs in linear mode - Application Note*. Online. 2022. Available at: https://assets.nexperia.com/documents/application-note/AN50006.pdf.

[3] SHUE, John and LEIDECKER, Henning. Power MOSFET Thermal Instability Operation Characterization Support. *NASA Technical Reports Server (NASA)*. 2010, p. 23. Online. Available at: https://ntrs.nasa.gov/api/citations/20100014777/downloads/20100014777.pdf.

[4] IXYS - LITTELFUSE. *MOSFETs Withstand Stress of Linear-Mode Operation*. Online. 2009. Available at: https://www.littelfuse.com/assetdocs/mosfets-withstand-stress-of-linear-mode-operation-en-white-paper?assetguid=7a45ee66-a951-486d-bd43-d3e19ebf3ce3.

[5] TEXAS INSTRUMENTS. *ADS131A04 Datasheet*. Online. 2016. Available at: https://www.ti.com/lit/ds/symlink/ads131a04.pdf.

[6] TEXAS INSTRUMENTS. *OPAx189 Datasheet*. Online. 2021. Available at: https://www.ti.com/lit/ds/symlink/opa189.pdf.

[7] ITEAD INTELLIGENT SYSTEMS CO., LTD. *NX8048P050-011C*. Online. Available at: https://nextion.tech/datasheets/nx8048p050-011c/.

Influence of Rotor Flux Barrier Geometry on Torque Ripple in a Synchronous Reluctance Motor for Electric Vehicle Applications

Hernán A. Bravo Urrea
Brno University of Technology
Brno, Czech Republic
281641@vut.cz

***Abstract*—Synchronous reluctance motors (SynRM) have recently attracted considerable attention as an alternative to permanent magnet machines in electric vehicle applications. Their simple rotor structure, absence of rare-earth materials, and high efficiency make them attractive candidates for traction applications. However, the main limitation of the SynRM is the presence of torque ripple caused by the anisotropic magnetic structure of the rotor and the spatial variation of magnetic reluctance. This paper analyzes the influence of rotor flux barrier geometry on the electromagnetic torque characteristics of a SynRM designed for electric vehicle applications. Two rotor configurations with three flux barriers are evaluated using finite element analysis (FEA). The results demonstrate that the geometry and orientation of the flux barriers strongly affect the saliency ratio, torque production and torque ripple behavior. In addition, skewing techniques are evaluated as a mitigation strategy for torque oscillations. The study provides a design approach for improving the electromagnetic performance of SynRM intended for traction applications.**

***Index Terms*—Synchronous reluctance motor, flux barriers, torque ripple, Finite Element Analysis, electric vehicles.**

I. Introduction

The rapid growth of electric mobility has significantly increased the demand for high–performance electric machines capable of delivering high torque density, high efficiency and reliable operation. Permanent magnet synchronous machines (PMSM) currently dominate the electric vehicle market due to their excellent electromagnetic performance. Nevertheless, their dependence on rare-earth materials such as neodymium and dysprosium introduces economic and geopolitical concerns associated with supply chain stability, manufacturing cost, and demagnetization [1].

In recent years, synchronous reluctance motors (SynRM) have emerged as a promising alternative for traction applications. Unlike PMSM machines, the SynRM rotor does not require permanent magnets or rotor windings, resulting in a simpler and more robust mechanical structure. In addition, rotor losses are significantly reduced, which can improve thermal behavior and overall efficiency [1] [2].

The torque production mechanism of the SynRM is fundamentally different from that of magnet-based machines. Instead of relying on magnetic excitation provided by permanent magnets, the synchronous reluctance machine produces torque as a result of the anisotropic rotor magnetic properties. This anisotropy is introduced through flux barriers embedded within the rotor lamination stack, which creates different magnetic reluctance along the direct d and quadrature q axes [3].

The electromagnetic torque of the SynRM is therefore proportional to the difference between the direct-axis and quadrature-axis inductance (L_{d}, L_{q}) [4] . Increasing this inductance difference improves the average torque capability of the machine. However, the geometric modifications required to achieve high saliency ratios often introduce spatial harmonics in the air-gap permeance distribution, which in turn leads to torque ripple [5].

Torque ripple represents one of the most important limitations of SynRM machines for traction applications, since it can generate vibrations, acoustic noise and additional mechanical stress on drivetrain components. Consequently, careful rotor design is required to balance torque capability and torque smoothness [5]. To reduce the torque ripple, several techniques have been investigated. In [6], asymmetric concentric-shaped barriers are used to analyze the spatial torque harmonics with respect to the permeance variation as a function of the rotor angular position. In addition, [7], the electromagnetic performance of the machine is investigated when permanent magnets are placed in the flux barriers. In [8] and [9], optimization algorithms are proposed to reduce torque ripple in machines with rectangular V-shaped flux barriers.

This paper investigates the influence of different flux barrier geometries on the torque characteristics of a synchronous reluctance motor designed for electric vehicle applications. two rotor configurations are analyzed through finite element simulations in order to evaluate their effect on torque production and torque ripple behavior.

II. Electromagnetic Torque and Torque Ripple in SynRM

It is well known that torque production in the SynRM is governed by the rotor´s magnetic anisotropy and depends on the saliency difference (L_{d}-L_{q}). Which defines the machine´s torque capability [10].

From a more general perspective, the torque developed by an electrical machine is defined as the vector cross product between the rotor and stator magnetic field. In the case of the

SynRM, the torque is produced by the interaction between the stator current vector i_s and the air-gap flux linkage ϕ_{m} [11].

$$T_e = \frac{m}{2} p(\phi_{\mathrm{m}} \times i_{\mathrm{s}}) \tag{1}$$

where m is the number of phases and p is the number of pole pairs.

The cross product $(\phi_{\mathrm{m}} \times i_{\mathrm{s}})$ leads to the following expression:

$$(\phi_{\mathrm{d}} i_{\mathrm{q}} - \phi_{\mathrm{q}} i_{\mathrm{d}}) \tag{2}$$

Based on (2), equation (1) can be rewritten as

$$T_e = \frac{3}{2} p(L_{\mathrm{d}} - L\mathrm{q}) i_{\mathrm{d}} i_{\mathrm{q}} \tag{3}$$

Equation (3) represents the developed electromagnetic torque of the synchronous reluctance motor.

This equation shows that the torque capability of the synchronous reluctance motor depends on the inductance difference between the direct and quadrature axes, known as reluctance torque.

Furthermore, the SynRM inductance varies with rotor position due to the spatial variation of magnetic reluctance introduced by the rotor flux barriers. The permeance of the air gap can therefore be expressed as a spatial periodic function [11].

$$P(\theta) = P_0 + \sum_{n=1}^{\infty} P_{\mathrm{n}} \cos(np\theta) \tag{4}$$

where P_0 represents the average air-gap permeance, P_{n} denotes the amplitude of the n-th spatial harmonic component, p is the number of pole pairs, and θ is the rotor angular position.

When the stator magnetomotive force interacts with this spatially modulated permeance, harmonic components appear in the electromagnetic torque.

Consequently, the instantaneous torque can be expressed as

$$T_e(\theta) = T_{avg} + \sum_{n=1}^{\infty} T_n \cos(n\theta) \tag{5}$$

where T_{avg} represents the average torque and T_n corresponds to the harmonic components responsible for torque ripple.

From an energy perspective, the electromagnetic torque can also be expressed as the derivative of the magnetic co-energy with respect to the rotor position

$$T_{\mathrm{e}} = \frac{\partial W'(\theta)}{\partial \theta} \tag{6}$$

where W is the stored magnetic energy and θ is the load angle (the angle between rotor axis and flux). W is expressed as follows:

$$W = \frac{1}{2} \int_{V_r} B_\delta H_\delta \, dV_{\mathrm{r}} \tag{7}$$

In a synchronous reluctance motor, the rotor tends to align with the stator magnetic field to minimize the magnetic reluctance along the direct axis. Any change in θ modifies the stored magnetic energy. If the stored magnetic energy varies periodically with the rotor position due to permeance modulation, oscillations naturally appear in the electromagnetic torque.

Hence, Equation (3) can be rewritten as follows:

$$T_e = \frac{m}{2} p(L_{\mathrm{d}} - L_{\mathrm{q}}) i_{\mathrm{d}} i_{\mathrm{q}} + \frac{1}{2} \left(i_{\mathrm{d}} \frac{\partial \lambda_{\mathrm{d}}}{\partial \theta} + i_{\mathrm{q}} \frac{\partial \lambda_{\mathrm{q}}}{\partial \theta} \right) \tag{8}$$

The first term represents the developed electromagnetic torque as previously discussed. The second term accounts for the variation of the flux linkages with respect to the rotor position and is associated with the angular modulation of L_{d} and L_{q}, which produces torque harmonic components (torque ripple) [12].

III. Finite Element Model

Due to the nonlinear magnetic behavior of electrical machines and the complex rotor geometry, analytical models are often insufficient to accurately evaluate torque ripple.

Therefore, finite element analysis was used to evaluate the electromagnetic behavior of the synchronous reluctance motor.

The finite element model allows accurate computation of the magnetic flux distribution, electromagnetic torque and spatial harmonics in the air gap.

Transient simulations were performed in order to obtain the instantaneous torque waveform as the rotor rotates. From these simulations, both the average torque and the torque ripple amplitude were obtained.

Transient simulations were performed in order to obtain the instantaneous torque waveform as the rotor rotates. From these simulations, both the average torque and the torque ripple amplitude were obtained. The simulations were carried out at the nominal machine speed of 1500 rpm. A total of 100 samples per electrical period were used, and the simulation was performed over 5 electrical cycles to ensure steady-state conditions.

Fig 1 and Fig 2 illustrate the SynRM geometries used to analyze the influence of rotor flux barrier on torque ripple.

IV. Results and Discussion

A. Magnetic Flux Distribution

The magnetic flux density distribution provides insight into how the rotor geometry affects the magnetic paths along the d and q axes. By the distribution of the magnetic flux, we can observe if d-axis is aligned with the magnetic field at θ (rotor position). Fig 3 and Fig 4 present the flux density maps obtained from finite element analysis.

For the first rotor configuration, the magnetic flux tends to concentrate in the iron ribs that separate the flux barriers. As the flux crosses these narrow regions, local saturation appears, which modifies the magnetic permeance distribution along the

air gap. This effect introduces spatial variations in the air-gap flux density.

For the second rotor configuration, the magnetic flux follows a smoother path along the direct axis of the rotor. The wider magnetic paths reduce the concentration of flux in the ribs and mitigate the appearance of localized saturation. As a result, the spatial variation of the air-gap flux density becomes less pronounced, contributing to a reduction in torque ripple, which can be seen further in this section.

B. Electromagnetic Torque Analysis

The instantaneous electromagnetic torque obtained from the finite element transient simulations is presented in Fig 5 and Fig 6. Both Figures illustrate that the machines produce different average torque values at the same current magnitude of 100 A. The amplitude of the torque oscillations differs significantly between the two rotor geometries. The first configuration presents stronger torque oscillations due to the abrupt variation of magnetic reluctance introduced by the flux barrier geometry. These variations introduce spatial harmonics in the air-gap permeance.The second configuration produces a smoother torque waveform since the permeance variation along the rotor circumference is more gradual.

Torque ripple was evaluated by comparing the peak-to-peak variation of the instantaneous torque in steady-state operation. The results indicate that the amplitude of the torque ripple is strongly influenced by the spatial distribution of the flux barriers. Although both machines exhibit a comparable number of flux barriers, diameter, number of slots, and use the same materials, the arc configuration significantly reduces torque ripple due to the smoother variation of the air-gap permeance. It is important to note that torque ripple is directly related to the variation of the inductance as a function of the rotor position angle, $L(\theta)$. In an ideal scenario, if $L(\theta)$ exhibits a purely sinusoidal behavior without harmonic components, the resulting electromagnetic torque would be practically constant.

In the second design, the geometry of the flux barriers presents a smoother arc-shaped transition, which leads to a more gradual variation of the magnetic permeance with

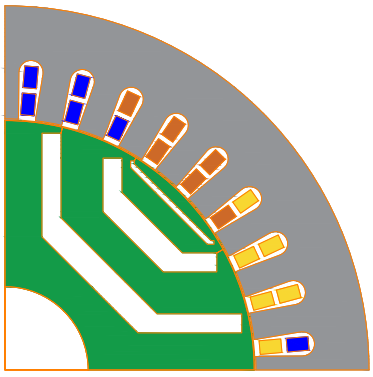

Fig. 1. Finite element model of the synchronous reluctance motor implemented in ANSYS.

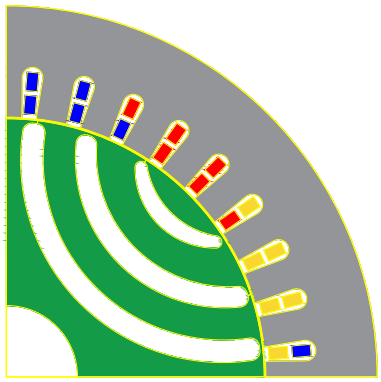

Fig. 2. Finite element model of the synchronous reluctance motor implemented in ANSYS.

respect to the angular position θ. This geometric smoothness attenuates the harmonic components of $L(\theta)$, thereby reducing torque oscillations and, consequently, the torque ripple. The main simulation results are summarized in Table I presents the simulation results.

On the other hand, Torque ripple components associated with the stator slotting effect and higher-order spatial harmonics can be reduced by applying skew to the stator slots. In order to further investigate this aspect, an additional analysis is carried out to evaluate how the torque ripple is affected when stator slot skewing is introduced in the synchronous reluctance machine, as investigated in [13], and [14]. The main simulation results after the skew are applied and summarized in Table II.

V. Conclusion

This paper investigated the influence of rotor flux barrier geometry on the torque characteristics of two synchronous reluctance motors. Finite element simulations were used to evaluate two rotor configurations with three flux barriers. The results demonstrate that the geometry of the barriers

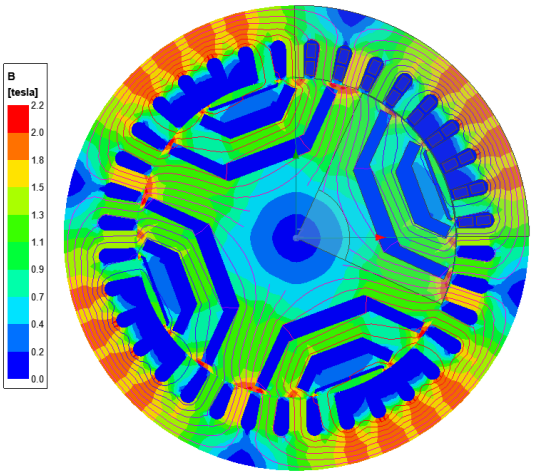

Fig. 3. Magnetic field distribution SynRM model 1.

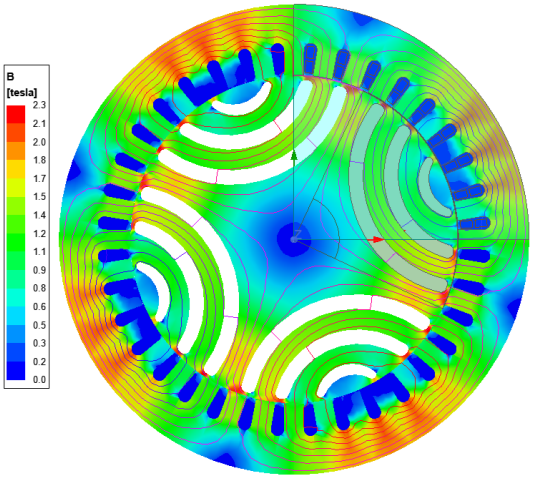

Fig. 4. Magnetic field distribution SynRM model 2.

strongly influences both the average torque and torque ripple behavior. Among the evaluated designs, the optimized rotor geometry combined with skewing techniques achieved the best compromise between torque production and ripple reduction.

Future work will focus on the experimental validation of the proposed machine and on the analysis of control strategies aimed at further improving torque smoothness.

REFERENCES

[1] N. Bianchi and M. Degano, "Electric vehicle traction based on synchronous reluctance motors," *IEEE Transactions on Industry Applications*, vol. 52, no. 6, pp. 4762–4769, 2016.

[2] Z. Zekic, "Modeling and analysis of synrm for hybrid electric vehicle application," Master's thesis, Chalmers University of Technology, Gothenburg, Sweden, 2016.

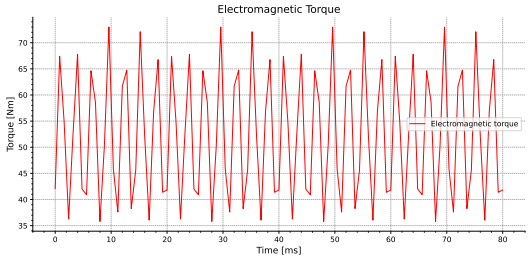

Fig. 5. Electromagnetic Torque SynRM Model 1 .

TABLE I
PERFORMANCE COMPARISON BETWEEN SYNRM MODELS

Parameter	SynRM Model I	SynRM Model II	Unit
Average Torque	52.2	109.18	Nm
Max Torque	73.06	129.58	Nm
Min Torque	35.79	92.12	Nm
L_d	5.84	10.09	mH
L_q	2.34	2.65	mH
Saliency (L_d/L_q)	2.50	3.81	–
Power Factor	0.50	0.55	–
Torque Ripple	71.40	34.31	%

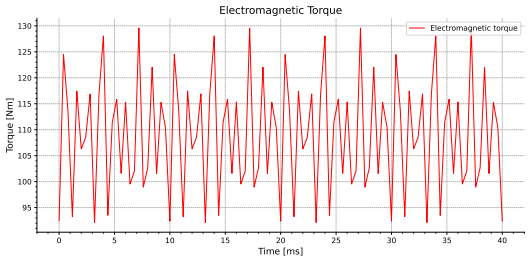

Fig. 6. Electromagnetic Torque SynRM Model 2

[3] H. D. Chi, "Modeling for the design of synchronous reluctance motors," in *2022 International Conference on Communications, Information, Electronic and Energy Systems (CIEES)*, Veliko Tarnovo, Bulgaria, 2022.

[4] M. Deepak, "Analysis of rotor flux barriers in synrm for performance enhancement in traction applications," in *2025 IEEE International Students' Conference on Electrical, Electronics, and Computer Science*, Bhopal, India, 2025.

[5] P. r. s. M. D. C.Bharatiraja, "A design on rotor flux barrier for line start synchronous reluctance machines for enhancing efficiency," in *2023 Second International Conference on Trends in Electrical, Electronics, and Computer Engineering (TEECCON)*, Kattigenahalli, India, 2023.

[6] Y. D. R. L. S. Z. X. Chai, "Design and optimization of asymmetric eccentric circular shaped barriers for synchronous reluctance machines," *IEEE Transactions on Industry Applications*, vol. 2, pp. 1–11, 2025.

[7] S. S. N. S. Kumar, "Flux-barrier based rotor design of permanent magnet assisted synchronous reluctance motor with power density and torque ripple improvement for e-mobility," in *2023 IEEE International Transportation Electrification Conference*, Chennai, India, 2023.

[8] S. E. E. Bekiroglu, "Design and double-stage optimization of synchronous reluctance motor for electric vehicles," *Electric Power Components and Systems*, vol. 51, pp. 2557–2572, 2023.

[9] Y. D. X. Chai, "Rotor design and optimization of synchronous reluctance machine with low torque ripple," *Journal of Power Electronics*, vol. 25, pp. 360–369, 2025.

[10] R. Panwar, "Field oriented control of synchronous reluctance motor using finite element co-simulation," in *2024 IEEE International Conference on Electronics, Computing and Communication Technologies (CONECCT)*, Bangalore, India, 2024.

[11] R. Moghaddam, "Synchronous reluctance machine (synrm) in variable speed drives (vsd) applications," Ph.D. dissertation, KTH School of Electrical Engineering, Stockholm, Sweden, 2011.

[12] M. P. Nikhila, "Torque ripple reduction of synrm using rotor structure modification," in *2024 International Conference on E-mobility, Power Control and Smart Systems (ICEMPS)*, Thiruvananthapuram, India, 2024.

[13] J. A. T. César Gallardo, Carlos Madariaga, "Impact of rotor step skew on the performance of synchronous reluctance machines," in *2023 IEEE Chilean Conference on Electrical, Electronics Engineering, Information and Communication Technologies (CHILECON)*, Valdivia, Chile, 2023.

[14] C. Gallardo and C. M. J. A. Tapia, "A method to determine the torque ripple harmonic reduction in skewed synchronous reluctance machines," *Applied Sciences*, pp. 1–19, 2023.

TABLE II
PERFORMANCE COMPARISON BETWEEN SYNRM MODELS WITH SKEWING TECHNIQUE

Parameter	SynRM Model I	SynRM Model II	Unit
Average Torque	51.5	107.01	Nm
Max Torque	55.88	111.4	Nm
Min Torque	47.42	103.05	Nm
L_d	5.79	9.95	mH
L_q	2.34	2.79	mH
Saliency (L_d/L_q)	2.47	3.57	–
Power Factor	0.50	0.54	–
Torque Ripple	16.43	7.80	%

Thermal System Modeling and Validation

1st Radim Vancl
Department of Control and Instrumentation
Brno University of Technology
Brno, Czech Republic
Radim.Vancl@vut.cz

2nd Michal Husák
Department of Control and Instrumentation
Brno University of Technology
Brno, Czech Republic
Michal.Husak@vut.cz

***Abstract*—The paper focuses on the design and validation of a heat exchanger station digital model. The main objective is to model the non-linear dynamics of the thermal system primarily for predictive maintenance, as well as to enable safe testing of control scenarios without interfering with the physical process. Two distinct modeling approaches are implemented and compared: a semi-physical discrete grey-box model based on differential equations and a data-driven black-box model utilizing Long Short-Term Memory (LSTM) neural networks. The parameters of the grey-box model were identified using the fminsearch and nlgreyest functions in MATLAB, while the LSTM network parameters were optimized via grid search. The models were validated using real-world data measured on a physical training station through Leave-One-Out Cross-Validation (LOOCV). The results demonstrate that the LSTM network achieved superior accuracy with an average Root Mean Square Error (RMSE) of 2.3 °C compared to 3.3 °C and 3.6 °C for the grey-box approaches.**

***Index Terms*—PLC, model, identification, MATLAB, heat exchanger, thermal system**

I. Introduction

With the advent of Industry 4.0 and increasing demands for production digitization, the use of digital representations of physical objects is becoming more frequent. These digital counterparts are primarily used to streamline the design, simulation, monitoring, and control of real systems. Based on the level of data integration between the physical and digital worlds, we distinguish three basic levels: digital model, digital shadow, and digital twin [1]–[4]. This paper focuses on the creation of a digital model, where there is not automatic data flow from the physical object to the digital one. We chose the thermodynamic model of the heat exchange station to test and validate the digital modeling on real hardware.

Heat exchangers are central elements of heating systems, serving to distribute heat or cold within buildings or factories. However, due to the non-linear character and time-varying nature of the processes occurring within the heat exchanger, designing optimal control is often a complex task. This complexity can be most effectively addressed using mathematical modeling, which allows for the safe testing of various scenarios without interfering with real operation.

The use of digital models extends beyond design and simulation. They are increasingly applied in the optimization of control systems, enabling better energy efficiency. From an ecological perspective, digital models can help reduce energy consumption and leakage by predicting system behaviors more accurately. They can also more dynamically respond to the varying heat demand in distribution systems, particularly for more complex systems. Moreover, digital models play a crucial role in predictive maintenance, as they monitor system performance and predict failures before they occur, minimizing downtime and maintenance costs. Virtual sensors, derived from digital models, can replace physical sensors, thus reducing the costs of hardware installation and maintenance. Furthermore, digital models contribute to cost reduction in production by improving the overall efficiency of manufacturing processes and identifying areas for resource optimization [5], [6].

The goal of this work is to design a discrete grey-box model of a heat exchanger station and identify its parameters using experimental data. Furthermore, the paper aims to compare this semi-physical approach with a black-box modeling approach utilizing a LSTM neural network without any prior knowledge of the physical structures. The comparison is based on the accuracy of the model outputs validated against real-world data measured on a physical training station.

II. Physical Model Description and Data Acquisition

The physical heat exchanger station modeled in this study is located at the Department of Control and Instrumentation, Brno University of Technology (BUT). The system consists of two main circuits: a primary circuit representing the heat distributor and a secondary circuit acting as the heat consumer. These circuits are thermally coupled via a counter-flow plate heat exchanger.

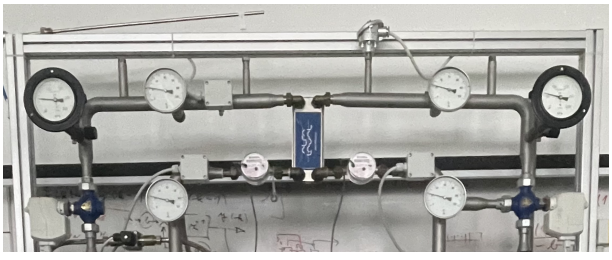

Fig. 1. Physical model

A. Physical Model

The piping and instrumentation diagram of the modeled system is depicted in Fig. 2.

The primary circuit is equipped with a flow heater (2 kW) serving as the heat source and a circulation pump. It includes a three-way control valve V1 and a two-way control valve V3, both operated by servomotors with a 0–10 V control signal. The secondary circuit contains a radiator that dissipates heat to the environment, simulating a thermal load, and is equipped with a single three-way mixing valve V2 and a circulation pump.

Instrumentation is critical for both control and identification. The system is equipped with resistance temperature detectors (RTD) with 4–20 mA current outputs placed at the inlet and outlet of the heat exchanger in both circuits (T1, T2, T3, T4) and an ambient temperature sensor ($\mathrm{T_{amb}}$).

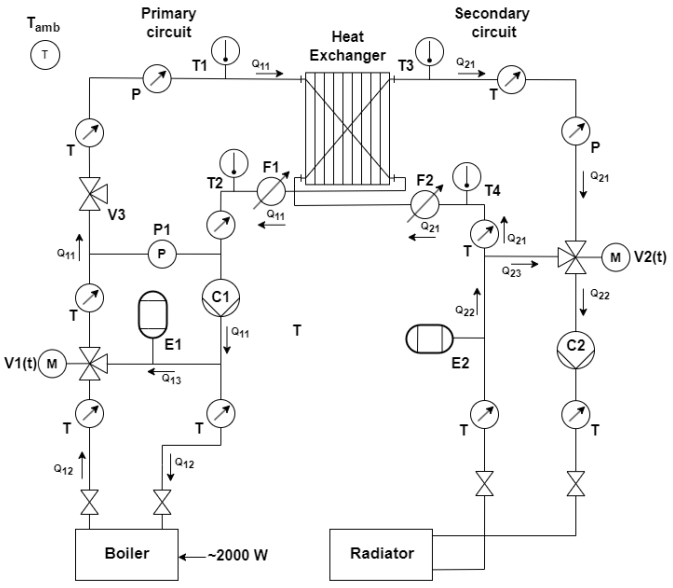

Fig. 2. Piping and instrumentation diagram

B. Control system

The entire technology is controlled by Siemens SIMATIC S7-1500 Programmable Logic Controller (PLC) (CPU 1512C-1 PN), which handle data acquisition and actuator control. The PLC also communicate with a Human-Machine Interface (HMI). The HMI is used to control the physical model, display its current status, and configure and execute measurement scenarios. The analog inputs for temperature and pressure sensors utilize a 4–20 mA current loop to minimize noise, while the actuators (valves) are controlled via 0–10 V voltage signals. The control logic and data handling were implemented in the TIA Portal.

C. Data Acquisition and Identification Scenarios

A large amount of high-quality data is necessary for identifying Grey-box parameters and for training the LSTM network. Data acquisition was performed using the *Traces* function within the TIA Portal, allowing for real-time signal recording without external measuring devices. The sampling period was set to $T_\mathrm{s} = 5$ seconds, which was determined to be sufficient to capture the transient response. Prior to each measurement, it is important to ensure that the pressures are equal in both the primary and secondary circuits. The operating pressure was set to 50 kPa. Before starting the measurement scenario, it is necessary to heat the water in the flow heater to the working temperature and turn on the circulation pumps. The boiler working temperature was set to 75 °C.

To identify the system dynamics and train the neural network, specific identification scenarios were designed. These scenarios involved randomized changes in the opening of valves V1 and V2, generated using a uniform distribution using the MATLAB *randi* function. This approach ensures that the dataset covers various operating points and transient states, which is crucial for the generalization capability of the data-driven LSTM model. A total of eight extensive measurement scenarios were recorded, providing a robust dataset for training and cross-validation. Each scenario consists of 11,111 samples, corresponding to a total duration of approximately 15.4 hours per scenario.

III. Mathematical Model

Two distinct modeling approaches were implemented to create the digital model: a semi-physical discrete Grey-box model and a data-driven Black-box model based on LSTM neural network.

A. Discrete Grey-Box Model

The continuous-time thermodynamic equations of the modeled heat exchange station, based on calorimetric principles and the law of conservation of energy, were comprehensively derived and validated in our previous work [7]. To enable the deployment of the digital shadow onto an industrial PLC for real-time execution, these continuous models had to be discretized.

The discretization was performed using the forward finite difference method (Euler method). The continuous derivative of a general state variable, such as temperature $T(t)$, is approximated as:

$$T'(t) \approx \lim_{\tau \to T_s} \frac{T(t+\tau) - T(t)}{\tau} \approx \frac{T(kT_\mathrm{s} + T_\mathrm{s}) - T(kT_\mathrm{s})}{T_\mathrm{s}} \tag{1}$$

where T_s is the sampling period (in seconds) and k represents the discrete time step.

Applying this approximation principle to the general continuous calorimetric equation yields its discrete recurrence form:

$$T_\mathrm{out}(kT_\mathrm{s} + T_\mathrm{s}) = \frac{T_\mathrm{s}}{m\,c}\left[p_\mathrm{u}(kT_\mathrm{s}) + p_\mathrm{p}(kT_\mathrm{s}) + p_\mathrm{v}(kT_\mathrm{s})\right] + T_\mathrm{out}(kT_\mathrm{s}) \tag{2}$$

where p_u, p_p, and p_v represent the discrete thermal powers (heat dissipation, input heating power, and flowing fluid

power), m is the mass of the water, and $c = 4180$ J·kg^{-1}·K^{-1} is the specific heat capacity.

To demonstrate the procedure, the heat exchanger is presented as the core non-linear component. The discrete state-space equations for the primary ($T_{1\text{out}}$) and secondary ($T_{2\text{out}}$) outlet temperatures are:

$$T_{1\text{out}}(kT_\text{s} + T_\text{s}) = \frac{T_\text{s}}{m_\text{x}\, c}\big\{c\, Q_{1i}\,[T_{1\text{in}}(kT_\text{s}) - T_{1\text{out}}(kT_\text{s})] + \\ + b_\text{x}\,[T_{2\text{out}}(kT_\text{s}) - T_{1\text{out}}(kT_\text{s})]\big\} + T_{1out}(kT_\text{s}) \tag{3a}$$

$$T_{2\text{out}}(kT_\text{s} + T_\text{s}) = \frac{T_\text{s}}{m_\text{x}\, c}\big\{c\, Q_{2o}\,[T_{2\text{in}}(kT_\text{s}) - T_{2\text{out}}(kT_\text{s})] + \\ + b_\text{x}\,[T_{1\text{out}}(kT_\text{s}) - T_{2\text{out}}(kT_\text{s})]\big\} + T_{2out}(kT_\text{s}) \tag{3b}$$

where m_x is the effective mass and b_x is the heat transfer coefficient.

The dynamic equations for the remaining physical components were discretized using the same approach.

B. LSTM Network Model

As a data-driven alternative, a LSTM network was employed. LSTM is a type of Recurrent Neural Network (RNN) capable of learning long-term dependencies, making it suitable for thermal systems [8]–[10]. The network architecture was implemented in MATLAB using the Deep Learning Toolbox [11] and consists of:

- **Sequence Input Layer:** 4 inputs (V1 and V2 opening values, boiler state, ambient temperature).
- **LSTM Layer:** Composed of memory cells with input, forget, and output gates to regulate information flow. The number of hidden units was optimized via Grid Search.
- **Fully Connected Layer:** Maps the hidden states to the 4 outputs (T1, T2, T3, T4).

IV. Identification Process

To assess the generalization capability of the proposed models and optimize their parameters, a Leave-One-Out Cross-Validation (LOOCV) approach was employed.

A. Leave-One-Out Cross-Validation (LOOCV)

The performance of the models was evaluated using LOOCV [12]. The dataset comprised eight independent measurement scenarios, resulting in eight validation folds. In each iteration, seven scenarios formed the training set for parameter identification and network training, while the single remaining scenario was held out for validation.

The accuracy of the prediction was quantified using the Root Mean Square Error (RMSE) metric:

$$\text{RMSE}_f = \sqrt{\frac{1}{KN}\sum_{i=1}^{N}\sum_{k=1}^{K}[\hat{y}_i(k) - y_i(k)]^2} \tag{4}$$

where $N = 4$ is the number of measured signals (four RTD temperatures), K is the number of samples in the scenario, $y_i(k)$ is the measured reference temperature, and $\hat{y}_i(k)$ is the model prediction. The overall performance of each method is determined by the average RMSE across all eight folds.

B. Grey-Box Parameter Identification

The unknown parameters of the semi-physical model were estimated using two distinct MATLAB optimization algorithms:

- **fminsearch:** A derivative-free method utilizing the Nelder-Mead simplex algorithm [13]. To accelerate convergence and ensure parameter positivity, the search space was transformed using a logarithmic scale ($\mathbf{x} = \ln \boldsymbol{\theta}$, where $\boldsymbol{\theta}$ is a vector containing all model parameters). The objective function was set to minimize the overall RMSE. This was calculated as the square root of the averaged Mean Squared Error (MSE) across all individual training scenarios. The maximum number of iterations was limited to 2500.
- **nlgreyest:** A specialized MATLAB function for non-linear grey-box model estimation. The identification was performed using the Levenberg-Marquardt algorithm [14], [15] due to its robust convergence properties. The maximum number of iterations was set to 300.

Both methods employed a discrete mathematical model defined by difference equations, implemented via the MATLAB *idnlgrey* function.

C. LSTM Network Training

The LSTM network was trained using the MATLAB *trainnet* function, with the MSE specified as the loss function. The training parameters, including the Adam optimizer and a maximum limit of 3000 epochs, were configured via the *trainingOptions*. To determine the optimal hyperparameters, a grid search was performed, systematically testing all possible combinations of hidden units and learning rates as outlined in Table I.

TABLE I
LSTM Grid Search Parameters

Hidden Units	Learning Rate
1, 2, 5, 10, 20, 50, 100, 200	0.001, 0.002, 0.005, 0.010

The configuration that yielded the lowest average validation RMSE was ultimately selected for the final comparison against the grey-box models.

To eliminate the effect of varying feature scales, the dataset was standardized using Z-scores prior to training:

$$Z = \frac{x - \mu}{\sigma} \tag{5}$$

where Z is the standardized value, x is the original (denormalized) value, μ is the mean value, and σ is the standard deviation.

To calculate the RMSE in the original physical units (°C), the network predictions were subsequently destandardized:

$$x = Z \cdot \sigma + \mu \tag{6}$$

V. Results and Discussion

The performance of the three modeling approaches was evaluated based on their prediction accuracy (RMSE) and computational demands during the identification or training phase. The parameter estimation process was executed on a system featuring a 16-core *AMD Ryzen 9 5950X* CPU and an *NVIDIA GeForce RTX 3090* GPU.

A. LSTM Hyperparameter Optimization

The Grid Search method revealed significant differences in network performance based on the chosen configuration (see Fig. 3). Evaluating all combinations from Table I, the optimal network architecture was identified as 100 hidden units and learning rate of 0.01. This configuration provided the best balance between capturing the complexity of non-linear dynamic and preventing overfitting, yielding the lowest validation error.

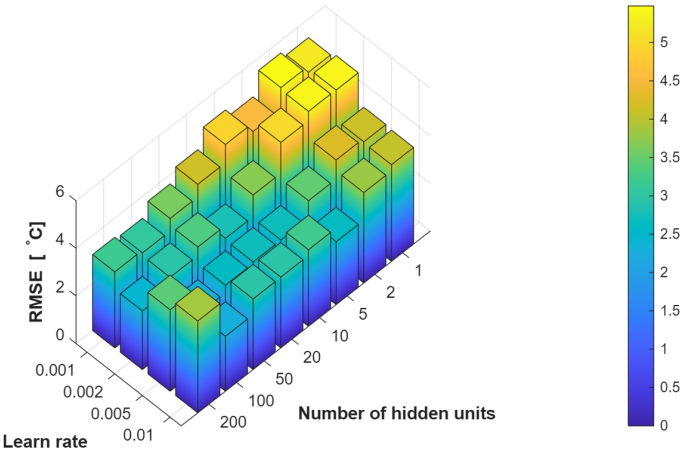

Fig. 3. Grid search results for the LSTM network based on average RMSE

B. Performance Comparison (RMSE)

The overall prediction accuracy was assessed using the average RMSE across all 8 validation folds. The results clearly indicate the superiority of the data-driven approach (see Fig. 4 and Fig. 5).

The **LSTM network** achieved the highest accuracy with an average RMSE of 2.3 °C, with Fold 8/8 yielding the lowest individual RMSE of 1.1 °C. Notably, Fold 1/8 was the only instance where the LSTM did not outperform the grey-box approaches.

In contrast, the semi-physical grey-box models exhibited higher prediction errors. The **fminsearch** method achieved an average RMSE of 3.3 °C (best performance in Fold 8/8 with 2.7 °C). The **nlgreyest** function yielded the highest average error of 3.6 °C (best performance in Fold 6/8 with 3.0 °C).

The lower accuracy of the grey-box models can be attributed to unmodeled parasitic phenomena and simplifying assumptions (e.g., constant heat capacity, ideal mixing, negligible heat losses in certain components), which the purely data-driven LSTM model are able to compensate.

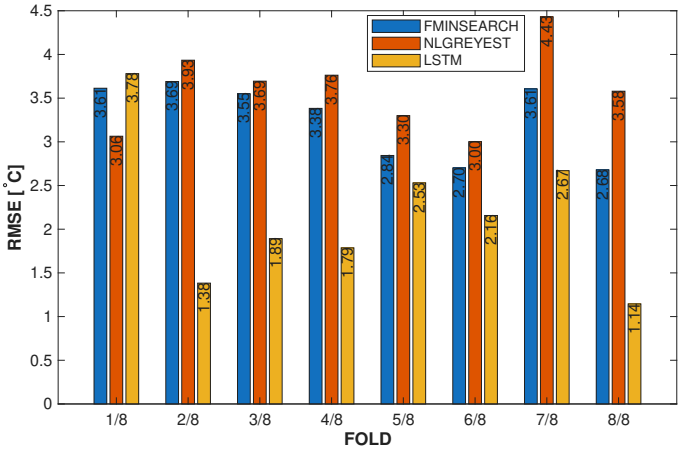

Fig. 4. RMSE comparison of the modeling approaches

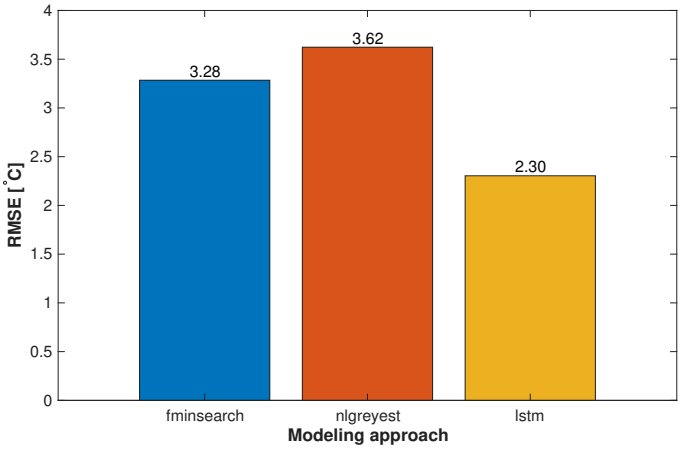

Fig. 5. Comparison of average RMSE across the modeling approaches

C. Computational Cost

An important aspect of practical deployment is the time required for parameter identification or network training (see Fig. 6). The training time for the best-performing LSTM configuration (100 hidden units and a learning rate of 0.01) remained relatively constant across all folds, taking approximately 49 minutes per fold. This consistency is due to the fixed number of epochs (3000) and the uniform size of the training sets.

Conversely, the grey-box optimization methods (*fminsearch* and *nlgreyest*) exhibited significant variability in computational time. The time heavily depended on the specific data fold and the initial parameter estimates, as these strongly influence the convergence speed of the underlying numerical solvers. Consequently, the process often converges and terminates before the predefined iteration limit is reached. Although this approach is computationally more efficient than the LSTM-based method, it is susceptible to becoming trapped in local minima unless the initial parameter estimates are sufficiently accurate.

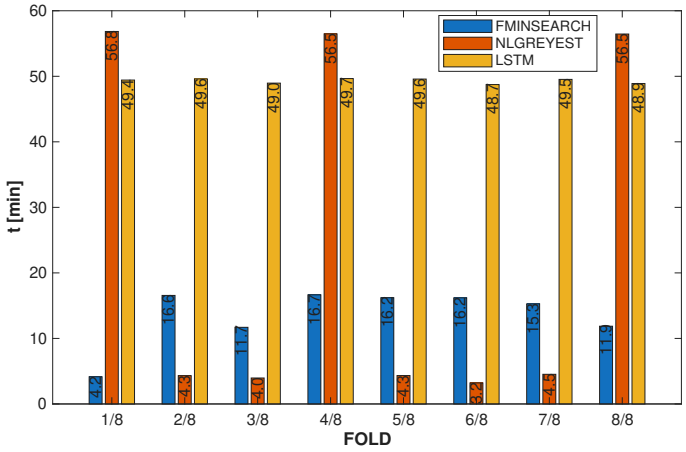

Fig. 6. Comparison of computational time for parameter identification across the modeling approaches

VI. Conclusion

This paper presented the design and validation of a digital model for a laboratory heat exchanger station. To capture the complex and non-linear dynamics of the thermal system, two distinct modeling paradigms were investigated: a discrete semi-physical grey-box model and a purely data-driven LSTM neural network. Both approaches were validated using LOOCV on datasets acquired from the physical process.

The results demonstrated that the LSTM network significantly outperformed the traditional grey-box approaches in terms of prediction accuracy, achieving an average RMSE of 2.3 °C compared to 3.3 °C and 3.6 °C for *fminsearch* and *nlgreyest*, respectively. Furthermore, the LSTM approach exhibited highly predictable and consistent training times.

While the grey-box models resulted in lower accuracy and variable identification times, they retain the advantage of physical interpretability. Nevertheless, for the purpose of creating a highly accurate digital model suitable for advanced control scenario testing, the LSTM-based data-driven approach proved to be the more effective solution.

Acknowledgment

The completion of this paper was made possible by the grant No. FEKT-S-26-8988 – “Advanced Methods in Cybernetics, Robotics, and Artificial Intelligence” financially supported by the Internal science fund of Brno University of Technology.

References

[1] R. SARACCO, “Digital Twins: Bridging Physical Space and Cyberspace”, IEEE Xplore, vol. 52, no. 12, 2019, doi: 10.1109/MC.2019.2942803.

[2] L. Li, S. Aslam, A. Wileman, and S. Perinpanayagam, “Digital Twin in Aerospace Industry: A Gentle Introduction”, IEEE Access, vol. 10, no. 10, 2022, doi: 10.1109/ACCESS.2021.3136458.

[3] M. Sjarov et al., “The Digital Twin Concept in Industry – A Review and Systematization”, in 2020 25th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA), IEEE, 2020, pp. 1789-1796. doi: 10.1109/ETFA46521.2020.9212089.

[4] W. M. P. van der Aalst, O. Hinz, and C. Weinhardt, “Resilient Digital Twins: Organizations Need to Prepare for the Unexpected”, Business & Information Systems Engineering, vol. 63, no. 6, Sep. 2021, doi: 10.1007/s12599-021-00721-z.

[5] KANTOR, M.; HUSÁK, M. Predictive maintenance with digital model. In Proceedings I of the 31st Conference STUDENT EEICT 2025: General papers. Brno: VUT, 2025. p. 87-90. ISBN: 978-80-214-6321-9.

[6] R. van Dinter, B. Tekinerdogan, and C. Catal, “Predictive maintenance using digital twins: A systematic literature review”, Information and Software Technology, vol. 151, 2022, doi: 10.1016/j.infsof.2022.107008.

[7] T. Benešl, M. Husák, O. Mihálik, R. Vancl, and Z. Bradáč, “Digital twin of heat exchange station”, IFAC-PapersOnLine, vol. 58, no. 9, 2024, doi: 10.1016/j.ifacol.2024.07.415.

[8] S. Hochreiter and J. Schmidhuber, “Long Short-Term Memory”, Neural Computation, vol. 9, no. 8, Nov. 1997, doi: 10.1162/neco.1997.9.8.1735.

[9] Y. Bengio, P. Simard, and P. Frasconi, “Learning long-term dependencies with gradient descent is difficult”, IEEE Transactions on Neural Networks, vol. 5, no. 2, 1994, doi: 10.1109/72.279181.

[10] F. A. Gers, J. Schmidhuber, and F. Cummins, “Learning to Forget: Continual Prediction with LSTM”, Neural Computation, vol. 12, no. 10, Oct. 2000, doi: 10.1162/089976600300015015.

[11] Deep Learning Toolbox”, MATLAB. Accessed: Mar. 07, 2026. [Online]. Available: https://www.mathworks.com/products/deep-learning.html

[12] G. James, D. Witten, T. Hastie, and R. Tibshirani, An Introduction to Statistical Learning with Applications in R. Springer, 2013. doi: 10.1007/978-1-4614-7138-7.

[13] J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright, “Convergence Properties of the Nelder–Mead Simplex Method in Low Dimensions”, SIAM Journal on Optimization, vol. 9, no. 1, 1998, doi: 10.1137/s1052623496303470.

[14] K. Levenberg, “A method for the solution of certain non-linear problems in least squares”, Quarterly of Applied Mathematics, vol. 2, no. 2, 1944, doi: 10.1090/qam/10666.

[15] D. W. Marquardt, “An Algorithm for Least-Squares Estimation of Nonlinear Parameters”, Journal of the Society for Industrial and Applied Mathematics, vol. 11, no. 2, 1963, doi: 10.1137/0111030.

ThermoFisher SCIENTIFIC

www.thermofisher.jobs.cz

Thermo Fisher Scientific is the world leader in scientific products and services. Our company employs more than 125,000 people worldwide and our turnover exceeds $44 billion. We supply our customers with a wide range of scientific products, from basic laboratory equipment to state-of-the-art analytical instruments and help our customers create a healthier, cleaner, and safer world.

The Brno branch focuses on the development and manufacture of unique scientific instruments – **electron microscopes and spectrometers**. Our customers include world-renowned universities, leading electronics manufacturers, and pharmaceutical companies. They also include development centers in the space, semiconductor, mining, automotive, and aviation industries.

The largest exporter in the South Moravian Region

6th largest exporter in the Czech Republic (2024)

TOP Employer in the South Moravian Region

More than 2,200 employees at the Brno branch



More about us

The world leader in serving science

Design and Implementation of a Framework for Automatic Generation of Attack Graphs Using Open Sources

1st Róbert Ruman
Department of Telecommunications
FEEC, Brno University of Technology
Brno, Czechia
xruman02@vutbr.cz

2nd Viet Anh Phan
Department of Telecommunications
FEEC, Brno University of Technology
Brno, Czechia
243760@vutbr.cz

***Abstract*—From the inserted identifier of one of the cybersecurity datasets (Common Vulnerabilities and Exposures, Common Weaknesses Enumeration, Common Attack Pattern Enumeration and Classification, MITRE ATTACK, MITRE DEFEND and Common Platform Enumeration), or a brief description of these entities, detailed information is fetched. Where these datasets do not provide links, the system suggests additional connections using a Natural Language Processing model. These steps are repeated for every item, creating a complete roadmap that shows how a vulnerability can lead through multiple stages of exploitation and ultimately to its mitigation techniques. The final output is displayed in a web-based interactive graph that allows users to display metadata such as the description, identifier and source of these entities. The displayed items are grouped based on their relevancy within the layer, and they can be moved by the user and upon interacting with them, a JSON file can opened displaying more detailed information about the specific node type. Testing with various types of input showed that the program handles all input accordingly, producing accurate results. Future improvements include boosting the accuracy of the NLP model.**

***Index Terms*—CVE, CWE, CAPEC, MITRE, ATT&CK, D3FEND, CPE, vulnerability, weakness, taxonomy mapping, severity, mitigation, penetration testing, NLP, Cybersecurity knowledge graphs, Vulnerability exploitation pathways, Mitre att&ck and defend integration, Natural language processing linking, Interactive visualization of security metadata**

I. Introduction

Analyzing vulnerabilities in isolation is no longer sufficient; understanding their relationships to underlying weaknesses, attack patterns, and defensive techniques is essential [1]. However, this is hindered by the fragmentation of cybersecurity knowledge across heterogeneous taxonomies [2]. Although authoritative datasets (e.g., CVE, CWE, CAPEC, ATT&CK, D3FEND) are deeply integrated into industry workflows, they evolve independently and provide only partial, unidirectional cross-references [3], [4]. Consequently, analysts must manually correlate disparate sources to reconstruct attack scenarios [5].

While attack graphs offer a promising approach to model these paths, existing generation tools typically construct static knowledge graphs limited to a subset of taxonomies. Furthermore, they restrict analyst workflows by requiring a single type of identifier as input and relying exclusively on predefined, often incomplete mappings [6], [7]. To address these limitations, this work introduces **Veacor**, a framework capable of accepting diverse identifiers or textual descriptions from multiple major cybersecurity taxonomies. Veacor automatically enriches and unifies this retrieved information—leveraging Natural Language Processing (NLP) to bridge missing links [8]—into a single, coherent attack roadmap.

This work is supported by Technology Agency of the Czech Republic under Grant FW11020057.

II. Related Work

While standardized taxonomies like CVE, CWE, CAPEC, ATT&CK, and D3FEND offer structured threat intelligence, their practical utility is hindered by sparse, unidirectional cross-mappings. This fragmentation makes it difficult to correlate vulnerability-centric data with adversarial behavior and defensive countermeasures. Consequently, prior research on attack graph generation and cybersecurity knowledge graphs typically relies on manually curated, static datasets restricted to a subset of available taxonomies [6]. Furthermore, many existing tools restrict analyst-driven exploratory workflows by requiring fixed input types, such as a specific CVE identifier.

In contrast to existing approaches, this work emphasizes flexible, identifier-agnostic roadmap generation across various domains. By combining explicit mappings with semantic inference, the proposed system aims to address the limitations observed in prior research while maintaining practical usability.

III. System Overview

The proposed system is designed to automatically generate an attack roadmap that captures relationships between vulnerabilities, weaknesses, attack patterns, adversary techniques, and defensive measures. The primary objective is to provide a unified, structured view of cyber-threat information starting from minimal user input.

The system accepts various input types, including identifiers, textual descriptions, vendors and their products. Upon

TABLE I
COMPARATIVE ANALYSIS OF AUTOMATED CYBERSECURITY MAPPING TOOLS

Metric	GraphKer	CVE2CAPEC	Veacor
Data Dependency	Official Links only	Manual RegEx	**Official + Semantic**
Link Discovery	Deterministic	Heuristic	**Inference-based**
Mapping Flow	Static/Historical	Linear Pipeline	**Relational Mesh**
Taxonomy Support	Fixed (4)	Sequential (5)	**Extensible (5+)**
Directionality	Unidirectional	Unidirectional	**Omnidirectional**

receiving an input identifier, the system first determines its type and initiates a guided expansion process based on predefined traversal rules. These rules define permissible transitions between taxonomies, such as mapping vulnerabilities to weaknesses, weaknesses to attack patterns, and attack techniques to defensive controls.

At a high level, the system follows a modular pipeline consisting of input classification, data retrieval, relationship expansion, and result consolidation. Each module operates independently but contributes to the construction of a single, coherent roadmap. The final output is represented as a structured graph-like dataset, where nodes correspond to cybersecurity entities and edges represent semantic or explicit relationships between them as shown on figure 1. By abstracting the underlying data sources and traversal logic, the system enables analysts to explore attack paths and defensive coverage without requiring deep familiarity with the individual taxonomies or manual correlation efforts.



Fig. 1. The layout of a visually represented

A. Data Sources and Knowledge Integration

The system integrates data from authoritative cybersecurity knowledge bases. The system integrates data from CVE, CWE, CAPEC, MITRE ATT&CK, and D3FEND using official APIs and repositories [9]–[13].

Knowledge integration is achieved through a hybrid approach that combines explicit relationships provided by the source datasets with inferred associations derived from natural language descriptions. Explicit mappings are used whenever available. However, because these mappings are often incomplete or unidirectional, the system supplements them using semantic similarity techniques applied to textual descriptions.

To ensure consistency and avoid redundancy, the system maintains a global registry of visited entities during roadmap generation. This mechanism prevents cyclic expansions and duplicate node retrieval. Additionally, caching is employed to reduce repeated data fetching and improve performance during large expansions.

By combining structured data retrieval with semantic inference, the system provides more comprehensive coverage of potential attack and defense relationships than approaches relying solely on predefined mappings.

B. Roadmap Generation Methodology

The roadmap generation process is centered on the systematic expansion of cybersecurity entities across multiple taxonomies while preserving semantic consistency and limiting combinatorial growth. Once the input identifier type is detected, the system applies a predefined set of expansion rules that govern permissible transitions between domains. These rules encode domain knowledge regarding how vulnerabilities, weaknesses, attack patterns, adversary techniques, and defensive measures are related.

For example, when starting from a CVE identifier, the system expands toward associated CWE entries, which are subsequently linked to CAPEC attack patterns, MITRE ATT&CK techniques, and finally D3FEND defensive measures. Conversely, when starting from a defensive technique, the expansion proceeds upward through associated attack techniques and inferred weaknesses and vulnerabilities. This bidirectional traversal enables flexible analysis from both offensive and defensive perspectives.

To address incomplete mappings across taxonomies, the system employs a hybrid linking strategy. Explicit links provided by authoritative sources are prioritized, while additional relationships are inferred using a Natural Language Processing model that levarages semantic similarity and keyword overlap. Only links exceeding a predefined confidence threshold – which is set to 0.85 equivalent to an 85% match – are retained. The scores for each link are calculated using cosine similarity which calculates the angle of vectors the sentence transformer created from the two textual inputs. This approach allows the system to infer plausible associations while minimizing false positives.

During expansion, the system tracks visited entities to prevent infinite loops and redundant retrievals. Optional aggregation limits are applied to control the maximum number of entities per category, ensuring that the resulting roadmap remains interpretable and suitable for analyst-driven exploration.

To evaluate the efficacy of this NLP-driven semantic inference, a meta-analysis has been conducted on a dataset of inferred links. Validation was performed using a dual LLM-as-a-Judge methodology to establish consensus-adjusted precision, comparing a semantically flexible model (Model A) against a strict, framework-bound model (Model B). [14]

As detailed in Table II, the underlying NLP tool demonstrates robust, highly practical accuracy for mapping complex, multi-stage cybersecurity relationships across disparate frameworks. Evaluating the system using strict structural matching against official taxonomies yields an artificially low base-

TABLE II
COMPARATIVE VALIDATION METRICS FOR NLP-INFERRED MAPPINGS

Metric	Model A (Semantic)	Model B (Conservative)	Consensus
Total Samples Evaluated	240	240	240
'VALID' Judgments	198	115	96
Precision	82.5%	47.9%	**40.0%**

line accuracy of approximately 9.3%. However, this metric is fundamentally misleading; the low score stems from the pervasive lack of explicit, officially documented links across existing taxonomies rather than inaccurate model inferences. Consequently, comparing dynamically generated relationships solely against this sparse ground truth is methodologically flawed, as it inherently penalizes logically sound, operationally valid mappings simply because they are not yet documented in official repositories.

To address this limitation and properly assess the unmapped connections, a dual LLM-as-a-Judge evaluation framework was deployed to independently verify the technical validity of the inferred links. By anchoring the dataset against these expert evaluations, the final "True Accuracy" is shown to fall between 47.9% (when evaluated under strict, conservative framework definitions) and 82.5% (when evaluated under flexible, semantic real-world applicability). Furthermore, the dual models established a reinforced, unassailable consensus baseline of 40.0% where both strictly agreed on the validity of the undocumented links. Ultimately, this demonstrates that the NLP-driven automated mapping capability is a powerful, production-ready triage asset capable of bridging the connectivity gaps in modern threat intelligence.

C. Proposed framework

The system is implemented as a modular Python-based framework, with separate components responsible for input handling, data retrieval, semantic inference, and roadmap construction. Each taxonomy is accessed through a dedicated retrieval module that encapsulates source-specific logic, such as API queries, web scraping, and data normalization.

Natural language processing is employed to support semantic inference between entities lacking explicit mappings. Textual descriptions are cleaned to remove noise such as citations and URLs, tokenized, and embedded using a sentence-transformer model. The sentence-transformer used in this project is the *multi-qa-mpnet-base-dot-v1* due to its higher classification performance [8]. A hybrid scoring mechanism combines semantic similarity, keyword overlap, and description length normalization to rank candidate relationships.

To improve performance and scalability, the system implements caching at multiple levels, including fetched entity data and computed embeddings. A global visitation registry ensures consistency during recursive expansion and avoids duplicate processing. The final roadmap is serialized into structured output files, allowing downstream visualization or integration with external analysis tools.

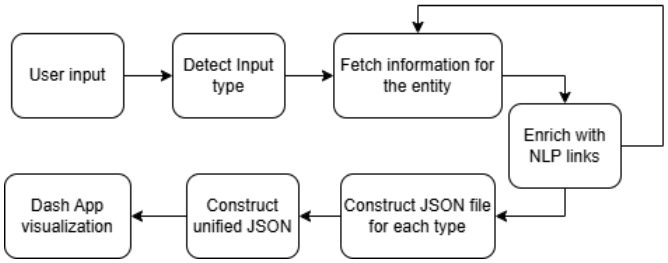

Fig. 2. Workflow of the proposed framework

D. Case Study

Testing with the CVE-2024-3094 identifier resulted in the generation of a roadmap containing 117 nodes across five taxonomic layers, as shown in Fig. 3. While traditional databases only provide a fraction of these links, the system's semantic inference successfully mapped the vulnerability to previously unlinked attack patterns and defensive strategies. By automatically bridging the gap between these isolated taxonomies, the generated roadmap allows a security analyst to identify critical defensive countermeasures (D3FEND) and attack patterns (CAPEC) that are not currently mapped in official databases. This enables proactive prioritization of mitigations for sophisticated supply-chain vulnerabilities, reducing the window of exposure by eliminating the need for manual, cross-dataset research during an active incident response.

Input Type	Avg. Runtime	Avg. Nodes	Processing time per Node
CVE	217.69s	73.6	2.96s
CWE	278.73s	108.77	2.56s
CAPEC	234.72s	80.71	2.90s
ATT&CK	198.13s	75.22	2.63s
D3FEND	169.23s	19.83	8.53s

TABLE III
AVERAGE RUNTIME AND NODE STATISTICS FOR EACH INPUT TYPE.

The generated roadmap (Fig. 3) illustrates the relationships between vulnerabilities, attack techniques, and mitigations. The system successfully identified relevant links even in the absence of explicit mappings.

A second test using a D3FEND identifier confirmed the system's ability to construct roadmaps starting from defensive techniques.

These results demonstrate that the proposed system can effectively transform fragmented cybersecurity data into a coherent attack roadmap for analysis and defensive planning.

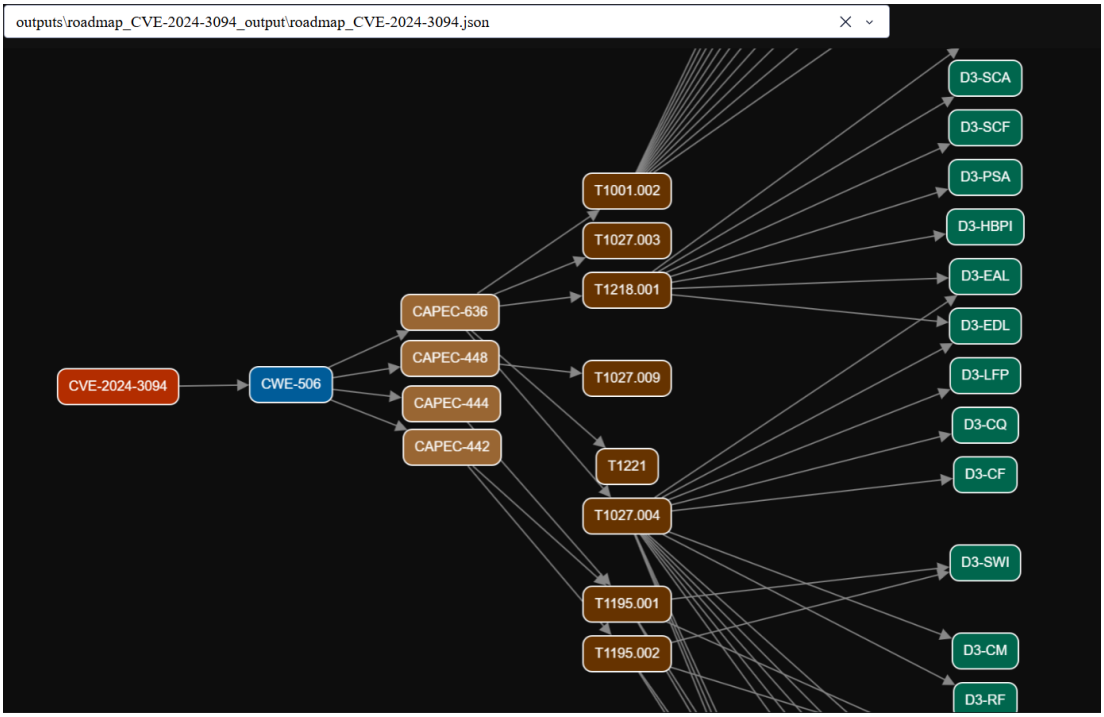

Fig. 3. Displayed roadmap by the system upon entering CVE-2024-3094 during testing

IV. Conclusion and Future Work

This paper presented a system for automatic generation of attack roadmaps by integrating multiple cybersecurity taxonomies into a unified representation. By supporting flexible input types and combining explicit mappings with semantic inference, the proposed approach reduces manual correlation effort and improves visibility across vulnerability, attack, and defense domains. Future work will focus on optimizing time- and resource-efficiency while further refining the accuracy of the NLP module. Despite high semantic precision, it is important to acknowledge the inherent limitations of NLP-based mapping. Semantic inference introduces the risk of false positives in cases where linguistic contexts are similar but underlying technical mechanisms differ. Given the observed 40.0% consensus among evaluation models, a human-in-the-loop triage approach is recommended to verify inferred links before they are integrated into production-level defensive playbooks. Ultimately, the proposed system aims to serve as a practical foundation for automated threat modeling and decision support in complex cybersecurity environments.

References

[1] J. Simola and T. Leppänen, "Identification of the Emerging Sources of Cybersecurity Threats," in *Proc. 24th Eur. Conf. Cyber Warfare Secur. (ECCWS)*, Reading, U.K., 2025. [Online]. Available: https://papers.academic-conferences.org/index.php/eccws/article/view/3569

[2] T.-W. Chen, E. Nunes, and P. Chan, "Challenges in Cyber Threat Intelligence Correlation," *IEEE Trans. Dependable Secure Comput.*, 2021. [Online]. Available: https://ieeexplore.ieee.org/document/9140492

[3] National Institute of Standards and Technology, "NVD Announces Need for Community Assistance to Improve CVE–CWE Mappings," 2023. [Online]. Available: https://nvd.nist.gov/general/news/need-user-assistance

[4] S. Patton, B. Miller, and T. Opferman, "Linking MITRE ATT&CK and CAPEC Using Natural Language Processing," in *Proc. ACM Workshop Cyber-Physical Syst. Secur.*, 2021. [Online]. Available: https://dl.acm.org/doi/10.1145/3461708.3461714

[5] M. Ali, H. Zhang, and L. Wang, "A Survey of Attack Graph Analysis Methods," *IEEE Access*, vol. 7, pp. 110361–110383, 2019. [Online]. Available: https://ieeexplore.ieee.org/document/8781911

[6] A.-M. Berzovitis, "GraphKer: Cybersecurity Graph Database in Neo4j," 2021. [Online]. Available: https://github.com/amberzovitis/GraphKer

[7] Z. Shi, N. Matyunin, K. Graffi, and D. Starobinski, "Uncovering CWE-CVE-CPE relations with threat knowledge graphs," 2023. [Online]. Available: https://arxiv.org/abs/2305.00632

[8] R. Othman, D. Rimawi, B. Rossi, and B. Russo, "From attack descriptions to vulnerabilities: A sentence transformer-based approach," *J. Syst. Softw.*, 2025, Art. no. 112615. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0164121225002845

[9] "The CVE Program," [Online]. Available: https://www.cve.org/About/Overview#AbouttheCVEProgram

[10] "Common Weakness Enumeration (CWE)," [Online]. Available: https://cwe.mitre.org/about/index.html

[11] "Common Attack Pattern Enumeration and Classification (CAPEC)," [Online]. Available: https://capec.mitre.org/about/index.html

[12] "MITRE ATT&CK Framework," 2023. [Online]. Available: https://www.crowdstrike.com/en-us/cybersecurity-101/cyberattacks/mitre-attack-framework/

[13] "MITRE D3FEND," [Online]. Available: https://d3fend.mitre.org/about/

[14] L. Zheng, W.-L. Chiang, Y. Sheng, S. Zhuang, Z. Wu, Y. Zhuang, Z. Lin, Z. Li, D. Li, E. Xing, and H. Zhang, "Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena," in *Advances in Neural Information Processing Systems*, vol. 36, pp. 46595–46623, 2023. [Online]. Available: https://arxiv.org/abs/2306.05685

Post-Quantum Encryption on FPGA Platform: Czech-Estonian Pilot

Jan Havlín
FEEC, Brno University of Technology
Brno, Czech Republic
241030@vut.cz

Adam Fridrich
FEEC, Brno University of Technology
Brno, Czech Republic
247536@vut.cz

Tomáš Lieskovan
FEEC, Brno University of Technology
Brno, Czech Republic
lieskovan@vut.cz

Patrik Dobiáš
FEEC, Brno University of Technology
Brno, Czech Republic
203169@vut.cz

Jan Hajný
FEEC, Brno University of Technology
Brno, Czech Republic
hajny@vut.cz

***Abstract*—Purpose – The aim of this paper is to present the results of testing a hardware encryptor developed at FEEC BUT that creates a quantum-safe encrypted communication channel. The purpose of this testing was to simulate a normal operational situation, observe the behavior of the encryptor, and record the properties of the connection. Furthermore, the encryptor was subjected to stress testing in order to observe its behavior and identify any shortcomings and optimization possibilities. All tests were also evaluated on unencrypted channels so that the results could be compared with those obtained on encrypted channels.**

Design – Testing was carried out manually using programs for testing and diagnosing network connections. The testing process was also subsequently analyzed in the Wireshark program. The testing design was also devised to allow comparison with the results of measurements conducted in a laboratory environment.

Findings – The results show that the encryptor can operate in a deployment very similar to its potential real-world use. At the same time, they highlight specific shortcomings and optimization opportunities, providing valuable insights for the further development of the encryptor.

Contribution – This paper offers a new perspective on the state of development of the encryptor and provides valuable information about the device's functionality. The paper presents insights that can be used for further development and practical application.

Project - This activity was a part of the Horizon Europe CHESS #101087529 project piloting activities between Czechia and Estonia. Two sectors (academia - BUT and industry - Cybernetica) were involved.

***Index Terms*—encryptor, iperf3, MikroTik, router, post-quantum cryptography**

I. Introduction

The aim of this paper is to provide a summary of the testing of a hardware encryptor based on an FPGA platform (hereinafter referred to as the encryptor). The testing was carried out during an internship at Cybernetica AS. During this time, the encryptor was not only tested, but also optimized and updated. The tests were selected to cover the properties of the network and hardware when deploying the encryptor. The goal of the testing was to measure the characteristics of the encryptor in an environment that approximates as closely as possible the conditions of a potential real-world deployment. To this end, an encrypted channel was established between server infrastructures in the Czech Republic and Estonia.

A. Related work

This work builds on previous papers published on the architecture and development of the hardware encryptor [1]. This project was enabled due to rapid developments in post-quantum cryptography, hardware acceleration, and Internet of Things post-quantum security, such as [2] [3], [4], [5].

II. Encryptor architecture

This section briefly summarizes the architecture and development of the hardware encryptor to date. The algorithms used in the project are the core for the security of the created communication channel. Numerous different algorithms are being used. For the creation of encryption key, the Elliptic Curve Diffie-Hellman, Module-Lattice-Based Key-Encapsulation Mechanism (ML-KEM) algorithms, and, optionally, a Quantum-Key Distribution (QKD) algorithms are used. The resulting key is then fetched as an input to the AES-256-GCM algorithm, which is responsible for encrypting the network traffic itself.

A. ECDH

The Elliptic Curve Diffie-Hellman [6] is a well known key agreement algorithm which allows establishing a shared secret between two parties over and untrusted channel. ECDH was selected as one of the three sources of keys because it is a standardized classical cryptography key establishment algorithm.

B. ML-KEM

The Crystals-Kyber (renamed after standardization as ML-KEM) [7] was chosen as the source of the post-quantum cryptography key. As the only NIST standardized key establishment mechanism, it provides safe source of secrets. Algorithm security is based on the Learning with errors lattice problem. Three variants can be used, varying in security levels and key lengths.

C. QKD

Quantum-Key Distribution is the third and final method through which the encryptor can gather the keys. However, this source is kept as optional because of the low rates and high adoption costs of this technology. Another problem, and the reason the QKD is not used in this testing as well, is the very limited functional range. For example, the distance between the encryptors in our testing scenario was approximately 5 times larger than the range of our laboratory-based QKD system Quantique CLAVIS 3. This system uses the Coherent One-Way QKD (COW-QKD) protocol [8]. COW-QKD enables work with weak coherent pulses at high bit rates. The basic version of this protocol uses pairs of coherent states, used for transmitting 0 and 1, and decoy sequences, designed to detect any eavesdropping attempts. As the result, the protocol is relatively easy to deploy and test, while maintaining high level security thanks to eavesdropping detection.

D. AES

The Rijndael algorithm, used in the Advanced Encryption Standard [9] is a symmetric, block cipher that can be used for the PQC-safe encryption. The main advantage is its speed. It is a very universal algorithm thanks to numerous working modes. In the encryptor, it is the Galois-Counter Mode. To ensure quantum resistance, a 256 bit key length is used.

E. Encryption key derivation

The key combiner implementation is very minimalistic with a resource optimization in mind. It uses two (or three) 256 bit keys and 512 bit transactions for the ciphertext and public parameters. Inside the combiner, those are split into a 64 bit transactions. The keys are then fetched to its respective HMAC functions. Results are then XORed and hashed using the SHA3-512, producing the final key. The complete combination scheme can be observed in the Figure 1.

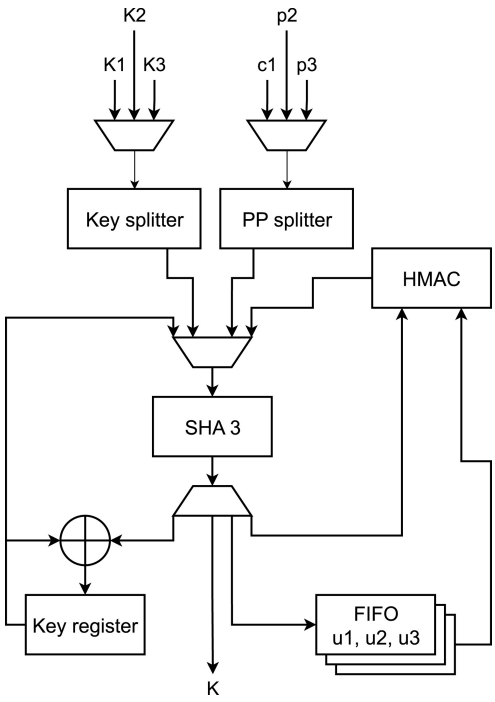

Fig. 1. Key derivation mechanism

III. Method

Before the testing phase, it was first necessary to test various connection topologies of the encryptor and select the most suitable one for testing. The selection process and the different tested topologies are presented in Chapter 3. After selecting and verifying the functionality of the connectivity, the actual testing was carried out. The ping tool was used to measure link functionality and latency. For measuring more detailed information about throughput and packet loss, the iperf3 tool was used. An advantage of this tool is also the possibility of detailed scalability of the computing resources available for testing (in this case, the option to set the number of CPU cores).

The testing itself was conducted according to a pre-prepared plan that reflected the plan of previously performed testing in a laboratory environment. This enables comparison of the individual results. The communication speed with TCP and UDP protocols was tested, as well as a long-term stress test aimed at verifying the stability of the connection. To generate traffic, a virtual machine with the Linux operating system, specifically the Debian 12 distribution, was used. On the other side of the connection, a machine with the Windows 11 operating system was used.

IV. Topology

Prior to the actual testing, a substantial focus was placed on identifying and testing an appropriate topology. The individual components used were a MikroTik CCR2004-16G-2S router, two Omada converters that convert traffic from 10 Gbps to 1 Gbps and vice versa, the encryptor itself, and a Lenovo ThinkBook G3 laptop used for generating and receiving traffic. An internet connection with a maximum speed of 1 Gbps was used to connect to the internet.

The initial version of the setup envisaged using the converters to translate traffic before and after the encryptor. However, this plan required a functional DHCP server provided directly by the end device in order to create a configuration SSH tunnel. Two different Windows utilities capable of turning a computer into a DHCP server were tested, but although they worked on the laptop side, they were unable to assign an IP address to the encryptor. Therefore, it was necessary to insert the router as a bridge between the encryptor and the PC. A DHCP server was then run on this bridge. The encryptor was subsequently successfully configured and started. During an attempt to connect, one of the converters did not work. The solution was to again use the router as a converter via a bridge. The converter was later brought back into operation, but using the router as the central element for all connections proved to be the most effective configuration, and thus the topology was no longer changed. The final setup on the Estonian side is illustrated in Figure 2.

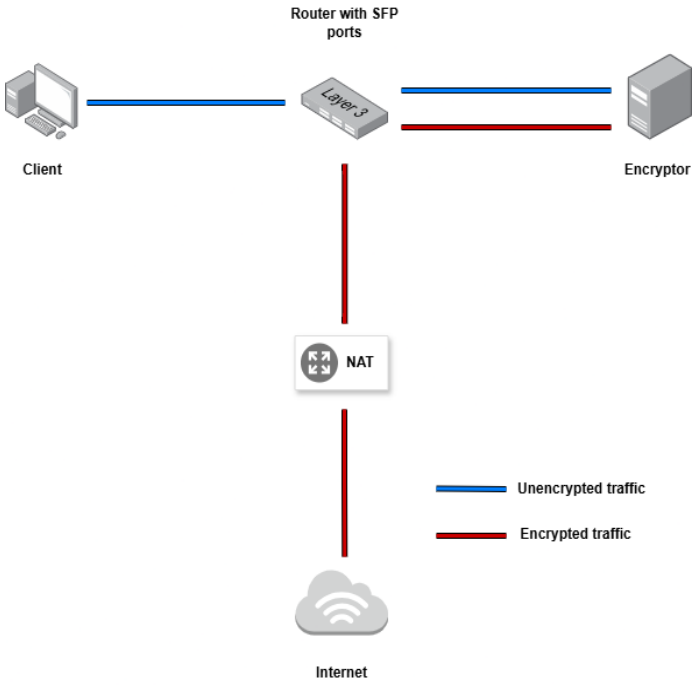

Fig. 2. Network Topology

A. Solved problems

With this setup, it was possible to establish a connection. However, a problem arose at the moment when the keys were to be established between the encryptors. It was later found that the cause was incorrect configuration in the encryptor's firmware. Before this error was fixed, an EoIP (Ethernet over IP) tunnel was created using MikroTik devices, through which it was possible to establish a working encrypted connection. The speed of this connection reached 400 Mbps.

After uploading a new version of the encryptor firmware, it was possible to disable the EoIP tunnel and configure src and dst NAT (Network Address Translation), thereby placing the encryptors and other devices into a private network. This solution is thus independent of the manufacturer of the network devices used (EoIP is a proprietary protocol of MikroTik).

V. Testing of the Czech-Estonian connection

This part of the paper discusses the testing of the established physical encrypted connection between the Czech Republic and Estonia. It presents the results as well as a description of the problems that were encountered during testing and their solutions.

A. Solved encryption issues

Practically throughout the entire testing period, it was necessary to solve the problem of unreliable data transmission over the encrypted connection. Although it was possible to keep the connection itself very stable, as demonstrated by testing with the ping tool, any communication using the UDP and TCP protocols fell far short of the target quality. The unreliability lay mainly in the high ratio of invalid, discarded, or never-delivered packets. It was therefore necessary to subject the communication chain to a detailed analysis. For this purpose, the Wireshark tool and the router's diagnostic modules were used. One of the corrective steps was to disable the Connection Tracking service, which MikroTik devices provide to simplify communication at the network and transport layers. It was also necessary to deal with firewall settings on both sides of the connection. Another problem was the specific nature of communication on devices running Windows OS. Windows uses TCP and UDP Offloading, a service that speeds up communication by processing packet checksums and segmentation directly in the network card hardware. Fortunately, this service can be disabled relatively easily, and the last remaining problem was thus the dropping of seemingly valid packets directly at the encryptors' interfaces.

The main issue with TCP and UDP traffic was the size of some packets. As the minimum Ethernet frame size is 64 bytes, IP packets with a smaller size than 46 bytes are padded with trailing before leaving the network card. However, the encryptor does not recognize this padding and encrypts it with the remaining data and extends the packet size by encapsulating it in UDP and including an authentication tag. Due to this, the ethernet frame exceeds 64 bytes and on some router in the internet network the ethernet trailing is removed. Because of this, verification of the authentication tag fails during decryption and the packet is discarded. For the TCP protocol, we solved this problem by adding a timestamp, which increased the packet to the required size. However, for the UDP protocol, there is no similar option. A solution in the form of modifying the firmware and testing it could not be implemented within the remaining testing time. Encrypted communication over the UDP protocol therefore remained non-functional for the entire duration of testing. However, for comparison purposes, we managed to measure the characteristics of unencrypted UDP communication. Communication over the TCP protocol was fully tested according to plan and, therefore, it was possible to measure the properties of the encrypted connection, as well as the reference values of the unencrypted connection.

B. Test results for the link between Czechia and Estonia

The actual testing was carried out using the iperf3 and ping tools. Bandwidth and latency were tested. For each test, both the properties of the encrypted connection and the properties of the unencrypted connection were measured. The connection was tested for different MTU sizes and different numbers of CPU cores used. It was necessary to adjust the test plan; however, the measurements performed correspond to the testing conducted in the laboratory environment and thus allow for direct comparison of the results. Among the mentioned adjustments was, for example, the omission of testing with active QKD servers, since it was not possible to obtain a QKD key over such a distance. The tables I and II shows the results of TCP testing. It is clear that the encryptor does not introduce any slowdown. The results of a 2-hour performance testing is

TABLE I
COMPARISON OF SPEEDS FOR AN ENCRYPTED TCP CONNECTION

Number of CPU cores	MTU [B]		
	576	1300	1500
1	237 Mbps	294 Mbps	296 Mbps
10	238 Mbps	296 Mbps	303 Mbps

TABLE II
COMPARISON OF SPEEDS FOR AN UNENCRYPTED TCP CONNECTION

Number of CPU cores	MTU [B]			
	576	1300	1500	9000
1	242 Mbps	297 Mbps	296 Mbps	299 Mbps
2	243 Mbps	295 Mbps	293 Mbps	300 Mbps
4	242 Mbps	297 Mbps	298 Mbps	303 Mbps
6	242 Mbps	295 Mbps	302 Mbps	300 Mbps
8	241 Mbps	296 Mbps	298 Mbps	300 Mbps
10	243 Mbps	297 Mbps	296 Mbps	298 Mbps

presented in III. Table IV shows the results of UDP protocol testing on an unencrypted channel.

VI. TESTED FIRMWARE VERSIONS

During testing of the encryptor, certain shortcomings were discovered in the design of communication between encryptors, which prevented direct connection over the Internet, and so changes were designed and tested that attempted to eliminate these shortcomings. In the end, three main firmware versions were developed and tested.

The first version, which was the original one and worked only over EoIP, uses IP packets with the IP protocol field in the header set to experimental values for key establishment. Furthermore, after the key is established, all encrypted traffic is handled by removing the IP header, encrypting the payload, and then reattaching a modified IP header. This solution had two issues that caused it to work only over EoIP. First, the experimental protocol values in the IP header likely caused some device on the path through the Internet to drop the packet, so it never reached the other side. Second, for encrypted packets, the original transport protocol designation remained in the IP header, and in the case of TCP packets, routers along the path attempted to verify the TCP checksum, which was incorrect because the data were encrypted at that point.

The second version was modified to more closely resemble the standard IPSec usage, where packets are encapsulated in the UDP protocol at the transport layer, encapsulating either the key-establishment data or the encrypted payload. With this firmware version, it was already possible to establish keys between both sides without using EoIP encapsulation. However, for reconstruction of the original protocol behind the IP header, the indication was implemented using the source port in the UDP packet, which turned out to be problematic, since NAT on routers in some cases changed this port number. This problem was resolved in the third firmware version, where this information was stored in the destination port instead of the source port.

TABLE III
RESULTS OF THE TWO-HOUR RELIABILITY TEST

Type of connection	Test duration
	2 hours
Without QKD	295 Mbps

TABLE IV
TABLE OF RESULTS OF AN UNENCRYPTED UDP CONNECTION

Number of CPU cores	MTU [B]			
	576	1300	1500	9000
1	97 Mbps	84 Mbps	87 Mbps	86 Mbps
2	104 Mbps	85 Mbps	87 Mbps	86 Mbps
4	107 Mbps	86 Mbps	85 Mbps	88 Mbps
6	106 Mbps	96 Mbps	87 Mbps	86 Mbps
8	107 Mbps	106 Mbps	85 Mbps	85 Mbps
10	98 Mbps	104 Mbps	85 Mbps	88 Mbps

A. Unresolved issues

In the final firmware version, encrypted traffic was already passing through and being correctly decrypted on the other side; however, there were problems with packets shorter than 64 B (for example SYN packets in TCP). Since according to the standard the minimum Ethernet frame length is 64 B, shorter packets were padded to this length with filler bits. The encryptor, however, cannot distinguish this and encrypts the filler bits as well. These bits are then discarded along the way, because the encrypted packet is larger and therefore these bits are no longer needed to maintain the minimum size of 64 B. Because of this, during decryption the packet is dropped, as the GCM tag cannot be verified.

Another issue is packet loss, where losing a packet used for key establishment causes the establishment of new keys to be interrupted, so the encryptors continue operating only with the keys that have already been established.

Furthermore, parts were identified that can be worked around by appropriately configuring the other devices, but modifying them in the firmware would simplify future testing. This concerns passing ARP messages without encryption, or even directly responding to these messages, which had to be worked around by setting static ARP tables on the end devices. Another such point is recalculating the TCP/UDP checksum after decryption, which is necessary when using the encryptor behind a NAT. In our testing this was worked around by configuring the NATs on both sides so that the source and destination IP addresses remained the same on both sides.

B. Conclusions

The testing provided new insights into the network requirements of the cipher implementations. All implemented solutions were created solely for verification within this testing, and it will be necessary to verify their correctness and optimize them for higher transmission speeds as well.

VII. Stress tests at the FEKT laboratory

In connection with testing between the Czech Republic and Estonia, stress tests were also carried out at the FEKT VUT laboratory. The aim of these tests was to verify the stability of the encryptors. The test results were then compared with the values for operation at greater distances between end stations. Tables V and VI present the results of tests of encrypted operation on a 10 G line.

TABLE V
10 Gbps connection without routing

TCP			
Number of CPU cores	MTU [B]		
	576	1300	1500
1	8,0 Gbps	9,1 Gbps	9,2 Gbps
10	8,1 Gbps	9,1 Gbps	9,2 Gbps
UDP			
Number of CPU cores	MTU [B]		
	576	1300	1500
1	1,2 Gbps	2,9 Gbps	3,3 Gbps
10	3,5 Gbps	7,9 Gbps	8,9 Gbps

TABLE VI
Connection with routing (limited to 1 Gbps)

TCP			
Number of CPU cores	MTU [B]		
	576	1300	1500
1	806 Mbps	902 Mbps	912 Mbps
10	806 Mbps	902 Mbps	912 Mbps
UDP			
Number of CPU cores	MTU [B]		
	576	1300	1500
1	840 Mbps	919 Mbps	926 Mbps
10	840 Mbps	919 Mbps	926 Mbps

The table of traffic values without routing shows that throughput limitations increase with smaller packet sizes. The limitation is also higher for UDP compared to TCP. These phenomena are caused by higher packet processing frequency, which puts more load on the CPU. The lower throughput of UDP traffic is caused by the iperf3 tool, where the process of generating UDP traffic itself is more computationally demanding than TCP traffic. Outside of these limitations, speeds approached the maximum line speed of 10 Gbps.

The results of traffic through routers were limited to a maximum of 1 Gbps.

VIII. Conclusion

A series of tests was carried out during the testing phase, yielding valuable insights for further development. From the results, it can be observed that the highest achieved speed was around 303 Mbps. This speed clearly does not indicate utilization of the cipher device's full potential. On the other hand, when compared with the reference values, we can see that the cipher device does not introduce any noticeable reduction in speed associated with the act of encryption itself. It is also worth mentioning the stability of the connection latency, which averaged 43 ms. The cipher device also passed a two-hour stress test. It should be noted, however, that this was only achieved on the third attempt.

Even more beneficial than the test results themselves were the discoveries and insights gained through the analysis and resolution of individual issues encountered during commissioning. These relate not only to the cipher device itself, but also to the configuration of network components and a deeper understanding of computer networking. Firmware development for the cipher device was also very intensive during this period and responded directly to the needs of testing. Identifying root causes and fixing them consumed most of the time allocated for testing; nevertheless, it was possible to acquire very valuable knowledge that will help the cipher device achieve higher reliability.

Acknowledgment

Funded by the European Union under Grant Agreement No. 101087529. Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

APPENDIX A

Figures 3 – 7 show actual devices during testing.

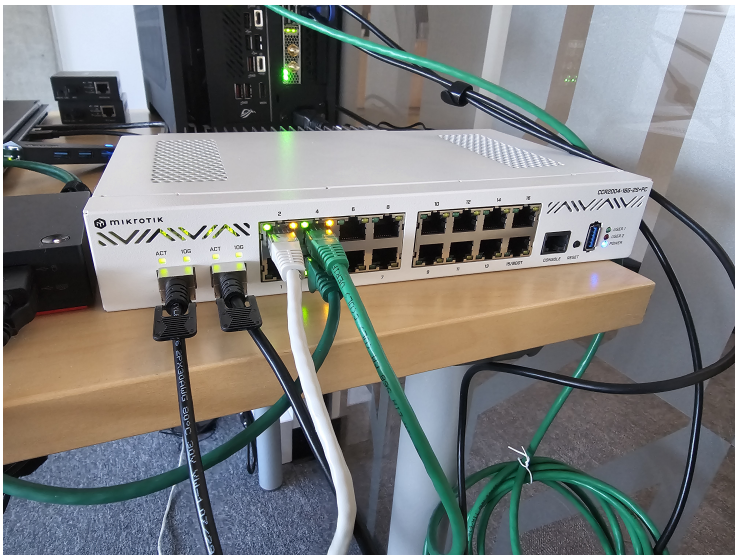

Fig. 3. Used components and final assembly



Fig. 4. Used components and final assembly



Fig. 5. Used components and final assembly



Fig. 6. Used components and final assembly



Fig. 7. Used components and final assembly

REFERENCES

[1] Peter Cíbik, Sara Ricci, Patrik Dobiáš, Jan Hajný, Lukáš Malina, and Jan Havlín. Quantum-resistant hardware-accelerated iot traffic encryptor, 2025.

[2] Peter Cibik, Patrik Dobias, Sara Ricci, Jan Hajny, Lukas Malina, Petr Jedlicka, and David Smekal. Pushing aes-256-gcm to its limits: Design, implementation and real fpga tests, 2024.

[3] A. Hodjat and I. Verbauwhede. Area-throughput trade-offs for fully pipelined 30 to 70 gbits/s aes processors, 2006.

[4] Maran van Heesch, Niels van Adrichem, Thomas Attema, and Thijs Veugen. Towards quantum-safe vpns and internet, 2019.

[5] Lukas Malina, Petr Dzurenda, Sara Ricci, Jan Hajny, Gautam Srivastava, Raimundas Matulevičius, Abasi-Amefon O Affia, Maryline Laurent, Nazatul Haque Sultan, and Qiang Tang. Post-quantum era privacy protection for intelligent infrastructures, 2021.

[6] Certicom Research. Sec 1: Elliptic curve cryptography. 2009.

[7] Federal Information Processing Standards Publication (FIPS) NIST FIPS 203. National institute of standards and technology module-lattice-based keyencapsulation mechanism standard. 2024.

[8] Rui-Qi Gao, Yuan-Mei Xie, Jie Gu, Wen-Bo Liu, Chen-Xun Weng, Bing-Hong Li, Hua-Lei Yin, and Zeng-Bing Chen. Simple security proof of coherent-one-way quantum key distribution, Jun 2022.

[9] 2023 Federal Information Processing Standards Publication (FIPS) NIST FIPS 197-upd1, updated May 9. National institute of standards and technology (2001) advanced encryption standard (aes). 2003.

Post-Quantum Migration of the EST Protocol: Evaluation on Constrained Microcontrollers

1st Jiří Hrachovina
Department of Telecommunication
Brno University of Technology
Brno, Czech Republic
240623@vutbr.cz

2nd Sara Ricci
Department of Telecommunication
Brno University of Technology
Brno, Czech Republic
0000-0003-0842-4951

***Abstract*—Quantum computers threaten widely used public key cryptographic algorithms such as RSA, ECDSA, and elliptic-curve Diffie–Hellman, which are currently employed in many public key infrastructure protocols. One such protocol is EST, used to automate certificate provisioning for embedded and IoT devices. Migrating EST to post-quantum cryptography is therefore necessary to ensure long-term security. In this work, we evaluate the feasibility of such a migration on constrained embedded systems. Several post-quantum key encapsulation and signature schemes were implemented and benchmarked on a Cortex-M4 microcontroller platform, measuring execution time, memory consumption, stack usage, and communication overhead. The results indicate that ML-KEM and ML-DSA provide the most practical performance among the evaluated schemes on the tested platform. Based on these findings, we propose a post-quantum migration strategy for EST in which ML-KEM is used for TLS key establishment and ML-DSA for certificate-related signatures, with parameter sets selected according to device capabilities.**

***Index Terms*—Post-Quantum Cryptography, Enrollment over Secure Transport (EST), Public Key Infrastructure (PKI), Constrained Devices, Embedded Systems.**

I. Introduction

Public key cryptography forms the basis of authentication and key-establishment mechanisms used in network security protocols and Public Key Infrastructure (PKI). Widely deployed schemes such as RSA and elliptic-curve cryptography rely on mathematical problems that could be efficiently solved by a sufficiently powerful quantum computer using Shor's algorithm, which would render current public-key systems insecure [1]. One such protocol is Enrollment over Secure Transport (EST) [2], which enables devices to automatically obtain digital certificates from a certification authority. EST operates on Transport Layer Security (TLS), which typically relies on classical public-key algorithms such as Rivest–Shamir–Adleman (RSA) and Elliptic Curve Digital Signature Algorithm (ECDSA) for authentication and Elliptic Curve Diffie-Hellman (ECDH) for key establishment. Since these algorithms are vulnerable to quantum attacks, EST in its current form cannot be considered quantum-resistant, motivating the integration of post-quantum cryptographic algorithms into certificate enrollment protocols.

A. Contribution

In this paper, we propose a post-quantum migration strategy for the EST protocol. To evaluate the feasibility of such a migration on constrained devices, we benchmark post-quantum standards from the National Institute of Standards and Technology (NIST), specifically the Module-Lattice-Based Key-Encapsulation Mechanism (ML-KEM) [3], the Module-Lattice-Based Digital Signature Algorithm (ML-DSA) [4], the Falcon Digital Signature Algorithm (FN-DSA) [5], the Stateless Hash-Based Digital Signature Algorithm (SLH-DSA) [6], and the Hamming Quasi-Cyclic (HQC) [7], all evaluated on a constrained microcontroller platform, using optimized implementations where available. For comparison, classical cryptographic schemes used in EST, namely X25519 and ECDSA, are also evaluated. The analysis compares computational performance, memory consumption, and communication overhead.

II. EST protocol

EST is a certificate enrollment protocol used in PKI to automate certificate provisioning for clients. The protocol builds on standard PKI components such as X.509 certificates, Public Key Cryptography Standards #10 (PKCS#10) Certificate Signing Requests (CSRs) and Cryptographic Message Syntax (CMS) message structures, integrating them into a modern architecture based on TLS. As illustrated in Fig. 1, an EST client establishes a TLS connection to the server, retrieves Certification Authority (CA) certificates, generates a key pair, and submits a CSR for certificate issuance.

The EST workflow relies on both asymmetric and symmetric cryptographic primitives. Specifically, TLS authentication and key exchange rely on RSA, ECDSA, and ECDH, while CSRs and certificates are signed using RSA or ECDSA. Moreover, AES and hash functions from the SHA family are used within TLS and CMS to provide confidentiality and integrity. Asymmetric algorithms such as RSA, ECDSA, and ECDH rely on mathematical problems that can be efficiently solved by Shor's algorithm on a sufficiently powerful quantum computer, and therefore require replacement by post-quantum alternatives. In contrast, symmetric primitives such as AES and SHA are less affected, as Grover's algorithm provides only a quadratic speed-up, effectively reducing their security strength

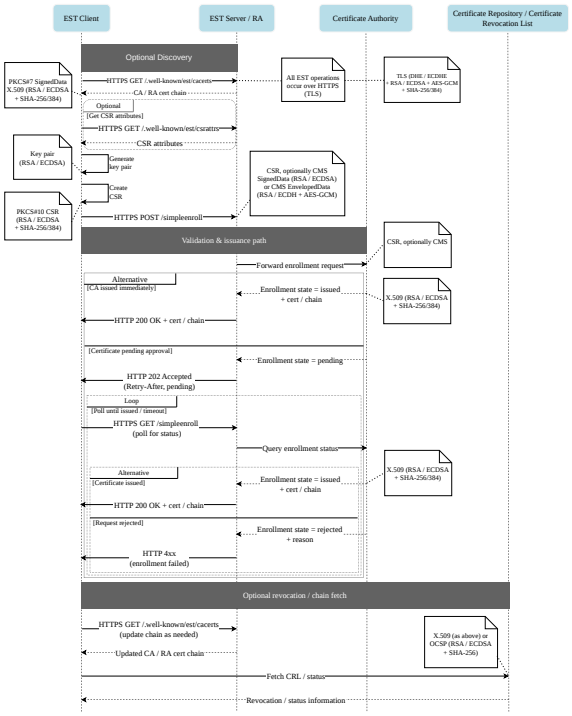

Fig. 1. Typical workflow of the EST protocol and associated cryptographic operations.

by roughly a factor of two [1]. Consequently, they can remain in use with appropriately increased parameters.

III. Experimental Results

All measurements were obtained using the same benchmarking framework to ensure consistent comparison across implementations. Randomness required by the evaluated algorithms was generated using the hardware random number generator of the STM32WB55RG platform. The main characteristics of the device are summarized in Table I.

Optimized implementations were used whenever available. The implementations of ML-KEM, ML-DSA, and FN-DSA were taken from the PQM4 project, which provides optimized implementations for the ARM Cortex-M4 platform [8]. For HQC and SLH-DSA, reference implementations from PQClean were used, as optimized Cortex-M4 implementations are not currently available. The classical algorithms X25519 and ECDSA (secp256r1) were implemented using optimized Cortex-M4 implementations developed by Emil Lenngren. All parameter sets were evaluated for the post-quantum schemes, while the elliptic-curve primitives were tested using their single available configuration.

Execution time was measured using the Data Watchpoint and Trace (DWT) cycle counter of the Cortex-M4 processor. Each operation was executed ten times and the reported metrics correspond to the average values obtained from these runs.

TABLE I
Device Specification

Parameter	Value
Microcontroller	STM32WB55RG
CPU	ARM Cortex-M4
Clock frequency	64 MHz
Flash memory	1 MB
SRAM	256 kB

The evaluation considers performance and resource metrics relevant for constrained devices. For key establishment mechanisms, key generation and shared-secret derivation times were measured, corresponding to encapsulation and decapsulation operations for KEM schemes. For digital signatures, signing and verification times were evaluated. Resource requirements were assessed by measuring flash memory usage, static RAM consumption, and peak stack usage. The reported RAM usage corresponds to memory allocated by the program at initialization, while peak stack usage captures additional memory used dynamically during execution. As these represent different aspects of memory consumption, peak stack usage may exceed the reported RAM values. Communication overhead was analyzed by comparing the sizes of public keys, ciphertexts, and signatures produced by the evaluated algorithms.

A. KEM results

The performance of the evaluated key encapsulation mechanisms is summarized in Tables II–IV, which report execution time, static memory usage, and peak stack consumption.

Table II shows that the ML-KEM family provides the best computational performance among the evaluated post-quantum schemes, with practical execution times on the evaluated microcontroller platform. In contrast, HQC implementations exhibit substantially higher computational cost, and HQC-256 could not be evaluated due to memory limitations on the target platform.

Table III summarizes the static memory requirements. ML-KEM requires moderate flash and RAM, while HQC demands significantly larger memory allocations.

Peak stack usage is shown in Table IV. ML-KEM requires only a few kilobytes of stack space, whereas HQC requires substantially larger allocations, exceeding feasible limits for some parameter sets. These results highlight the challenges of deploying code-based KEM schemes on constrained microcontrollers.

TABLE II
KEM Performance

Scheme	KeyGen (cycles)	Enc (cycles)	Dec (cycles)
X25519	599,461	598,701	-
ML-KEM-512	914,147	849,530	890,586
ML-KEM-768	1,409,180	1,437,206	1,491,470
ML-KEM-1024	2,253,315	2,271,792	2,338,610
HQC-128	50,740,583	102,320,691	153,999,265
HQC-192	154,678,880	310,854,788	466,581,861
HQC-256	283,236,485	568,505,738	-

TABLE III
KEM MEMORY USAGE

Scheme	Flash (B)	RAM (B)
X25519	18,016	1,092
ML-KEM-512	34,524	4,164
ML-KEM-768	34,492	5,636
ML-KEM-1024	34,860	7,268
HQC-128	39,604	10,024
HQC-192	42,068	19,120
HQC-256	47,220	30,020

Overall, ML-KEM provides a more practical balance between computational cost and memory requirements than HQC on the evaluated Cortex-M4 platform.

B. Signature Results

The results for the evaluated digital signature schemes are presented in Tables V–VII, which summarize computational performance, memory requirements, and stack usage.

Table V shows that the classical ECDSA implementation remains significantly more efficient than all evaluated post-quantum schemes. Among the post-quantum algorithms, ML-DSA provides the lowest computational cost, while higher security parameter sets incur increased overhead. FN-DSA exhibits a different performance profile, with relatively fast verification but significantly more expensive key generation and signing. In contrast, the stateless hash-based SLH-DSA schemes show very high signing costs, limiting their practicality on constrained devices.

Table VI summarizes the memory requirements. ML-DSA requires moderate flash and RAM, while FN-DSA requires substantially larger flash memory. SLH-DSA uses relatively small flash but significantly more RAM.

Peak stack usage is shown in Table VII. ECDSA has the lowest stack requirements, ML-DSA requires moderate stack space, and FN-DSA incurs significantly higher stack usage during signing.

Overall, ML-DSA achieves the most balanced trade-off between computational performance and memory usage among the evaluated post-quantum signature schemes on the tested Cortex-M4 platform.

IV. IMPACT ON EST DEPLOYMENT

The benchmarks presented provide information on the feasibility of deploying post-quantum cryptography within the EST protocol on constrained devices. The results show that while several post-quantum schemes can be executed on embedded microcontrollers, their performance and memory requirements differ significantly.

Among the evaluated KEM schemes, ML-KEM provides the most favorable balance between performance and memory consumption. In contrast, HQC exhibits substantially higher computational cost and memory usage, which may limit its deployment on constrained platforms. For digital signatures, ML-DSA offers the most practical trade-off between execution time and memory usage among the evaluated schemes. These results suggest that ML-KEM and ML-DSA are suitable candidates for post-quantum migration of EST in constrained environments. Depending on available computational resources and security requirements, different parameter sets may be selected. Table VIII summarizes recommended configurations for several device classes.

The impact on EST latency is dominated by the increased cost of post-quantum digital signatures, particularly during CSR generation and certificate verification on the client. Based on the measured results, a simplified client-side enrollment flow is approximately 11-13$\times$ slower than in the classical setting, with the majority of the overhead stemming from signing and verification. However, as certificate enrollment is typically infrequent, this increase remains acceptable for most IoT deployments.

TABLE V
SIGNATURE PERFORMANCE

Scheme	KeyGen (cycles)	Sign (cycles)	Verify (cycles)
ECDSA	385,381	432,115	1,171,950
ML-DSA-44	3,915,028	18,930,640	4,663,146
ML-DSA-65	7,510,451	43,577,655	8,177,324
ML-DSA-87	12,765,054	44,849,740	14,211,881
FN-DSA-512	83,815,681	26,728,558	495,958
FN-DSA-1024	312,785,980	184,567,401	962,898
SLH-DSA-128f	22,663,177	530,371,512	31,710,494
SLH-DSA-192f	33,905,868	986,393,016	52,885,044
SLH-DSA-256f	91,728,511	2,072,558,492	57,819,436

V. PROPOSED EST PQ MIGRATION

Migrating EST to a post-quantum setting requires replacing the classical asymmetric primitives used in TLS key establishment, certificate signing, and certificate request generation with post-quantum alternatives. Fig. 2 illustrates the proposed migration strategy. We consider a deployment scenario with

TABLE IV
KEM STACK USAGE

Scheme	KeyGen (B)	Enc (B)	Dec (B)
X25519	656	648	656
ML-KEM-512	2,736	2,784	2,768
ML-KEM-768	3,256	3,296	3,280
ML-KEM-1024	3,768	3,808	3,792
HQC-128	33,720	51,512	55,968
HQC-192	65,744	101,712	110,376
HQC-256	103,832	131,584	-

TABLE VI
SIGNATURE MEMORY USAGE

Scheme	Flash (B)	RAM (B)
ECDSA	32,288	1,092
ML-DSA-44	42,176	7,228
ML-DSA-65	41,464	10,232
ML-DSA-87	41,760	13,052
FN-DSA-512	113,924	3,872
FN-DSA-1024	113,924	6,404
SLH-DSA-128f	27,828	18,120
SLH-DSA-192f	60,124	36,744
SLH-DSA-256f	60,424	50,984

TABLE VII
SIGNATURE STACK USAGE

Scheme	KeyGen (B)	Sign (B)	Verify (B)
ECDSA	1,808	1,944	2,208
ML-DSA-44	4,840	5,480	3,104
ML-DSA-65	4,840	7,008	3,104
ML-DSA-87	4,840	8,544	3,104
FN-DSA-512	14,536	42,168	3,192
FN-DSA-1024	27,848	82,208	5,240
SLH-DSA-128f	3,032	2,472	2,784
SLH-DSA-192f	5,592	4,360	4,432
SLH-DSA-256f	8,592	6,368	6,056

heterogeneous capabilities, where the EST client is a constrained microcontroller-based IoT device, the EST server runs on a gateway or enterprise platform, and the CA operates on cloud or enterprise infrastructure.

In the proposed design, TLS key establishment between the EST client and server relies on ML-KEM. Since key establishment is performed jointly by both endpoints, both parties must support the same parameters. Consequently, parameter selection is determined by the most constrained participant, typically the client. For microcontroller-based devices, ML-KEM-512 represents a practical choice.

Digital signatures for certificate signing requests and certificate issuance are based on ML-DSA. Constrained clients generate CSRs using lightweight parameter sets such as ML-DSA-44. To ensure interoperability, certificates issued by the CA must use parameter sets supported by the client, and therefore ML-DSA-44 is used in constrained deployments. Higher-security parameter sets (ML-DSA-65 or ML-DSA-87) and higher-security ML-KEM variants (ML-KEM-768 or ML-KEM-1024) may be employed for communication between the EST server and the CA, where client-side constraints do not apply.

Symmetric primitives used in TLS and CMS, such as AES-GCM and SHA-2, are less affected by quantum attacks and do not require replacement. However, due to Grover's algorithm, their effective security is reduced, requiring larger parameters such as AES-256 and SHA-384 or SHA-512.

VI. CONCLUSION

This paper evaluates the feasibility of migrating the EST protocol to post-quantum cryptography in constrained embedded environments. Several post-quantum key encapsulation and signature schemes were implemented and benchmarked on a Cortex-M4 platform, measuring execution time, memory consumption, stack usage, and communication overhead.

The results indicate that ML-KEM provides a more practical balance between computational cost and memory requirements

TABLE VIII
RECOMMENDED PQC ALGORITHMS FOR EST DEPLOYMENT

Device Class	KEM	Signature
Microcontroller	ML-KEM-512	ML-DSA-44
Medium IoT Device	ML-KEM-768	ML-DSA-65
Cloud / Server	ML-KEM-1024	ML-DSA-87

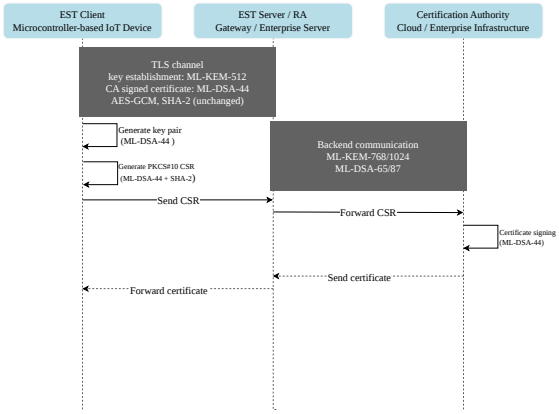

Fig. 2. Proposed post-quantum migration of the EST certificate enrollment workflow.

than HQC, while ML-DSA achieves the most favorable trade-off among the evaluated signature schemes. These findings show that post-quantum cryptography can be deployed within EST even on constrained microcontroller-based devices. Accordingly, ML-KEM and ML-DSA are identified as suitable candidates, with lightweight parameter sets (ML-KEM-512 and ML-DSA-44) used by constrained clients and stronger parameters by more capable infrastructure components.

REFERENCES

[1] BERNSTEIN, Daniel J. Introduction to post-quantum cryptography. In: BUCHMANN, Johannes and Erik DAHMEN, eds. Post-Quantum Cryptography. Berlin: Springer, 2009, pp. 1–14. ISBN 978-3-540-88701-0.

[2] Pritikin, M., Yee, P., and Harkins, D., Enrollment over Secure Transport (EST), RFC 7030, IETF, 2013.

[3] NIST. FIPS 203: Module-Lattice-Based Key-Encapsulation Mechanism Standard (ML-KEM). Gaithersburg, MD: National Institute of Standards and Technology, 2024. DOI: 10.6028/NIST.FIPS.203.

[4] NIST. FIPS 204: Module-Lattice-Based Digital Signature Standard (ML-DSA). Gaithersburg, MD: National Institute of Standards and Technology, 2024. DOI: 10.6028/NIST.FIPS.204.

[5] Fouque, P.-A., Kirchner, P., Tibouchi, M., and Prest, T., Falcon: Fast-Fourier Lattice-Based Compact Signatures over NTRU, NIST PQC submission, 2018.

[6] NIST. FIPS 205: Stateless Hash-Based Digital Signature Standard (SLH-DSA). Gaithersburg, MD: National Institute of Standards and Technology, 2024.

[7] C. A. Melchor et al., HQC: Hamming Quasi-Cyclic Key Encapsulation Mechanism, NIST Post-Quantum Cryptography Standardization Project, 2018.

[8] Kannwischer, M. J., Rijneveld, J., Schwabe, P., and Stoffelen, K., pqm4: Testing and Benchmarking NIST PQC on ARM Cortex-M4, AFRICACRYPT 2020.

[9] Kannwischer, M. J., et al., PQClean: Clean, Portable, and Tested Implementations of Post-Quantum Cryptography, 2022.

A PDMS stopper-based cell migration assay as an alternative to the Scratch assay

1st Inna Zumberg
Department of Biomedical Engineering
Brno University of Technology
Brno, Czech Republic
zumberg@vutbr.cz

2nd Johannes Jun Wei Low
Center for Molecular and Cellular Bioengineering
Dresden University of Technology
Dresden, Germany
johannes.low_jun_wei@tu-dresden.de

3rd Anna Taubenberger
Center for Molecular and Cellular Bioengineering
Dresden University of Technology
Dresden, Germany
anna.taubenberger@tu-dresden.de

4th Vratislav Cmiel
Department of Biomedical Engineering
Brno University of Technology
Brno, Czech Republic
cmiel@vut.cz

***Abstract*—Wound healing assays are fundamental tools for studying collective cell motility, tissue repair and cancer cell invasiveness. The conventional Scratch assay is simple and widely used, but it produces variable wound geometry and mechanically damages cells and the underlying matrix at the wound edge, confounding measurements of migration and local viability. Exclusion-based methods avoid mechanical injury and yield more reproducible gaps, but commercial inserts can be costly or inflexible with respect to different culture formats. Here, we present a low-cost method using biopsy-punched PDMS stoppers (5×1 mm) to generate well-defined, reproducible exclusion zones and compare its performance against p200 pipette tip scratching. We evaluate stopper geometry and between-replicate variability, quantify migration under standard and reduced FBS conditions in NIH-3T3 and SaOS-2 cells, and assess local cytotoxicity at the wound edge using Live/Dead staining. We aim to determine whether PDMS stoppers provide a practical, low-perturbation alternative to mechanical wounding for quantitative migration studies.**

Keywords—PDMS, Stopper, Cell migration, SaOS-2, NIH-3T3, Scratch assay

I. Introduction

Cell migration is a fundamental process in both physiological processes, such as tissue formation and immune responses [1], and pathological conditions, including cancer metastasis [2]. To study these phenomena, the wound healing assay serves as a widely used tool [3]. The core principle of this assay involves creating an artificial, cell-free gap within a confluent cell monolayer and monitoring the subsequent migration of cells from the wound edges until the gap is closed and new cell-cell contacts are established. This method allows researchers to quantify migration rates and observe the morphological characteristics of migrating cells, such as leading-edge polarization.

To initiate the assay, a cell-free area must be generated using techniques that can be broadly categorized into two main approaches: cell wounding and cell exclusion. The first approach involves mechanical disruption of the monolayer (e.g., scraping with a pipette tip, a spatula, a specialized scratching tool, or stamps [4]). Modern advancements have introduced diverse wounding principles, including chemical methods using lytic agents, electrical methods that electropermeabilize cells, and optical wounding with laser beam. This "scratch" results in the immediate removal of cells and the release of cellular debris and damage to the substrate, which can influence the behavior of neighboring cells. The second approach relies on exclusion (barrier) methods, in which a physical barrier or insert (e.g., ibidi culture inserts [5] or custom stoppers [6,7]) is placed on the culture surface prior to cell seeding. Cells grow around the barrier until a confluent monolayer is formed. Once the barrier is removed, a well-defined cell-free zone remains. This method avoids mechanical damage to the cells at the wound edge and provides a more consistent gap width.

Both approaches have advantages and limitations. Mechanical scratching is simple and inexpensive but strongly depends on manual execution, often resulting in variable wound widths and irregular edges. Moreover, the mechanical damage to cells and the extracellular matrix can trigger signaling events that influence the migration process. Barrier-based methods, including commercial inserts and custom devices, produce highly defined and reproducible cell-free areas but may require specialized materials and incur higher costs. Therefore, there remains a need for a simple, low-cost, and reproducible method to generate defined cell-free regions while minimizing mechanical damage to surrounding cells. In this work, we present laboratory-fabricated polydimethylsiloxane (PDMS) stoppers as a practical and economical solution for creating cell-free areas in *in vitro* cell migration assays.

II. Materials and Methods

A. Fabrication of PDMS stoppers

PDMS (SYLGARD™ 184 Silicone Elastomer Kit) was prepared following the manufacturer's instructions. Briefly, the Elastomer Base and Curing Agent were mixed at a 10:1 (w/w) ratio and mixed thoroughly. The mixture was degassed under vacuum until visible bubbles were removed, then poured into 70 mm polystyrene culture dishes to produce a 5 mm thick layer (≈20,2 g). The PDMS was left to cure at room temperature for 7 days. After curing, a 2×2 cm block of PDMS was cut from the layer with a scalpel, peeled from the dish, and transferred to a clean 100 mm glass Petri dish. Care was taken to keep the

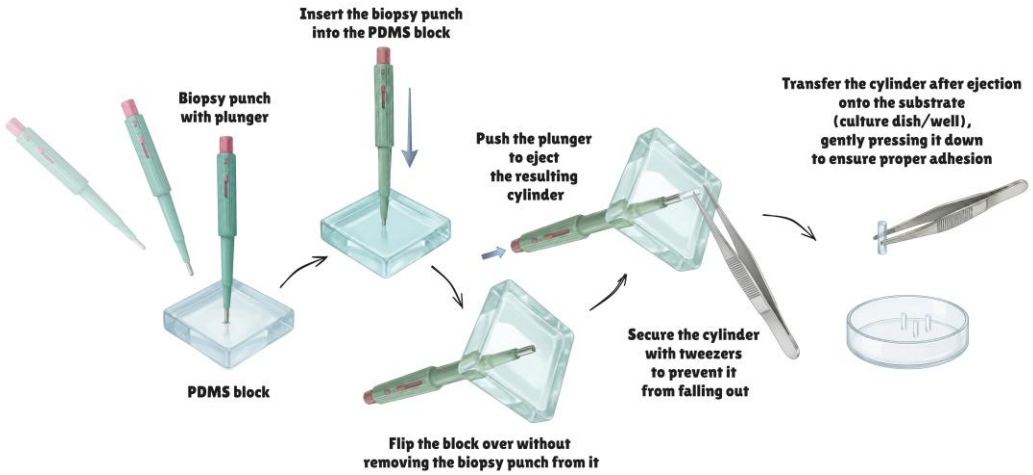

Fig. 1. Schematic representation of PDMS stopper fabrication. *Note: Image generated using Illustrae, 2026, https://illustrae.co/*

bottom (dish-contacting) surface free of dust and particles, as contamination reduces surface tack and can impair stopper adhesion. When necessary, the bottom surface was gently wiped with a tissue soaked in 70% ethanol to remove particles.

A biopsy punch (1 mm diameter biopsy punch with plunger, BPP-10F, Kai Medical, Japan) was placed vertically on the PDMS block and pressed down with moderate force until the punch reached the dish surface. The PDMS block was then lifted off the dish while leaving the punch in place; the plunger was slowly pushed to eject the PDMS cylinder, which was grasped with fine tweezers to prevent loss. Each stopper was transferred to a new dish preserving its orientation relative to the original block (the side that faced the culture dish during curing is the tacky surface and should face the culture surface when stoppers are applied). A schematic overview of the fabrication procedure is shown in Figure 1. All tools and surfaces were handled with clean gloves and forceps to minimize contamination. Finished stoppers can be sterilized (e.g., by ethanol rinse and UV exposure) prior to use, as required by the specific application.

B. Cell culture

Human osteosarcoma SaOS-2 cells were obtained from CLS (Cell Lines Service GmbH) and maintained in high glucose Dulbecco's Modified Eagle Medium (Sigma-Aldrich) and Ham's F-12 medium with L-Glutamine (Serana) (DMEM:F12, 1:1) supplemented with 10% fetal bovine serum (FBS, Biosera) and 1% of penicillin-streptomycin (Sigma-Aldrich) at 37 ºC and 5% CO_2. The medium was changed every 2-3 days, and cells were subcultured after reaching 90% confluence.

NIH-3T3 mouse embryonic fibroblasts were obtained from CLS and maintained in high glucose DMEM supplemented with 5% FBS and 1% of penicillin-streptomycin at 37 ºC and 5% CO_2. The medium was changed every 2 days, and cells were subcultured after reaching 90% confluence.

C. Wound healing assay using PDMS stoppers

Five PDMS stoppers were placed with fine tweezers in the central region of 8.5 cm^2 tissue culture dishes (Deltalab) with ≈2 mm spacing between adjacent stoppers. To prevent air bubbles from forming on the stopper surfaces, the dishes with stoppers were incubated in phosphate-buffered saline (PBS) for 10 minutes at 37 ºC and 5% CO_2 before cell seeding. Two dishes were prepared for each cell line, and $5 \cdot 10^5$ cells were seeded into each dish. Then, dishes were cultured under standard conditions for 3-4 days until a confluent cell monolayer was formed.

After confluence was reached, the PDMS stoppers were removed by gentle vertical peeling with tweezers. Culture medium was aspirated and the dishes were rinsed with PBS to remove detached cells and debris. Each dish was then filled with 2 ml of the appropriate culture medium. For NIH-3T3 cells, the control dish was maintained in medium containing 5% FBS, while the test dish was supplied with medium containing 0.5% FBS to induce serum starvation and reduce the migration rate. The same procedure was applied to SaOS-2 cells, except that the control medium contained 10% FBS and the test medium contained 1% FBS.

The microscopic images (1920×1440 pixels, pixel size 1.297×1.297 µm) of cell-free area were acquired using Leica DMi8 phase contrast microscope equipped with a CCD camera (LEICA DFC7000 T) immediately after wounding (0 hours), then in 24 and 48 hours using a 5× objective magnification. The image acquisition parameters were the same in all groups.

D. Live/Dead staining

Two culture dishes with PDMS stoppers were prepared as described in Section 2C, and $5 \cdot 10^5$ cells (SaOS-2 or NIH-3T3) were seeded into each dish. Cells were cultured under standard conditions until a confluent monolayer was formed. After reaching confluence, PDMS stoppers were gently removed from each dish using sterile tweezers. Additionally, two linear "scratches" were made in each dish using a sterile P200 pipette tip, positioned near the dish edges. Culture medium was then aspirated and the dishes were rinsed once with PBS to remove detached cells and debris.

For cell viability assessment, a fresh staining solution consisting of 1 µM calcein AM and 2 µg/ml propidium iodide (PI) (both from Sigma-Aldrich) in PBS was prepared. Then, 2 ml of the staining solution were added to each dish and cells were incubated for 15 min at 37 °C in the dark. After incubation,

the staining solution was removed, the dishes were rinsed with PBS, and fresh PBS was added prior to fluorescence imaging. The microscopic images (1024×1024 pixels, pixel size 1.136×1.136 µm) of cell-free area were acquired using Leica TCS SP8 X confocal laser scanning microscope equipped with a 10× objective. Excitation and emission wavelengths were set to 490 nm and 500–530 nm for live cells visualization, while 535 nm and 550–620 nm were used for dead cells imaging.

E. Statistical analyses

Statistical analyses were performed using two-sample t-test in MATLAB software (R2025b). Results are based on three independent experiments. Nonparametric Mann–Whitney U test was applied if data failed the Shapiro–Wilk test of normality. Statistical significance was set at $p<0.05$.

III. Results and Discussion

A. Geometry of PDMS stoppers and resulting cell-free area

Figure 2A–E summarizes the PDMS stopper-based wound healing workflow: PDMS stoppers were placed on 8.5 cm^2 tissue culture dish (Fig. 2A), $5 \cdot 10^5$ SaOS-2 cells (pre-labelled with 5 µM CellTracker Green CMFDA (Thermo Fisher Scientific) for 30 min) were seeded (Fig. 2B), imaged at 4 h and 24 h post-seeding (Fig. 2C–D), and stoppers were removed with tweezers (Fig. 2E). Qualitatively, stoppers remained adherent to the substrate during incubation and produced well-defined, cell-free areas corresponding directly to the stopper contact footprint.

To quantify the stopper footprint and the resulting geometry of the cell-free area, 30 stoppers were prepared, placed on culture dish and imaged at 5× phase contrast. Images were analysed in ImageJ software [8]: the oval selection tool was used to determine width and height (major and minor axes of the fitted ellipse), and the same selection provided contact area and perimeter measurements. The resulting dimensions were highly reproducible: width = 979±10 µm, height = 982±11 µm (mean±SD), with mean contact area 0.755 mm^2 and perimeter 3.08 mm, confirming that the PDMS stoppers produce a consistent, near-circular exclusion footprint.

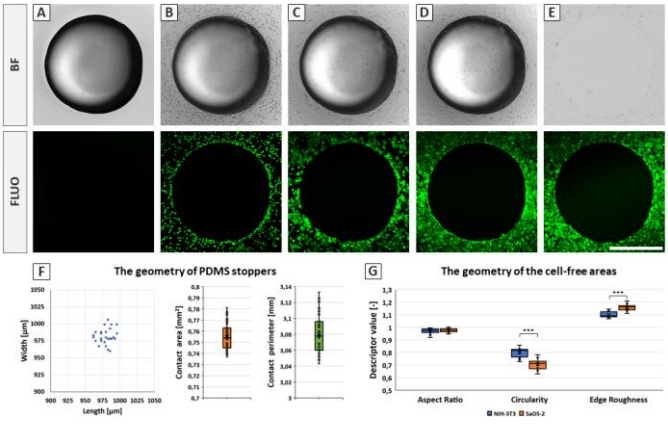

Fig. 2. PDMS stopper-based cell migration assay: (A) PDMS stopper placed on a culture dish. (B) Seeding of pre-labelled SaOS-2 cells. (C) Cell attachment and spreading after 4 h. (D) Confluent cell monolayer formed after 24 h. (E) Removal of stoppers with tweezers producing well-defined cell-free area corresponding to the stopper contact area. Scale bar 500 µm. (F-G) The geometry of PDMS stopper footprints and generated cell-free boundary. The significance levels are indicated as $p<0.05$ (*), $p<0.01$ (**), $p<0.001$ (***).

We next evaluated how the geometry of the cell-populated border surrounding each cell-free area varies with cell type. Three metrics were used (computed from segmented cell-free areas using Wound Healing size tool [9]):

$$AR\ (Aspect\ Ratio) = \frac{\min\ (width,height)}{\max\ (width,height)}, \quad (1)$$

$$C\ (Circularity) = \frac{4\pi A}{P^2}, \quad (2)$$

where A is segmented cell-free area in mm^2 and P is measured perimeter of cell-free area in mm; $C = 1$ for a perfect circle, decreases with elongation or irregularity.

$$R\ (Edge\ Roughness) = \frac{P}{P_{fe}}, \quad (3)$$

where P_{fe} is a perimeter of a fitted ellipse; $R = 1$ for a smooth circular edge, $R>1$ with increasing boundary irregularity.

Aspect ratio was similar for both cell lines (NIH-3T3: 0.970±0.022, SaOS-2: 0.975±0.017; $P>0.05$), consistent with AR reflecting the stopper geometry rather than cell-dependent modulation. In contrast, metrics characterizing the cell border differed significantly between cell types: circularity was 0.799±0.040 for NIH-3T3 versus 0.708±0.047 for SaOS-2 ($P<0.001$), and edge roughness was 1.100±0.026 for NIH-3T3 versus 1.160±0.029 for SaOS-2 ($P<0.001$). Thus, although the gross stopper footprint is reproducible, the fine morphology of the cell-free edge is cell-type dependent: SaOS-2 cells form less circular and rougher boundaries than NIH-3T3, suggesting a more irregular boundary after stopper removal. This difference may reflect distinct morphological and adhesion properties of the two cell types. In addition, we observed that SaOS-2 cells can partially grow along or over the stopper edges during culture, which may increase their susceptibility to mechanical disturbance upon stopper removal. Consequently, the removal process itself may contribute to local perturbation of edge cells and increased irregularity of the wound margin, representing a potential limitation of the method.

B. Serum-dependent migration of NIH-3T3 and SaOS-2 cells assessed using PDMS stoppers

PDMS stoppers were utilized to compare the migration under control and reduced-serum conditions for SaOS-2 and NIH-3T3 cells. Quantitatively (Fig. 3), lowering the FBS concentration resulted in a statistically significant decrease in cell-free area closure for both cell lines relative to their matched controls, with a markedly greater effect in SaOS-2 cells (–50.5% after 48 hours compared to –21.3% in NIH-3T3 cells). Gap closure in this assay may reflect a combination of migration and proliferation, although reduced serum conditions are expected to limit proliferative activity. The kinetics of cell-free area closure differed by cell type: SaOS-2 displayed an approximately linear, steady decrease in the exclusion area, whereas NIH-3T3 showed a faster, non-linear (accelerating) closure profile. These contrasting responses likely reflect differences in biological dependence on serum: serum provides both motogenic and

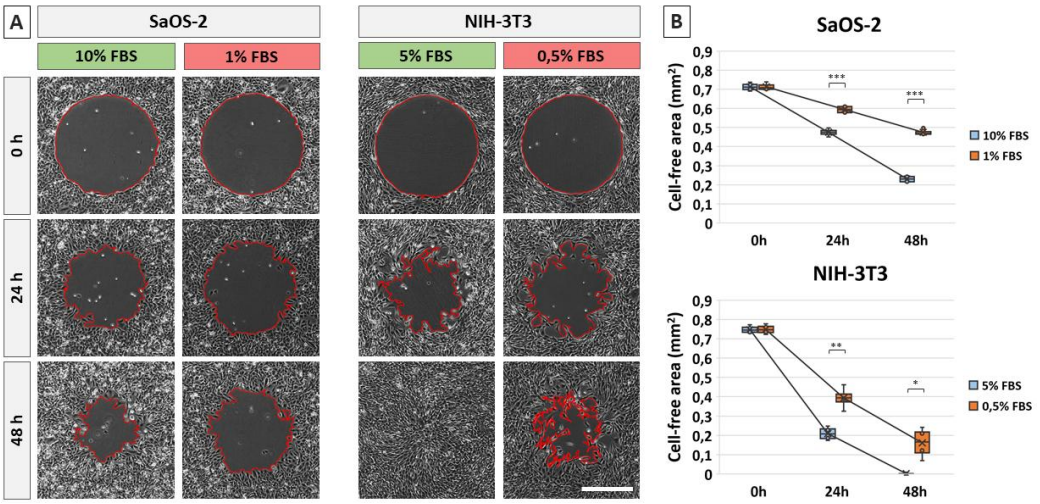

Fig. 3. Assessment of SaOS-2 and NIH-3T3 cell migration under reduced FBS concentration using PDMS stoppers: (A) Microscopic images acquired 0, 24 and 48 hours after PDMS stopper removal. The boundaries of the cell-free area were identified using Wound healing size tool in ImageJ. (B) Quantitative analysis of cell-free area. The significance levels are indicated as $p<0.05$ (*), $p<0.01$ (**), $p<0.001$ (***).

mitogenic cues, and SaOS-2 migration appears to be particularly serum-dependent. Cell morphology and edge behavior could also contribute, as SaOS-2 produced rougher, less circular wound margins (and tend to grow onto stopper edges and detach during removal), which may dissipate collective traction forces and slow coherent front advance; whereas NIH-3T3 cells, with smoother edges and fibroblast-like spreading, close gaps more rapidly.

C. Cell viability evaluation at the border of the cell-free area

The potential cytotoxicity introduced by cell wounding (Scratch assay) versus cell exclusion (PDMS stopper) approaches was assessed by quantifying cell viability in the immediate vicinity of the cell-free area. After performing either a conventional pipette tip scratch or PDMS stopper exclusion as described above, cells were stained with calcein AM and PI and imaged by confocal microscopy using identical acquisition settings for all samples. The resulting images were analyzed using a custom MATLAB algorithm (Fig. 4). First, the boundary of the cell-free area was determined by generating a binary mask using the Wound Healing Tool in ImageJ. This mask was then used to define a band-shaped region of interest (ROI) encompassing cells located within a ≈40 µm distance (35 pixels) from the wound edge. Only cells within this border region were included in the analysis, as they are most likely to be directly affected by mechanical interaction with the pipette tip or by physical contact with the PDMS stopper. For each image, fluorescence intensities were quantified to calculate normalized green fluorescence (NGF) and normalized red fluorescence (NRF) corresponding to calcein and PI signal per pixel, respectively. In addition, the normalized green/red ratio (NGF/NRF) was calculated as an indicator of cell viability. A higher green/red fluorescence ratio indicates a greater proportion of viable cells with intact plasma membranes, whereas elevated red fluorescence reflects increased cell death or membrane damage.

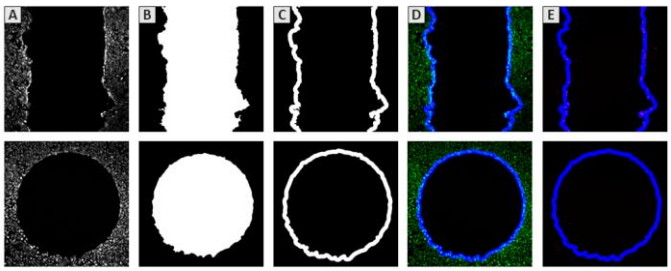

Fig. 4. Individual steps of the algorithm: (A) Original grayscale image (B) Binary mask created using Wound healing assay tool (C) Band-shaped ROI (D-E) Visualization of analyzed cells by overlaying a band-shaped ROI on the fluorescence images of (D) live and (E) dead cells.

The Live/Dead staining results (Fig. 5A) demonstrate that the conventional Scratch assay induces substantially greater local cytotoxicity than the PDMS stopper exclusion method. The normalized green/red ratio in the stopper group was significantly higher than in the scratch group, corresponding to a ≈2.6-fold increase for NIH-3T3 and a ≈2.0-fold increase for SaOS-2. These findings indicate a larger proportion of cells with intact membranes adjacent to PDMS-generated exclusion zones, consistent with the expectation that mechanical wounding both kills cells directly and provokes secondary damage in neighboring cells. The markedly larger fold-change observed for NIH-3T3 suggests that this cell type may experience greater collateral effects from pipette tip scratching, potentially related to stronger intercellular and cell–matrix mechanical coupling. Mechanical disruption of the monolayer may transmit forces through these connections to adjacent cells, which could contribute to secondary membrane damage and reduced viability near the wound edge. In contrast, the epithelial-like SaOS-2 cells appear less susceptible to this indirect mechanical stress, possibly due to differences in cell–cell junction organization or adhesion dynamics. In addition, scratching may disrupt the underlying extracellular matrix and release intracellular contents, potentially inducing local cytotoxic signaling. By contrast, PDMS stopper-based exclusion avoids direct mechanical removal of cells and instead prevents cell adhesion within a defined region, which may help preserve the integrity of surrounding cells and the extracellular matrix.

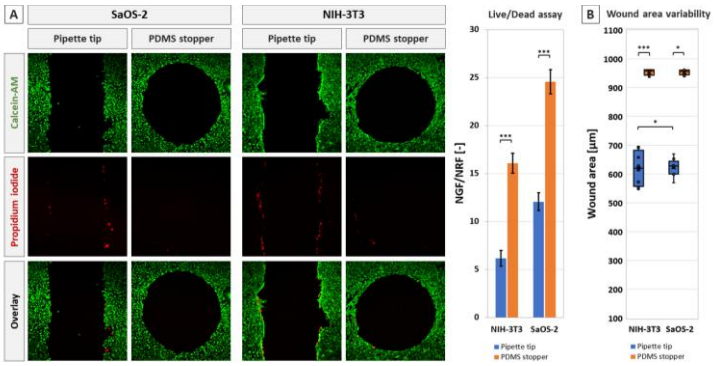

Fig. 5. (A) PDMS stoppers preserve neighboring-cell viability compared with pipette tip scratching (B) Variability of initial gap width generated by pipette tip scratching and PDMS stoppers. The significance levels are indicated as $p<0.05$ (*), $p<0.001$ (***).

The variability analysis of initial gap width further reinforces the practical advantages of the PDMS stopper approach (Fig. 5B). For Scratch assay, gap width was measured at five randomly selected positions per image, whereas for stopper exclusion, the diameter of the cell-free area was measured at five random orientations. For NIH-3T3, pipette tip scratching produced a mean width of 635.11±73.23 µm (mean±SD), whereas PDMS stoppers yielded 953.03±9.53 µm; for SaOS-2 the corresponding values were 630.29±38.48 µm for pipette and 952.49±8.69 µm for stoppers. Wound area variability was evaluated using Levene's test and expressed as the coefficient of variation (CV). Stopper-defined regions exhibited low and comparable variability between cell lines ($p = 0.8851$; CV≈1.0 for NIH-3T3 and 0.91 for SaOS-2), whereas pipette-generated scratches showed substantially higher variability (CV = 11.53 for NIH-3T3 and 4.82 for SaOS-2), with a significant difference between cell lines ($p = 0.019$). Variability was also significantly lower for stopper-based regions compared to scratches ($p<0.001$ for NIH-3T3 and $p=0.03$ for SaOS-2). Thus, mechanical scratching produced substantially greater between-replicate variability in initial wound size than stopper-based exclusion, particularly for NIH-3T3. This operator dependence of the Scratch assay is well documented [3,4,6,7] and arises because manual scraping removes cells and disturbs the substrate to varying degrees depending on tip angle, applied force, and local monolayer heterogeneity. The relatively high variability observed for NIH-3T3 cells may reflect their stronger mechanical coupling, whereby identical scraping actions produce uneven tearing and inconsistent gap widths across experiments. By contrast, the PDMS stoppers produced highly consistent contact footprints, and this reproducibility can be further improved by standardizing punch technique (angle, applied force) and by pre-screening stopper dimensions prior to seeding. From an experimental-design perspective, the lower variability in cell-free area eliminates a major source of confounding in migration assays, since inconsistent initial wound geometry alters closure kinetics and complicates comparisons across conditions and cell types.

CONCLUSION

Compared to conventional exclusion assays using commercial inserts or more complex fabrication methods (e.g., 3D printing), the PDMS stoppers provide a simple, cost-effective, and reproducible method for generating well-defined cell-free areas for *in vitro* migration assays. The stoppers can be easily fabricated using a biopsy punch from widely available PDMS, and can be reused after standard cleaning (e.g., ethanol washing). Live/Dead analysis confirmed markedly higher cell viability adjacent to stopper-generated gaps compared with conventional scratching. Taken together, these results demonstrate that the PDMS stopper assay offers clear advantages over the traditional Scratch assay in terms of preserving local cell viability and minimizing operator-dependent variability.

REFERENCES

[1] X. Fu, G. Liu, A. Halim, Y. Ju, Q. Luo, and G. Song, "Mesenchymal Stem Cell Migration and Tissue Repair," Cells, vol. 8, no. 8, 784, 2019.

[2] D. Wirtz, K. Konstantopoulos, and P. Searson, "The physics of cancer: the role of physical interactions and mechanical forces in metastasis," Nat Rev Cancer, vol. 11, pp. 512–522, 2011.

[3] J. E. N. Jonkman, J. A. Cathcart, F. Xu, M. E. Bartolini, J. E. Amon, K. M. Stevens, and P. Colarusso, "An introduction to the wound healing assay using live-cell microscopy," Cell Adhesion & Migration, vol. 8, no. 5, pp. 440–451, 2014.

[4] R. Riahi, Y. Yang, D. D. Zhang, and P. K. Wong, "Advances in Wound-Healing Assays for Probing Collective Cell Migration," Journal of Laboratory Automation, vol. 17, no. 1, pp. 59–65, 2012.

[5] S. Elgass, A. Cooper, and M. Chopra, "Lycopene treatment of prostate cancer cell lines inhibits adhesion and migration properties of the cells," Int J Med Sci., vol. 11, no. 9, pp. 948–954, 2014.

[6] A. S. Joshi, M. Madhusudanan, and I. Mijakovic, "3D printed inserts for reproducible high throughput screening of cell migration," Front. Cell Dev. Biol., vol. 11, 1256250, 2023.

[7] W. Gough, K. I. Hulkower, R. Lynch, P. Mcglynn, M. Uhlik, L. Yan, and J. A. Lee, "A Quantitative, Facile, and High-Throughput Image-Based Cell Migration Method Is a Robust Alternative to the Scratch Assay," Journal of Biomolecular Screening, vol. 16, no. 2, pp. 155–163, 2011.

[8] T. J. Collins, "ImageJ for Microscopy," BioTechniques, vol. 43. no. 1, pp. 25–30, 2007.

[9] A. Suarez-Arnedo et al., "An image J plugin for the high throughput image analysis of in vitro scratch wound healing assays", PLOS ONE, vol. 15, no. 7, 2020.

Electroactive Double-Layer Janus Biohybrid Hydrogel for Cardiac Tissue Engineering

1st Romana Ščotková
Department of Biomedical Engineering
Faculty of Electrical Engineering and Communication
Brno University of Technology
Brno, Czech Republic
romana.scotkova@vut.cz

2nd Matyáš Dostalík
Department of Biomedical Engineering
Faculty of Electrical Engineering and Communication
Brno University of Technology
Brno, Cech Republic
256273@vut.cz

3rd Zdenka Fohlerová
Department of Biomedical Engineering
Faculty of Electrical Engineering and Communication
Brno University of Technology
Brno, Czech Republic, Country
fohlerova@vut.cz

***Abstract*— The development of biomimetic scaffolds for cardiac repair requires a sophisticated integration of mechanical resilience, electrical coupling, and spatial functionality. Conventional symmetric hydrogels often struggle to balance the contradictory requirements of robust tissue adhesion and the prevention of post-operative visceral fouling. This work describes the fabrication and characterization of a double-layer electroconductive Janus hydrogel based on a poly(vinyl alcohol) (PVA) and collagen (Coll) biohybrid matrix. By utilizing a charge-balance switching mechanism, the Janus architecture achieves asymmetric properties: a bottom layer optimized for rapid adhesion to wet cardiac tissue under acidic conditions, and an anti-adhesive top layer created under neutral conditions to prevent unintended tissue tethering. To bridge the electrical insulation typical of polymer networks, PEDOT:PSS was incorporated as a conductive filler. Results demonstrated that the integration of 5 v/v % PEDOT:PSS established a robust percolation network, significantly reducing charge-transfer resistance and lowering impedance within the physiologically relevant frequency range (0 to 200 Hz). Degradation profiles confirmed that the conductive filler provides mechanical reinforcement, revealing a stable yet transient support system over a 21-day period. These findings suggest that the PVA/Coll/PEDOT:PSS Janus hydrogel represents a promising stimuli-responsive platform for advanced cardiac patches, offering directional functionality and enhanced bioelectrical integration.**

Keywords*— *electroconductive hydrogel, janus hydrogel, polyvinyl alcohol, collagen, regenerative medicine

I. Introduction

The use of advanced materials is becoming prevalent in a wide range of applications in various industries. Innovative functional materials, offering effective physicochemical performance, are building a broader conceptual picture of a new generation of polymeric materials and are making their way into the field of therapeutic devices like drug carriers, scaffolds, or bandages. Tissue engineering, as a fast-moving interdisciplinary field, is transitioning from biologically passive methods for damaged tissue repair to stimuli-responsive platforms. One of them being scaffolds offering a niche tissue mimicry with necessary physicochemical support to the dynamic bioelectrical environment of the human body [1, 2].

The human physiological environment is inherently dynamic and bioelectrical; to achieve true tissue mimicry, scaffolds must bridge the gap between synthetic architectural support and biological signalling. By integrating smart materials that respond to environmental cues, a niche environment can be designed that not only supports cell attachment but actively directs cellular behavior through biophysical pathways. A promising strategy to achieve this is the development of dual-function scaffolds combining piezoelectric and electroconductive polymers, which creates a biomimetic environment that effectively translates mechanical movement into the bioelectrical signals necessary for regeneration. To establish the structural foundation of such a system, specific fabrication techniques and material selections must be considered [2, 3].

To move these hydrogels from passive material to active components, electrical conductivity must be introduced into the network by incorporating a conductive filler. Electroconductive hydrogels represent a class of biomaterials designed to overcome the electrical insulation typical of standard polymer networks in hydrogel matrix. By incorporating conductive composites, such as electroconductive polymers with conjugated π-bond systems. While many conductive fillers such as graphene or carbon nanotubes offer high performance, PEDOT:PSS [poly(2,3-dihydrothieno-1,4-dioxin)poly(styrene sulfonate)] known for its dual ionic and electroconductivity, high chemical stability, mechanical flexibility and biocompatibility offers a suitable alternative in regenerative medicine applications [4-6].

Traditional hydrogels, while widely used, are often limited by their homogenous structure and symmetric properties, which in some cases fail to mimic the complexity of natural tissue enviroment. In contrast, Janus hydrogels are new type of promising materials due to their dual architecture, characterized by asymmetric chemical compositions or surface properties on opposing sides [7]. This asymmetric design is particularly

advantageous for internal wound healing and tissue patches, as double-layer bioadhesives can be engineered to exhibit opposing adhesive and anti-adhesive properties [7, 8].

Typically, these systems comprise two distinct layers where the adhesive layer is often functionalized with high densities of free carboxyl (–COOH) groups, while the anti-adhesive layer is constructed using antifouling polymers to prevent unwanted adhesions to surrounding tissue while regeneration process. This overcomes the inherent contradiction in tissue healing, where strategies to prevent adhesions, such as superhydrophobic surfaces or dense pores and often inhibit the migration of important repair cells like fibroblasts , thereby affecting regeneration efficiency [7, 9]. For cardiac applications, these hydrogels must also function as flexible sensors, demanding high mechanical stability and low hysteresis to ensure reliability and durability under the cyclic mechanical strain of the heartbeat. These properties are vital for reducing energy dissipation and preventing fatigue during repeated deformation; however, a significant material science challenge exists in the trade-off between achieving high toughness and low hysteresis. Despite these difficulties, systems that simultaneously achieve both high tensile strain and low hysteresis remains rare but essential for biomimetic cardiac patches [10].

To achieve these multifaceted goals, researchers often utilize collagen (Coll), the most abundant structural protein in the extracellular matrix, favored for its exceptional biocompatibility. However, collagen-based materials frequently suffer from inadequate mechanical strength compared to native tissues, leading researchers to employ polymer blending the association of natural and synthetic polymers to modulate physical properties through compositional changes [11]. Building on this asymmetric and functionalized framework, Janus hydrogels offer directional functionality where each region can be independently engineered to perform a specific role, such as providing spatially resolved responsiveness to external stimuli like pH or temperature [7].

The performance of these hydrogels is further dictated by their internal architecture and conductivity. The crosslinking method used for PVA-based hydrogels significantly impacts their final biological and mechanical properties due to variations in the interaction forces between polymer chains. Consequently, PVA hydrogels are typically crosslinked using physical methods, where molecular linkages consist primarily of physical entanglement, hydrogen bonds, and hydrophobic or electrostatic interactions. The most common process for PVA-based hydrogels is freezing/thawing (F-T). This method forms crystalline regions via internal intramolecular hydrogen bonds, while intermolecular regions remain limited. During freezing, PVA chains are displaced by water molecule crystallization, leading to the formation of high-concentration aggregates. As the material thaws, the loose water molecules between the chains are expelled, increasing the number of physical entanglements. The final properties of the hydrogel can be controlled by adjusting the number of F-T cycles, the freezing temperature, and the processing time [7].

The aim of this research was to fabricate and characterize a double-layer electroconductive PVA/Collagen Janus hydrogel with the conductive polymer PEDOT:PSS. The bottom hydrogel is constructed under acidic conditions, featuring abundant -COOH and -NH3+ residues for adhesion to diverse wet tissues. Conversely, the top hydrogel is created under neutral conditions, achieving a balanced charge, thus offering a possible prevention of postoperative tissue adhesion. This biohybrid hydrogel aims to integrate directional stimuli-responsiveness and anisotropic deformation with the necessary electrical conductivity for possible cardiac tissue implants.

II. Methods and Materials

A. Hydrogel Fabrication

Aqueous solutions of 10 wt % PVA was prepared by dissolving preweighed quantities of PVA in distilled water and heating for 1 h at 90 °C. The PVA solution was cooled slowly to room temperature. For the adhesive layer of the hydrogel, Col (4 wt. %, dissolved in 0,1 HCl) was mixed with the PVA emulsion by stirring at room temperature to forma PVA /Col solution to form a 1:1 ratio. Second, 1,5 ml of PVA and 500 µl of PVA/Col solutions were casted into each well of a 12 well culture plate until it was spread smoothly to form a coating. In the final stage of preparation, PEDOT:PSS was incorporated into the hydrogel precursor at concentrations of 1, 3, and 5% (v/v). The resulting mixture were stirred continuously to ensure complete homogenization before further processing. A freeze-thaw (F-T) treatment was used to physically cross-link the coated hydrogel. Briefly, coated plates were frozen at 20 C overnight and slowly thawed in a sterile hood with the UV light on. The sample underwent freeze-thaw cycles three times to form an hydrogel bilayer conduit film scaffold. and to ensure stable crosslinked network of hydrogel.

B. Mechanical Properties Characterization

Rheological measurements of the both hydrogel layers (PVA/PEDOT:PSS and PVA/Coll/PEDOT:PSS) samples were performed on an oscillatory rheometer (DHR-2, TA Instruments) using a 20 mm parallel-plate with crosshatched surface and a gap of 2000 µm with Peltier Plate tempered 25 °C. The whole area of geometry was covered by hydrogel sample. Firstly, the samples were conditioned designated temperature of 25 °C for 180 s. Secondly, amplitude sweeps at a constant frequency of oscillation (1 Hz) from 0.01 to 1 000 % (logarithmic scale, 8 points per decade) were performed to determine differences in the viscoelastic behaviour of the hydrogels. Results are reported as the average of at least two independent measurements in the Results and Discussion sections. The software TRIOS (TA Instruments) was used for analysis.

The in vitro degradation of double-layer hydrogel was monitored by measuring the weight loss over time. Hydrogel underwent full drying, and the initial weight was recorded followed by immersion in 2 ml PBS media with a weekly change of the solution. A weight comparison was performed at each predetermined time interval and then weighed. The degradation was calculated by equation:

$$\Delta w = (w_0 - w_t)/w_0 \times 100\ \% \quad (1)$$

where w_0 and w_t are the weights before and after degradation

C. Electrochemical Impedance Spectroscopy

Electrochemical impedance spectroscopy (EIS) measurements were conducted at room temperature using an electrochemical workstation (Metrohm Autolab). The experimental arrangement comprised gold interdigitated electrodes (conductance constant: 9.450 µm.) placed in the bottom of homemade chambre filled with final hydrogel. The cell was properly closed to avoid hydrogel drying during the measurement. EIS measurements were performed at the frequency range from 0.1 Hz to 10 000 kHz, using an amplitude of ±10 mV.

III. Results and Disscussions

The double layer electroactive Janus hydrogel has a bilayer structure composed of Polyvinyl Alcohol and Collagen, respectively. PEDOT:PSS, as both ionically and electrically conductive polymer, was added to both layers to improve the hydrogel conductivity. The introduction of conductive polymer improved the electrical performance of the hydrogel. Because Collagen enhanced adhesive layer was constructed on the surface of the high concentration PVA-only layer, an interpenetrating layer appeared between the layers.

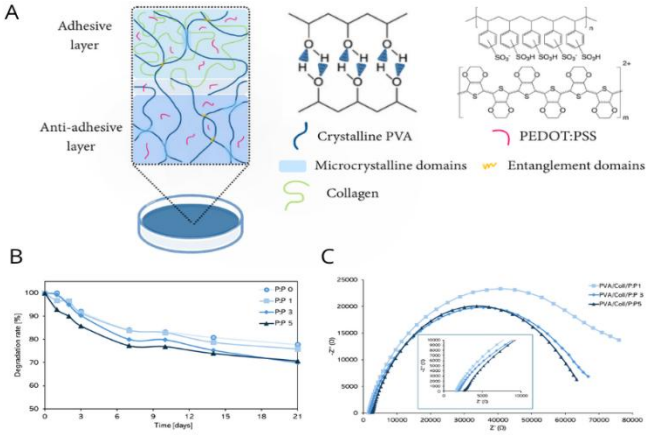

Fig. 1. A) Structure of double-layer hydrogel consisting of non-adhesive PVA and adhesive PVA/Coll substrates with dispersed PEDOT:PSS particles, B) Degradation profile of hydrogel samples with PEDOT:PSS concentration from 0 to 5 v/v %, C) Nyquist diagram of Janus hydrogel with various concentrations of PEDOT:PSS incorporated (1, 3 and 5 v/v %).

Dynamic rheological tests were conducted to evaluate the softness of the hydrogel layer and the mechanical stability of the hydrogel layers. Fig. 1A presents the viscoelastic moduli of the non-adhesive hydrogel layer at, highlighting its distinct soft matter characteristics, with G' not greater than G". The trend however changes for the adhesive layer of hydrogel (Fig. 1B) with Storage modulus exceeding the loss modulus indicating greater mechanical strength and later deformation with higher strain applied. Across varying strain amplitudes, the hydrogel exhibited nearly constant G' and G" values, indicating the reversibility of the network structure. As a result, the hydrogel demonstrated strong mechanical robustness and stability, capable of accommodating the deformations of the hydrogel layer. Across all tested frequency ranges, the hydrogel consistently exhibited higher G' than G" values, indicating their stability and viscoelastic behavior. Additionally, the oscillatory frequency sweep results of the analysis showed that G' values consistently exceeded G" values, indicating the dominant elastic of the material for typical gel-like structures.

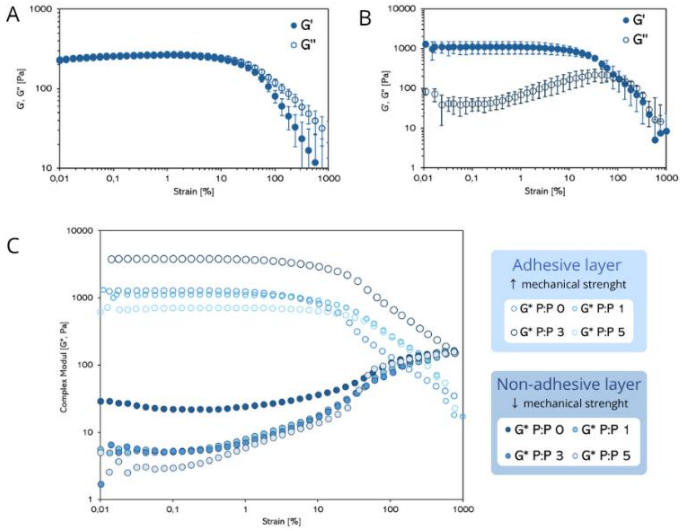

Fig. 2. Evaluation of rheological performance with amplitude sweep tests. A) Viscoelastic moduli of adhesive PVA/Coll/PEDOT:PSS layer, B) Viscoelastic moduli of anti-adhesive PVA/PEDOT:PSS layer of the hydrogel, C) Comparison of Complex moduli G* of anti-adhesive PVA and adhesive PVA/Coll layers showing significant decrease of mechanical strength after Collagen incorporation.

As a viscoelastic material, PEDOT:PSS-PVA hydrogel also exhibits a wide linear viscoelastic region at 25 °C from the rheological characterization of energy storage modulus (G') and loss modulus (G'') (Fig. *2*). Compared to pure PVA hydrogel, the higher G* of PEDOT:PSS-PVA hydrogel indicates a stronger mechanical influence due to the introduction of the PEDOT:PSS network. Conversely, the higher G'' value of PEDOT:PSS-PVA hydrogel reveals a microphase semi-separation network of PEDOT:PSS.

The degradation profiles (Fig. *1*, B) of the fabricated hydrogel scaffold revealed distinct material stability over a 21 day period and further confirmed the degradability of the hydrogels under conditions close to physiological ones (PBS, pH 7.4). During the first 24 hours, percentual mass loss remained minimal, likely to initial stages of swelling of the polymer network and the gradual PBS media absorption. The most noticeable reduction in mass was during the day 3 to 5, after which the hydrogel mass stabilized. The incorporation of the conductive filler appeared to modulate the degradation rate. Hydrogels with PEDOT:PSS (mainly samples with higher concentrations, 3 or 5 v/v %) exhibited a more rapid mass loss compared to hydrogel without conductive filler. The reason for leveled out degradation rates for hydrogels containing PEDOT:PSS might be caused by PEDOT:PSS particles simply dispersed in the matrix of PVA or PVA/Coll causing slightly lower degree of crosslinking, causing un-crosslinked PVA or Collagen to

diffuse from the matrix. This degradation behavior suggests balanced architecture the resulting scaffold maintains its integrity for the required therapeutic timeframe.

Electrochemical impedance spectroscopy (EIS; Fig. *1*, C) was employed to analyse the charge transport inside the biohybrid network. The results demonstrated that the PVA/Coll hydrogel containing the highest concentration of PEDOT:PSS (5 v/v %) possessed the lowest impedance across the measured spectrum. This enhancement was particularly pronounced at lower frequencies, similar those with the physiologically relevant range of 0–200 Hz. To further characterize the conductive mechanisms, Nyquist plots were utilized, which revealed that an increased concentration of the PEDOT:PSS conductive filler resulted in a reduced semicircle radius. This reduction represents the lowest charge-transfer resistance and indicates a high-velocity charge-transfer capability within the matrix. These findings suggest that the incorporation of PEDOT:PSS and its integration through both covalent and physical bonding with the PVA/Collagen polymer network effectively introduces the electrical performance to the scaffold.

IV. Conclusion

In this study, a double-layer Janus hydrogel was successfully developed, addressing the complex mechanical and electrical demands for cardiac tissue engineering implants such as cardiac patches. Heterogenous architecture effectively addresses the adhesive and anti-adhesive requirements of a cardiac patch. The synergy between natural biopolymer (Collagen) and synthetic polyvinyl alcohol, processed through physical freezing/thawing cycles, resulted in a biohybrid matrix with the necessary structural integrity and mechanical support. The incorporation of PEDOT:PSS transformed the passive scaffold into an electroactive platform, achieving optimal charge-transfer kinetics and low impedance at lower Overall, this electroconductive Janus system provides a versatile foundation for cardiac patches application, combining directional stimuli-responsiveness with the electrical connectivity required for synchronized heart function.

References

[1] Andrey, V. et al., 2024. Piezoelectric PVDF and its copolymers in biomedicine: innovations and applications. *Biomaterials Science*, 12(20), pp.5164-5185.

[2] Samadi, A. et al., 2022. Comparative review of piezoelectric biomaterials approach for bone tissue engineering. *Journal of Biomaterials Science, Polymer Edition*, 33(12), pp.1555-1594

[3] Havlíková, T. et al., 2025. Adaptability of Electrospun PVDF Nanofibers in Bone Tissue Engineering. *Polymers*, 17(3), p.330.

[4] Paulsen, B.D. et al., 2020. Organic mixed ionic–electronic conductors. *Nature Materials*, 19(1), pp.13-26

[5] Gao, C. et al., 2022. Recent Development of Conductive Hydrogels for Tissue Engineering: Review and Perspective. Macromolecular Bioscience, 22(8).

[6] Ghasemi-Mobarakeh, L. et al., 2009. Electrical Stimulation of Nerve Cells Using Conductive Nanofibrous Scaffolds for Nerve Tissue Engineering. Tissue Engineering Part A, 15(11), pp.3605-3619.

[7] Yan, T. et al., 2025. Multifunctional Janus Hydrogels: Surface Design Strategies for Next-Generation Clinical Solutions. *Gels*, 11(5), p.343.

[8] Peng, W. et al., 2025. Charge balance transition enabled Janus hydrogel for robust wet-tissue adhesion and anti-postoperative adhesion. *Bioactive Materials*, 52(5), pp.123-138.

[9] Wang, H. et al., 2023. An Integrally Formed Janus Hydrogel for Robust Wet-Tissue Adhesive and Anti-Postoperative Adhesion. *Advanced Materials*, 35(23), pp.123-138.

[10] Song, Z. et al., 2025. Janus adhesion hydrogels with low hysteresis and high sensitivity toward precision flexible electronics underwater. *Chemical Engineering Journal*, 521(23), p.166386.

[11] B. Sarti, M. Scandola, Viscoelastic and thermal properties of collagen/poly(vinyl alcohol) blends, Biomaterials 16 (1995) 785-792.

Segmentation of Cervical Spinal Rootlets in MRI Scans of Patients with Cord Compression

Kateřina Krejčí
Department of Biomedical Engineering
FEEC, Brno University of Technology
Brno, Czech Republic
0009-0009-5817-4840

Michal Nohel
Department of Biomedical Engineering
FEEC, Brno University of Technology
Brno, Czech Republic
0000-0002-2679-2160

Roman Jakubíček
Department of Biomedical Engineering
FEEC, Brno University of Technology
Brno, Czech Republic
0000-0003-4293-260X

Jiří Chmelík
Department of Biomedical Engineering
FEEC, Brno University of Technology
Brno, Czech Republic
0000-0001-9950-6279

Jan Valošek
University Hospital Olomouc
Palacký University Olomouc
Olomouc, Czech Republic
0000-0002-7398-4990

***Abstract*—Spinal rootlets segmentation and spinal level estimation are relevant steps within applications such as spinal cord lesion classification and functional magnetic resonance imaging group analyses. Current automatic segmentation methods were developed solely on magnetic resonance imaging (MRI) scans of healthy adults and the performance on patients' MRI scans remains unexplored. The goal of this study was to develop three segmentation models (two U-Net-based and one transformer-based) on MRI scans of both healthy subjects and patients with degenerative cervical myelopathy (DCM). The performance of proposed models were compared with existing segmentation model developed solely on healthy subjects. Preliminary results indicate similar performance across the U-Net–based models, with comparable Dice scores and Hausdorff distances, outperforming the transformer-based model. Relative to the existing model, the U-Net–based Clinic-RootletSeg model showed slightly improved overall performance; however, the existing model also performed well on MRI scans of DCM patients, even though it was trained solely on MRI scans of healthy subjects.**

***Index Terms*—Spinal Rootlets Segmentation, Degenerative Cervical Myelopathy, Magnetic Resonance Imaging, Deep Learning, U-Net, Transformer**

I. Introduction

Spinal rootlets are bundles of nerve fibers connecting the spinal cord to peripheral parts of the body. Accurate spinal rootlets segmentation from magnetic resonance imaging (MRI) scans enables the estimation of neurological spinal levels, which is relevant for spinal cord lesion classification [1], [2], neuromodulation therapy [3], [4], and functional magnetic resonance imaging (fMRI) group analyses [5], [6]. Existing deep-learning-based spinal rootlets segmentation methods [7], [8] were developed exclusively on MRI scans of healthy subjects, and their applicability to clinical data (e.g., subjects with degenerative spinal cord compression) remains unexplored. Degenerative cervical myelopathy (DCM) typically affects only part of the spinal region, mostly at the C5–C7 levels [9]. Segmentation in those regions is particularly challenging due to changes of the anatomy, reduced tissue contrast, flow void artifacts, and motion-related artifacts. The aim of this study is to compare the performance of the currently available model [8] with three (two U-Net-based and one transformer-based) newly trained models on clinical MRI scans of healthy subjects and DCM patients.

Supported by the Ministry of Health of the Czech Republic, grant number NW26J-04-00101.

II. Data and Methods

A. Dataset description

This study included MRI scans from three datasets of healthy subjects: (i) 3T T2w scans from the open-access OpenNeuro ds004507 dataset [10], (ii) 3T T2w scans from the open-access spine-generic multi-subject dataset [11], (iii) 7T MP2RAGE scans from a private MP2RAGE single-site dataset [12], [13]; and a private DCM single-site dataset of 3T T2w scans of patients with degenerative cervical myelopathy. Detailed characteristics of the datasets are included in Table I. Figure 1 presents a representative comparison of T2w MRI scans between a healthy subject and a patient with degenerative cervical myelopathy. In the DCM dataset, the pathology occurs mainly at the C5–C7 spinal levels.

B. Ground truth labels obtaining

For the OpenNeuro, spine-generic, and MP2RAGE datasets, the existing ground truth labels were used as in the development of the RootletSeg model [8]. For the DCM dataset, where ground truth labels were not available, the RootletSeg model was applied. The output segmentations were manually verified and corrected using FSLeyes [14] if necessary. To avoid data leakage and ensure an objective comparison of models, ground truth labels for the test set of the DCM dataset were created manually from scratch (i.e., without using the RootletSeg model). All ground truth segmentations were level-specific, with each spinal level assigned a unique label (e.g., label 2 for C2, label 3 for C3).

TABLE I
CHARACTERISTICS OF STUDY DATASETS (TSE - TURBO SPIN ECHO, MP2RAGE - MAGNETIZATION PREPARED 2 RAPID ACQUISITION GRADIENT ECHOES)

Dataset	OpenNeuro	spine-generic	MP2RAGE	DCM
Total participants (male/female)	7 (5/2)	24 (12/12)	19 (11/8)	19 (3/16)
Age (mean ± std)	22.57 ± 0.53	29.58 ± 6.53	29.84 ± 6.66	53.16 ± 9.80
Total MRI scans (train/test)	12 (10/2)	24 (21/3)	57 (45/12)	43 (34/9)
Manufacturer (Siemens/GE)	12/0	20/4	57/0	43/0
Field strength (3T/7T)	12/0	24/0	0/57	43/0
Sequence	TSE	TSE	MP2RAGE	TSE
Voxel size (mm)	0.6 × 0.6 × 0.6	0.8 × 0.8 × 0.8	0.7 × 0.7 × 0.7	0.8 × 0.8 × 0.8

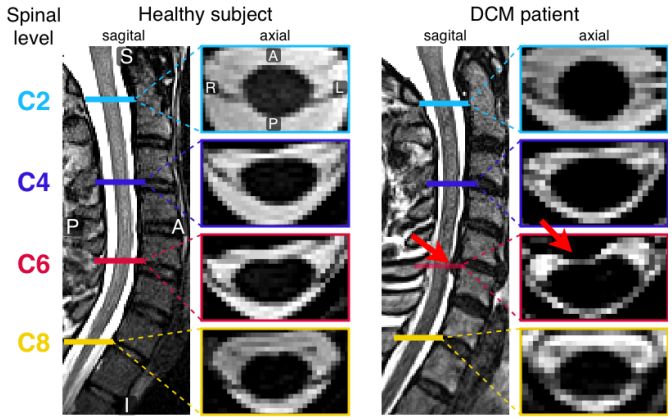

Fig. 1. Comparison of T2w MRI scan between a healthy subject and a degenerative cervical myelopathy (DCM) patient; red arrows show the point of the spinal compression (A – anterior, P – posterior, S – superior, I – inferior, R – right, L – left)

C. Deep-learning model training

The development of models for spinal rootlet segmentation was performed using the nnUNetv2 framework [15]. The MRI scans of healthy subjects were split following the same way as in the RootletSeg model [8] (trained solely on MRI scans of healthy subjects). Additionally, MRI scans of DCM patients were randomly divided into 80 % for training and 20 % for testing (Table II). The training inputs of all models were MRI scans and corresponding level-specific segmentation of spinal rootlets. Within this study, we developed three models with different 3D architectures: Clinic-RootletSeg (classical U-Net [16]), Clinic-ResEncM (Residual Encoder M U-Net [17]), and Clinic-TransformerM (Transformer Primus M [18]).

Preprocessing with the nnUNetv2 framework involved re-orienting both the input scans and their corresponding ground truth labels to the Right-to-Left × Posterior-to-Anterior × Inferior-to-Superior (RPI) orientation, intensity z-score normalization and resampling into isotropic 0.7 mm resolution before training.

The nnUNetv2 framework proposed the same configuration of these parameters for all architectures: 6 stages, instance normalization technique, Leaky ReLU activation function, batch size 2, patch size 48 × 224 × 224, learning rate 0.01, Dice Loss and Cross-Entropy loss function and Stochastic Gradient Descent with Nesterov Momentum optimizer. Random spatial transformation (translation, rotation, scaling), mirroring, Gaussian noise and Gaussian blur, brightness and contrast adjustment, Gamma transformation and low-resolution simulation were used as data augmentation. Trained models were applied on the testing dataset and their performance was evaluated by the Dice score and 95th percentile Hausdorff Distance (HD95).

III. RESULTS

Within this study, we developed three models with different 3D architectures: Clinic-RootletSeg (classical U-Net), Clinic-ResEncM (Residual Encoder M U-Net), and Clinic-TransformerM (Transformer Primus M). A comparison of the model parameters with the existing RootletSeg model is shown in Table II.

TABLE II
COMPARISON OF DIFFERENT MODEL SETTINGS (DCM - DEGENERATIVE CERVICAL MYELOPATHY)

Model	Orig-RootletSeg	Clinic-RootletSeg	Clinic-ResEncM	Clinic-TransformerM
Architecture	U-Net	U-Net	Residual Encoder M U-Net	Residual Encoder M U-Net
Patch size	352 × 96 × 128	48 × 224 × 224	48 × 224 × 224	48 × 224 × 224
Number of parameters	88,217,234	87,889,554	382,740,818	150,014,410
Subjects (Healthy/DCM)	50 (50/0)	69 (50/19)	69 (50/19)	69 (50/19)
Training scans (Healthy/DCM)	76 (76/0)	110 (76/34)	110 (76/34)	110 (76/34)
Testing scans (Healthy/DCM)	26 (17/9)	26 (17/9)	26 (17/9)	26 (17/9)

The models were tested on a total of 26 MRI testing scans, including 17 scans of healthy subjects and 9 scans of DCM patients. A visual comparison of model predictions and manual ground truth labels is presented in Figure 2 for a healthy subject and in Figure 3 for a DCM patient.

The overall Dice scores across all spinal levels, MRI contrasts, and subjects were: Orig-RootletSeg (61.47 ± 12.17 %), Clinic-ResEncM (61.20 ± 11.90 %), Clinic-RootletSeg (60.95 ± 11.81 %), and Clinic-TransformerM (53.59 ± 11.6 %). The 95th percentile Hausdorff distances were: Clinic-RootletSeg (5.16 ± 2.27 mm), Orig-RootletSeg (5.21 ± 2.19 mm), Clinic-TransformerM (5.81 ± 2.21 mm) and Clinic-ResEncM (6.08 ± 13.44 mm). The results across C2–T1 spinal levels are shown in Table III and the corresponding box plots are in the Figure 4. The performance of U-Net-based models is comparable and the performance of the transformer-based model Clinic- TransformerM is lower (lowest Dice score and highest HD95). Based on the Dice score and HD95 across individual spinal levels, the Clinic-RootletSeg model achieved slightly better performance than other U-Net-based models (Orig-RootletSeg, Clinic ResEncM). The lowest performance was observed for the Clinic-TransformerM model.

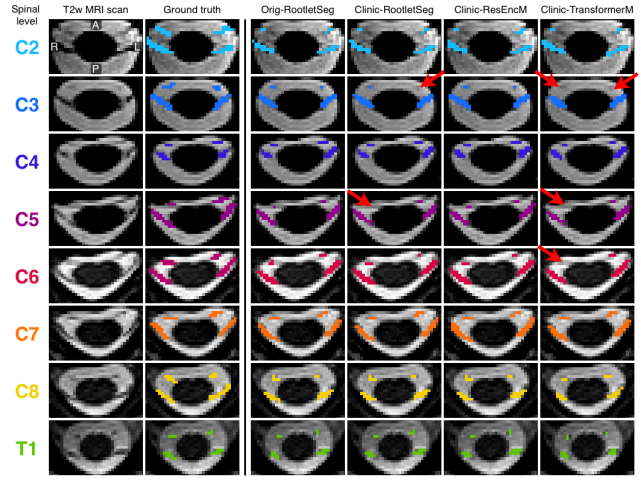

Fig. 2. Visual comparison of model predictions and manual ground truth labels of healthy subject MRI scan (red arrows show missing rootlets segmentation in predictions)

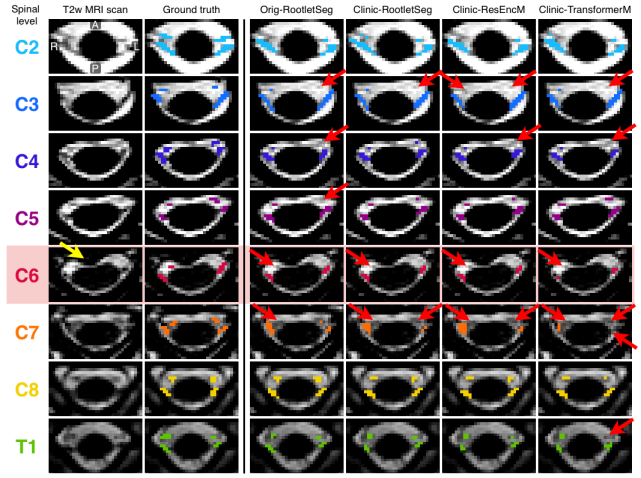

Fig. 3. Visual comparison of model predictions and manual ground truth labels of MRI scan of DCM patient at C6 spinal level (red arrows show missing rootlets segmentation in predictions and yellow arrow shows the point of spinal compression)

The performance of the Clinic-RootletSeg and Orig-RootletSeg models was compared on T2w MRI scans of healthy subjects (5 scans) and DCM patients (9 scans). For healthy subjects, the Dice scores across all spinal levels were 64.46 ± 9.66 % for Orig-RootletSeg and 64.30 ± 8.91 % for Clinic-RootletSeg, with HD95 values of 5.31 ± 2.14 mm and 5.19 ± 2.14 mm, respectively. For DCM patients, the Dice scores were 55.83 ± 14.32 % for Orig-RootletSeg and 54.62 ± 13.91 % for Clinic-RootletSeg, and the HD95 values were 5.04 ± 2.30 mm and 5.21 ± 2.93 mm. The comparison across spinal levels is presented in Table IV and illustrated in Figure 5. Both models showed lower Dice scores in the group with DCM patients, particularly at spinal levels C4–C7, but comparable Hausdorff distances at those levels. Detailed comparison with visual illustration of one healthy subject and one DCM patient is shown in Figure 6.

TABLE III
OVERALL COMPARISON OF DICE SCORE (A) AND 95TH HAUSDORFF DISTANCE (B) OF MODELS ACROSS SPINAL LEVELS (THE BEST RESULT PER EACH SPINAL LEVEL IS IN BOLD)

(a) Dice score [%]

Spinal level	Orig-RootletSeg	Clinic-RootletSeg	Clinic-ResEncM	Clinic-TransformerM
C2	63.98 ± 9.63	**64.74 ± 9.08**	64.52 ± 9.36	59.09 ± 10.51
C3	63.83 ± 12.64	**64.68 ± 11.94**	64.20 ± 11.77	53.28 ± 13.45
C4	59.60 ± 8.38	58.97 ± 8.35	**60.13 ± 8.28**	49.96 ± 8.80
C5	59.82 ± 10.10	**60.25 ± 10.88**	60.22 ± 10.62	53.67 ± 9.43
C6	**60.24 ± 13.67**	58.15 ± 14.91	60.19 ± 13.41	55.69 ± 11.86
C7	59.24 ± 17.12	**59.38 ± 15.60**	59.33 ± 16.19	51.82 ± 15.19
C8	**62.25 ± 14.96**	61.29 ± 13.12	61.59 ± 14.58	54.6 ± 11.25
T1	**62.82 ± 8.39**	60.17 ± 7.65	59.43 ± 8.90	50.6 ± 9.71

(b) 95th percentile Hausdorff distance [mm]

Spinal level	Orig-RootletSeg	Clinic-RootletSeg	Clinic-ResEncM	Clinic-TransformerM
C2	5.29 ± 1.71	**4.68 ± 1.54**	4.92 ± 1.64	5.51 ± 2.14
C3	5.02 ± 1.84	**4.90 ± 1.75**	4.94 ± 1.93	5.65 ± 1.62
C4	5.26 ± 1.66	5.12 ± 1.74	**5.07 ± 1.71**	5.52 ± 1.45
C5	4.34 ± 1.67	**4.24 ± 1.25**	4.48 ± 1.43	4.56 ± 1.29
C6	**4.54 ± 1.36**	4.71 ± 1.22	4.48 ± 1.34	5.43 ± 1.85
C7	**5.36 ± 2.23**	6.08 ± 3.8	12.37 ± 37.56	6.22 ± 2.57
C8	5.81 ± 3.08	**5.75 ± 2.9**	5.95 ± 3.37	6.43 ± 3.12
T1	6.1 ± 3.03	**5.83 ± 2.28**	6.39 ± 3.03	7.13 ± 2.31

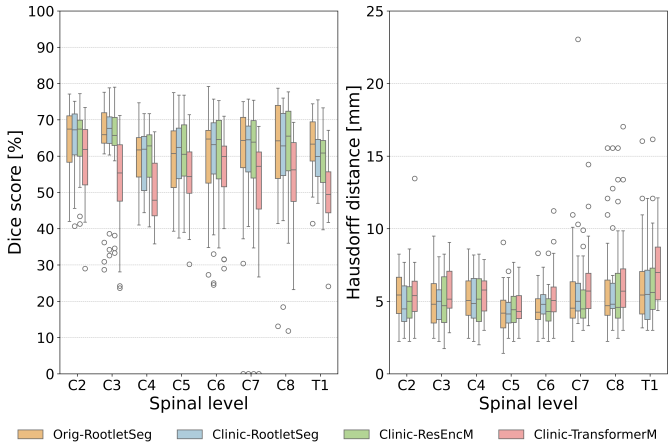

Fig. 4. Box plots of the Dice score and 95th percentile Hausdorff distance of models across C2-T1 spinal levels (26 testing MRI scans), showing comparable performance between U-Net-based architectures (Orig-RootletSeg, Clinic-RootletSeg and Clinic-ResEncM models) and lower performance of transformer-based architecture (Clinic-TransformerM)

IV. DISCUSSION AND CONCLUSION

In this work, we proposed three models (two U-Net-based: Clinic-RootletSeg and Clinic ResEncM model; and one transformer-based Clinic-TransformerM model) for an automatic spinal rootlets segmentation from T2w and MP2RAGE MRI scans of healthy subjects and DCM patients and compared their performance with existing (Orig-RootletSeg) model [8], which was trained solely on MRI scans of healthy adults.

The results indicate that the original RootletSeg model generalizes well to MRI data of DCM patients, despite not

TABLE IV
COMPARISON OF DICE SCORE AND 95TH HAUSDORFF DISTANCE OF ORIG-ROOTLETSEG AND CLINIC-ROOTLETSEG MODELS SEPARATELY FOR TESTING MRI SCANS OF HEALTHY SUBJECTS AND DCM PATIENTS ACROSS SPINAL LEVELS (THE BEST RESULT PER EACH SPINAL LEVEL IS IN BOLD)

(a) Dice (healthy subjects) [%]

Spinal level	Orig-RootletSeg	Clinic-RootletSeg
C2	61.99 ± 11.10	**62.87 ± 10.46**
C3	62.26 ± 15.33	**62.95 ± 14.18**
C4	61.14 ± 6.62	61.26 ± 6.89
C5	64.15 ± 8.39	**65.89 ± 7.04**
C6	**68.32 ± 5.28**	66.92 ± 5.61
C7	**67.34 ± 5.44**	66.56 ± 5.56
C8	**67.54 ± 11.88**	65.28 ± 10.79
T1	62.95 ± 7.49	62.69 ± 7.13

(c) HD95 (healthy subjects) [mm]

Spinal level	Orig-RootletSeg	Clinic-RootletSeg
C2	5.95 ± 1.39	**5.16 ± 1.53**
C3	5.34 ± 2.15	**5.09 ± 2.07**
C4	5.57 ± 1.60	**5.46 ± 1.79**
C5	4.41 ± 1.70	**4.16 ± 1.01**
C6	**4.51 ± 1.42**	4.79 ± 1.28
C7	**4.74 ± 1.55**	5.22 ± 1.64
C8	**5.18 ± 2.46**	5.40 ± 2.18
T1	6.75 ± 3.41	**5.84 ± 2.57**

(b) Dice (DCM patients) [%]

Spinal level	Orig-RootletSeg	Clinic-RootletSeg
C2	67.76 ± 4.38	**68.27 ± 4.17**
C3	66.80 ± 3.75	**67.96 ± 4.99**
C4	**56.68 ± 10.82**	54.63 ± 9.51
C5	**51.62 ± 7.88**	49.58 ± 8.68
C6	**44.99 ± 11.31**	41.57 ± 12.59
C7	43.94 ± 21.32	**45.82 ± 19.55**
C8	52.27 ± 15.66	**53.76 ± 14.40**
T1	**62.57 ± 10.38**	55.39 ± 6.48

(d) HD95 (DCM patients) [mm]

Spinal level	Orig-RootletSeg	Clinic-RootletSeg
C2	4.04 ± 1.62	**3.78 ± 1.14**
C3	**4.44 ± 0.85**	4.52 ± 0.85
C4	4.68 ± 1.70	**4.47 ± 1.52**
C5	**4.19 ± 1.72**	4.41 ± 1.66
C6	4.58 ± 1.32	**4.58 ± 1.15**
C7	**6.51 ± 2.91**	7.71 ± 5.93
C8	7.01 ± 3.88	**6.41 ± 4.02**
T1	**4.86 ± 1.71**	5.80 ± 1.73

being trained on such data. Its performance is comparable to that of the Clinic-RootletSeg model, which was trained with these data included. The Clinic-RootletSeg model achieved the best overall performance in terms of Dice score and 95th percentile Hausdorff distance (HD95) across all MRI scans and spinal levels from C2 to T1 (Table III).

Visual comparison shows that all segmentation models struggled to accurately segment some ventral rootlets, which were missing within the model predictions (red arrows in Figure 2 and Figure 3). This limitation was most evident in the Clinic-TransformerM model, which also produced the lowest Dice scores and highest HD95 values across all spinal levels.

The results show that the U-Net–based models (Orig-RootletSeg, Clinic-RootletSeg) achieve better performance than the transformer-based Clinic-TransformerM model, while using fewer parameters (∼88 million vs. ∼150 million parameters) and therefore require lower computational complexity. The Residual Encoder M U-Net-based architecture did not improve the segmentation model's performance, and the results of the Clinic-ResEncM model were comparable to or slightly worse than those of the Orig-RootletSeg and Clinic-RootletSeg models. In addition, Clinic-ResEncM has the most parameters (∼382 million) and is therefore the most computationally demanding.

The performance of Orig-RootletSeg and Clinic-RootletSeg models between healthy subjects and DCM patients was compared on T2w MRI scans (Table IV and Figure 6). The large Dice score variation of the rootlets segmentation performance between groups and small HD95 performance difference can be explained by the small size of the segmented structures.

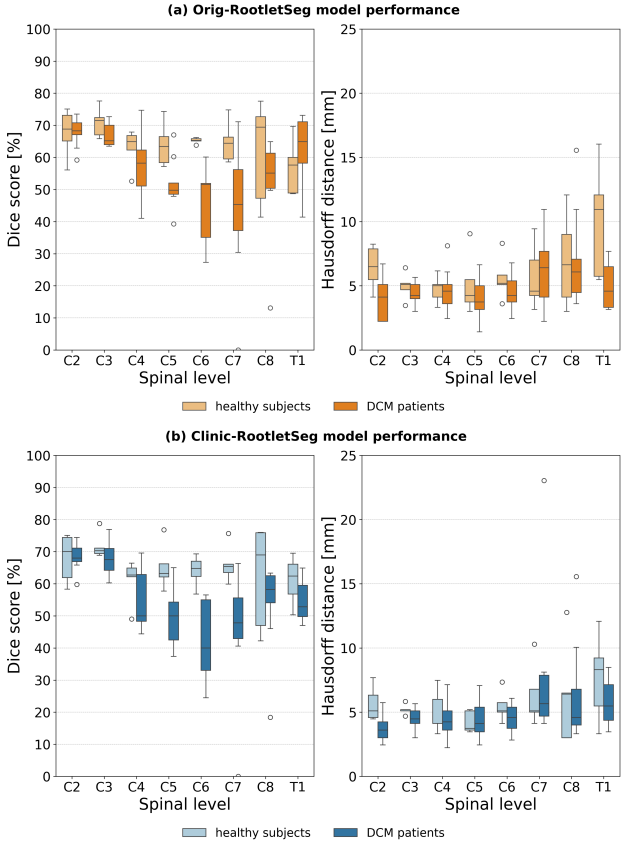

Fig. 5. Box plots visualisation of the Dice score and 95th percentile Hausdorff distance of models on testing T2w MRI scans of healthy subjects (n=5) and DCM patients (n=9) separately across C2-T1 spinal levels (showing larger difference in Dice score between groups at C4-C7 spinal levels, but lower difference in HD95)

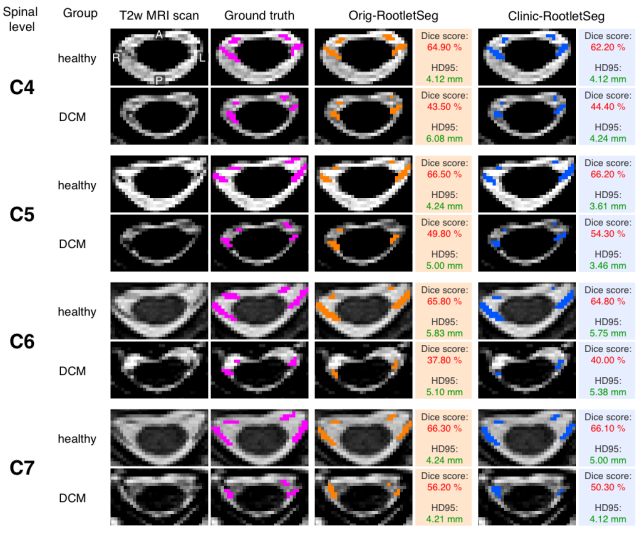

Fig. 6. Comparison of the Dice score and 95th percentile Hausdorff distance of Orig-RootletSeg and Clinic-RootletSeg models across C4-C7 spinal levels of healthy subject and DCM patient, showing larger difference between Dice scores and lower difference between Hausdorff distances (DCM - degenerative cervical myelopathy)

Spinal rootlets cover only a few voxels in the MRI scan and minor differences in overlap between the predicted and ground truth segmentation can lead to relatively large changes in Dice scores. In contrast, HD95 measures boundary distances and is therefore less sensitive to small volumetric differences when the predicted and ground truth segmentation remain spatially close. As a result, variations in Dice occur, while HD95 remains relatively stable across subjects. In addition, MRI scans of DCM patients are more affected by changes of the anatomy, reduced tissue contrast, and motion-related artefacts resulting with reduced image quality and lower visibility of rootlets, which could lead to incomplete segmentations (red arrows in the Figure 3). In such cases, the models localize the main rootlet structures correctly, resulting in similar HD95 values, while missing parts of the rootlets lead to lower Dice scores. Overall, the HD95 results confirm the visual comparison (Figure 6), which shows that the models are stable in predictions for healthy subjects and also DCM patients with spinal cord compression.

In a future study, other possible extensions could involve including MRI scans of subjects with other spinal pathologies (such as spinal lesions) in the training process, with an evaluation of performance across different pathologies. To improve accessibility for future research, the Clinic-RootletSeg model could be implemented within the freely available Spinal Cord Toolbox [19].

ACKNOWLEDGMENT

We thank Lumír Hrabálek, Vlastimil Novák, Přemysl Stejskal, Klára Balážová, Julien Cohen-Adad, Sandrine Bédard, Falk Eippert, Ulrike Horn and Virginie Callot for their assistance with the data acquisition of the datasets. This study is supported by the Ministry of Health of the Czech Republic, grant number NW26J-04-00101. Computational resources were provided by the e-INFRA CZ project (ID:90254), supported by the Ministry of Education, Youth and Sports of the Czech Republic.

REFERENCES

[1] S. Mohajeri Moghaddam and A. A. Bhatt, "Location, length, and enhancement: systematic approach to differentiating intramedullary spinal cord lesions", *Insights into Imaging*, vol. 9, no. 4, Jun. 2018, doi: 10.1007/s13244-018-0608-3.

[2] C. S. Ahuja *et al.*, "Traumatic spinal cord injury", *Nature Reviews Disease Primers*, vol. 3, no. 1, Apr. 2017, doi: 10.1038/nrdp.2017.18.

[3] R. Vallejo, "Neuromodulation of the Cervical Spinal Cord in the Treatment of Chronic Intractable Neck and Upper Extremity Pain: A Case Series and Review of the Literature", *Pain Physician*, vol. 2;10, no. 3;2, Mar. 2007, doi: 10.36076/ppj.2007/10/305.

[4] A. Rowald *et al.*, "Activity-dependent spinal cord neuromodulation rapidly restores trunk and leg motor functions after complete paralysis", *Nature Medicine*, vol. 28, no. 2, 2022, doi: 10.1038/s41591-021-01663-5.

[5] S. Bédard, J. Valošek, V. Oliva, K. A. Weber II, and J. Cohen-Adad, "Rootlets-based registration to the PAM50 spinal cord template", *Imaging Neuroscience*, vol. 3, 2025, doi: 10.1162/imag.a.123.

[6] S. Bédard *et al.*, "Spatial distribution of spinal cord fMRI activity with electrocutaneous stimulation", Mar. 2026, doi: 10.64898/2026.02.26.26347215.

[7] J. Valošek, T. Mathieu, R. Schlienger, O. S. Kowalczyk, and J. Cohen-Adad, "Automatic segmentation of the spinal cord nerve rootlets", *Imaging Neuroscience*, vol. 2, 2024, doi: 10.1162/imag_a_00218.

[8] K. Krejci *et al.*, "RootletSeg: Deep learning method for spinal rootlets segmentation across MRI contrasts", *ArXiv (Cornell University)*, Sep. 2025, doi: 10.48550/arxiv.2509.16255.

[9] S. Bédard *et al.*, "Normalizing spinal cord compression measures in degenerative cervical myelopathy", *The Spine Journal*, vol. 25, no. 9, 2025, doi: 10.1016/j.spinee.2025.03.012.

[10] S. Bédard, M. Bouthillier, and J. Cohen-Adad, "Pontomedullary junction as a reference for spinal cord cross-sectional area: validation across neck positions", *Scientific Reports*, vol. 13, no. 1, Aug. 2023, doi: 10.1038/s41598-023-40731-3.

[11] J. Cohen-Adad *et al.*, "Open-access quantitative MRI data of the spinal cord and reproducibility across participants, sites and manufacturers", *Scientific Data*, vol. 8, no. 1, Aug. 2021, doi: 10.1038/s41597-021-00941-8.

[12] U. Horn *et al.*, "Ultra-high-field fMRI reveals layer-specific responses in the human spinal cord", Jul. 2025, doi: 10.1101/2025.07.17.665316.

[13] A. Massire, M. Taso, P. Besson, M. Guye, J. -P. Ranjeva, and V. Callot, "High-resolution multi-parametric quantitative magnetic resonance imaging of the human cervical spinal cord at 7T", *NeuroImage*, vol. 143, 2016, doi: 10.1016/j.neuroimage.2016.08.055.

[14] M. Jenkinson, C. F. Beckmann, T. E. J. Behrens, M. W. Woolrich, and S. M. Smith, "FSL", *NeuroImage*, vol. 62, no. 2, 2012, doi: 10.1016/j.neuroimage.2011.09.015.

[15] F. Isensee, P. F. Jaeger, S. A. A. Kohl, J. Petersen, and K. H. Maier-Hein, "NnU-Net: a self-configuring method for deep learning-based biomedical image segmentation", *Nature Methods*, vol. 18, no. 2, Dec. 2020, doi: 10.1038/s41592-020-01008-z.

[16] O. Ronneberger, "U-Net: Convolutional Networks for Biomedical Image Segmentation", *ArXiv (Cornell University)*, May 2015, doi: 10.48550/arxiv.1505.04597.

[17] F. Isensee and K. H. Maier-Hein, "An attempt at beating the 3D U-Net", 2019, doi: 10.48550/arxiv.1908.02182.

[18] T. Wald *et al.*, "Primus: Enforcing Attention Usage for 3D Medical Image Segmentation", *ArXiv (Cornell University)*, Mar. 2025, doi: 10.48550/arxiv.2503.01835.

[19] B. De Leener *et al.*, "SCT: Spinal Cord Toolbox, an open-source software for processing spinal cord MRI data", *NeuroImage*, vol. 145, 2017, doi: 10.1016/j.neuroimage.2016.10.009.

Garrett
ADVANCING MOTION



TRANSFORMATIVE TECHNOLOGY IS IN OUR NATURE

DISCOVER THE JOB OPENINGS AT
GARRETT MOTION R&D CENTER
IN BRNO

Contact: kariera@garrettmotion.com

JOB PORTAL

CAREERS SITE

Turbo & Hybrid Technology

Zero Emission

Audio Inpainting with Endless Sounds

Michal Švento
Department of Telecommunications
Brno University of Technology
Brno, Czech Republic
michal.svento@vut.cz

***Abstract*—Seamlessly reconstructing missing or corrupted segments of an audio recording remains a challenging problem in both audio restoration and creative sound design. This paper proposes a novel application of the endless sounds technique, a spectral phase-randomisation approach that generates arbitrarily long continuations of a stationary signal, to audio inpainting and creative segment stitching. By independently synthesising endless continuations from both sides of a gap and blending them via energy-normalised, time-aligned crossfading, the proposed pipeline achieves perceptually seamless transitions at negligible computational cost, requiring only a small number of FFT operations and a cross-correlation. The method is well-suited to stationary and quasi-stationary material such as ambient textures and sustained instrument tones, and naturally extends to the creative merging of unrelated audio clips without requiring any shared spectral content. In contrast to deep learning methods, the approach demands no training data and is suitable for real-time and resource-constrained applications. We discuss the trade-offs in context-length selection and identify adaptive context selection and perceptual evaluation as directions for future work.**

***Index Terms*—endless sounds, FFT, audio inpainting, audio restoration**

I. Introduction

The extension of audio signals from a limited context is an important problem in audio signal processing. It enables generation of long sound textures from short recordings and is useful in applications such as sound design in cinematography and games, and ambience generation for audiovisual performance. Traditional approaches to audio synthesis rely on signal-processing techniques such as granular synthesis [1] or concatenative methods [2]. More recently, data-driven approaches based on statistical modeling [3] or deep neural networks have been proposed for audio synthesis [4]–[6]. Despite their differences, both families of methods tend to be either computationally demanding or reliant on large amounts of training data.

In contrast, spectral-domain methods offer a simple and computationally efficient alternative for generating stationary sounds. A notable example is the method proposed in [7], which generates arbitrarily long audio signals by preserving the spectral magnitude of a short recording while randomizing the phase in the Fourier domain. The resulting signal maintains the spectral characteristics of the original sound while producing a perceptually plausible endless continuation.

The work was supported by the Czech Science Foundation (GAČR) Project No. 23-07294S. The author is grateful to NVIDIA for donation of the Titan XP graphic card, which has been used in this research. PhD study of Michal Švento is supervised by Pavel Rajmic. The author thanks Pavel Rajmic, Ondřej Mokrý and Eloi Moliner for their valuable comments and improvements.

While such a principle was originally demonstrated for extending stationary environmental and ambient sounds, the underlying mechanism of decoupling spectral magnitude from phase is not inherently limited to this setting. In this work, we investigate whether the same approach can be applied to tasks that involve bridging two audio segments rather than extending a single one. Specifically, we investigate its use in audio inpainting, where a missing segment must be reconstructed between two known contexts, and in the creative merging of unrelated audio clips. These applications demonstrate that simple spectral-domain manipulation can produce perceptually convincing transitions at a negligible computational cost.

This paper is organized as follows. Sec. II reviews the method for generating endless audio signals from a limited context. Sec. III introduces the proposed application scenarios and describes how the principle extends to each task. Finally, Sec. V and VI provide a discussion of the approach and conclude the paper.

II. Endless Sounds

Assume a single-channel discrete-time signal $\mathbf{x} \in \mathbb{R}^L$ of length L. Using the Discrete Fourier Transform (DFT), the spectrum of the signal can be analysed and its significant frequency bins identified. The spectrum $\mathbf{X} \in \mathbb{C}^L$ for audio signals comes typically in the polar form:

$$\mathbf{X}[k] = |\mathbf{X}[k]|\, \mathrm{e}^{\mathrm{j}\boldsymbol{\theta}[k]}, \tag{1}$$

where $|\mathbf{X}[k]|$ is the magnitude and $\boldsymbol{\theta}[k]$ is the phase at bin k.

In [7], a computationally efficient method was proposed for extending a stationary portion of an audio signal by exploiting the Fast Fourier Transform (FFT) algorithm, denoted as $\mathcal{F}\{\cdot\}$. Their approach proceeds as follows. The original signal $\mathbf{x}$ is zero-padded to a new target length M. If M is not a power of two it is rounded up to the nearest such value, since the FFT operates most efficiently on power-of-two lengths [8]. The spectrum $\hat{\mathbf{X}}$, computed from the zero-padded signal, $\hat{\mathbf{x}}$, is then split into its magnitude $|\hat{\mathbf{X}}|$ and phase $\boldsymbol{\theta}_{\hat{\mathbf{x}}}$ components. To generate the endless sound, the magnitude $|\hat{\mathbf{X}}|$ is preserved while the original phase $\boldsymbol{\theta}$ is replaced by a random phase $\boldsymbol{\phi}$ drawn uniformly from $[-\pi, \pi)$:

$$\hat{\mathbf{Y}}[k] = |\hat{\mathbf{X}}[k]|\, \mathrm{e}^{\mathrm{j}\boldsymbol{\phi}[k]}, \qquad \boldsymbol{\phi}[k] \sim \mathcal{U}(-\pi, \pi). \tag{2}$$

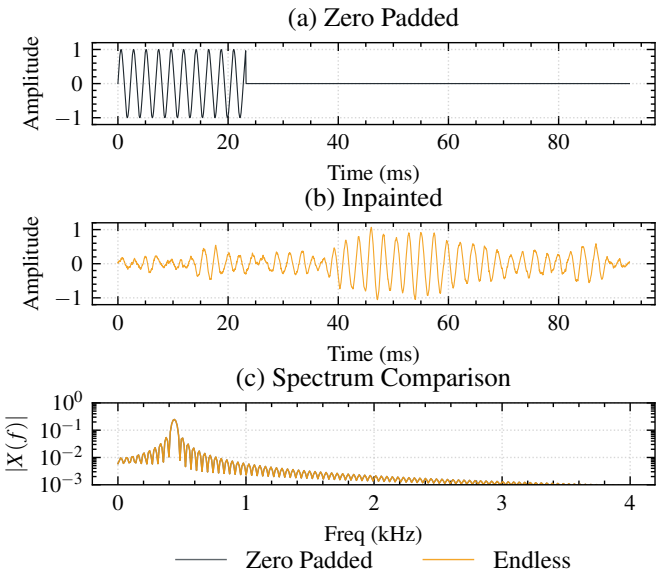

Fig. 1. Creating endless sounds. (a) Zero-padded signal, (b) Extended signal obtained by phase randomisation, (c) identical spectra of zero-padded and endless signals.

The output signal $\mathbf{y}$ is finally obtained by applying the inverse FFT to $\hat{\mathbf{Y}}$. A comparison of the original and extended signals together with their spectra is shown in Fig. 1.

This method was originally proposed for non-tonal sounds. A key property of the DFT is that it implicitly assumes the analysed signal to be periodic. When applied to finite-length real-world signals, this assumption introduces a discontinuity at the signal boundaries during periodisation, known as the *wrap-around effect* [9].

For noise-like signals, this artefact is usually less noticeable because their periodisation does not produce a strong transient. In contrast, harmonic and generally periodic signals, such as music, make the discontinuity more perceptible. A phase mismatch at the signal boundaries produces a click that, after zero-padding, spreads across the spectrum and becomes audible as a high-frequency noise.

This effect can be substantially reduced by tapering the signal prior to zero-padding using a standard analysis window, such as the Hann or Hamming window [10]. Windowing smooths the signal boundaries and suppresses the periodisation artefact, but attenuates the signal energy near the edges. When the extended signal is used in subsequent processing stages, this energy loss must therefore be compensated, as discussed in Section III.

III. Audio Inpainting via Endless Sounds

Audio inpainting is the task of reconstructing a missing or corrupted segment within a longer recording [11]–[13]. We propose a signal inpainting algorithm that achieves this by synthesising two *endless* continuations, one grown forward from the available left context and one grown backward from the right context and blending them inside the gap.

Given a signal $\mathbf{x}$ of length N with a gap of length G samples, left and right context segments are extracted of length $S \leq N$ adjacent to the gap. The pipeline consists of three main stages: (i) statistically consistent signal extension using a spectral phase-randomisation technique, (ii) energy normalisation and time-alignment of each extension to its neighbouring context, and (iii) crossfading to produce a perceptually seamless reconstruction. This pipeline is shown on the test signal in Fig. 2 and as a diagram in Fig. 3.

A. Endless Sounds for Signal Extension

Using the method described in Section II, the context on each side of the gap is independently extended. Here, $\mathbf{x}_c \in \mathbb{R}^S$ denotes the context segment cropped from the original signal adjacent to the gap boundary. It is windowed by a Hamming window, zero-padded to the nearest power-of-two length $\hat{M} = 2^{\lceil \log_2 M \rceil}$, and extended via Eq. (2). The resulting endless extension is denoted $\hat{\mathbf{x}}_L$ for the left side and $\hat{\mathbf{x}}_R$ for the right side, where each $\hat{\mathbf{x}}$ is thus a synthesised continuation of its respective $\mathbf{x}_c$. Because the phases of each extension are drawn independently and uniformly from $[-\pi, \pi)$, the two signals are mutually uncorrelated. The choice of context length S directly affects the spectral quality of each extension and is discussed in Section V.

B. Energy Normalisation and Time Alignment

A practical limitation of spectral randomisation arises from the windowing step. It reduces the contribution of samples near the gap boundary, precisely where accurate reconstruction is most critical. Furthermore, phase randomisation does not strictly preserve short-time energy. Therefore, an energy-normalisation step is applied before crossfading.

To compensate the windowing effect, S context samples from $\mathbf{x}_c$ nearest to the boundary are selected. A scalar gain σ is then derived as:

$$\sigma = \sqrt{\frac{\sum_{s=0}^{S-1} \hat{\mathbf{x}}^2[s]}{\sum_{s=0}^{S-1} \mathbf{v}[s]\, \mathbf{x}_c^2[s]}}, \tag{3}$$

where $\mathbf{v}$ is a linearly increasing ramp that places greater weight on samples close to the gap boundary. The synthesised extension $\hat{\mathbf{x}}$ is subsequently normalised by this gain via $\hat{\mathbf{x}} = \hat{\mathbf{x}}/\sigma$.

Because the extension is generated with an arbitrary phase offset, a coarse time-alignment step identifies the shift τ^* that maximises the cross-correlation between the end of the available context and the beginning of the extension:

$$\tau^* = \arg\max_{\tau} \left[\mathcal{F}^{-1}\{\mathcal{F}\{\mathbf{x}_c\} \cdot \overline{\mathcal{F}\{\hat{\mathbf{x}}\}}\}\right](\tau). \tag{4}$$

The extension is time-shifted by τ^* prior to merging to avoid phasing artefacts.

Analogous operations are applied to the right context, with the direction of alignment and the ramp window reversed, finally yielding a left extension $\mathbf{y}_L$ and a right extension $\mathbf{y}_R$.

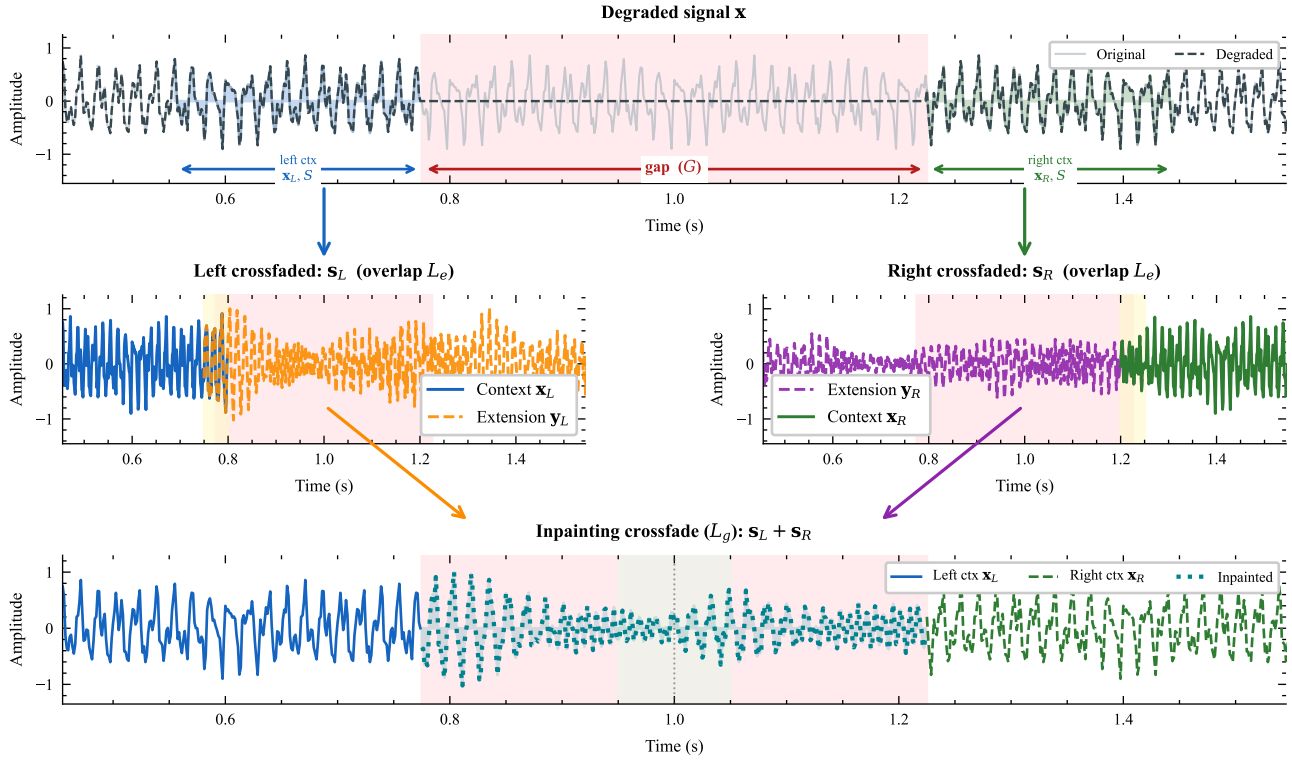

Fig. 2. Audio inpainting pipeline applied to a sustained tone with a gap of length G. (Top) The degraded signal with the left context $\mathbf{x}_L$ and right context $\mathbf{x}_R$ adjacent to the gap. (Middle left) Left-side crossfade $\mathbf{s}_L$, where the left context transitions smoothly into the endless extension $\mathbf{y}_L$ over L_e samples. (Middle right) Right-side crossfade $\mathbf{s}_R$, where the right context and extension $\mathbf{y}_R$ are blended symmetrically in reversed time. (Bottom) Final inpainting crossfade combining $\mathbf{s}_L$ and $\mathbf{s}_R$ over L_g samples inside the gap (shaded region). The inpainted signal closely follows the spectral envelope of the original.

C. Crossfading and Final Reconstruction

With the two normalised and aligned extensions ($\mathbf{y}_L$ and $\mathbf{y}_R$) at hand, the reconstruction proceeds via two successive crossfading phases. Both steps rely on the same underlying principle: an *equal-power* fade, in which the crossfade weights $\mathbf{w}_1[n]$ and $\mathbf{w}_2[n]$ satisfy $\mathbf{w}_1[n]^2 + \mathbf{w}_2[n]^2 = 1$ at every sample n. This condition ensures that the total signal power remains constant throughout the transition, preserving perceived loudness even when the two signals being blended are uncorrelated [14]. Specifically, the weights are drawn from a cosine window,

$$\mathbf{w}_1[n] = \cos\Big(\tfrac{\pi}{2}\,\tfrac{n}{L}\Big), \qquad \mathbf{w}_2[n] = \sin\Big(\tfrac{\pi}{2}\,\tfrac{n}{L}\Big), \tag{5}$$

over a transition of L samples, where one signal fades out while the other fades in.

a) Context–Extension Blending: The first pass applies this equal-power fade over L_e samples to blend each generated extension with its neighbouring reliable context. For the left side, this produces a signal $\mathbf{s}_L$ that begins with the original waveform $\mathbf{x}_L$ and transitions smoothly into the left extension $\mathbf{y}_L$. The right side is handled symmetrically in the reversed time direction to yield $\mathbf{s}_R$. Because the fade is equal-power, the blend remains perceptually seamless even at the boundary where synthesised content meets recorded audio.

b) Inpainting Crossfade: The second pass combines $\mathbf{s}_L$ and $\mathbf{s}_R$ inside the gap itself, again using the equal-power cosine window, now over an overlap of L_g samples:

$$\hat{\mathbf{x}}[n] = \mathbf{w}_1[n]\,\mathbf{s}_L[n] + \mathbf{w}_2[n]\,\mathbf{s}_R[n], \quad n = 0, \ldots, L_g - 1. \tag{6}$$

That results as inpainted signal from two, left and right, contexts.

IV. Creative Segment Stitching

The proposed pipeline extends naturally to the creative merging of two unrelated audio clips that share no spectral content. The goal is to synthetise a perceptually plausible transition between two arbitrary sounds $\mathbf{x}_A$ and $\mathbf{x}_B$.

The procedure is identical to the inpainting case. An endless continuation is grown forward from the tail of $\mathbf{x}_A$ and backward from the head of $\mathbf{x}_B$, both energy-normalised. This two extensions are blended inside a transition region of length G via the equal-power crossfade as described in Section III-C. Fig. 4 shows spectra of two unrelated clips and the synthesised transition between them.

V. Discussion

The quality of the reconstruction depends heavily on the stationarity of the signal within the context window and therefore on the choice of context length S. For quasi-stationary signals such as sustained instrument tones or ambient noise,

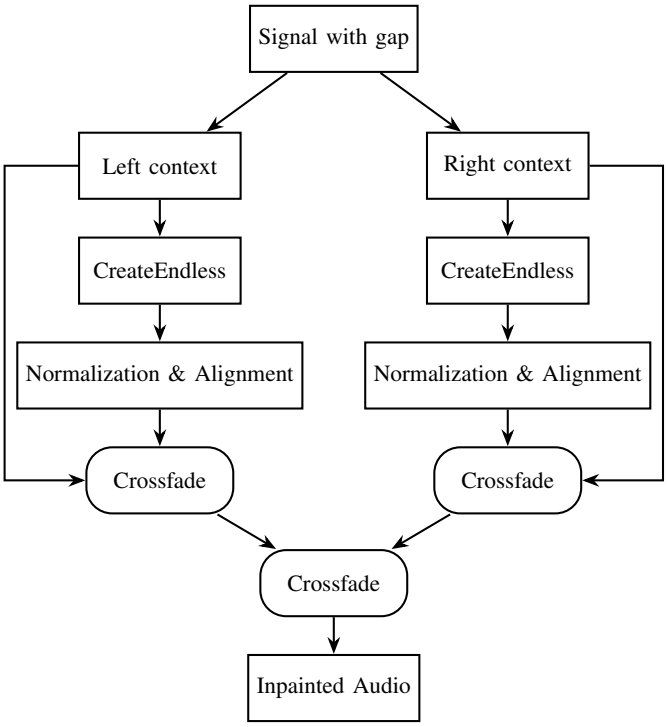

Fig. 3. Audio inpainting pipeline illustrated for two audio segments (one gap in the signal).

the method performs well, as the power spectrum remains stable over the context duration. Conversely, for highly non-stationary material with rapid transients, speech with strong pitch modulation, or music with fast harmonic changes, the spectral envelope is a mixture of all components.

The choice of context length S involves a further trade-off. A longer context provides a more representative spectral envelope, but increases the risk of including non-stationary content near the gap boundary, such as transients or harmonic changes, which would blur the estimated spectrum. A shorter context avoids this but may underrepresent low-frequency components that require many cycles to be reliably captured. An adaptive, signal-dependent choice of S is therefore a natural direction for future work.

A practical limitation is that perceptual quality degrades as the gap length grows relative to the context length S. For very long gaps, the two independently synthesised extensions share no temporal structure, and the crossfade at the centre of the gap may produce a subtle energy dip or textural discontinuity. As a practical guideline, the method is most reliable when the gap length does not exceed S.

A key advantage of the proposed approach is its computational cost. The entire pipeline reduces to a small number of FFT operations, a cross-correlation (again only 3 FFTs), and a pointwise multiply-accumulate for the crossfades. This makes it orders of magnitude cheaper than data-driven methods that require extensive training and high inference times, making it suitable for real-time or resource-constrained applications such as embedded audio or live performance tools, where deep learning inference is usually impractical.

VI. Conclusion and Future work

In this paper, a computationally efficient method for audio inpainting and segment stitching was proposed, building on the concept of endless sounds. By synthesising continuations from both sides of a gap and blending them via energy-normalised crossfading, the approach produces perceptually seamless transitions at negligible computational cost. The method is best suited for stationary and quasi-stationary signals such as ambient textures and sustained tones. Future work will explore adaptive context selection, perceptual evaluation against deep learning baselines, and extensions to multi-channel audio. The source code, audio examples, and setup instructions are publicly available at https://gitlab.utko.feec.vutbr.cz/xsvent00/endless_inpainting.

References

[1] M. Fröjd and A. Horner, "Sound texture synthesis using an overlap-add/granular synthesis approach," *AES: Journal of the Audio Engineering Society*, vol. 57, pp. 29–37, 01 2009.

[2] D. Schwarz, "State of the art in sound texture synthesis," in *Proceedings of the 14th International Conference on Digital Audio Effects (DAFx-11)*, Sep. 2011.

[3] E. Gutiérrez, F. Font, X. Serra, and L. Wyse, "A statistics-driven differentiable approach for sound texture synthesis and analysis," in *Proceedings of the 28th International Conference on Digital Audio Effects (DAFx25)*, Ancona, Italy, September 2025, 2–5 September 2025.

[4] C. Hawthorne, A. Jaegle, C. Cangea, S. Borgeaud, C. Nash, M. Malinowski, S. Dieleman, O. Vinyals, M. Botvinick, I. Simon *et al.*, "General-purpose, long-context autoregressive modeling with perceiver ar," in *Proceedings of the 39th International Conference on Machine Learning*, 2022. [Online]. Available: https://arxiv.org/abs/2202.07765

[5] J. Engel, C. Resnick, A. Roberts, S. Dieleman, M. Norouzi, D. Eck, and K. Simonyan, "Neural audio synthesis of musical notes with wavenet autoencoders," in *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ser. ICML'17. JMLR.org, 2017, p. 1068–1077.

[6] J. M. Antognini, M. Hoffman, and R. J. Weiss, "Audio texture synthesis with random neural networks: Improving diversity and quality," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 3587–3591.

[7] V. Välimäki, J. Rämö, and F. Esqueda, "Creating endless sounds," *Proceedings of the International Conference on Digital Audio Effects, DAFx*, p. 32 – 39, 2018.

[8] J. W. Cooley and J. W. Tukey, "An algorithm for the machine calculation of complex Fourier series," *Mathematics of Computation*, vol. 19, no. 90, pp. 297–301, 1965. [Online]. Available: https://www.ams.org/mcom/1965-19-090/S0025-5718-1965-0178586-1/

[9] J. D. Reiss and A. P. McPherson, *Audio Effects: Theory, Implementation and Application*. Routledge, Jan. 2026. [Online]. Available: http://dx.doi.org/10.1201/9781003593942

[10] F. J. Harris, "On the use of windows for harmonic analysis with the discrete Fourier transform," *Proceedings of the IEEE*, vol. 66, no. 1, pp. 51–83, 1978.

[11] E. Moliner and V. Välimäki, "Diffusion-based audio inpainting," *Journal of the Audio Engineering Society*, vol. 72, no. 3, pp. 100–113, Mar. 2024.

[12] O. Mokrý and P. Rajmic, "Audio inpainting: Revisited and reweighted," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2906–2918, 2020.

[13] O. Mokrý, P. Záviška, P. Rajmic, and V. Veselý, "Introducing SPAIN (SParse Audio INpainter)," in *2019 27th European Signal Processing Conference (EUSIPCO)*. IEEE, 2019.

[14] V. Pulkki and M. Karjalainen, *Communication Acoustics*. Wiley, Jan. 2015.

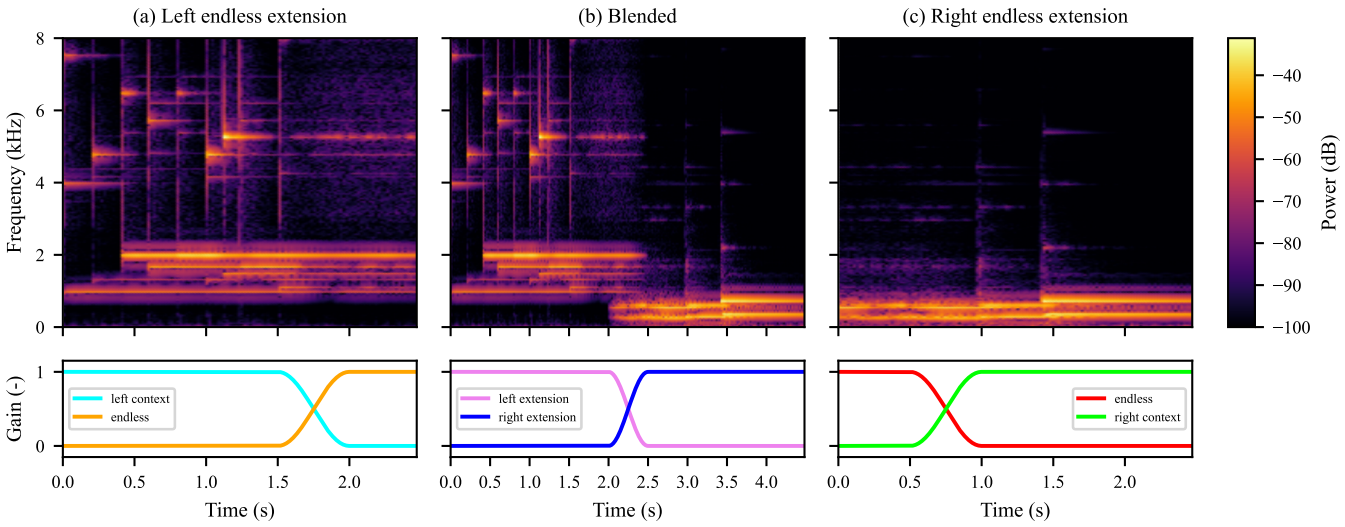

Fig. 4. Spectrograms and gain envelopes illustrating the stitching of two unrelated audio clips. Panel **(a)** shows $\mathbf{x}_A$ smoothly fading into its endless extension, and panel **(c)** shows $\mathbf{x}_B$ fading into its extension. Panel **(b)** shows the final stitched transition obtained by equal-power crossfading between the two generated extensions over a window of length G. The bottom plots show the corresponding mixing gains over the transition window.

Gray-Box Modeling of Flanger Effect

David Leitgeb
Department of Telecommunications
Brno University of Technology, Technicka 12
616 00 Brno, Czech Republic
david.leitgeb@vut.cz

***Abstract*—This paper presents a gray-box neural network model of the flanger audio effect. The proposed architecture estimates the time-variant transfer function of the flanger effect and applies it to the Short-Time Fourier Transform (STFT) of the input signal to produce the output signal. The model was implemented in the PyTorch framework and evaluated using the error-to-signal ratio (ESR) metric, with results compared to a previously proposed black-box approach. Experimental results show that the model is capable of reproducing the typical comb-filter character of the flanger effect while achieving competitive performance for the considered parameter settings.**

***Index Terms*—gray-box modeling, audio effects modeling, flanger effect, neural networks, digital signal processing**

I. Introduction

Machine learning (ML) has seen a significant increase in use in recent years across various fields, including digital audio processing. Several areas of audio processing have explored ML-based approaches, most notably nonlinear audio effects modeling, virtual analog (VA) modeling, and musical instrument modeling. A key advantage of ML-based methods is their ability to model complex audio effects that are difficult to implement using traditional signal processing techniques. However, achieving sufficiently accurate results often requires a large amount of training data.

Generally, the use of ML in audio modeling can be divided into three main approaches: white-box, black-box, and gray-box. White-box models allow accurate recreation of various audio effects. However, detailed knowledge of the effect and its internal structure is required. In contrast, black-box models typically rely solely on the input-output relationship and do not incorporate any information about the internal structure of the effect. Finally, gray-box models represent a combination of these two approaches and require partial knowledge of the system structure together with input-output audio data pairs. Depending on the audio effect being modeled, an appropriate neural network architecture must be designed in order to minimize the required dataset size and computational complexity while maintaining accurate results.

Various neural network architectures have been used to model different types of audio effects. These effects and their underlying signal processing algorithms are described in [1]. Depending on the type of effect being modeled, different challenges may arise. This paper focuses on modeling time-varying audio effects, such as phasers and flangers, which are commonly used in music production. Their distinguishing characteristic is the modulation of specific parameters using a low-frequency oscillator (LFO) [2]. This modulation can occur over time intervals of several seconds, representing a significant challenge during training. To address this issue, a frame-based approach can be employed [3].

The primary goal of this paper is to implement a model architecture that can serve as a baseline for future work and experiments. The proposed model focuses specifically on the flanger effect. The architecture is largely based on the model presented in [4]. However, while that model follows a black-box approach and predicts a transfer function in the form of a complex matrix, the architecture proposed in this paper adopts a gray-box design by estimating the parameters of the underlying flanger structure. As the authors of [4] do not provide publicly available source code, the implementation was carried out independently. Certain aspects of the proposed model were also adopted from [3]. The model was implemented using the PyTorch framework[1]. Audio examples corresponding to the experiments presented in this paper are provided on a dedicated GitHub page[2], where the target and predicted signals can be directly compared by listening.

This paper is organized as follows. Section II describes the methodology and implementation of the proposed model. Section III presents selected experiments conducted using the trained model. Section IV summarizes the results and concludes the paper.

II. Methods

This section describes the methodology used in this work. It provides a detailed overview of the flanger effect used in the experiments, as well as the training signal and dataset structure. The proposed model architecture and the training procedure are described in detail.

A. Flanger Implementation

In this work, the proposed model focuses on the flanger effect. Both flangers and phasers share a characteristic sound that is recognizable by a comb-filter effect in the spectrum of the processed signal. A phaser achieves this effect using a cascade of notch filters, whereas a flanger forms a comb filter by summing a delayed copy of the input signal with the original signal. Compared to a phaser, this results in a significantly larger number of minima in the audible frequency spectrum.

[1]https://pytorch.org/
[2]https://david-leitgeb.github.io/graybox-flanger/

The training dataset was generated using a simple flanger digital signal processing (DSP) implementation according to [1]:

$$\begin{aligned} x_{\mathrm{d}}(n) &= x(n) + g_{\mathrm{fb}}\, x_{\mathrm{d}}\big(n - M(n)\big), \\ y(n) &= (1-g)\, x(n) + g\, x_{\mathrm{d}}\big(n - M(n)\big), \end{aligned} \tag{1}$$

where $x(n)$ denotes the input signal, g_{fb} is the feedback gain, g is the mix parameter, $x_{\mathrm{d}}(n)$ represents the internal delay-line signal, $M(n)$ is the delay modulation signal generated by the LFO, and $y(n)$ is the output signal. Since the delay parameter is modulated by the LFO, the delay line requires fractional-delay processing in order to avoid audible artifacts. Therefore, the DSP implementation also includes second-order polynomial sample interpolation according to [2].

This DSP implementation was chosen in order to automate the dataset preparation process. Tools such as [5] could further automate this process by allowing audio effects implemented as plug-ins in the VST format to be used for dataset generation.

B. Training Signal

The training signal was constructed according to [3], [4]. It consists of a chirp signal with a burst repetition period of 40 ms. Originally, [3] proposed a period of 30 ms, but [4] increased this value to 40 ms without further explanation. In order to match and compare the results, the same value was used. The total length of the training signal is 6.67 seconds.

The signal is then processed by a cascade of 64 first-order all-pass filters. The transfer function of each filter is given by [3]:

$$A(z) = \frac{p - z^{-1}}{1 - pz^{-1}}, \tag{2}$$

where $p = 0.9$ is the filter coefficient. According to [3], only a few seconds of this signal should be sufficient to train the model due to its spectral flatness. More details about this type of signal can be found in [6].

C. Dataset

One of the main advantages of the proposed gray-box approach is that the model estimates the time-varying transfer function of the system rather than learning a direct input-output mapping, reducing the amount of data required for training. The dataset used for training consists of 12 processed signals and a single unprocessed signal. The only parameters that are varied between the individual audio samples are the LFO frequency (also commonly referred to as the rate) and the feedback gain. The parameter values used to generate the dataset are shown in Table I.

TABLE I
FLANGER PARAMETER VALUES USED FOR THE DATASET.

LFO rate [Hz]	0.23, 0.73, 1.13
Feedback [%]	0, 25, 50, 75
Delay [ms]	1.5
Mix [%]	70

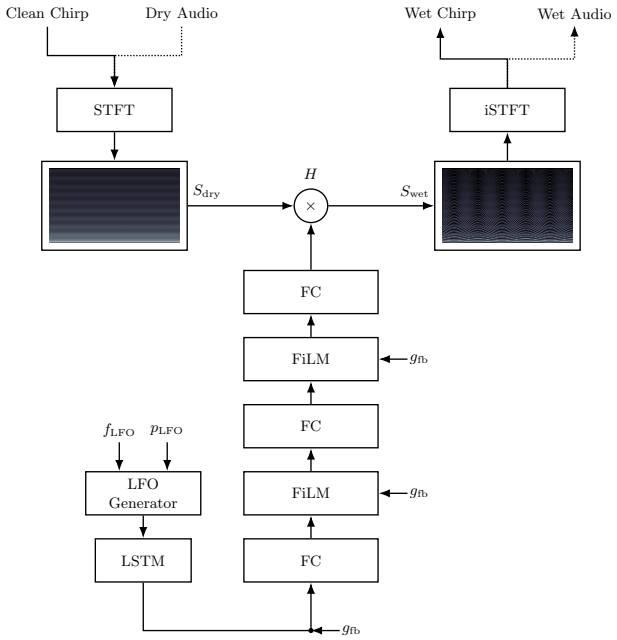

Fig. 1. Proposed model architecture based on [4]. Solid and dotted lines indicate STFT input and iSTFT output data used during training and inference, respectively. Parameters f_{LFO}, p_{LFO}, and g_{fb} are used in both stages.

D. Model Architecture

As stated earlier, the model architecture proposed in this work is based on the architecture presented in [4], where a black-box neural network model, referred to as CONMOD, is introduced. It utilizes the Feature-wise Linear Modulation (FiLM) layer proposed by [7]. In [4], the FiLM layers were used to enable effect embedding and allow the model to be trained on multiple effects simultaneously. However, this functionality is not used in the model proposed in this paper.

The structure of the model can be visually divided into three parts, see Fig. 1. The first part performs the Short-Time Fourier Transform (STFT) of the unprocessed signal and the inverse Short-Time Fourier Transform (iSTFT) of the processed signal. Between these two operations, a multiplication with the predicted transfer function H is performed:

$$S_{\mathrm{wet}}(f,t) = H(f,t)\, S_{\mathrm{dry}}(f,t), \tag{3}$$

where S_{dry} is the STFT matrix of the input signal, H is the predicted transfer function, and S_{wet} is the STFT matrix of the predicted output signal.

The transfer function is predicted by an MLP block consisting of fully connected (FC) and FiLM layers. The final part of the architecture consists of an LFO generator and an LSTM module. The LFO generator takes the LFO rate (f_{LFO}) and phase (p_{LFO}) as inputs, and the generated LFO signal is then used as input to the LSTM. The output of the LSTM is subsequently passed to the first FC layer.

During the implementation process, the feedback value (g_{fb}) was also added as a concatenated input to the first FC layer in order to stabilize training, which previously failed

TABLE II
TRAINING PARAMETERS.

Training:	
Epochs	10,000
Batch size	1
Optimizer	Adam
LR exponential decay	0.9997
Dataset:	
Sample rate	44.1 kHz
STFT:	
FFT size	4096
Frame size	1764
Hop size	441
Model:	
LSTM hidden size	32
MLP hidden size	512

to converge. This behavior was unexpected, since the model proposed in [4] uses FiLM layers to apply conditioning based on the feedback parameter. The most likely explanation is a discrepancy between the implementations, although this was not fully investigated at the time of writing this paper. Since the feedback parameter is provided directly as an input to the FC layer, it was also considered whether the FiLM layers were necessary. However, experiments with the FiLM layers disabled showed that they contribute significantly to the learning process, and without them, the model failed to converge.

E. Model Training

The parameter values used during training are listed in Table II. These values were selected according to [3], [4].

The model was trained on an NVIDIA RTX 4070 GPU, and training took approximately 40 minutes. During the implementation testing phase, 5000 epochs were sufficient to achieve acceptable accuracy. However, the number of epochs was increased in order to enable comparison with the results reported in [4] in Section III-B.

After implementing the model, an appropriate loss function had to be selected for training. As a starting point, the training objective was chosen according to [3], [4]. The first paper uses exclusively the error-to-signal ratio (ESR), while the second additionally employs the multi-resolution spectral loss (MRSL). These loss functions are defined as follows:

$$\mathcal{L}_{\mathrm{ESR}}(y, \hat{y}) = \frac{\sum_{l=0}^{L-1}(y[l] - \hat{y}[l])^2}{\sum_{l=0}^{L-1} y[l]^2}, \tag{4}$$

where $y[l]$ is the target signal, $\hat{y}[l]$ is the predicted signal, and L is the length of the training data in samples.

$$\mathcal{L}_{\mathrm{MRSL}} = \sum_i \|S_i - \hat{S}_i\|_1 + \|\log S_i - \log \hat{S}_i\|_1, \tag{5}$$

where S_i is the STFT of the target signal, $\hat{S}_i$ is the STFT of the predicted signal, and i indexes the different STFT resolutions corresponding to FFT sizes of 512, 1024, and 2048. The STFT overlap was set to 75 %. The final training objective is defined as a weighted combination of these two loss terms:

$$\mathcal{L}(y, \hat{y}) = \lambda \mathcal{L}_{\mathrm{ESR}} + \mathcal{L}_{\mathrm{MRSL}}, \tag{6}$$

where λ is the weight of the ESR loss term. During training, several values of this parameter were tested, and the value $\lambda = 2$ was selected for the model evaluated in this paper. Other loss functions were also explored during the implementation process, but the final model was trained using the objective defined above. With this configuration, the model achieved sufficient accuracy while maintaining relatively low computational requirements.

F. Transfer Function

The model proposed in [4] predicts the transfer function in the form of a complex matrix that is multiplied element-wise with the STFT matrix of the dry signal. Following the gray-box approach, the model proposed in this paper instead predicts the parameters of the internal structure of the flanger effect, corresponding to the following transfer function [1]:

$$H(e^{j\omega}, n) = (1 - g) + g \frac{e^{-j\omega M(n)}}{1 - g_{\mathrm{fb}} e^{j\omega M(n)}}, \tag{7}$$

where g is the mix parameter, $M(n)$ is the delay modulation signal generated by the LFO, and g_{fb} is the feedback gain. The predicted transfer function is then multiplied by the STFT of the dry signal as described in (3). After applying the inverse STFT, the predicted output signal is obtained.

III. EXPERIMENTS

The following section presents a selection of experiments conducted using the proposed model. These include an evaluation of the comb-filter character reproduced by the model, an evaluation based on ESR, and an analysis of the influence of sequence length on the ESR metric. Audio files corresponding to these experiments are available on the dedicated GitHub page.

A. Comb Filter

The first experiment evaluates whether the proposed model correctly reproduces the comb-filter character, which represents the defining characteristic of the flanger effect. The input signal used during evaluation was white noise. The LFO rate and feedback gain were set to 0.23 Hz and 75 %, respectively.

Fig. 2 and Fig. 3 show the spectrograms of the target and predicted signals. The predicted signal closely resembles the target signal, indicating that the model successfully captures the characteristic spectral character of the effect.

B. ESR Evaluation

One of the metrics used to evaluate the proposed model was ESR. Fig. 4 shows the ESR values for different combinations of LFO rate and feedback gain. The vertical axis corresponds to the LFO rate, while the horizontal axis represents the feedback gain. Axis labels in black correspond to parameter configurations used during training, while red labels represent unseen parameter combinations.

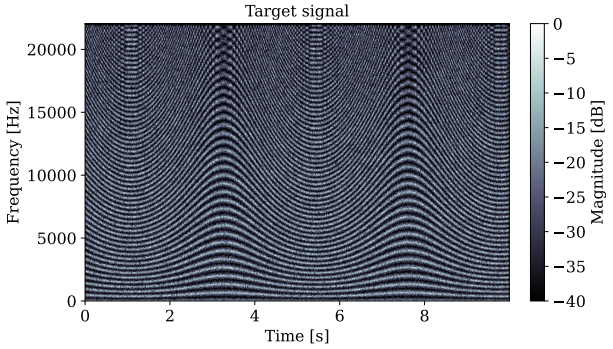

Fig. 2. Spectrogram of the target signal.

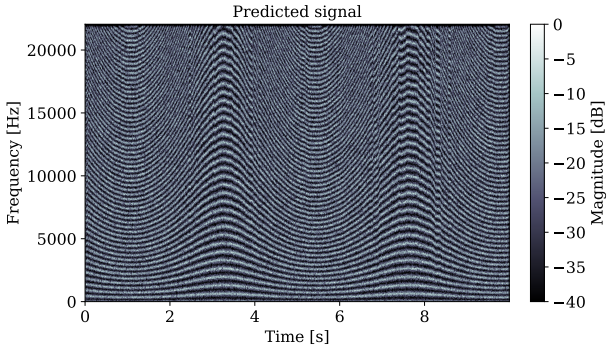

Fig. 3. Spectrogram of the predicted signal.

The evaluation was performed using the same 10 s DI guitar recording used in [4]. However, it should be noted that the flanger DSP implementation used to generate the dataset differs from the one used in the reference work. While [4] employs the Studio One[3] flanger effect, the dataset used in this work was generated using the simplified flanger implementation described in Section II-A. Therefore, the comparison should be interpreted with this difference in mind.

As shown in Fig. 4, the ESR increases for higher LFO rates and feedback gains. This behavior is expected, as stronger feedback and faster modulation lead to greater spectral variation, which makes estimation of the flanger parameters more challenging. It should also be noted that the LFO rate of 10 Hz lies far outside the training range, as the highest rate used during training was 1.13 Hz.

Table III compares the ESR values of the CONMOD model and the proposed method for selected LFO rate and feedback configurations. The proposed model achieves lower ESR than CONMOD for low-feedback settings across all considered LFO rates. As the feedback increases, the ESR rises for both methods, with a larger increase observed for the proposed model, especially at high feedback. Overall, the results indicate that the proposed gray-box approach performs well for the evaluated parameter values, while strong feedback combined with fast modulation remains the most challenging scenario.

[3]http://www.presonus.com/studioone

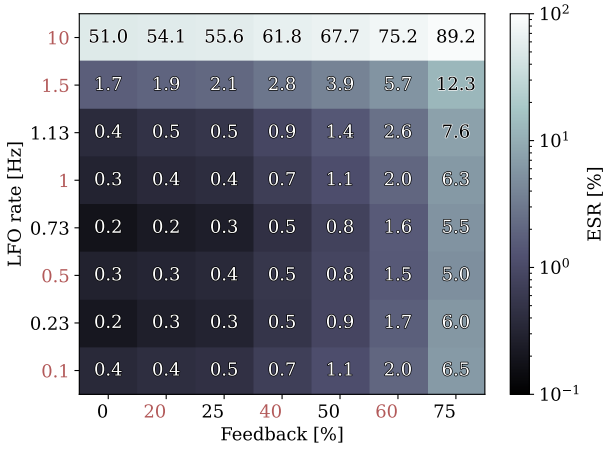

Fig. 4. ESR values for different LFO rate and feedback gain combinations. Axis labels indicate seen (black) and unseen (red) parameter settings.

TABLE III
COMPARISON OF ESR RESULTS FOR THE PROPOSED MODEL AND THE CONMOD MODEL.

LFO rate [Hz]	Feedback [%]	ESR [%] CONMOD [4]	ESR [%] Gray-box
0.23	0	1.7	0.2
	25	1.0	0.3
	50	1.1	0.9
	75	4.9	6.0
0.73	0	1.2	0.2
	25	0.6	0.3
	50	0.7	0.8
	75	4.3	5.5
1.13	0	1.1	0.4
	25	0.6	0.5
	50	1.0	1.4
	75	5.7	7.6

C. ESR Dependence on Sequence Length

Another experiment evaluated how the ESR changes with the length of the evaluated signal for different parameter configurations. The results are shown in Fig. 5. The test signal was a one-minute guitar recording obtained from [3].

As expected, higher feedback values result in larger ESR. For both LFO rates, the configurations with feedback = 75 % exhibit significantly higher ESR than the corresponding configurations with feedback = 0 %. This behavior is consistent with the results presented in the previous experiment.

For configurations with feedback = 0 %, the ESR remains relatively low for different sequence lengths. In particular, the configuration with LFO rate = 0.23 Hz and feedback = 0 % shows a gradual decrease of ESR as the sequence length increases. This behavior may be related to the long period of the modulation signal. By comparison, the configuration with LFO rate = 1.13 Hz and feedback = 0 % shows a slight increase of ESR as the sequence length increases, although the change remains relatively small.

In contrast, the configuration with LFO rate = 1.13 Hz and feedback = 75 % remains relatively stable as the sequence length increases. This suggests that the error is dominated by

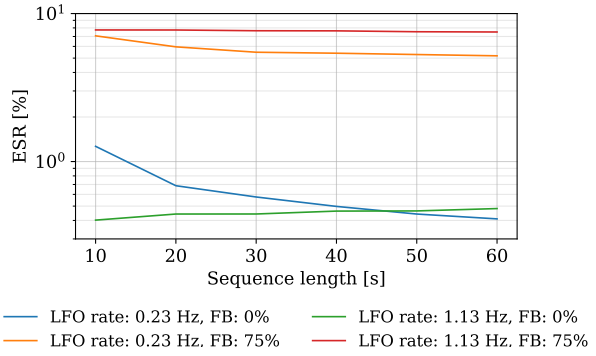

Fig. 5. ESR dependence on sequence length for different LFO rate and feedback configurations.

the difficulty of accurately modeling the feedback behavior rather than by the length of the evaluated signal.

IV. Conclusion

The main outcome of this work is a gray-box neural network model of the flanger effect. The proposed architecture estimates the time-varying transfer function of the flanger effect, which is then applied to the STFT of the dry signal to obtain the predicted output signal. The model was evaluated using the ESR metric and compared with the black-box approach proposed in [4].

The experimental results show that the proposed model is capable of reproducing the characteristic comb-filter structure of the flanger effect and achieves competitive performance for the evaluated parameter values. However, higher feedback values and faster modulation rates remain more challenging, which is reflected in increased ESR values in these configurations.

Future work will focus on incorporating perceptual evaluation of the proposed model. While objective metrics such as PEMO-Q [8] could be considered, they are not specifically designed for this type of task and may not provide reliable results. Therefore, the focus will be on subjective listening tests, such as ABC tests [9] and MUSHRA [10]. Additional work may include improving the robustness of the model for strong feedback settings, extending the approach to additional flanger implementations, and investigating potential adjustments to the training parameters.

Audio files corresponding to the experiments presented in this paper are available on the dedicated GitHub page.

Acknowledgment

I would like to thank my supervisor doc. Ing. Jiří Schimmel, Ph.D. for guidance during the preparation of this paper.

References

[1] U. Zölzer, *DAFX: Digital Audio Effects*, 2nd ed. Chichester, West Sussex, England: Wiley, 2011.

[2] J. D. Reiss and A. P. McPherson, *Audio Effects: Theory, Implementation and Application*. Hoboken: CRC Press, 2015.

[3] A. Carson, C. Valentini-Botinhao, S. King, and S. Bilbao, "Differentiable grey-box modelling of phaser effects using frame-based spectral processing," in *Proc. of the 26th Int. Conf. on Digital Audio Effects (DAFx23)*, Copenhagen, Denmark, 2023.

[4] G. Lee, H. Kim, J. Lee, and J. Nam, "CONMOD: Controllable Neural Frame-Based Modulation Effects," in *Proc. of the 27th Int. Conf. on Digital Audio Effects (DAFx24)*, E. De Sena and J. Mannall, Eds., Sept. 2024, pp. 9–16.

[5] A. Nizamutdinov, "Generování zvukových trénovacích dat pro neuronové sítě," Bachelor's thesis, Vysoké učení technické v Brně. Fakulta elektrotechniky a komunikačních technologií. Ústav telekomunikací, Brno, Czech Republic, 2025, accessed: Mar. 5, 2026. [Online]. Available: http://hdl.handle.net/11012/254318

[6] A. Wright and V. Välimäki, "Neural Modeling of Phaser and Flanging Effects," *Journal of the Audio Engineering Society*, vol. 69, no. 7/8, pp. 517–529, Jul. 2021.

[7] E. Perez, F. Strub, H. de Vries, V. Dumoulin, and A. Courville, "FiLM: Visual reasoning with a general conditioning layer," *Proc. of the AAAI Conf. on Artificial Intelligence*, vol. 32, no. 1, Apr. 2018.

[8] R. Huber and B. Kollmeier, "PEMO-Q—A new method for objective audio quality assessment using a model of auditory perception," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 6, pp. 1902–1911, 2006.

[9] ITU-R, "Methods for the subjective assessment of small impairments in audio systems (ABC)," International Telecommunication Union, Tech. Rep. BS.1116-3, 2015.

[10] ——, "Method for the subjective assessment of intermediate quality level of audio systems (MUSHRA)," International Telecommunication Union, Tech. Rep. BS.1534-3, 2015.

Defect-Aware Approach to Validation of MEMS Microphones in Spherical Arrays

Matej Liska
Department of Telecommunications
Brno University of Technology
Brno, Czech Republic
xliska19@vut.cz

***Abstract*—This paper presents defect-aware method for the systematic validation of MEMS-based spherical microphone arrays. Addressing the inherent discrepancies between idealized spatial models and physical hardware realizations, we propose a diagnostic framework designed to mitigate manufacturing stochasticity and electromechanical tolerances. The method employs a reference-calibrated and compensated technique, where each array element is normalized against a high-precision microphone to eliminate spectral bias from the measurement chain. An automated analytical pipeline evaluates these compensated responses to identify non-conforming units across critical spectral regions. By benchmarking individual microphone performance against an aggregate array-relative reference, the system isolates anomalies such as acoustic port occlusions, mechanical placement and sensitivity change. This approach facilitates precise sensitivity matching and provides a deterministic criterion for finding defective microphones, ensuring high precision in direction of acoustic axis and spatial sampling stability for advanced sound field analysis.**

***Index Terms*—defect detection, spherical microphone array, MEMS microphone arrays, manufacturing tolerances, anechoic measurement**

I. Introduction

The design and implementation of spherical microphone arrays (SMAs) constitute a critical research area across diverse disciplines, including architectural acoustics, adaptive beamforming, and high-order spatial audio rendering [1]. While the spatial configuration of transducers significantly dictates the accuracy of sound field reconstruction, traditional research predominantly focuses on identifying optimal sampling topologies, such as spherical t-designs, to achieve high-precision spherical harmonic analysis while minimizing system complexity [2] [3]. However, the transition from an idealized sampling model to physical hardware introduces significant challenges that are often overlooked in theoretical frameworks.

Although sampling theorems extended to the spherical domain [4] provide a mathematical basis for order-limited reconstructions [5], these models assume uniform microphone performance and deterministic geometric alignment. In practice, the assembly of MEMS-based arrays is susceptible to manufacturing stochasticity, including soldering inconsistencies and mechanical integration tolerances [6]. Such discrepancies lead to deviations in individual microphone responses, fundamentally compromising the spatial integrity of the captured sound field. Specifically, phenomena such as acoustic port occlusion—frequently caused by flux residue or mechanical misalignment during the integration process—result in substantial magnitude attenuation and phase aberrations.

Standard compensation routines [7], which typically address only uniform gain offsets, are insufficient for identifying these localized hardware failures. Furthermore, attempting to compensate for a critically defective microphone through traditional equalization can severely degrade the signal-to-noise ratio (SNR) by over-amplifying corrupted channels [8]. This paper proposes a defect-aware validation methodology that extends beyond conventional sensitivity matching. Utilizing measurements within an anechoic environment, an automated diagnostic sequence characterizes each transducer through reference-compensated analysis. The proposed framework systematically identifies and isolates hardware outliers, ensuring that the physical implementation maintains the spatial sampling performance intended by its theoretical design.

The paper is organized as follows: Section II describes the measurement setup and hardware implementation. Section III details the diagnostic and defect identification algorithm. Section IV presents the experimental results and the resulting calibration strategy, and Section V concludes the work.

II. System Architecture and Measurement Methodology

The experimental validation was conducted in an anechoic chamber using a logarithmic sine sweep from 20 Hz to 20 kHz. This configuration provides the baseline for the multi-channel acquisition and the subsequent diagnostic analysis.

A. Hardware Architecture

The SMA is based on configuration of TDM (Time-Division Multiplexed) MEMS microphones connected to Matlab software [9].

In comparison to previous research, the hardware was modified to utilizes a flexible-rigid PCB to daisy-chain the microphones. While this architecture allows for a compact 15 mm thick spherical shell, the high component density and the 23.2 mm maximum spacing between capsules introduce significant assembly challenges. Manual soldering of the MEMS microphone connectors within these spatial constraints is particularly prone to manufacturing imperfections, such

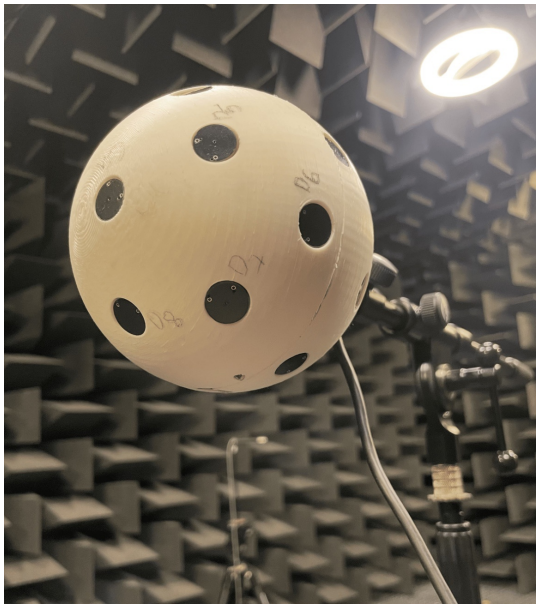

Fig. 1. Spherical microphone array setup.

as acoustic port occlusion or inconsistent mechanical stress, which necessitate the robust validation procedures analyzed in the subsequent sections.

B. Spherical Topology and Mechanical Design

The spatial distribution follows an optimized 25-point configuration, selected through a comparative analysis of spherical sampling techniques [10].

This topology demonstrates superior Euclidean distance distribution and spherical uniformity, which are critical for high-order sound field decomposition. The physical implementation consists of a 3D-printed spherical shell with an outer diameter of 112 mm as shown in Fig. 1.

The internal volume is optimized to house the flexible PCB while maintaining structural rigidity, ensuring that each microphone is mounted in a precise acoustic port corresponding to the theoretical sampling coordinates.

C. Calibration and Reference Compensation

To eliminate the spectral bias of the acoustic source and the measurement environment, a two-stage calibration procedure was implemented:

- Reference Calibration: A Brüel & Kjær Type 2671 reference microphone was calibrated using acoustics calibrator generating pressure of 1 Pa. The loudspeaker output was subsequently adjusted to maintain a constant 1 Pa pressure level at the reference position 1 m on axis from calibrated microphone.
- Source Compensation: The signal of the reference microphone were sampled across the entire frequency range from which the RMS level was calculated. The output of measurement was stored as the reference magnitude response $M_{ref}(f)$.

D. Spectral Analysis and Array Normalization

The diagnostic process concludes with a spectral normalization pipeline designed to isolate response of each microphone. Utilizing the previously recorded reference magnitude response $M_{ref}(f)$, the raw acoustic signals from the array are processed using a Fast Fourier Transform (FFT) with a window length N_{FFT} corresponding to the nearest power of two of the sample count to ensure computational efficiency.

To facilitate a direct comparison with the reference baseline, the magnitude spectrum $M_{raw}(f)$ of each microphone channel is resampled using linear interpolation to align with the reference frequency vector f_{ref}. The final compensated magnitude response $H_i(f)$ for the i-th microphone is derived by subtracting the reference magnitude $M_{ref}(f)$, effectively de-convolving the source [11] and environment characteristics:

$$H_i(f) = [M_{raw}(f) - M_{ref}(f)] + O_{cal}. \tag{1}$$

In this context, O_{cal} represents a localized calibration offset calculated at 1 kHz to normalize each microphone's sensitivity level. To mitigate stochastic measurement noise while preserving critical electromechanical features, a quadratic regression over each window smoothing (LOESS) algorithm [12] with a smoothing factor of 0.025 is applied to the compensated data.

This processed dataset serves as the primary input for the subsequent defect identification and statistical anomaly detection phase.

III. Diagnostic and Defect Identification Algorithm

The proposed diagnostic framework transitions from the normalized magnitude responses $H_i(f)$ to a validated microphone population through a multi-stage analytical pipeline. The core of the methodology lies in statistical evaluation of each microphone against the aggregate performance of microphone array, ensuring that the detection logic is adaptive to the specific acoustic loading and diffraction effects of the 112 mm spherical housing.

A. Aggregate Array Reference and Statistical Benchmarking

To establish an objective baseline for a functional microphone integrated into the 3D-printed shell, the algorithm calculates the mean magnitude response $\bar{H}(f)$ across the entire population. Since each microphone was sequentially positioned on-axis at a distance of 1 m from the acoustic source, the resulting $H_i(f)$ values are spatially invariant. Unlike idealized manufacturer specifications, this array-relative mean accounts for systemic deviations, such as the scattering and frequency-dependent shadowing inherent to the rigid spherical geometry. The reference baseline is defined as:

$$\bar{H}(f) = \frac{1}{N} \sum_{i=1}^{N} H_i(f) \tag{2}$$

To further quantify the manufacturing consistency, the algorithm computes the frequency-dependent standard deviation

$\sigma(f)$ [13], providing a metric for the spectral dispersion of the microphones:

$$\sigma(f) = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N} |H_i(f) - \bar{H}(f)|^2} \tag{3}$$

This measurement configuration allows $\sigma(f)$ to strictly represent manufacturing tolerances and mechanical integration inconsistencies, such as soldering defects or acoustic port misalignments, rather than spatial shadowing or angular sensitivity variations. The magnitude of $\sigma(f)$ typically increases at higher frequencies, where mechanical tolerances and the precision of the 15 mm acoustic port geometry have a more pronounced impact on the microphone's acoustic impedance.

B. Physical Defect Analysis

The identification of defective microphone is based on the spectral deviation $\Delta_i(f) = \bar{H}(f) - H_i(f)$. The pipeline evaluates $\Delta_i(f)$ across critical diagnostic windows, targeting distinct physical failure modes identified during the manual assembly of the high-density flexible PCB.

Detection of Acoustic Port Occlusion

The most critical defect identified in MEMS-based arrays is the partial or total occlusion of the acoustic inlet. Due to the high component density and the 23.2 mm maximum spacing between capsules, manual soldering is prone to depositing flux residue or soldering paste into the microphone port as shown in Fig. 2.

This obstruction acts as a first-order low-pass filter, where the effective diameter of the port is reduced. The diagnostic target is set at $f_{high} = 17$ kHz, where even minor occlusion result in a significant magnitude roll-off. A deterministic threshold $\Gamma_{HF} = 3$ dB is applied:

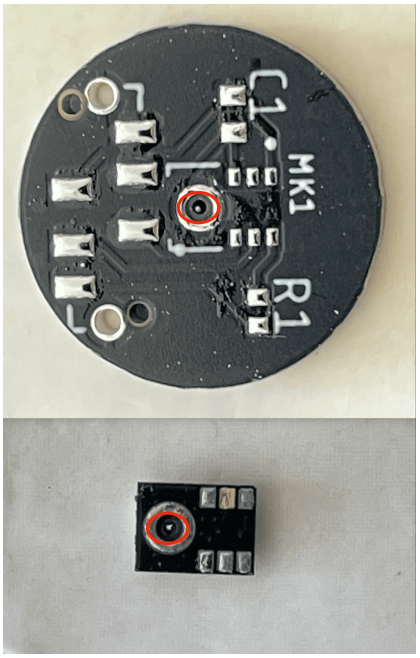

Fig. 2. Acoustic port occlusion.

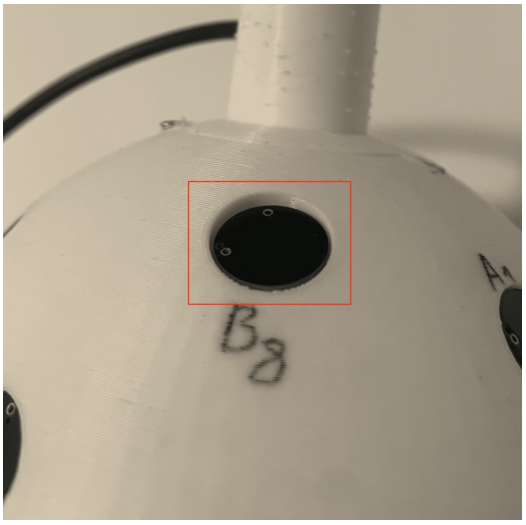

Fig. 3. Insufficiently sealed microphone PCB.

$$\text{Status}_{i,HF} = \begin{cases} \text{Defective}, & \text{if } \bar{H}(f_h) - H_i(f_h) > \Gamma_{HF} \\ \text{Functional}, & \text{otherwise} \end{cases} \tag{4}$$

Detection of Internal Cavity Resonances

The second defect analysis relates to acoustic coupling with internal standing waves of sphere. With an internal diameter of $D_{in} = 82$ mm, the air volume encapsulated by the 15 mm rigid shell acts as a spherical cavity resonator. According to the wave equation in spherical coordinates, the characteristic frequencies for these internal cavity modes are defined by the roots of the spherical Bessel functions. For the given geometry, the first and second modes occur at $f_{res} \approx 3.8$–4.8 kHz. If a microphone is not perfectly sealed from the internal volume, these standing wave patterns introduce localized spectral spikes and distortions, providing a deterministic indicator of mechanical assembly failure.

Ideally, each microphone is hermetically sealed from the internal volume by the PCB and mechanical seating. However, if a microphone is incorrectly oriented on the flexible PCB or insufficiently sealed as shown in Fig. 3, it becomes sensitive to internal pressure oscillations. Microphones exhibiting significant magnitude deviations in this band are identified as "leaky," as they respond to the internal resonance of the hollow sphere instead of the external incident sound field. This diagnostic window is critical for verifying the mechanical integrity of the 3D-printed mounting ports.

Sensitivity Validation

To establish a baseline for electromechanical consistency, the algorithm evaluates the array response at 500 Hz. For a sphere diameter of 112 mm, the wavelength ($\lambda \approx 68$ cm) ensures that the acoustic field is relatively uniform across all transducers, with minimal diffraction effects. Deviations in this band ($\Delta_{i,500\,Hz} > 2$ dB) are used to identify manufacturing

tolerances of the MEMS microphones or gain discrepancies within the digital acquisition chain. Validating the array at this frequency is a prerequisite for subsequent high-frequency analysis, as it isolates electronic gain offsets from mechanical integration defects such as port occlusion or internal resonances.

C. Algorithm Implementation and Classification

The output of the diagnostic pipeline is a classification matrix used for the subsequent procedure. Table I summarizes the detection criteria and their physical significance.

TABLE I
DIAGNOSTIC DETECTION CRITERIA AND IDENTIFIED FAILURE MODES

Frequency Band	Threshold	Identified Physical Issue
High (17 kHz)	$>$ 3 dB drop	Port occlusion
Mid (3.8-4.8 kHz)	$>$ 3 dB dev.	Cavity resonance
Low (500 Hz)	$>$ 2 dB dev.	Sensitivity mismatch

By isolating these non-conforming units, the framework prevents the degradation of the spatial sampling integrity, ensuring that only high-fidelity microphones contribute to the spherical harmonic decomposition.

IV. Experimental Results and Discussion

This section presents the performance evaluation of the 25-element spherical microphone array and the validation of the automated diagnostic pipeline. The analysis focuses on the spectral dispersion across the microphone population and the identification of specific hardware non-conformities that compromise the spatial sampling integrity.

A. Array Dispersion and Reference Mean Analysis

The measured frequency responses of all 25 microphones, following the reference-based compensation and LOESS smoothing, are shown in Fig. 4. The individual magnitude responses (grey curves) exhibit high consistency in the mid-frequency range (400 Hz to 3 kHz), where the standard deviation $\sigma(f)$ remains below 0.7 dB. However, as the frequency increases towards the Nyquist limit of the 112 mm sphere, the spectral dispersion significantly expands, particularly above 15 kHz.

The array-relative mean $\bar{H}(f)$ (black solid line) serves as the baseline for the diagnostic logic. By utilizing this adaptive reference instead of an idealized datasheet curve, the algorithm successfully accounts for the 3D-printed housing's diffraction patterns and the specific acoustic impedance of the 15 mm rigid ports.

B. Identification of Critical Hardware Defects

Applying the multi-band diagnostic thresholds revealed five critical hardware failures within the prototype. The automated reporting from the MATLAB pipeline is summarized in Table II.

The most significant failure was observed in **Mic 08**, which exhibited an extreme magnitude drop of **22.03 dB** at 17 kHz

TABLE II
IDENTIFIED DEFECTIVE MICROPHONES AND MAGNITUDE DEVIATIONS

Mic ID	Target [kHz]	Dev. [dB]	Identified Defect Type
08	17	22.03	Total Port Occlusion
22	17	11.30	Partial Port Occlusion
21	17	8.25	Partial Port Occlusion
13	3,8	3.28	Internal Cavity Leakage
04	3,8	2.97	Mechanical Misalignment

(red dashed line in Fig. 4). Physical inspection as shown in Fig. 2 confirmed that the acoustic inlet was almost entirely occluded by soldering paste. This defect originated during the manual assembly of the high-density flexible PCB within the 15 mm spherical shell constraints. Similarly, **Mic 21** and **Mic 22** showed significant high-frequency attenuation ($>$ 8 dB), indicating partial blockages that would introduce severe errors in high-order spherical harmonic decomposition.

C. Cavity Mode Interference and Misalignment

The diagnostic window at 3.8 kHz successfully isolated **Mic 04** and **Mic 13**, which exhibited deviations exceeding the 3 dB threshold. As analyzed in Section III, these frequencies correspond to the internal cavity modes of the 82 mm hollow core. The deviations suggest that these specific microphones were insufficiently sealed from the internal air volume as shown in Fig. 3. Consequently, these microphones should respond to internal pressure oscillations rather than the external sound field, leading to "leaky" characteristics that degrade the array's directional stability.

V. Conclusion

This paper presented a defect-aware methodology for the validation and calibration of MEMS-based spherical microphone arrays, addressing the inherent discrepancy between idealized spatial sampling models and physical hardware realizations. By implementing a reference-compensated diagnostic pipeline, we successfully identified and characterized manufacturing non-conformities that are frequently overlooked in conventional calibration routines.

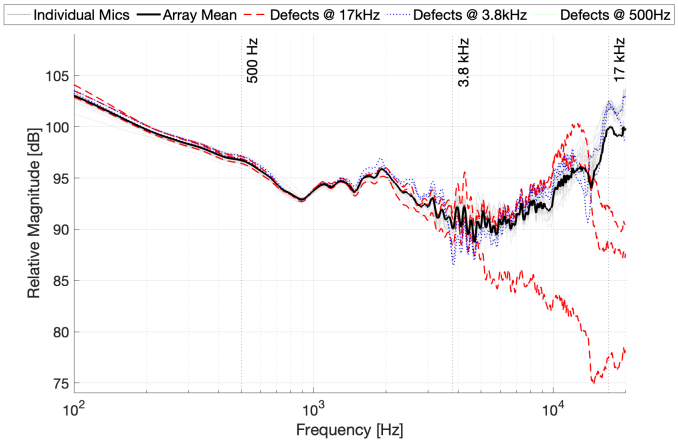

Fig. 4. Multiband defect analysis graph.

Experimental results demonstrated that manufacturing stochasticity, specifically acoustic port occlusion and mechanical misalignment within the 15 mm rigid shell, leads to significant spectral deviations. The proposed automated identification system isolated critical failures, including a significant 22.03 dB magnitude drop in a single microphone unit caused by soldering paste obstruction. Furthermore, the analysis of spectral anomalies in the 3.8–4.8 kHz range provided evidence of internal cavity resonances, indicating localized sealing inconsistencies in the 3D-printed shell.

The implementation of this defect identification procedure allows for the systematic isolation of compromised channels before further spatial processing. By identifying hardware-specific anomalies—such as acoustic leakage at low frequencies, cavity resonances at mid-frequencies, and port occlusion at high frequencies—this approach maintains the spatial sampling integrity of the 25-element array. This framework ensures high precision in the acoustic axis alignment and provides a reliable basis for transitioning from theoretical spherical harmonic designs to physical prototypes.

Future research will focus on the real-time integration of this diagnostic layer into adaptive beamforming algorithms to evaluate the long-term spatial stability of the calibrated array in complex acoustic environments.

REFERENCES

[1] B. Rafaely, "Analysis and Design of Spherical Microphone Arrays," in IEEE Transactions on Speech and Audio Processing, vol. 13, no. 1, pp. 135-143. 2005. doi: 10.1109/TSA.2004.839244.

[2] R. H. Hardin and N. J. A. Sloane, "McLaren's improved snub cube and other new spherical designs in three dimensions," Discrete and Computational Geometry, vol. 15, no. 4, pp. 429–441, 1995.

[3] H. Groemer, "Geometric Applications of Fourier Series and Spherical Harmonics," Cambridge University Press, New York, 1996.

[4] K. E. Atkinson and W. Han, Spherical Harmonics and Approximations on the Unit Sphere: An Introduction. Springer-Verlag, Berlin, Heidelberg, 2012.

[5] J. D. McEwen, G. Puy, J.-P. Thiran, P. Vandergheynst, D. Van De Ville and Y. Wiaux, "Sampling theorems and compressive sensing on the sphere," in Proc. SPIE 8138, Wavelets and Sparsity XIV, San Diego, CA, USA, 2011, doi: 10.1117/12.895318

[6] P. E. Kladitis, V. M. Bright, and J. P. Kharoufeh, "Uncertainty in Manufacture and Assembly of Multiple-Joint Solder Self-Assembled Microelectromechanical Systems (MEMS)," J. Manuf. Process., vol. 6, no. 1, pp. 32–50, 2004, doi: 10.1016/S1526-6125(04)70058-4.

[7] A. Politis and H. Gamper, "Comparing modeled and measurement-based spherical harmonic encoding filters for spherical microphone arrays," in Proc. IEEE Workshop Appl. Signal Process. Audio Acoust. (WASPAA), New Paltz, NY, USA, pp. 224–228, 2017, doi: 10.1109/WASPAA.2017.8170028.

[8] P. Wijnings, S. Stuijk, R. Scholte, and H. Corporaal, "Characterization of Mems Microphone Sensitivity and Phase Distributions with Applications in Array Processing," in Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP), Toronto, ON, Canada, 2021, pp. 4480–4484, doi: 10.1109/ICASSP39728.2021.9413352.

[9] M. Liska and J. Tomešek, "FPGA-based sound acquisition prototype from MEMS microphone array," in Proceedings of the 30th Conference STUDENT EEICT 2024, Brno University of Technology, Faculty of Electrical Engineering and Communication, 2024.

[10] M. Liska, "Sampling Configuration Analysis for the Design of Spherical Microphone Array," in Proceedings of the 31st Conference STUDENT EEICT 2025, Brno University of Technology, Faculty of Electrical Engineering and Communication, 2025.

[11] S. Müller and P. Massarani, "Transfer-function measurement with sweeps," J. Audio Eng. Soc., vol. 49, no. 6, pp. 443–471, 2001.

[12] MathWorks, Inc., "Filtering and Smoothing Data," MATLAB & Simulink Documentation. [Online]. Available: https://www.mathworks.com/help/curvefit/smoothing-data.html (accessed Mar. 1, 2026).

[13] J. Meyer and G. W. Elko, "A highly scalable spherical microphone array based on an orthonormal decomposition of the soundfield," in Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP), Orlando, FL, USA, 2002. doi: 10.1109/ICASSP.2002.5744968.

[14] T. Li and Y. Zhang, "Mode Analysis in a Spherical Resonator," Open Access Library Journal, vol. 10, no. 1, pp. 1-9, 2023. doi: 10.4236/oalib.1109690.

Design and Implementation of a Modular IoT Gateway for Heterogeneous Networks

1st Pavel Palurik
Department of Telecommunications
Brno University of Technology
Brno, Czech Republic
pavel.palurik@vut.cz

2nd Antonin Cerny
Institute of Automation and Computer Science
Brno University of Technology
Brno, Czech Republic
229309@vut.cz

3rd Juraj Goryl
Department of Telecommunications
Brno University of Technology
Brno, Czech Republic
240913@vut.cz

4th Dmitri Moltchanov
Department of Electrical Engineering
Tampere University
Tampere, Finland
dmitri.moltchanov@tuni.fi

5th Radim Dvorak
Department of Telecommunications
Brno University of Technology
Brno, Czech Republic
xdvora2g@vut.cz

6th Pavel Masek
Department of Telecommunications
Brno University of Technology
Brno, Czech Republic
masekpavel@vut.cz

***Abstract*—With the introduction of new Internet of Things (IoT) services, such as energy grid maintenance and remote control, the reliability requirements for conventional massive machine-type communications (mMTC) have become significantly stricter. One of the primary methods to improve the reliability of such services in heterogeneous Industrial IoT (IIoT) environments is to utilize Multiple Radio Access Technologies (Multi-RAT) within a single gateway device. To address this challenge, this work introduces the design and physical implementation of a novel modular hardware platform based on a high-performance Single Board Computer (Raspberry Pi CM5). The core of the novelty lies in the integration of dual M.2 interface slots, enabling the simultaneous operation of licensed cellular technologies (e.g., NB-IoT, LTE Cat-M) for wide-area backhaul and unlicensed local area networks (e.g., 6LoWPAN, DECT-2020, ESP-NOW) for sensor data aggregation. The development process involves custom hardware and multi-layer PCB design, complemented by proprietary routing firmware.**

***Index Terms*—modular IoT, hybrid connectivity, heterogeneous networks, DECT-2020, M.2 interface**

I. Introduction

The Internet of Things (IoT) has undergone significant advancements during the last decade, marked by a transition from isolated wireless sensor networks (WSN) to ubiquitous infrastructure that supports modern cyber-physical systems (CPS) and Industry 4.0 applications [1]. As Industrial IoT (IIoT) applications grow in complexity, there is a fundamental shift toward flexible, universal wireless solutions capable of seamlessly fusing diverse data streams [2]. Progressing toward the sixth-generation (6G) era, these IIoT networks must satisfy stringent Key Performance Indicators (KPIs) for massive machine-type communications (mMTC) [3], such as a packet loss rate (PLR) below 1% and an end-to-end latency under 10 seconds for 99% of packets, as defined by ITU-R M.2410 guidelines [4], [5].

However, as the IoT ecosystem evolves, new applications are appearing—such as energy grid maintenance, remote control, and advanced industrial automation that demand much stricter guaranties in terms of packet delivery and delay. Satisfying these emerging near Ultra-Reliable Low-Latency Communications (near-URLLC) requirements using conventional cellular IoT solutions, such as eMTC and NB-IoT, is becoming increasingly difficult due to their bandwidth and latency limitations [4]. Therefore, realizing a resilient and multi-service IIoT ecosystem requires overcoming the fragmentation between licensed cellular networks and local unlicensed communication standards.

No single wireless technology is capable of simultaneously satisfying all these requirements, as different technologies inherently involve design trade-offs [6]. IIoT infrastructures are heterogeneous environments where multiple communication technologies have to coexist to support diverse application requirements. For example, long-range low-power networks typically offer limited data rates, while high-throughput technologies often require higher energy consumption or shorter communication ranges. One of the possible solutions is to use multi-RAT nodes, which improves reliability, latence and power consumption [7], [8]. Multi-RAT nodes enable modern system design by allowing communication system capabilities to be tailored to deployment-specific requirements.

To address these challenges, this paper presents the design and physical implementation of a novel, modular multi-RAT IoT gateway. By leveraging a high-performance edge-computing core and a customized dual M.2 interface architecture, the proposed platform bridges the current connectivity gap. It facilitates the simultaneous and seamless operation of licensed long-range cellular backhaul networks and highly efficient, localized unlicensed sensor networks.

The remainder of this paper is organized as follows. Section II defines the fundamental system requirements for the proposed multi-RAT IoT architecture. Section III details the communication technology selection, focusing on the integration of licensed cellular and unlicensed local area networks. Section IV presents the hardware design and physical im-

DOI 10.13164/eeict.2026.172

plementation of the gateway prototype. Finally, Section V concludes the paper and outlines future research directions.

II. System Requirements for Hybrid IoT

This section defines the key functional and technical parameters that a modern hybrid gateway must meet for deployment in heterogeneous IIoT environments, such as critical transportation and industrial infrastructure. Given the massive expansion of IoT devices and the complexity of modern network topologies, it is essential that the platform effectively interconnects different communication layers while maintaining high security, reliability, and low latency [9]. The following requirements form the basis for the design of a modular IoT gateway:

- Emerging Near-URLLC Requirements
- Network Resilience and Connectivity Stability
- Heterogeneous Connectivity (Multi-RAT)
- Modular Expandability and Scalability
- Real-time Performance.

The following sections cover these topics.

A. Emerging Near-URLLC Requirements

The modern IIoT market is driven by novel applications such as factory automation, UAV connectivity, and energy grid monitoring—whose performance requirements fall between traditional mMTC and full URLLC. These emerging near-URLLC scenarios demand stricter latency, reliability, and throughput than baseline mMTC metrics [5] and are evolving faster than current standardization cycles [4]. Consequently, relying solely on single-RAT mMTC solutions (e.g., LTE-M, NB-IoT, or standalone DECT-2020 NR) often fails to concurrently satisfy all these constraints. Furthermore, while the upcoming 3GPP 5G RedCap standard addresses higher throughput demands, it may still lack the deterministic latency and ultra-reliability necessary for critical operations. Therefore, transitional solutions are essential to reliably support these demanding services today.

B. Network Resilience and Connectivity Stability

In the domain of IIoT networks, ensuring reliable and continuous connectivity is paramount for the operational stability of the entire ecosystem and its constituent components. Specifically, sensor segments are often organized into a hierarchical cluster-tree (multi-tree) topologies consisting of leaf nodes and superior edge gateways which improves scalability and resource utilization [10]. Within this architecture, edge gateways play a vital role in protocol conversion, routing selection, and data exchange between fundamentally different physical and MAC layers [9]. To maintain high availability, the implementation of failure detection and multi-modal routing mechanisms is essential.

The system architecture must support communication forwarding and the seamless switching of communication technologies when connection parameters deteriorate (e.g., due to severe channel fading, interference, or mobility) [11]. Robustness is further reinforced through data delivery acknowledgment protocols and efficient energy-balancing strategies to prevent premature node failures [12]. Any localized increase in undelivered or unsent information significantly degrades the overall resilience and reliability of the monitoring infrastructure.

C. Heterogeneous Connectivity (Multi-RAT)

The inherent fragmentation across various IoT ecosystems necessitates hardware solutions capable of bridging diverse communication standards. As highlighted in the previous section, utilizing a multi-RAT approach is essential to significantly improve Quality of Service (QoS) specifically connectivity reliability and network availability across heterogeneous environments. Such a multi-RAT configuration enables the convergence of WSN, utilizing unlicensed bands (e.g., Wi-Fi, Bluetooth Low Energy), and Mobile Cellular Networks (MCN) operating in licensed spectrums (e.g., NB-IoT, LTE-M, 5G) [13]. This is possible due to the presence of a modular system equipped with two submodules based on distinct wireless technologies. This strategic combination leverages both spectrum diversity and economic advantages related to deployment costs and infrastructure availability. Such hybrid setups are particularly essential for industrial and transport infrastructure, where radio coverage is frequently compromised by harsh environmental conditions, structural shielding (e.g., by metals and thick walls), and electromagnetic interference. By employing a modular interface, the gateway ensures the necessary flexibility to adapt to evolving technological standards and diverse sensor inputs. Furthermore, the system incorporates a layer dedicated to local data processing and for forwarding of information to specific communication interfaces [9].

D. Modular Expandability and Scalability

A modular architecture was chosen to ensure maximum scalability and adaptability across diverse industrial and transportation deployment scenarios. The integration of standardized interfaces, specifically dual M.2 connectors, facilitates a fundamental ”plug-and-play” approach for swapping communication submodules. This modularity is particularly vital for long-lifecycle infrastructure, such as railway systems, as it enables hardware upgrades to emerging standards including 5G RedCap, NB-IoT NTN, or DECT-2020 without requiring a complete system overhaul. By decoupling the communication layer from the core processing unit, conceptually similar to the separation of control and forwarding planes in Software-Defined Networks (SDN) [9], the gateway can be flexibly customized to meet specific sensory requirements through various modular board configurations. This approach significantly reduces long-term maintenance and installation costs and allows for seamless deployment upgrades. Furthermore, the physical architecture supports adaptable installation configurations, facilitating seamless integration into a wide range of industrial enclosures and critical infrastructure nodes. The hardware layout is optimized for future-proofing, providing

a versatile foundation for the integration of new sensing technologies and enhanced edge computing capabilities.

E. Real-time Performance

Real-time performance represents a critical requirement for the proposed gateway, particularly regarding its deployment in time-critical IIoT environments, such as smart factories and other industrial complexes. To achieve deterministic communication, the gateway must efficiently address latency arising from data processing and the fusion of heterogeneous networks. Industrial applications require varying levels of latency, ranging from stringent cycles of 0.1–2 ms for real-time synchronous data and precision motion control, to 10–100 ms for periodic data transmission and general monitoring [9]. The gateway's modular architecture is designed to target Time-Sensitive Networking (TSN) standards, which provide deterministic transmission services with bounded low delay and traffic scheduling across diverse interfaces (e.g., Ethernet, Wi-Fi 6, and 5G). Through the potential implementation of Edge Computing, the gateway aims to minimize Round-Trip Time (RTT) by processing data locally and transmitting only essential metadata to the cloud. This distributed approach effectively mitigates not only node congestion issues (so-called "hotspots" or energy holes), a common problem where nodes or gateways near the base station rapidly deplete their resources due to heavy relay traffic [10], but also ensures the stability of time-critical control loops, even under high network loads.

III. Communication Technology Selection

To satisfy the system requirements defined in Section II, communication architecture is structurally divided into licensed and unlicensed tiers. This section details the selection of specific multi-RAT communication standards for both long-range backhaul and local sensor data collection, alongside the available hardware modules utilized for their physical implementation.

A. Licensed Cellular IoT (Long-Range Backhaul)

To satisfy the stringent network resilience and wide-area connectivity requirements defined in Section II, the gateway's backhaul architecture relies on MCNs operating within the licensed frequency spectrum. Serving as the critical edge-to-cloud uplink, the platform integrates 3GPP-standardized Low-Power Wide-Area Network (LPWAN) technologies, specifically Narrowband IoT (NB-IoT) and LTE Category-M (LTE-M/Cat-M1). These protocols offer complementary characteristics essential for heterogeneous IIoT deployments [14]. NB-IoT is optimized for massive mMTC, providing an extended link budget for deep indoor penetration in structurally shielded industrial or underground environments, alongside maximized energy efficiency for intermittent, low-throughput transmissions [15]. Conversely, LTE-M caters to dynamic deployment scenarios by offering natively supported seamless mobility, higher data throughput, and significantly reduced transmission latency [16]. By leveraging this dual-stack cellular capability, the gateway can effectively adapt to varying QoS demands.

To physically realize this dual-stack cellular capability, the gateway utilizes an M.2 module featuring the Quectel BG77. Furthermore, to ensure long-term hardware viability and scalability within the proposed modular dual M.2 framework, this licensed communication tier is designed to seamlessly accommodate emerging 3GPP standards, including 5G Reduced Capability (RedCap) and Non-Terrestrial Networks (NB-IoT NTN).

B. Unlicensed Local Area Networks (Sensor Collection)

To complement the licensed edge-to-cloud backhaul, local sensor data acquisition is driven by WSNs operating in unlicensed spectrums. The primary technology targeted for this localized tier is DECT-2020 NR. This choice is backed by the standardized nature of the technology, its fast evolution, emerging multi-vendor availability, and its inherent capability to fully satisfy the ITU-R M.2410 requirements for 5G mMTC and URLLC systems. It operates in the dedicated 1.9 GHz license-exempt frequency band, which inherently mitigates the severe interference prevalent in crowded 2.4 GHz ISM bands [17], [18].

DECT-2020 NR is uniquely designed to satisfy mMTC and URLLC requirements [4], [19]. At the physical and MAC layers, it leverages Cyclic Prefix Orthogonal Frequency-Division Multiplexing (CP-OFDM), Hybrid ARQ (HARQ) with incremental redundancy, and Listen-Before-Talk (LBT) channel access to ensure timely data transmission [20]. Whenever needed, scheduled-based regime can be utilized to ensure highly reliable and deterministic delivery. Additionally, the physical and data link layers allow for flexible configurations tailored for fine-grained tuning of quality-of-service parameters [21].

Furthermore, the standard employs an autonomous, infrastructure-less cluster-tree mesh topology. Devices dynamically alternate between Fixed Termination (FT) and Portable Termination (PT) modes, enabling flexible multi-hop routing and local radio resource management without the need for centralized base stations [4]. This hierarchical architecture prevents single points of failure [19] and supports ultra-dense deployment, scalable up to millions of devices per square kilometer, which is critical for pervasive IIoT monitoring [22].

While the dedicated DECT-2020 NR hardware ecosystem reaches full market maturity [17], the current gateway prototype bridges this gap using a customized M.2 adapter equipped with an ESP32-C6 microcontroller. This module runs the ESP-NOW peer-to-peer protocol, effectively simulating the low-latency cluster-tree data aggregation [19].However, it is important to note that this transitional ESP-NOW is inherently limited by its operation in the 2.4 GHz band, making the local tier more susceptible to electromagnetic interference and less resilient compared to the dedicated 1.9 GHz spectrum utilized by native DECT-2020 NR deployments.

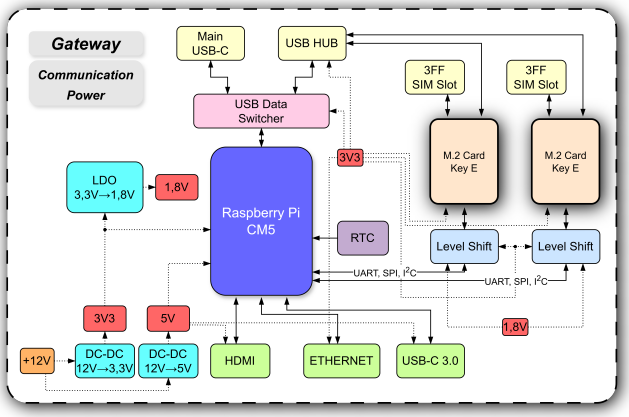

Fig. 1. Communication and power routing block diagram.

IV. Proposed Gateway Prototype

Unlike standard mass-produced M.2 carrier boards, which lack industrial-grade power distribution and integrated local processing, the proposed modular gateway prototype is designed as a highly integrated hardware development platform. The novelty lies in the simultaneous integration of dual M.2 Key E slots directly alongside the Raspberry Pi Compute Module 5 (RP-CM5) edge-computing core on a custom multi-layer PCB. This unique architecture enables the concurrent operation of heterogeneous communication technologies discussed in Section III, with specific efforts to mitigate EMI through optimized signal routing and dedicated ground planes.

A significant advantage of the RP-CM5 module is its direct integration onto the printed circuit board (PCB) using two 100-pin board-to-board connectors. This mounting method allows the processor module to be easily replaced in the event of damage while maintaining the integrity of the main gateway board. Such modularity is particularly beneficial in experimental and educational environments, where hardware components may be frequently replaced or upgraded.

For the implemented prototype, the RP-CM5 variant equipped with 8 GB of RAM and 64 GB of internal storage was selected in order to provide sufficient computational resources for local data processing and edge computing tasks.

A. Communication and Signal Routing

To physically realize the multi-RAT concept, the communication architecture and signal routing are built around two independent M.2 slots (Key E). The overall signal routing and interconnection between the individual functional blocks are detailed in the communication block diagram, see Fig. 1. This diagram describes the data paths between the RP-CM5 processor module, the M.2 communication interfaces, peripheral interfaces, and the dual SIM card modules.

This dual-slot configuration directly supports the simultaneous operation of the licensed Quectel BG77 cellular module and the unlicensed ESP32-C6 local aggregation module. Each M.2 slot is interfaced with the RP-CM5 via high-speed and standard communication busses, including UART, SPI, I²C, and USB. To enable communication with multiple devices through a single USB differential pair available on the RP-CM5, the architecture incorporates a USB hub that distributes the communication channel between the connected modules. Additionally, several general-purpose input/output (GPIO) signals are routed from the RP-CM5 to the M.2 connectors. Since the RP-CM5 supports multiple logic voltage levels, a dedicated level-translation circuit is integrated within the communication paths to convert logic levels between 3.3 V and 1.8 V, ensuring reliable signal routing across peripheral modules that operate at different voltage domains.

B. Power Distribution and PCB Design

To accommodate the high peak power demands of the RP-CM5 and active LPWAN modules, standard USB power delivery is insufficient. Consequently, the platform is driven by an external 12 V DC industrial input. As illustrated in Fig. 1, the internal power distribution network utilizes high-efficiency DC-DC step-down converters (5 V, 3.3 V) and a linear LDO regulator (1.8 V) to reliably supply the core processing and communication subsystems.

The complete hardware layout and schematic design were developed using the Altium Designer environment on a custom 4-layer PCB. The physical implementation of the modular gateway, illustrated by the 3D PCB render in Fig. 2, is structured into strictly separated functional blocks: the communication subsystem, the power distribution network, and peripheral interfaces. High-speed communication lines and differential signal pairs were carefully routed to minimize electromagnetic interference (EMI) between the cellular uplinks and local sensor networks, providing a robust foundation for the multi-RAT IoT gateway.

V. Conclusion

This paper addresses the growing demand for near-URLLC performance in the modern IIoT market by presenting the design and physical implementation of a novel multi-RAT gateway. Because current single-RAT solutions, relying on baseline ITU-R M.2410 metrics (e.g., packet loss rate $< 1\%$, latency < 10 s), struggle to accommodate requirements that fall between traditional mMTC and full URLLC, adopting a multi-RAT approach has become essential to ensure network reliability and availability. To seamlessly bridge this connectivity gap, the proposed modular platform leverages a high-performance Raspberry Pi Compute Module 5 and a dual M.2 interface architecture.

The validated prototype successfully demonstrates the fusion of a licensed cellular backhaul (NB-IoT/LTE-M) with a low-latency unlicensed sensor tier (ESP-NOW simulating DECT-2020 NR) on a custom PCB. Future work will focus on integrating native DECT-2020 NR and 3GPP 5G RedCap hardware modules, alongside developing machine-learning-driven algorithms. By dynamically evaluating real-time channel conditions and QoS requirements, these algorithms will autonomously switch between available multi-RAT interfaces, further maximizing the resilience and energy efficiency of the hybrid IIoT infrastructure.

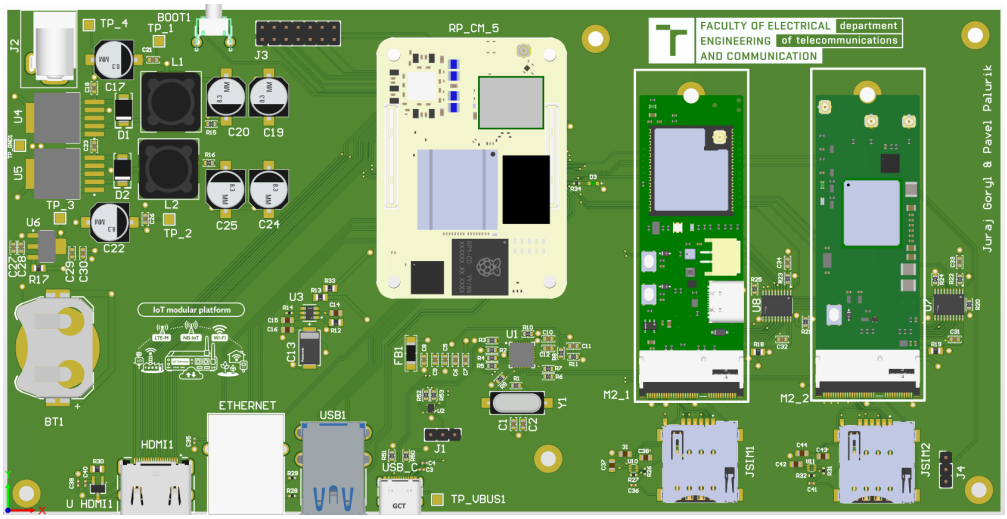

Fig. 2. 3D render of the implemented PCB for the proposed gateway.

ACKNOWLEDGMENT

This publication is co-financed with state support from the Technology Agency of the Czech Republic within the framework of the National Competence Centers Program. This publication is financed within the framework of the National Recovery Plan from the European Recovery and Resilience Instrument. Project FEIM, TN02000067.

REFERENCES

[1] J. Shehu Yalli, M. Hilmi Hasan, and A. Abubakar Badawi, "Internet of Things (IoT): Origins, Embedded Technologies, Smart Applications, and Its Growth in the Last Decade," *IEEE Access*, vol. 12, pp. 91 357–91 382, 2024.

[2] N. H. Mahmood, G. Berardinelli, E. J. Khatib, R. Hashemi, C. De Lima, and M. Latva-aho, "A Functional Architecture for 6G Special-Purpose Industrial IoT Networks," *IEEE Transactions on Industrial Informatics*, vol. 19, no. 3, pp. 2530–2540, 2023.

[3] S. S. Sefati, A. U. Haq, Nidhi, R. Craciunescu, S. Halunga, A. Mihovska, and O. Fratu, "A Comprehensive Survey on Resource Management in 6G Network Based on Internet of Things," *IEEE Access*, vol. 12, pp. 113 741–113 784, 2024.

[4] A. Samuylov, D. Moltchanov, A. Gaydamaka, E. Lyczkowski, A. Frotzscher, F. von Schoettler, J. Pirskanen, J. Numminen, J. Salokannel, E. Llaguno, R. Pires, D. Bayar, and M. Valkama, "Empowering Near-URLLC IoT with 5G DECT-2020 NR: Current State and the Road Ahead," *IEEE Communications Magazine*, vol. 63, no. 4, pp. 130–136, 2025.

[5] ITU, "Minimum Requirements Related to Technical Performance for IMT-2020 Radio Interface," ITU-R, Report ITU-R M.2410-0, July 2017, accessed: 2026-03-07. [Online]. Available: https://www.itu.int/pub/R-REP-M.2410-2017

[6] N. Lethaby, "Wireless Connectivity for the Internet of Things: One Size Does Not Fit All," Texas Instruments White Paper, October 2017.

[7] R. Mozny, P. Masek, D. Moltchanov, M. Stusek, P. Mlynek, Y. Koucheryavy, and J. Hosek, "Characterizing Optimal LPWAN Access Delay in Massive Multi-RAT Smart Grid Deployments," *Internet of Things*, vol. 25, p. 101001, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2542660523003244

[8] M. Stusek, P. Masek, D. Moltchanov, N. Stepanov, J. Hosek, and Y. Koucheryavy, "Performance Assessment of Reinforcement Learning Policies for Battery Lifetime Extension in Mobile Multi-RAT LPWAN Scenarios," *IEEE Internet of Things Journal*, vol. 9, no. 24, pp. 25 581–25 595, 2022.

[9] D. Xia, C. Jiang, J. Wan, J. Jin, V. C. Leung, and M. Martínez-García, "Heterogeneous Network Access and Fusion in Smart Factory: A Survey," *ACM Computing Surveys*, vol. 55, no. 6, pp. 1–31, 2022.

[10] O. I. Khalaf, C. A. T. Romero, S. Hassan, and M. T. Iqbal, "Mitigating Hotspot Issues in Heterogeneous Wireless Sensor Networks," *Journal of Sensors*, vol. 2022, no. 1, p. 7909472, 2022.

[11] J. Kong, F. T. Dagefu, and T. J. Moore, "Covert Routing in Heterogeneous Networks," *IEEE Transactions on Information Forensics and Security*, vol. 19, pp. 7047–7059, 2024.

[12] F. Jibreel, E. Tuyishimire, and M. I. Daabo, "An Enhanced Heterogeneous Gateway-Based Energy-Aware Multi-Hop Routing Protocol for Wireless Sensor Networks," *Information*, vol. 13, no. 4, 2022. [Online]. Available: https://www.mdpi.com/2078-2489/13/4/166

[13] R. Djehaiche, S. Aidel, A. Sawalmeh, N. Saeed, and A. H. Alenezi, "Adaptive Control of IoT/M2M Devices in Smart Buildings Using Heterogeneous Wireless Networks," *IEEE Sensors Journal*, vol. 23, no. 7, pp. 7836–7849, 2023.

[14] A. Boiano *et al.*, "An Empirical Comparison of NB-IoT and eMTC Uplink Performance for IoT Applications," *IEEE Internet of Things Journal*, vol. 10, no. 14, pp. 12 345–12 356, 2023.

[15] D. Orlovs *et al.*, "LPWAN Technologies for IoT: Real-World Deployment, Performance and Comparative Analysis," *IoT*, vol. 6, no. 4, p. 77, 2025.

[16] M. S. Z. Khan, "Comparative Analysis of LTE-M, 4G Cat-1bis and NB-IoT for IoT Connectivity," *IEEE Access*, vol. 13, no. 1, pp. 45 678–45 692, 2025.

[17] F. Graf, T. Watteyne, and M. Villnow, "What to Expect When Using DECT NR+," *IEEE Communications Magazine*, pp. 1–7, 2025.

[18] T. Troccoli, H. J. Damsgaard, J. Pirskanen, E. S. Lohan, A. Ometov, J. Morte, J. Nurmi, and V. Kaseva, "Novel Direction-of-Arrival-Based Localization in Massive DECT-2020 5G NR Networks," *IEEE Internet of Things Journal*, vol. 12, no. 3, pp. 2690–2712, 2025.

[19] R. Kovalchukov, D. Moltchanov, J. Pirskanen, J. Säe, J. Numminen, Y. Koucheryavy, and M. Valkama, "DECT-2020 New Radio: The Next Step toward 5G Massive Machine-Type Communications," *IEEE Communications Magazine*, vol. 60, no. 6, pp. 58–64, 2022.

[20] T. Nihtilä and H. Berg, "Energy Consumption of DECT-2020 NR Mesh Networks," in *2022 Joint European Conference on Networks and Communications 6G Summit (EuCNC/6G Summit)*, 2022, pp. 196–201.

[21] R. Glazkov, A. Samuylov, A. Gaydamaka, D. Moltchanov, J. Pirskanen, J. Numminen, and M. Valkama, "Improving Near-URLLC Services Performance Over Multi-Hop Topologies in DECT-2020 NR Systems," in *2025 IEEE 101st Vehicular Technology Conference (VTC2025-Spring)*, 2025, pp. 1–7.

[22] A. Samuylov, D. Moltchanov, J. Pirskanen, J. Numminen, Y. Koucheryavy, and M. Valkama, "Performance Assessment of DECT-2020 NR and Classic DECT Coexistence Mechanisms," in *2023 IEEE 97th Vehicular Technology Conference (VTC2023-Spring)*, 2023, pp. 1–7.

Cognitive Radio implementation using GNU Radio

Martin Kousal
Department of Radio Electronics
FEEC, Brno University of Technology
Brno, Czech Republic
martin.kousal@vut.cz

Roman Maršálek
Department of Radio Electronics
FEEC, Brno University of Technology
Brno, Czech Republic
marsaler@vut.cz

***Abstract*—Cognitive Radio (CR) has emerged as one of the most promising solutions to address the spectrum utilization problem. The main idea behind CR is to design intelligent wireless communication systems capable of sensing their environment and dynamically adapting their transmission parameters to optimize given spectrum usage. The developed GNU Radio blocks implement spectrum sensing, occupancy detection, and dynamic allocation of Orthogonal Frequency-Division Multiplexing (OFDM) subcarriers based on occupancy information provided by the previously described energy detector block. The implemented systems can detect unused frequency bands and utilize them without causing harmful interference to licensed users. To verify correct operation of the cognitive radio components implemented in a practical environment, experiments were performed using real radio hardware. The objective of these experiments was to observe the resulting spectrum behavior during transmission. In the first stage of testing, the system was connected using a direct coaxial cable between the USRP transmitter output and the FSVR spectrum analyzer input. In the second stage, the system was tested with receiving and transmitting signals over the air and captured by spectum analyzer.**

***Index Terms*—Cognitive Radio, SDR, sensing, spectrum, GNU Radio**

I. Introduction

The rapid growth of wireless communication technologies has significantly increased the demand for radio spectrum resources. Modern applications such as fifth generation (5G) mobile networks or milimeter-wave systems require reliable and high-capacity wireless communication. However, the radio frequency (RF) spectrum is a limited resource and traditional spectrum allocation policies assign fixed frequency bands to licensed users, often resulting in ineffictient utilization of the spectrum. Study [1] has shown that many licensed frequency bands remains underutilized at different times and locations, even though the spectrum appears fully allocated from a regulatory perspective.

Cognitive Radio (CR) has emerged as one of the most promising solutions to address the spectrum utilization problem. The concept of CR was first introduced by Joseph Mitola in the late 1990s as an extension of Software-Defined Radio (SDR) [2]. The main idea behind CR is to design intelligent wireless communication systems capable of sensing their environment and dynamically adapting their transmission parameters to optimize given spectrum usage. Such systems can detect unused frequency bands and utilize them without causing harmful interference to licensed users. The aim of this paper is to describe functions of cognitive radio, including spectrum sensing, occupancy detection, and the cognitive radio transmitter. Then with the usage of the open-source project, GNU Radio[1] implements these functions as Out Of Tree Module (OOT) and verify their functionality.

The paper is organized as follows. Chapter II provides an overview of commonly used spectrum sensing techniques relevant to cognitive radio systems. Chapter III briefly introduces the architecture of GNU Radio and the principles of developing custom processing blocks. Chapter IV describes the implementation of the proposed Out Of Tree module and the individual blocks developed for spectrum sensing and dynamic transmission. The complete system flowgraph is presented in Chapter V. Chapter VI evaluates the proposed system using recorded 5G signals in a simulation environment, while Chapter VII presents experimental verification using real radio hardware.

II. Spectrum sensing techniques

Accurate sensing is essential to avoid harmful interference with other users while allowing access to unused parts of the wireless spectrum. Several spectrum sensing techniques have been proposed in the literature, each offering different trade-offs in terms of complexity, detection performance, and required prior knowledge of the primary signal. Among the most commonly studied approaches are energy (power) detection, cyclostationary feature detection, or matched filter detection [3], [4].

A. Energy detector

The energy detector is one of the simplest techniques due to its low computational complexity and minimal required knowledge about the received signal. The principle of this method is based on measuring the energy of the received signal within a specific frequency band and comparing it to a predefined threshold. If the measured energy exceeds the given threshold, the presence of a primary user is assumed. Otherwise, the channel is considered unused.

Energy detection offers several advantages, including simple implementation and no requirement for detailed information

The research has been funded by the project GNU Radio contributions under European Space Agency (ESA) contract No. 4000141041/23/UK/AL implemented by JIC.

[1]https://www.gnuradio.org/

about the received signal. However, it also has significant limitations. Its performance is highly sensitive to noise uncertainty, and it cannot distinguish between different signal types, which may lead to false detections when interference or noise level fluctuation occurs. These limitations make the energy detector less reliable in low signal-to-noise ratio (SNR) environments.

B. Cyclostationary detector

Cyclostationary feature detection exploits the periodic statistical properties present in many modulated wireless signals. Unlike noise, which is typically stationary, modulated signals often exhibit periodicity in their mean, autocorrelation, or spectral components due to the used modulation type, coding scheme, or pilot symbols [5].

This method analyzes the cyclic autocorrelation function or the spectral correlation density of the received signal to identify these periodic features. Because noise does not exhibit cyclostationary characteristics, the detector can effectively distinguish between noise and useful signals even under low SNR conditions. In addition, it can differentiate between different types of signals, which makes it useful for signal classification tasks. However, this approach requires higher computational complexity and longer observation time, which may limit its applicability in a real-time system. What is the most important thing, we must know the basic parameters of the signal that we are looking for in the received spectrum. This fact is a fundamental limitation if we want to observe an unknown received spectrum.

C. Matched filter detector

Matched filter detection is considered the optimal detection technique when the characteristics of the received signal are known. This method may be used to find repeated preambles or pilot symbols, but only if they are not changed during transmission and are all known at the detector.

In this approach, the receiver uses prior knowledge of the transmitted signal, such as modulation type or pilot sequences. By correlating the received signal with this known pattern, the matched filter can achieve very fast and reliable detection, even at relatively low SNR levels. In addition, implementing multiple matched filters may be necessary if different signal types must be detected, increasing the system complexity.

III. GNU Radio architecture

The architecture of GNU Radio is based on a modular design in which signal processing operations are implemented as blocks connected within a flowgraph. Each block performs a specific function, such as filtering, modulation, demodulation, or spectral analysis. Data flow between blocks through input and output ports in the form of streams or vectors. The flowgraph defines the overall processing chain that the runtime scheduler executes.

GNU Radio applications are typically developed using two main components. The first component is the runtime, written primarily in `C++`, which provides efficient execution of signal processing blocks and manages the data flow between them. Almost all native blocks are written as effective as they can be, so it makes no sense to implement, for example, our own Fast Fourier Transform (FFT). The second component is the Python interface, which allows users to construct and control flowgraphs using Python scripts. In addition, the GNU Radio Companion graphical tool provides an interface to design and configure flowgraphs without extensive programming. Signal sources for implemented systems can be either simulated signals or real radio hardware. Software-Defined Radio (SDR) devices such as USRP can be connected to GNU Radio to receive and transmit RF signals as drop-in native blocks.

A. Custom blocks

Although GNU Radio provides a rich set of predefined blocks, implementing a more complex signal processing chain using only native blocks can introduce several limitations. Many algorithms require specific mathematical operations or data manipulation that may not be available directly as native blocks. As a result, the flowgraph can become overly complex, consisting of a large number of interconnected blocks that reduce readability and maintainability.

The Embedded Python Block partially addresses this issue by allowing custom Python code to be inserted into the flowgraph. However, this method is suitable for small pieces of code or simple algorithms. The main issue is that if you want to reuse this piece of code, you must copy it out of your current flowgraph and insert it somewhere else, which is not an optimal way of maintaining code.

To resolve the described issue, there exist Out Of Tree Python Modules where it is possible to create custom blocks. These modules are then installed into GNU Radio and are available as other native blocks. You can maintain them outside of the GNU Radio and perform testbenches.

IV. Developed blocks

Using GNU Radio development tools[2], the *spectrumSharing* OOT module was created. Python language was used for the simplicity of its syntax and widely available mathematic libraries that are used for signal processing algorithms. If processing throughput will be a problem in the future, all created blocks could be rewritten using `C++` language. All custom-created blocks are located inside this module, and after proper installation they are available to the user for his flowgraph. Their source code consists of a Python file containing its own functionality and YAML file used for interface definition and description.

A. Energy detector

The first block created, named *energyDetector*, performs energy-based spectrum sensing on vectors of complex input samples. The block processes vectors of length N, corresponding to the FFT size, which represent frequency-domain samples obtained from a preceding FFT processing stage. The main purpose of the block is to determine whether individual frequency subbands are occupied or free. To achieve this, the

[2]https://wiki.gnuradio.org/index.php/ModtoolArchitecture

input FFT vector is divided into a configurable number of equally sized subbands. The energy within each subband is computed by averaging the squared magnitude of the complex samples belonging to that subband and converting them to a logarithmic scale as follows:

$$y = 10 \cdot \log_{10} \sum_{n=0}^{N-1} |x_n|^2, \tag{1}$$

where N is size of the subband and x is input sample.

Based on the calculated powers in each subband, a dynamic thresholding mechanism is applied. The threshold T is determined relative to the minimum and maximum power observed in the current vector, according to:

$$T = P_{min} + \alpha(P_{max} - P_{min}), \tag{2}$$

where P_{min} and P_{max} denote minimum and maximum subband power values and α is a configurable threshold parameter. Subbands, whose energy exceeds the threshold, are marked as occupied. The resulting binary decision is then expanded so that all FFT bins belonging to the detected subband share the same occupancy value, either occupied or free.

The output of the block is therefore a vector of values representing the occupancy state of each FFT bin. A value of 1 indicates that a signal has been detected, while 0 indicates that the subband is considered free.

B. Carrier allocator

To enable opportunistic transmission in available spectrum regions, a custom transmission block named *carrierAllocator* was developed. The block implements dynamic allocation of Orthogonal Frequency-Division Multiplexing (OFDM) subcarriers based on spectrum occupancy information provided by the previously described energy detector block.

The *carrierAllocator* block generates OFDM symbols in the frequency domain and places modulated data only on subcarriers that are marked as unoccupied. The block receives a vector with size of the FFT representing the occupancy state of individual FFT bins. For each input vector, the block determines the indices of free positions and generates modulated symbols on them. In the current implementation, Quadrature Phase Shift Keying (QPSK) symbols are randomly generated using pseudo-random number generator, but can be extended for higher-order modulations in the future. The amplitude of the generated symbols is controlled by a configurable transmission amplitude parameter. All other subcarriers remain set to zero, ensuring that no energy is transmitted in occupied spectral regions.

In addition to the occupancy input, the block includes a second input stream that acts as a transmission enable control signal. When the transmission enable control signal is set to zero, the block outputs a vector of zeros, effectively disabling the transmission.

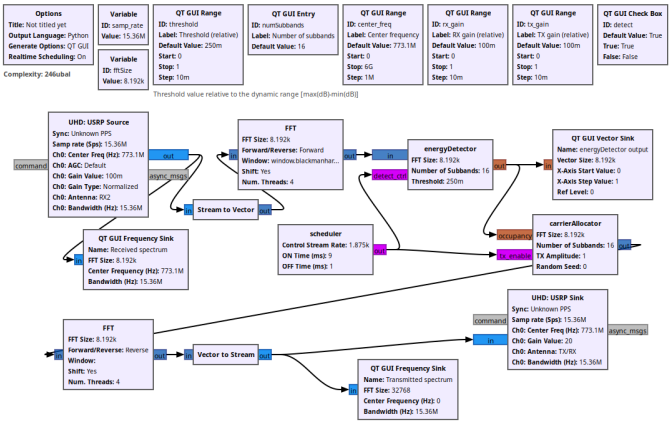

Fig. 1. GNU Radio flowgraph with custom blocks for spectrum sharing.

C. Scheduler

To coordinate the sensing and transmission phases, a dedicated control block named *scheduler* was implemented. This is done by periodically generating a control signal that enables or disables transmission and vice versa controls sensing. This sensing process is inspired by LTE-V2X systems [6], which also has some sensing period before it starts transmission. The principle can also be simply called Listen Before Talk.

The scheduler block produces a continuous stream of control values (either 1 or 0) without requiring any input signals. The transmission phase keeps the output value set to 1, and the sensing or idle phase keeps the output set as 0. These phases repeat periodically, creating a deterministic duty cycle. The duration of these periods is defined by configurable parameters expressed in milliseconds, where the default is 1 ms for sensing and 9 ms for transmission.

Internally, the block converts the configured time durations into the corresponding number of output samples using the specified control stream rate. This control stream rate should be given as a parameter and should be defined as the sample rate divided by FFT size.

V. Complete flowgraph

After developing all required blocks, the final flowgraph could be created using GUI. Two different flowgraphs were created, one for testing purposes with a recorded or synthetic signal, and the other for use with USRP for receiving and transmitting over the air.

The flowgraph for real-world usage is shown in Fig. 1. First of all, the block *UHD: USRP Source* is used for receiving data from USRP, where output streams of complex data are converted into vectors of the same size as the FFT length. This vector is then passed to block *FFT* where transformation from time-domain to frequency-domain is performed, and this result is then passed to *energyDetector* block. The output of the detection is passed to the *carrierAllocator* block, where the modulation is performed, and then the last processing step is the Inverse FFT calculation to transform from the frequency-domain to the time-domain representation. Before

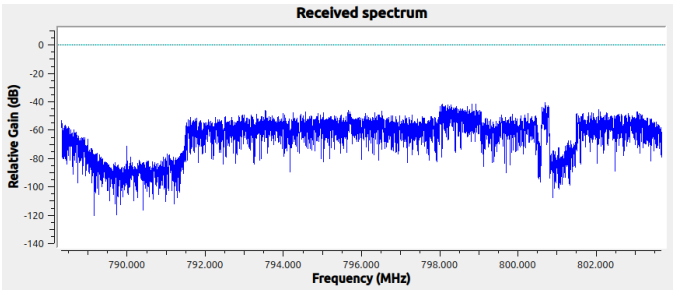

Fig. 2. Spectrum of received signal used for simulation.

the last block only the transformation from vectors to stream is performed and the generated signal is transmitted by the final *UHD: USRP Sink* block.

VI. Simulation Verification

To verify the functionality of the flowgraph and created custom blocks, the recorded signal loaded from a file can be used. The signal was captured with a sample rate of 15.36 MHz and center frequency of 796 MHz, where one of Czech local cellular operators is providing an LTE connectivity with a used bandwidth of 10 MHz. Due to the national restrictions and spectrum allocations to other operators, there is guard band (given by the difference between assigned and occupied bandwidth) aside of this used part of spectrum. The ”blank” space in the spectrum is ideal for testing cognitive radio, because it is exactly known where this space is, and from this knowledge, it could be easy to determine if our blocks are working correctly. This received signal can be seen at Fig. 2.

For the purpose of simulation, the parameters were set as follows. The size of the FFT was set to 8192, the number of subbands to 64, and the threshold α to 0.2. With these parameters provided, the flowgraph was run, and its output can be seen in Fig. 3. We can clearly see that the detector output is correctly reaching the space in the spectrum below and above the present signal from the cellular operator. The allocator, accordingly, is generating the signal to transmit in these spaces.

VII. Real Hardware Measurements

To validate the functionality of the developed GNU Radio blocks in a practical environment, experiments were performed using real radio hardware. The objective of these experiments was to verify the correct operation of the implemented cognitive radio components and to observe the resulting spectrum behavior during transmission. The testing setup consisted of an SDR transceiver implemented in USRP B210[3] with its coverage from 70 MHz to 6 GHz and a bandwidth of up to 56 MHz. A reference receiver was based on the Rohde & Schwarz FSVR Real-Time Spectrum Analyzer[4].

[3]https://www.ettus.com/all-products/ub210-kit/

[4]https://www.rohde-schwarz.com/us/products/test-and-measurement/benchtop-analyzers/rs-fsvr-real-time-spectrum-analyzer_63493-11047.html

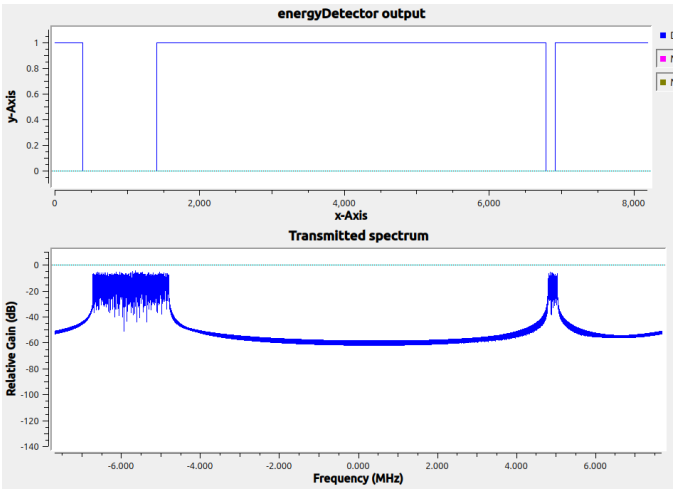

Fig. 3. Simulation output for detector and spectrum of transmitted signal.

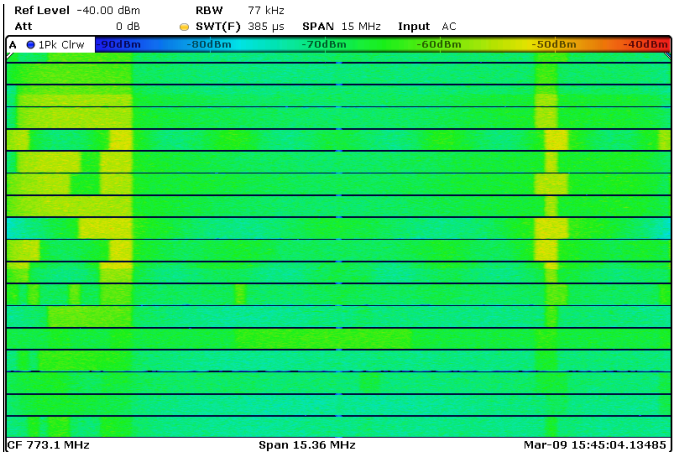

Fig. 4. Spectrogram of the cognitive radio transmitter connected via cable.

A. Cable-based verification

In the first stage of testing, the system was connected using a direct coaxial cable between the USRP transmitter output, and the FSVR spectrum analyzer. This configuration ensured a controlled environment without external radio interference or channel effects. The receiver of USRP is connected to the omnidirectional antenna and is receiving signals from the existing cellular operator network.

Operating the system in this configuration allowed direct observation of the transmitted signal in the frequency domain. This stage was used primarily to verify the correct operation of the carrier allocation block and to confirm that the transmitted spectrum matched the expected pattern. The output of the transmitted spectrum is shown in Fig. 4. The yellow color represents the OFDM signal transmitted in subbands that were considered free. We can clearly see the repeating pattern of the sensing and transmitting period maintained by the scheduler block.

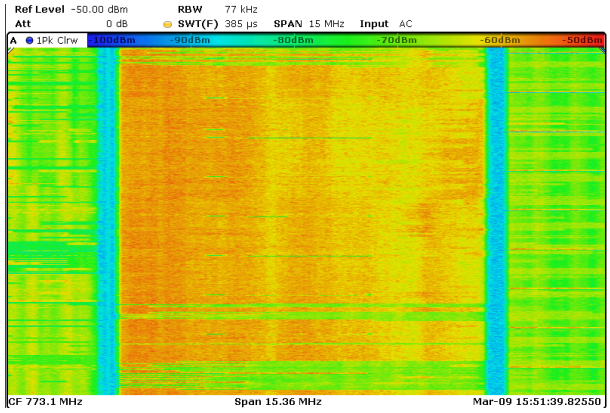

Fig. 5. Observed spectrogram of ambient environment.

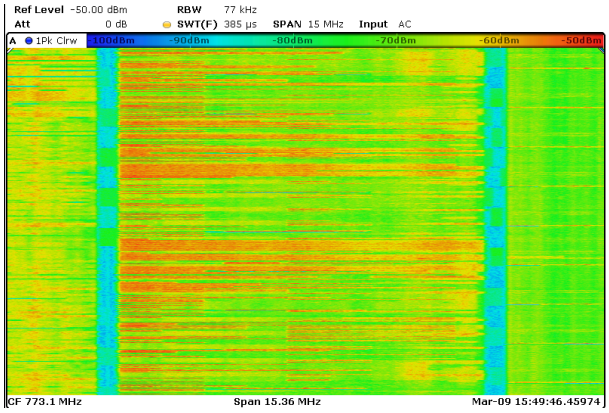

Fig. 6. Observed spectrogram with cognitive radio transmission.

B. Over-The-Air experiment

After verifying the system in a cable-connected configuration, the experiment was extended to a wireless scenario, in which the primary 5G network signal was replayed by the Rohde & Schwarz SMW generator in the anechoic chamber to avoid interference to operator network. The USRP transmitted the signal using an antenna at its RF frontend, while the FSVR spectrum analyzer observed the spectrum over the air.

The first measurement was performed with the cognitive radio turned off, to get insight into how the received spectrum really looks without any of our interference. The spectrogram of this measurement can be found in Fig. 5. The center frequency was set to 773.1 MHz, and the bandwidth was set to 15.36 MHz, which is the same as in the cognitive radio. At this given frequency, the 5G network of one of the local operators is present. From this spectrogram we can observe that the two main blank spaces, where we want to focus with our cognitive radio transmission, are placed beyond and above the 10 MHz occupied block at the center.

The second measurement was performed with the cognitive radio turned on with the following parameters: FFT size of 8192, number of subbands equal to 64, and dynamic threshold set to 0.2, i.e, 20 % of the dynamic range. These settings lead to the spectrogram received with FSVR shown in Fig. 6. From this figure, it is clear that our cognitive radio system is transmitting in the gaps of the spectrum. The transmission does not occur in every time slice, because the threshold and corresponding size of the subband were set up so that sometimes it is considered free and sometimes not, just to show the function of scheduling and not interfering with itself.

VIII. Conclusion

This paper presented the implementation and experimental verification of a cognitive radio system developed within the GNU Radio framework. The proposed solution combines spectrum sensing and opportunistic transmission in order to enable dynamic spectrum access without interfering with existing wireless users.

Several custom GNU Radio blocks were implemented as part of an Out Of Tree module. The *energyDetector* block performs energy-based spectrum sensing and identifies occupied frequency subbands. Based on this information, the *carrierAllocator* block dynamically assigns OFDM subcarriers only to spectral regions detected as free. A dedicated *scheduler* block coordinates the sensing and transmission phases using a periodic listen-before-talk mechanism.

The functionality of the proposed system was validated in both simulation and experiments using real radio hardware. Measurements performed with a USRP B210 transceiver and a Rohde & Schwarz FSVR spectrum analyzer confirmed that the system correctly detects unused spectral regions and transmits signals within these gaps. The results demonstrate the feasibility of implementing cognitive radio concepts using GNU Radio and software-defined radio hardware. Developed module and example flowgraph is available at GitHub[5].

Future work may focus on improving detection robustness, for example by employing more advanced sensing techniques or adaptive thresholding mechanisms, as well as extending the system toward a complete communication link with a receiver.

References

[1] V. Valenta, Z. Fedra, R. Marsalek, G. Baudoin and M. Villegas, "Towards cognitive radio networks: Spectrum utilization measurements in suburb environment," 2009 IEEE Radio and Wireless Symposium, San Diego, CA, USA, 2009, pp. 352-355, doi: 10.1109/RWS.2009.4957353

[2] J. Mitola and G. Q. Maguire, "Cognitive radio: making software radios more personal," in IEEE Personal Communications, vol. 6, no. 4, pp. 13-18, Aug. 1999, doi: 10.1109/98.788210

[3] T. Yucek and H. Arslan, "A survey of spectrum sensing algorithms for cognitive radio applications," in IEEE Communications Surveys & Tutorials, vol. 11, no. 1, pp. 116-130, First Quarter 2009, doi: 10.1109/SURV.2009.090109

[4] S. Haykin, "Cognitive radio: brain-empowered wireless communications," in IEEE Journal on Selected Areas in Communications, vol. 23, no. 2, pp. 201-220, Feb. 2005, doi: 10.1109/JSAC.2004.839380

[5] I. F. Akyildiz, W.-Y. Lee, M. C. Vuran, and S. Mohanty, "NeXt generation/dynamic spectrum access/cognitive radio wireless networks: A survey," Computer Networks, vol. 50, no. 13, pp. 2127–2159, 2006, doi: https://doi.org/10.1016/j.comnet.2006.05.001

[6] C. Ghodhbane, M. Kassab, S. Maaloul, H. Aniss, and M. Berbineau, "A Study of LTE-V2X Mode 4 Performances in a Multiapplication Context," IEEE Access, vol. 10, pp. pp63579-63591, Jun. 2022, doi: 10.1109/ACCESS.2022.3182508

[5] https://github.com/mkousal/GNU-Radio-Spectrum-Sharing

The Impact of Over-the-Air Updates on Cellular IoT Systems Performance

1st Radim Dvořák
Department of Telecommunications
Brno University of Technology
Brno, Czech Republic
xdvora2g@vutbr.cz

2nd Pavel Paluřík
Department of Telecommunications
Brno University of Technology
Brno, Czech Republic
xpalur01@vutbr.cz

3rd Thao Dinh Le
Department of Telecommunications
Brno University of Technology
Brno, Czech Republic
dinh.thao.le@vut.cz

4th Dmitri Moltchanov
Department of Electrical Engineering
Tampere University
Tampere, Finland
dmitri.moltchanov@tuni.fi

***Abstract*—Massive machine type communications (mMTC) is an integral part of 5G ecosystem that has recently attracted a lot of attention from the market. The appearance of new services not only requires stricter performance guarantees at the air interfaces but also introduces new functional requirements such as over-the-air updates (OTA). Unfortunately, differently from enhanced mobility broadband service (eMBB), mMTC technologies including Narrowband Internet-of-Things (NB-IoT) and LTE-M, have not received significant attention from the standardization bodies in OTA domain. The aim of this study is to implement and test the impact of over-the-air (OTA) updates on the performance of conventional traffic. To this aim, we composed a testbed to evaluate the total overhead generated by common industry protocols. The results indicate the UDP-based protocols having a significant advantage in terms of lowering the total system overhead. When combined with the streaming OTA method for 1 MB file, the total overhead can be reduced by up to 75% as opposed to the TCP variants depending on the scenario.**

***Index Terms*—CIoT, NB-IoT, OTA, protocols, implementation, performance**

I. Introduction

In recent years, the Cellular IoT (CIoT) technologies have seen a significant increase in adoption for the industrial and consumer low-power market. The two most distinct technologies of the 5G ecosystem are the LTE-M (LTE Cat-M) and NB-IoT (LTE Cat-NB), holding approximately 60.25% share of the global new Low Power Wide Area (LPWAN) deployments in 2025 [1].

Both NB-IoT and LTE-M technologies are tailored to provide significant device-cost reduction and increased coverage in exchange for reduction of advanced features and reduced throughput and as such are ideal for devices with low-power mass produced devices with limited data throughput and latency requirements where coverage is the main concern. The latter of the two, NB-IoT, provides the MCL (Maximum Coupling Loss) of 164 dB utilizing bandwidth of 180 kHz and achieving up to 64 kbps (LTE Cat-NB1) or 160 kbps (LTE Cat-NB2) theoretically in uplink and 26 kbps (LTE Cat-NB1) or 127 kbps (LTE Cat-NB2) in downlink throughput respectively, which fulfills the requirements for most low-power devices currently available on the market [1].

By today's industry standards, an end device is expected to include means of changing specific parameters or adjust the firmware to the customers' needs. While some devices are capable of direct user configuration (via mobile application or other means) some devices do not have such capability and need to be reprogrammed to fit the demands of the customer which requires an on-site personnel or a device hand-out to the manufacturer including additional costs.

Nowadays, the above mentioned reconfigurations are often handled via the over-the-air updates (OTA) to save costs of the on-site technician and labor. While the OTA implementation is simple in most cases, in the limited environment, such as NB-IoT deployment, it may prove challenging due to the network connection bottleneck in downlink direction. Even though, the NB-IoT is capable of reaching up to 127 kbps in downlink direction, it is often significantly lower, as the network is shared among other devices, and as such creates a significant challenge when transferring bigger amount of data such as firmware, which could reach up to several MB. Additionally, selected transfer protocol will also significantly contribute to the total amount of data transferred [1], [2].

The goal of this paper is to explore current IoT protocols commonly used on the market for data transmission and explore the options of the OTA updates via these protocols. In the first section the paper, an overview of the commonly used IoT protocols is layed out followed by a section describing different mechanisms used in OTA applications nowadays. In the fourth section of the paper a practical experiment is conducted and protocols are compared for their OTA applicability. Finally, in the last section, we conclude the paper.

The rest of the paper is organized as follows. We begin in Section II with overview of the state-of-the-art IoT protocols overview. OTA update mechanisms are provides in Section III. We evaluate the impact of OTA updates in Section IV. Finally,

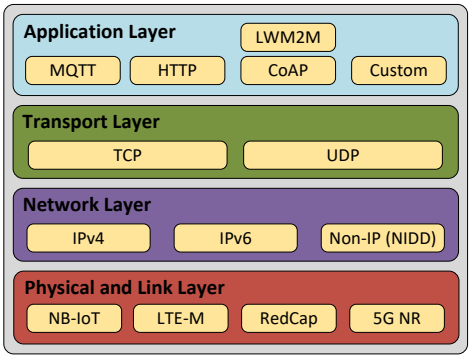

Fig. 1. Potential 5G-IoT stack configurations.

conclusions are stated in the last section.

II. IoT Protocols Overview

Nowadays, several protocols have established themselves as the de-facto standards in CIoT data transmission applications. From the conventional transport layer protocols such as UDP (User Datagram Protocol) and TCP (Transport Control Protocol) to specific applications protocols, namely, Message Queuing Telementry Transport (MQTT) and Constrained Application Protocol (CoAP) with a fairly new protocol presented by the Open Mobile Alliance (OMA) – Lightweight Machine to Machine (LWM2M). The selection of the protocol stack influences not only the communication structure but also the overhead data that is carried out by the communications system [3].

With the LPWAN technologies the target is to optimize the communication pattern for a particular technology as well as, in most cases, minimize the amount of data transmitted to save costs and energy. Potential CIoT stack configurations discussed below in this section are shown in Fig. 1.

A. Bare UDP/TCP Protocols

For simple devices implementing the NB-IoT technology, many manufacturers opt to use bare transport protocols for simplicity and overhead reduction. UDP and TCP are well-known in the industry as overhead controlled reliable stack. While the UDP offers a fixed overhead of 8 bytes with no delivery guarantee and congestion control, the TCP offers a minimum overhead of 20 bytes with delivery and transmission control [4], [5].

B. MQTT Protocol

MQTT is a lightweight, publish-subscribe application layer network protocol. It is designed for connections with remote locations, where a ”small code footprint” is required or the network bandwidth is limited. MQTT strictly requires a connection-oriented transport, almost exclusively TCP/IP stack, to ensure ordered and lossless delivery of data. Nowadays, the MQTT is a well-known, dominant protocol on the field of IoT devices either for consumer use or for industry [6].

MQTT relies on a central Broker (Server) that distributes messages to clients based on topic subscriptions and features a very small fixed header of just 2 bytes. However, due to the fact it utilizes the TCP protocol at the transport layer, it inherits the overhead of the three-way handshake and the persistent connection maintenance – keep-alive messages. In high-latency environments like NB-IoT, the cumulative overhead of TCP acknowledgments can become significant compared to the actual payload. Additionally, due to its connection-oriented nature, the application for battery-powered, constrained devices is limited as the connection has to be maintained via a keep-alive mechanism affecting the battery-life and also inducing additional overheads [6].

C. CoAP Protocol

CoAP is a specialized web transfer protocol for use with constrained nodes and networks, defined in RFC 7252. It is designed to easily interface with HTTP for integration with the web while meeting specialized requirements such as multicast support, very low overhead, and simplicity [7].

CoAP follows a Client-Server model but is logically divided into two layers: messaging (handling reliability and duplication) and request/response (using method codes such as GET, POST, PUT, DELETE). It is primarily designed to run over UDP, significantly reducing the connection setup overhead seen in TCP and also the fixed header is just 4 bytes. To compensate for UDP's lack of reliability, CoAP uses ”Confirmable” (CON) messages which require a lightweight 4-byte acknowledgment [7].

D. LWM2M Protocol

LWM2M is a device management protocol also designed by the OMA. It is not a standalone transport protocol but rather a service layer that typically runs on top of CoAP. LWM2M utilizes a hierarchical object-based resource model. Each device is represented as a collection of objects (e.g., device information, connectivity monitoring), which contain resources (e.g., manufacturer, radio signal strength) [8].

As LWM2M is implemented on top of CoAP/UDP, it maintains a low transport footprint. However, it introduces additional application-layer overhead due to its structured data formats, originally, Type-Length Value (TLV) or JavaScript Object Notation (JSON), now often optimized with Concise Binary Object Representation (CBOR). It is highly optimized for CIoT technologies, supporting features like ”Queue Mode” to allow devices to sleep for long periods without losing their registration status [8].

III. OTA Update Mechanisms in CIoT

The method by which a firmware binary is delivered from a remote server to a constrained device significantly impacts the total data overhead and the energy consumption of the node. In the context of NB-IoT, two primary architectural patterns are utilized: (i) discrete block polling and (ii) continuous stream delivery. We now proceed explaining these methods.

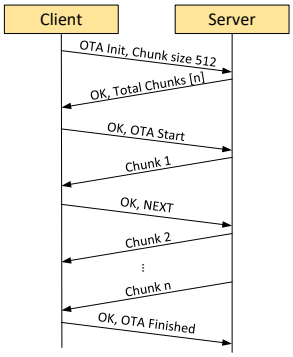

Fig. 2. Signaling diagram of the discrete block polling approach.

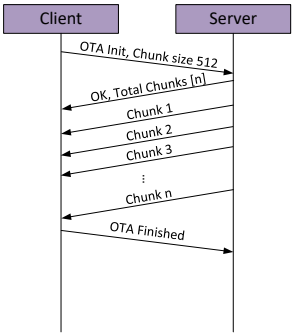

Fig. 3. Signaling diagram of the streaming-based approach.

A. Discrete Block Polling

In case of discrete block polling, the device maintains full control over the download process, see Fig. 2. The device initiates a request for a specific ”chunk” or ”block” of the firmware (e.g., 512 bytes or 1024 bytes), waits for the server's response, processes the data, and then sends a subsequent request for the next block. This provides high resilience in unstable network conditions. If a connection is lost, the device knows exactly which block was last received. It also allows the device to manage its internal buffer and flash-writing cycles without the risk of an overflow.

This is the more ”expensive” method in terms of transport layer overhead. To put things into perspective, for a 1 MB file, if the block size is 512 bytes, the device must send approximately 2,048 uplink requests. In NB-IoT, every uplink transmission requires a Random Access Procedure (RACH) if the device has entered a Radio Resource Control (RRC) protocol's IDLE state between blocks, leading to a massive accumulation of signaling overhead and energy drain [1].

B. Streaming Mechanism

In a streaming scenario, the device sends a single initial request to the server to start the update, see Fig. 3. The server then responds by sending the firmware file as a continuous flow of packets until the entire firmware is delivered or the connection is terminated. This method is significantly more efficient regarding the data-to-overhead ratio. By minimizing the number of uplink requests, the device stays in RRC Connected state, avoiding the repetitive signaling cost of the RACH procedure [1].

While being more efficient, this approach requires robust flow control typically handled by TCP in MQTT or block-wise transfer ”windowing” in advanced CoAP implementations. In the NB-IoT environment, a continuous stream can saturate the narrow downlink bandwidth, and if a packet is dropped, the retransmission of large segments can lead to ”head-of-line blocking,” where the entire stream stalls while waiting for the missing piece [7].

C. Chunk Size Consideration

In addition to the protocol and the overall OTA architecture approach, a careful selection of the chunk size has to be considered as it may influence the overall data exchange. Usually, the goal is to maximize the chunk size to the maximum MTU to effectively utilize the network and minimize the protocol overhead as it lowers the amount of individual messages that need to be transferred through the network.

However, such approach may prove counter-productive for the CIoT devices, in particular for those located in harsh radio conditions. The RF environment is characterized by non-negligible bit error rate (BER), which causes an increase in repetitions required for a successful data exchange with the probability of unsuccessful data exchange scaling up with the data size. In extreme cases where BER is high it may even result in a failure and data loss, forcing a retransmission from the higher layers and also possibly from the server, further increasing the overhead [9].

For battery-powered devices, the increased repetition impact may have a significant impact on the lifetime of the device. Therefore it is important do adapt the overall architecture with the end-device capabilities and conditions in mind [2].

IV. OTA Evaluation

In this section, we will evaluate the introduced OTA updates methods. We begin with the description of configuration and parameters and then proceed reporting the obtained results.

A. Configuration and Parameters

To evaluate various OTA protocols and approaches in the CIoT environment a testbed consisting of a Quectel BC660K and a ESP32-S3 was assembled, see Fig. 4. The Quectel BC660K is a NB-IoT communication module supporting both version of the technology (Cat-NB1 and Cat-NB2) providing a common AT interface over UART (Universal Asynchronous Receiver-Transmitter). The communication with the data server for the OTA updates was done solely through the module to mimic the real-world deployment. The ESP32-S3s role, as a main controller of the operation, is to command the module through the UART and process the OTA data. During the test, the ESP32-S3 was connected to the computer for monitoring and logs acquisition for later processing. The comparison was conducted in a stable environment with the RSRP of -85 dBm and +12 dB SINR.

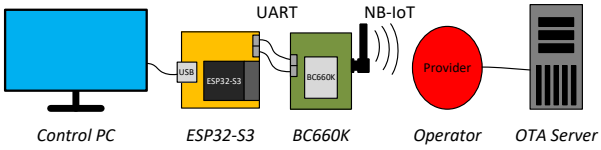

Fig. 4. The testbed for OTA evaluation.

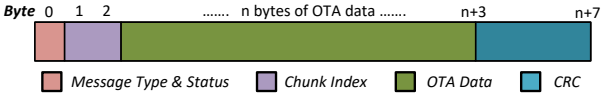

Fig. 5. Protocol used for transport layer evaluation.

The evaluation is based on downloading a **1 MB** firmware file utilizing various IoT protocols and OTA schemes through the NB-IoT technology. While such firmware size is not common among the simplest, low-power devices, some utilize a more capable MCUs/MPUs or complex firmware while utilizing low-bandwidth technology, such as NB-IoT, solely for purposes of status or accumulative data reporting, i.e., not live services.

B. Transport Layer Protocols Performance

For the transport layer protocols both OTA update methods, the streaming and polling mechanisms were evaluated. The serving server side was implemented via a simple custom python script providing a bare minimum protocol for the firmware transfer. The protocol implements a three byte header for traffic control, where one byte is reserved for the packet identification and type and the other two representing the current chunk of the data. At the end of each packet a 32 bits cyclic redundancy check (CRC) sequence is sent for data integrity purposes, as depicted in Fig. 5.

Depending on the transport protocol the application protocols may not be required, e.g., in cases of TCP, and due to the fixed, well-known size could be omitted from the results, however in cases of UDP it is essential to successfully update as some messages may get lost along the network chain and may lead, in worst cases, in catastrophic scenarios if the integrity is not verified.

The results depicted in Fig. 6 show the comparison of the underlying transport protocols for various chunk sizes and modes. It is evident that the transport-layer overhead is significantly higher for the TCP protocol due to its higher header size and a connection-oriented nature. Similar increase in overhead is present for the Polling architecture as additional request message has to be sent to the remote server, increasing the traffic volume. Further, the effect of a chunk size is evident where depending on it the overhead increases again as more messages are required to deliver the equal amount of useful data of the update binary.

C. Application Layer Protocols Performance

For the application protocols the measurement was conducted on custom-made servers tailored for maximum efficiency. As the amount of overhead is highly variable with protocols on the application layer, e.g. topic for the MQTT or Universal Resource Identified (URI) for the CoAP, the variables were optimized to minimize its effect for the comparison. However, as these variables are an integral part of the protocol and system integration, an experiment was conducted to compare the effect on the total overhead. For the OTA chunk distribution the inherent protocol mechanisms were utilized as efficiently as possible, e.g. a wildcard was used for the chunk order number for the MQTT protocol with client subscribed to it etc. For the CoAP and LWM2M protocols a CON with piggy-backed responses approach with 0-length token was chosen for maximum effectivity. Also, for MQTT a QoS 0 was utilized.

In Fig 7, showcasing the comparison results for each application protocol, a significant difference is presented. The MQTT protocol assumes the biggest data overhead for both Polling and Stream approaches due to the underlying TCP protocol which adds an unignorable amount of additional overhead (up to 146% for message size of 64 bytes). The UDP based protocols, CoAP and LWM2M show lower overhead where for the minimum message size an overhead of 48 % is achieved for the CoAP. The LWM2M shows additional overhead over the CoAP as it uses the CoAP protocol as an underlying protocol with well-defined URI.

Fig. 8 showcases the impact of the size of the variable fields in protocols MQTT and CoAP. The LWM2M is omitted as it

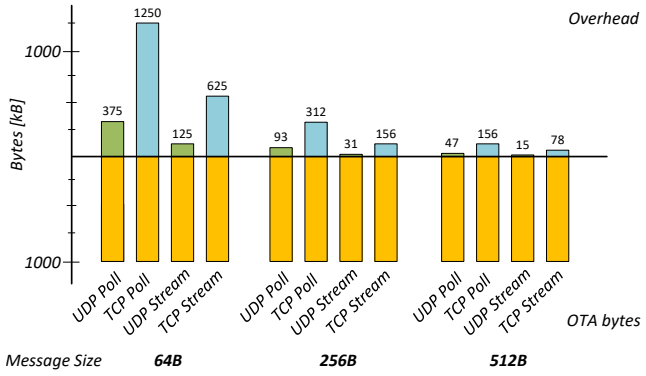

Fig. 6. Comparison of Transport protocols for OTA.

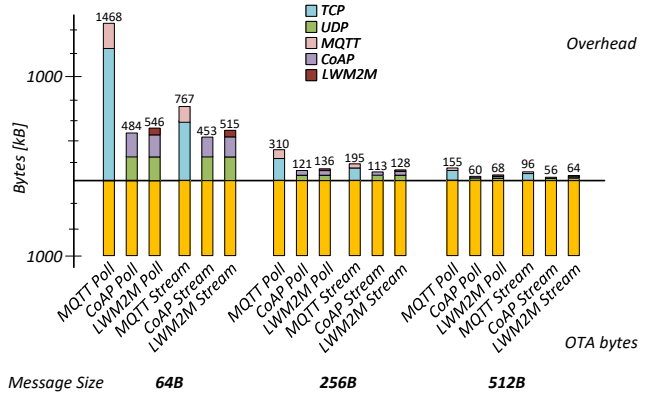

Fig. 7. Comparison of Application protocols for OTA.

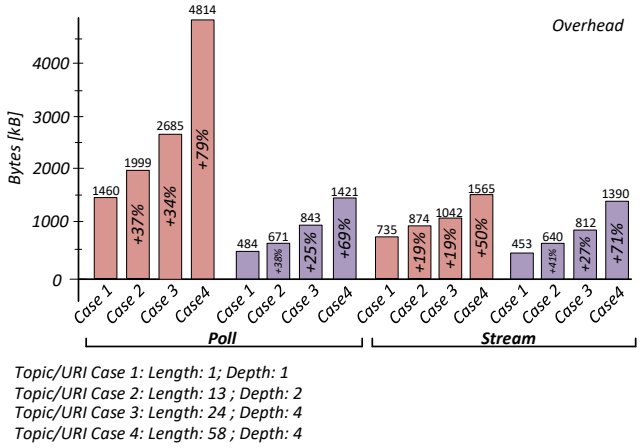

Fig. 8. Comparison of various topic/URI sizes.

is bound to a well-defined architecture designed by OMA.

Here the conducted evaluation consisted of 64B message (to enhance the visible difference) with various topic (MQTT) and URI (CoAP) lengths and depths to explore the significancy and effect on the total overhead.

Here, as per previous result the MQTT adds more total overhead over the CoAP. Interestingly, in cases of longer nested paths for URI and topics, the increase of CoAP overhead is comparatively higher than the MQTT for the streaming methid. In case of case 1 versus case 2, the MQTT added 50% of overhead due to the size change, however the CoAP protocol added 71% for the same length of the variable. This is caused by the fact that in the CoAP protocol header, the URI is not encoded as a single URI-path but rather a multiple individual URI-path options, each with designated option number and length, effectively increasing the header size.

V. Conclusion

Nowadays, the OTA updates are becoming increasingly important for CIoT technologies as more manufacturers enter the market and more use-cases are being added to the IoT managed/monitored assets. For CIoT technologies, specifically for NB-IoT, OTA updates not only creates additional significant traffic load potentially affecting performance guarantees of conventional mMTC services but may also incur additional costs for service providers. As a result, optimizing the overheads associated with OTA updates in CIoT systems is of special importance.

In this paper, several protocols common to the IoT industry, were selected for evaluation, specifically UDP, TCP, MQTT, CoAP and LWM2M. An experimental measurement was conducted to determine the typical transport and application overhead for each of the protocols utilizing the NB-IoT technology with different OTA approaches (i) polling and (ii) streaming. Our numerical results show that, the MQTT protocol, while being very popular, yielded the worst results in terms of overhead reaching up to 146 % of overhead vs. useful data for the smaller chunk sizes. The best results in terms of overhead were achieved with a bare UDP protocol, due to its small fixed 8 bytes header (averaging 12.5% of overhead for 64 bytes with the stream method). However the UDP does not provide any means of traffic steering and acknowledgements, therefore an additional protocol is required, which may not be easily integrable into the already existing platforms when custom protocol is used. From this standpoint the CoAP protocol seems to be a great all-rounded protocol, achieving about 46% of overhead for small, 64 byte size chunks and 6% for bigger size of 512B while providing the application with features such as confirmed messages etc. The chunk size also plays a crucial role in the total overhead. Bigger chunk size will be able to lower the amount of overhead, however the bigger the chunk size the lower the odds when delivering the data in harsh radio conditions. Therefore an adaptive approach based on the radio conditions is recommended.

In terms of incurred overheads, the streaming technique is superior however it is prone to failures where if no suitable recovery mechanism is implemented the process will have to be repeated potentialy causing even bigger overhead than the polling method.

Acknowledgment

This publication is co-financed with state support from the Technology Agency of the Czech Republic within the framework of the National Competence Centers Program. This publication is financed within the framework of the National Recovery Plan from the European Recovery and Resilience Instrument. Project FEIM, TN02000067.

References

[1] O. Liberg, M. Sundberg, Y.-P. E. Wang, J. Bergman, J. Sachs, and G. Wikström, *Cellular Internet of Things: From Massive Deployments to Critical 5G Applications*, second edition ed. Academic Press, 2020.

[2] F. Mahfoudhi, A. K. Sultania, and J. Famaey, "Over-the-air firmware updates for constrained nb-iot devices," *Sensors*, vol. 22, no. 19, p. 7572, 2022.

[3] P. Masek, J. Hosek, K. Zeman, M. Stusek, D. Kovac, P. Cika, J. Masek, S. Andreev, and F. Kröpfl, "Implementation of true iot vision: survey on enabling protocols and hands-on experience," *International Journal of Distributed Sensor Networks*, vol. 12, no. 4, p. 8160282, 2016.

[4] J. Postel, "User datagram protocol," Internet Requests for Comments, RFC Editor, RFC 768, August 1980. [Online]. Available: https://www.rfc-editor.org/info/rfc768

[5] F. Gont, P. Touch, R. Bush, and P. E. Tarjan, "Transmission control protocol (tcp) specification," Internet Requests for Comments, RFC Editor, RFC 9293, August 2022. [Online]. Available: https://www.rfc-editor.org/rfc/rfc9293.txt

[6] OASIS, "MQTT Version 5.0 OASIS Standard," https://docs.oasis-open.org/mqtt/mqtt/v5.0/mqtt-v5.0.pdf, 2019, Accessed: 2023-07-01.

[7] C. B. Z. Shelby, K. Hartke, "Common Format and MIME Type for Comma-Separated Values (CSV) Files," Internet Requests for Comments, RFC Editor, RFC 7252, 10 2014. [Online]. Available: https://datatracker.ietf.org/doc/html/rfc7252

[8] O. M. Alliance, "Lightweight Machine to Machine Technical Specification," *Approved Version*, vol. 1, no. 1, 2017.

[9] R. Dvorak, P. Palurik, T. D. Le, M. Stusek, P. Masek, D. Moltchanov, and J. Hosek, "Verification of communication parameters for ntn communication," in *2025 17th International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT)*. IEEE, 2025, pp. 87–93.

MESIT

GROUP OMNIPOL



Chceš být součástí týmu, který tvoří technologie pro budoucnost? Přidej se k nám!

POPIS SPOLEČNOSTI

MESIT (MEMBER OF OMNIPOL GROUP) je technologicky vyspělá společnost vyvíjející a vyrábějící **pokročilé systémy pro letectví, obranu a bezpečnost.** Naše technologie najdete v **letounech, obrněných vozidlech, ponorkách i v komunikačních systémech armád** po celém světě.
V našem vlastním **R&D** realizujeme projekty spojující **precizní inženýrství, inovace** a **vysoké nároky na kvalitu a spolehlivost** – vyvíjíme technologie, kde je klíčová odolnost, přesnost a dlouhodobá funkčnost.

HLEDÁME

- studenty a absolventy **technických oborů**
- zájemce o **letectví, vojenskou komunikaci, kybernetickou bezpečnost**
- motivované k práci na **reálných projektech**

Máte zájem vyzkoušet si moderní technologie pro výzkum včetně testovacích laboratoří?

NABÍZÍME

- **moderní technické zázemí** a nejnovější technologie
- práce na unikátních projektech
- možnost odborného rozvoje a kariérního růstu
- školení a jazykové kurzy
- **flexibilní pracovní podmínky**, HO, hybridní práce
- **5 týdnů dovolené**
- pracovní doba 37,5 hod/týden
- **příspěvek na stravování** a vlastní jídelna
- benefitní poukázky na **volnočasové aktivity**
- náborový příspěvek

PRŮBĚH VÝBĚROVÉHO ŘÍZENÍ

Zašlete nám svůj životopis a krátký motivační dopis a my se s vámi rádi sejdeme na osobní schůzce, kde probereme možnosti další spolupráce.

MESIT holding, a.s., Sokolovská 573, Mařatice, 686 01 Uherské Hradiště, Česká republika

Kontaktní osoba: Bc. Hana Bartošková
personalni@mesit.cz, 702 243 528

Nabízíme spolupráci pro obory:

- Strojní
- Elektro
- Prům. inženýrství
- Informatika
- Ekonomie
- Chemie

Ve formě:

- Bakalářka
- Diplomka
- Praxe/stáž

Možnosti během studia:

- praxe/stáž
- DP a BP

Nastartuj svou kariéru u nás!

www.mesit.cz

ML-Driven Real-Time In-Network Intrusion Classification on P4-Programmable Switches

1st Viet Anh Phan
Department of Telecommunications
Brno University of Technology
616 00 Brno, Czech Republic
243760@vut.cz

2nd Jan Jerabek
Department of Telecommunications
Brno University of Technology
616 00 Brno, Czech Republic
jan.jerabek@vut.cz

***Abstract*—Network intrusion detection is a cornerstone of modern cybersecurity, yet existing P4-based approaches have a critical blind spot. A lot of research relies on packet-level features like header fields and individual packet sizes, which lack the discriminative power to capture behavioral patterns across a flow's lifetime. Feature extraction at the flow level within the P4 data plane has been almost entirely overlooked, with only a single prior study attempting it. Moreover, prior work is limited to a narrow selection of ML models and relies on exact matching in Match-Action tables, which cannot handle the large value ranges inherent in real network flow features. This paper presents a unified framework that closes these gaps. We implement bidirectional flow feature extraction directly in the P4 data plane, computing per-flow statistical features in line with traffic processing. To meet P4 hardware's strict match-field constraints, we encode ML classifiers as range match tables, which accommodate continuous and wide-ranging feature values while substantially reducing the number of required table entries compared to exact matching. The framework supports three classifiers, namely Decision Tree, Random Forest, and XGBoost, and covers the full pipeline of the intrusion detection system.**

***Index Terms*—P4, programmable data plane, in-network machine learning, in-network feature extraction, attack detection, PCA, Decision Tree, Random Forest, XGBoost**

I. Introduction

The programmability of P4 switches has motivated a growing body of research on offloading security functions directly to the data plane. Early work demonstrated that real-time DDoS detection using entropy estimation and statistical thresholds can be performed entirely within P4 pipelines [1] [2], with subsequent systems achieving full in-network mitigation using information-theoretic analysis and novel P4-compatible logarithmic approximation algorithms [3] [4]. In parallel, researchers have explored embedding trained machine learning (ML) models into the match-action tables of programmable switches. Pioneering frameworks such as IIsy mapped Decision Tree (DT), SVM, and Naive Bayes classifiers to per-feature match-action tables [5], while SwitchTree demonstrated Random Forest (RF) inference at line speed for attack detection [6]. More recently, Planter has emerged as a modular, open-source toolchain supporting a wide range of ML algorithms across multiple hardware targets, including the Intel Tofino [7].

This work is supported by Ministry of Industry and Trade of the Czech Republic under Grant CZ.01.01.01/01/24_062/0007470.

Despite the promising results achieved by existing approaches, several critical challenges remain unresolved in P4-based attack detection research. First, the majority of current systems rely on packet-by-packet or unidirectional feature extraction, which captures only a partial view of network communication and fails to leverage the bidirectional flow-level statistics that are essential for distinguishing sophisticated attack patterns from benign traffic. Second, existing implementations predominantly employ exact-match tables to encode ML model logic, which leads to a rapid overflow of table entries as the feature space or model complexity grows, severely limiting scalability on hardware targets with constrained Static Random Access Memory (SRAM) and Ternary content-addressable memory (TCAM) resources. Third, the restricted number of pipeline stages in programmable switch Application-Specific Integrated Circuit (ASIC) inherently limits the depth of ML models that can be deployed, constraining the expressiveness and classification accuracy achievable within the data plane. Fourth, while pioneering frameworks have demonstrated the feasibility of embedding individual ML algorithms into match-action pipelines, the implementation of multiple diverse ML models following a unified, systematic approach remains largely unexplored, with most works targeting a single algorithm rather than providing a generalizable methodology.

This paper addresses these limitations by proposing a real-time in-network multiple attacks detection system on P4-programmable switches that introduces bidirectional flow feature extraction, employs range-match encoding to mitigate table entry overflow, and implements several ML models through a novel, unified inference pipeline, thereby advancing the state of the art toward practical, scalable, and accurate in-network security. The source code and implementation of the proposed system are available at: https://github.com/vafekt/p4sec.git

II. System Architecture

The proposed Intrusion detection system (IDS) is structured as a two-plane architecture illustrated conceptually in Fig. 1 and 2. The *control plane* handles all offline, compute-intensive work: PCAP-based feature extraction, model training (Decision Tree, Random Forest, or

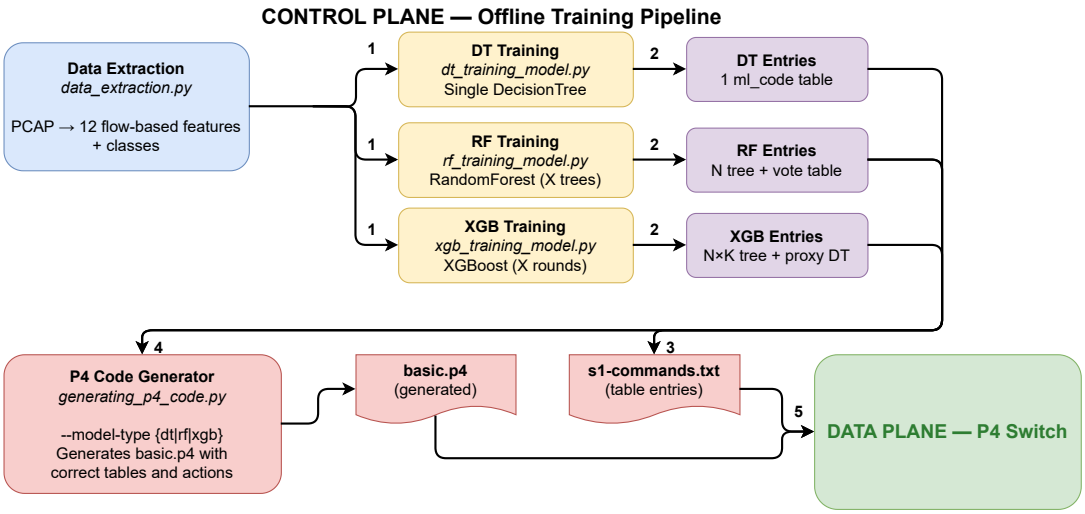

Fig. 1: In-Network ML-Based Intrusion Detection System Framework - Control Plane

XGBoost), and the translation of the trained model into P4 `table_add` commands. The *data plane* is implemented on a P4-programmable BMv2 software switch (v1model architecture); it performs real-time bidirectional flow tracking in register arrays, applies the pre-compiled ML classifier through range-match table lookups, and forwards per-flow digests to the P4Runtime controller for logging.

A central design requirement of the proposed framework is that the features used to train the ML model (in the control plane Fig. 1) and the features computed at classification time (in the data plane Fig. 2) are *identical by construction*. To enforce this, the same extraction logic is implemented twice: once in Python (control plane, offline) to build the labelled training dataset from PCAP captures, and also in P4 (data plane, online) to maintain live flow statistics in register arrays and emit them as digest messages. Both implementations share the same canonical direction convention, the same 5-tuple flow key, the same timeout value, and the same termination conditions, so there is no train/inference distribution mismatch arising from feature-computation differences.

A. Flow Identification: the Canonical 5-Tuple

Each network flow is uniquely identified by a 5-tuple:

$$\mathbf{k} = (\text{SrcIP},\ \text{SrcPort},\ \text{DstIP},\ \text{DstPort},\ \text{Protocol}). \quad (1)$$

Because a bidirectional flow produces packets in both directions, the 5-tuple must be *normalised* so that packets from A→B and B→A map to the same flow record. The canonical form assigns the endpoint with the numerically smaller IP address as the "source":

$$\text{canonical}(a,b) = \begin{cases} (a,b) & \text{if } \texttt{int}(a.\text{IP}) \le \texttt{int}(b.\text{IP}), \\ (b,a) & \text{otherwise.} \end{cases} \quad (2)$$

A packet is labelled *forward* when its source matches the canonical source, and *backward* otherwise. This direction flag drives the asymmetric feature counters (FwdPktCount, BwdPktCount, FwdBytes, BwdBytes) that capture communication asymmetry. This is a property that packet-level or unidirectional flow-level approaches cannot exploit.

B. Feature Extraction

The extracted features fall into 4 groups, summarised in Table I.

TABLE I: Flow Feature Set

#	Feature	Group	Description
Identifiers (not used as ML features)			
–	SrcIP, DstIP	identifier	Canonical endpoint addresses
–	SrcPort, DstPort	identifier	Associated transport ports
–	Protocol	identifier	IP protocol number (TCP=6, UDP=17)
ML features			
1	Duration	temporal	Time from first to last packet (ns)
2	MaxIAT	temporal	Maximum inter-arrival time (ns)
3	FwdPktCount	volume	Packet count, forward direction
4	BwdPktCount	volume	Packet count, backward direction
5	FwdBytes	volume	Total bytes, forward direction
6	BwdBytes	volume	Total bytes, backward direction
7	MaxWinSize	TCP/UDP	Maximum TCP window size observed
8	UrgCount	TCP/UDP	Packets with URG flag set
9	FlagsSyn	TCP/UDP	Packets with SYN flag set
10	FlagsAck	TCP/UDP	Packets with ACK set
11	FlagsFin	TCP/UDP	Packets with FIN flag set
12	FlagsRst	TCP/UDP	Packets with RST flag set

The combination of bidirectional volume counters, timing statistics, and per-flag counts is deliberately compact: it fits within the register-array and metadata constraints of a P4 target while remaining semantically rich enough to discriminate anomalous traffic patterns.

A flow record is initialised on the first packet matching its canonical 5-tuple. The record is finalized and its feature vector emitted, under three conditions:

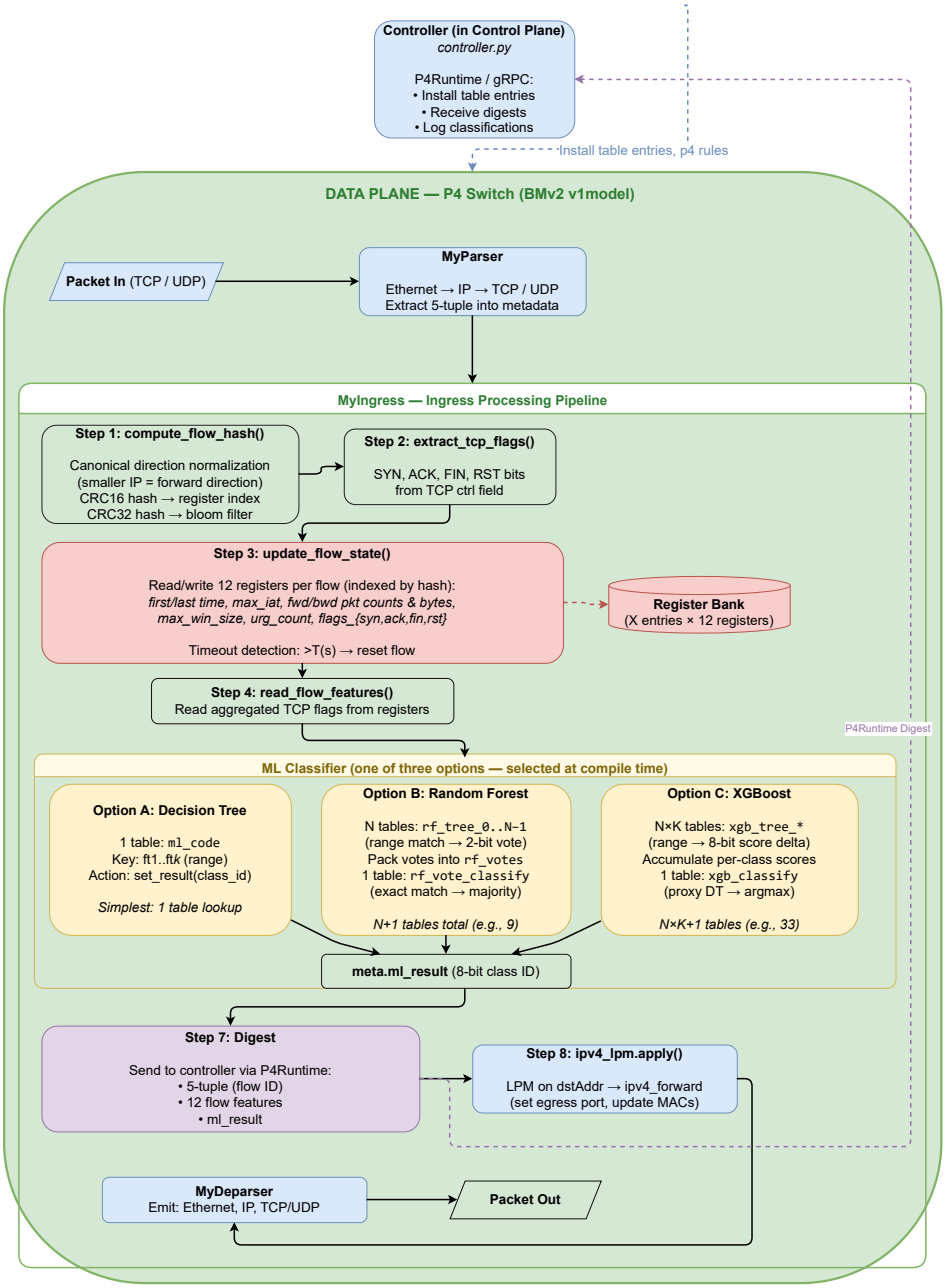

Fig. 2: In-Network ML-Based Intrusion Detection System Framework - Data Plane

1) **TCP close:** a packet carries a FIN or RST flag.
2) **Idle timeout:** no packet belonging to the flow has been observed for $T_{\text{idle}} = 120\,\text{s}$.
3) **End of capture:** all captured packets have been processed.

As shown in Table II, existing P4-based detection systems each address only a subset of the properties required for general-purpose in-network intrusion detection. Planter [7], P4Pir [8] and P4DDLe [9] operate at the packet level, making them structurally incapable of observing behavioural

TABLE II: Comparison of P4-Based In-Network Feature Extraction Approaches

Work	Granularity	Bidirectional	Num of Features	Coll. avoid.
Planter [7]	packet	×	var.	×
SwitchTree [6]	flow	×	7	✓
P4Pir [8]	packet	×	7	×
P4DDLe [9]	packet	×	var.	✓
Proposed	flow	✓	**12**	✓

var. = framework or use case dependent.

patterns that emerge only across the lifetime of a flow, such as byte asymmetry, inter-arrival timing, or TCP handshake anomalies. SwitchTree [6] is the closest in spirit to the proposed work, extracting 7 per-flow features (duration, jitter, inter-arrival of packets, round-trip time, bitrate, average packet size and loss rates) at line rate; however, it accumulates those features in a single traffic direction, so directional asymmetry between a client and a server remains unobservable. The proposed system is the only approach in this comparison that combines bidirectional per-flow feature extraction, a 12-feature set computed entirely inside the P4 dataplane, and explicit hash collision avoidance through a dual-hash Bloom filter, addressing the individual limitations of all four prior systems simultaneously.

C. Machine Learning Implementation

All three classifiers (DT, RF and XGB) share the same deployment strategy. A trained model is inspected in the control plane and every decision path is converted into a P4 `table_add` command carrying *range-match* keys on the quantised PC-code features. The resulting text file (`s1-commands.txt`) is replayed by the P4Runtime controller at startup, populating the switch's TCAM tables. At classification time the dataplane performs a single pass through the loaded tables, in which there are no floating-point arithmetic, no loops, no external calls and produces an 8-bit class identifier stored in `meta.ml_result`. The three models differ in the number of tables they require and in how the final class decision is formed.

A trained `Decision Tree Classifier` (performed in the control plane Fig. 1) is a binary tree $\mathcal{T}$ in which every internal node tests one feature against a threshold θ. Each root-to-leaf path defines a hyper-rectangular region of the feature space and a class label. Algorithm 1 converts each such path into one range-match rule.

The conversion from a floating-point threshold θ to integer P4 bounds follows the convention $(\ell, h] \to [\ell{+}1, h]$, consistent with "feature $\leq$ θ" split semantics. Priority is assigned inversely to the total range width so that the most specific rule wins in the event of overlap. The entire DT requires exactly **one P4 table** (`ml_code`).

The key enabler of the proposed in-network ML implementation is the use of *ternary range-match* tables in the P4 dataplane rather than exact-match hash tables (used by all other previous research). This section explains the difference using a concrete example derived from the trained DT.

Consider a DT with depth 3, 5 features, 4 classes, and 8 leaves, as illustrated in Fig. 3. Each root-to-leaf path defines a set of inequality constraints on individual features; features not tested along a given path are unconstrained. The range-match encoding maps each leaf directly to one P4 table entry: constrained features carry their derived lower and upper bounds, while unconstrained features are assigned the full wildcard range $[0, \texttt{MAX_VAL}]$. In this example, features are represented using 16-bit values, therefore $\texttt{MAX_VAL} = 2^{16} - 1 = 65\,535$, which defines the maximum representable value in the feature domain. All specific numeric bounds appearing in the branches are illustrative values produced by a particular decision tree instance trained by the DT model for the given dataset.

Algorithm 1 DT Path-to-P4-Rule Conversion

Require: Trained DT $\mathcal{T}$, feature list $\mathbf{f}$, $\texttt{MAX_VAL} = 2^B - 1$
Ensure: `table_add` entries for `MyIngress.ml_code`

rules $\leftarrow [\,]$
function TRAVERSE(node n, path π)
  if n is a leaf **then**
    $c \leftarrow \arg\max\big(\text{value}[n]\big)$
    $D \leftarrow$ MINIMISE(π) ▷ tighten per-feature bounds
    append (D, c) to rules
    return
  end if
  TRAVERSE(n_{left}, $\pi \cup \{(f_j, \leq, \theta)\}$)
  TRAVERSE(n_{right}, $\pi \cup \{(f_j, >, \theta)\}$)
end function
TRAVERSE(root, $\emptyset$)
Sort rules by ascending total range width ▷ specificity order
for each rule (D, c) with priority p **do**
  for each feature f_j **do**
    $\ell_j \leftarrow \lfloor D[f_j].\min \rfloor + 1$ (or 0 if unbounded below)
    $h_j \leftarrow \lfloor D[f_j].\max \rfloor$ (or `MAX_VAL` if unbounded above)
  end for
  Emit: `table_add MyIngress.ml_code set_result`
    $[\ell_1 \to h_1]\ [\ell_2 \to h_2]\ \cdots\ [\ell_k \to h_k]\ \Rightarrow\ c\ \ p$
end for

The result is that **8 leaves produce exactly 8 P4 table entries**, regardless of the quantisation precision B. Each entry covers an entire axis-aligned hyper-rectangle of the feature space in a single TCAM lookup, and unused feature dimensions are handled by wildcards at zero additional cost. An exact-match table cannot express inequalities or wildcards: every possible combination of feature values that should map to a given class must be enumerated as its own row. For k features each quantised to B bits, the number of entries required to cover the same feature space is:

$$|\mathcal{E}_{\text{exact}}| = 2^{k \cdot B}. \tag{3}$$

The principles of RF and XGB are inherited from DT-based algorithm and, due to the limited scope of this paper, their full implementation details cannot be described here; however, they are available in the accompanying GitHub project.

III. EVALUATION

The system is evaluated using the *Mu-IoT* dataset [10], a publicly available IoT network traffic benchmark collected at Manipur University. The dataset includes multiple attack types that are grouped into four traffic categories: *Benign* (normal

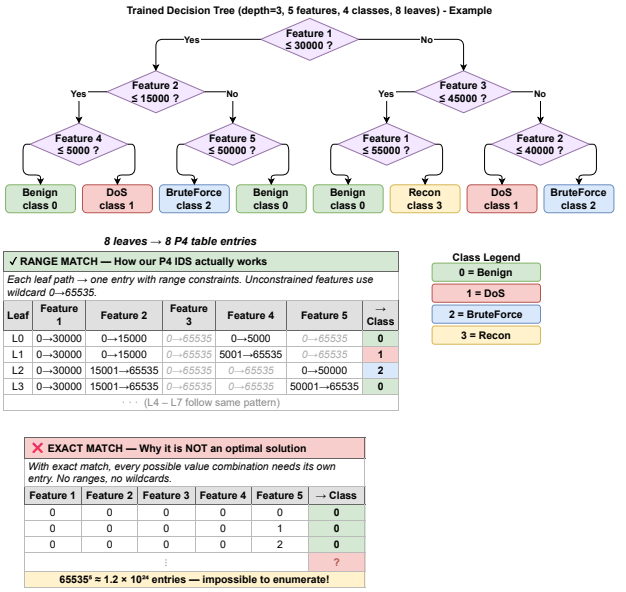

Fig. 3: Principle of Generating Entries to P4 Switch (DT model) and Approach comparison - Simplified Example

traffic) and three attack classes - *DDoS* (volumetric flooding attacks), *PasswordHacking* (brute-force credential attacks), and *Reconnaissance* (network scanning activities). In total, the dataset contains approximately 22.8 GB of raw network traffic.

All three classifiers use 16-bit quantization for range encoding in the P4 match tables. For Random Forest, the ensemble is fixed at 8 trees to keep the number of generated table entries within practical hardware limits. Table III reports precision, recall, and F1-score per class for each classifier on the held-out test set.

TABLE III: Per-Class Classification Results on the Mu-IoT Test Set

	Decision Tree			Random Forest		
Class	Precision	Recall	F1	Precision	Recall	F1
Benign	0.843	0.984	0.908	0.875	0.990	0.929
DDoS	0.993	0.989	0.991	0.994	0.995	0.994
PwdHacking	0.957	0.839	0.894	0.974	0.872	0.920
Reconnaissance	0.995	0.917	0.954	1.000	0.929	0.963
Accuracy		95.72%			96.24%	
Macro F1		0.937			0.952	

	XGBoost		
Class	Precision	Recall	F1
Benign	0.915	0.983	0.948
DDoS	0.991	0.983	0.987
PwdHacking	0.929	0.866	0.896
Reconnaissance	0.994	0.994	0.994
Accuracy		96.86%	
Macro F1		0.956	

All three models perform strongly overall. DDoS is the best detected class across every model, with F1 scores of 0.991, 0.994, and 0.987 for DT, RF, and XGBoost respectively, reflecting that volumetric attacks produce highly consistent flow patterns. Reconnaissance also classifies well, reaching perfect precision under RF (1.000) and near perfect F1 under XGBoost (0.994).

Across overall accuracy, XGBoost leads at 96.86%, followed by RF at 96.24% and DT at 95.72%. The macro F1 trend is consistent: XGBoost 0.956, RF 0.952, DT 0.937. The gap between DT and the two ensemble methods is modest but consistent, suggesting that even a single tree encodes sufficient decision boundaries given the flow level features, while the ensemble methods provide a small but measurable gain particularly on the harder classes.

Conclusion

This paper presented a fully in-network intrusion detection framework that compiles trained ML classifiers into P4 range match table entries executable at line rate on a programmable switch. Bidirectional flow features are extracted using identical logic in both the offline training pipeline and the live P4 dataplane, ensuring consistent feature distributions at inference time. Range match encoding maps each tree leaf to a set of table entries defined over the original feature space, accommodating continuous and wide ranging feature values without exact matching. Evaluation on the Mu-IoT dataset yielded overall accuracies of 95.72%, 96.24%, and 96.86% for Decision Tree, Random Forest, and XGBoost respectively. Future work will investigate richer feature representations, and deployment on Intel Tofino hardware.

References

[1] Lapolli A.C., Adilson Marques J., Gaspary L.P. Offloading real-time DDoS attack detection to programmable data planes (2019) 2019 IFIP/IEEE Symposium on Integrated Network and Service Management, IM 2019, art. no. 8717869, pp. 19 - 27

[2] Ilha A.D.S., Lapolli A.C., Marques J.A., Gaspary L.P. Euclid: A Fully In-Network, P4-Based Approach for Real-Time DDoS Attack Detection and Mitigation (2021) IEEE Transactions on Network and Service Management, 18 (3), pp. 3121 - 3139 DOI: 10.1109/TNSM.2020.3048265

[3] Ding D., Savi M., Siracusa D. Tracking Normalized Network Traffic Entropy to Detect DDoS Attacks in P4 (2022) IEEE Transactions on Dependable and Secure Computing, 19 (6), pp. 4019 - 4031 DOI: 10.1109/TDSC.2021.3116345

[4] Gao S., Handley M., Vissicchio S. Stats 101 in P4: Towards In-Switch Anomaly Detection (2021) HotNets 2021 - Proceedings of the 20th ACM Workshop on Hot Topics in Networks, pp. 84 - 90 DOI: 10.1145/3484266.3487370

[5] Xiong Z., Zilberman N. Do switches dream of machine learning?: Toward in-network classification (2019) HotNets 2019 - Proceedings of the 18th ACM Workshop on Hot Topics in Networks, pp. 25 - 33 DOI: 10.1145/3365609.3365864

[6] Lee J.-H., Singh K. SwitchTree: in-network computing and traffic analyses with Random Forests (2025) Neural Computing and Applications, 37 (28), pp. 23143 - 23154 DOI: 10.1007/s00521-020-05440-2

[7] Zheng C., Zang M., Hong X., Perreault L., Bensoussane R., Vargaftik S., Ben-Itzhak Y., Zilberman N. Planter: Rapid Prototyping of In-Network Machine Learning Inference (2024) Computer Communication Review, 54 (1), pp. 2 - 20 DOI: 10.1145/3687230.3687232

[8] M. Zang, C. Zheng, L. Dittmann and N. Zilberman, "Toward Continuous Threat Defense: in-Network Traffic Analysis for IoT Gateways," in IEEE Internet of Things Journal, vol. 11, no. 6, pp. 9244-9257, 15 March15, 2024, doi: 10.1109/JIOT.2023.3323771

[9] Doriguzzi-Corin R., Knob L.A.D., Mendozzi L., Siracusa D., Savi M. Introducing packet-level analysis in programmable data planes to advance Network Intrusion Detection (2024) Computer Networks, 239, art. no. 110162 DOI: 10.1016/j.comnet.2023.110162

[10] U. B. Clinton and N. Hoque, "MU-IoT: A New IoT Intrusion Dataset for Network and Application Layer Attacks Analysis," in IEEE Access, vol. 12, pp. 166068-166092, 2024, doi: 10.1109/ACCESS.2024.3494052

Interactive Graph-based Environment for Open Source Intelligence

1st Willi Lazarov
Department of Telecommunications
FEEC, Brno University of Technology
Brno, Czechia
lazarov@vut.cz

2nd Roman Klampar
Department of Telecommunications
FEEC, Brno University of Technology
Brno, Czechia
xklamp01@vut.cz

3rd Zdenek Martinasek
Department of Telecommunications
FEEC, Brno University of Technology
Brno, Czechia
martinasek@vut.cz

***Abstract*—With the ever-increasing volume of data along with the proliferation of digital services and their users, the information contained within cyberspace has become a valuable and integral commodity of the modern world. In a security context, these data represent a critical element for the detection, prevention, and investigation of threats across domains such as cybersecurity, criminal justice, national security, military intelligence, and others. However, collecting data from open sources is a particularly challenging process given their unstructured nature, the prevalence of duplicates, noise, and other adverse quality factors. The analysis of such large volumes of data poses further practical challenges, especially when it involves searching for and correlating information across multiple heterogeneous sources. These challenges are equally pertinent to Open Source Intelligence (OSINT) – a discipline focused on the systematic collection, analysis, and interpretation of publicly available data. This paper addresses the aforementioned challenges by proposing an interactive environment that enhances analytical accuracy through graph-based structures, enabling analysts to visualize, correlate, and verify intelligence information.**

***Index Terms*—OSINT, open-source data, cyberspace, graph-based analysis, intelligence gathering**

I. Introduction

The internet and all other subsets of cyberspace have become the single largest medium for information exchange on a global scale. Current estimates suggest that approximately 403×10^6 terabytes of data are generated every day, a figure that continues to grow each year [1]. These data take numerous forms, including text with semantic content, images, video, audio, and many others. However, the majority of these data remain unstructured and may contain duplicates, errors, or other quality deficiencies [2]. In the context of cybersecurity, the information embedded within these data plays a critical role in tasks such as cyberattack detection, threat actor deanonymization, criminal investigation, and beyond. A substantial portion of the overall data landscape consists of publicly accessible, open-source data and information. For investigative and preventive purposes, systematic collection and analysis of such data are essential for producing actionable intelligence toward a defined objective. The discipline that addresses this need is known as Open Source Intelligence (OSINT).

The research described in this paper was financially supported by the Technology Agency of the Czech Republic, project No. FW11020057.

As an intelligence discipline, OSINT encompasses the end-to-end process of gathering, processing, and deriving actionable meaning from publicly available data. Rather than collecting raw data alone, it more precisely denotes information that has undergone systematic identification, evaluation, and selection before being delivered to relevant stakeholders [3]. The applicability of OSINT spans a broad spectrum of fields, including intelligence agencies, law enforcement, cybersecurity, investigative journalism, and academic research, among others [4]. Despite its capabilities, OSINT faces several persistent challenges stemming from the large volume, varied structures, and noise present in open-source data. This paper addresses these challenges through the design, implementation, and verification of an interactive environment that integrates graph-based analysis to support OSINT investigations.

II. Background and Related Work

OSINT is one of several established intelligence disciplines, each defined by its primary data source and collection methodology. It involves systematic collection, analysis, and interpretation of publicly available data from digital sources such as social networks, websites, public archives, and metadata [3]. Other notable disciplines include Human Intelligence (HUMINT), deriving information from interpersonal sources; Social Media Intelligence (SOCMINT), focusing on gathering and analyzing data from social media platforms; and Geospatial Intelligence (GEOINT), targeting imagery and location-based data [5]. Although the boundaries of these disciplines are formally distinct, they frequently overlap in practice. For instance, SOCMINT analysis may simultaneously include OSINT, HUMINT, and GEOINT techniques, depending on the nature and objectives of the investigation.

The OSINT process distinguishes four core components according to their increasing intelligence value. At the base level, Open Source Data (OSD) represent raw, unprocessed data collected from open sources. Once these data are structured and contextualized, they become Open Source Information (OSIF). When OSIF is further discovered, selected, and exploited for its intelligence relevance, it constitutes OSINT in the proper sense. Finally, Validated OSINT (OSINT-V) denotes intelligence that has been corroborated through reputable sources, thus increasing its certainty and credibility [5].

In relation to the work presented in this paper, we provide a brief overview of existing research in the OSINT domain, with a particular focus on graph-based data analysis and visualization, as these represent the core functions of our proposed solution. In the academic literature, research utilizing graph theory to support OSINT investigations remains limited. The most relevant identified work includes knowledge graphs for fact-checking [6], graph-based visualization of OSINT data for analyzing homicide cases [7], and combining graphs with large language models to improve intelligence insights [8]. Despite the limited amount of research utilizing graph theory for OSINT investigations, we observed a larger number of papers (for example, [9]–[11]) using graph-based structures and analysis for cyber threat intelligence (CTI). The common objective observed was to aggregate, interconnect, and visualize data for modeling threat actors, campaigns, and attack paths to support CTI analysis and attribution.

III. Graph-based Environment

To address the research gap identified in Section II and streamline data analysis and visualization within OSINT investigations, this paper presents the design and implementation of an interactive environment based on graph structures. The primary objective of the graph-based environment is to visualize the OSIF entities and their relationships, further enhanced by fully interactive and customizable components. In the following subsections, we present our solution, from its architecture to integration into OSINT practice.

A. Formal Definition

To define the proposed analytical environment, we formalize its underlying data model. The environment operates over OSIF entities represented using a graph-based structure that captures two fundamental types of relationship: hierarchical connections reflecting parent-child decompositions and associative connections representing arbitrary semantic or contextual links between entities.

Let the OSINT analytical environment be defined as a directed attributed multigraph

$$\mathcal{E} = (V, E, \tau_E, \alpha), \tag{1}$$

where:

- $V = \{v_1, v_2, \ldots, v_n\}$ is a finite, non-empty set of nodes, each representing an individual OSIF entity;
- $E \subseteq V \times V$ is a set of directed edges representing relationships between entities;
- $\tau_E : E \to \mathcal{T}_E$ is an edge-typing function, where $\mathcal{T}_E = \{h, a\}$ is the set of edge types with h denoting a hierarchical (parent–child) edge and a denoting an associative (non-hierarchical) edge;
- $\alpha : V \to 2^A$ is an attribute function that maps each node $v \in V$ to a set of key–value pairs $\alpha(v) \subseteq A$, where $A = \{(k, val) \mid k \in \mathcal{K},\ val \in \mathcal{V}_k\}$ is the universe of admissible attributes defined over attribute keys $\mathcal{K}$ and their corresponding value domains $\mathcal{V}_k$.

The edge set is partitioned into hierarchical edges $E_h = \{e \in E \mid \tau_E(e) = h\}$ and associative edges $E_a = \{e \in E \mid \tau_E(e) = a\}$ such that $E = E_h \dot{\cup} E_a$. The subgraph (V, E_h) forms a directed forest, meaning each node has at most one hierarchical parent:

$$\forall\, v_i \in V : \big|\{v_k \in V \mid (v_k, v_i) \in E_h\}\big| \leq 1. \tag{2}$$

For any node $v_i \in V$, the set of its children is defined as $C(v_i) = \{v_j \in V \mid (v_i, v_j) \in E_h\}$. A node with $C(v_i) = \emptyset$ is a leaf, while a node without a hierarchical parent is a root. In contrast, associative edges do not impose structural constraints, which means that an edge $(v_i, v_j) \in E_a$ can exist between any two distinct nodes regardless of their hierarchical relationship, allowing the graph environment to capture semantic connections without disrupting tree decomposition.

Figure 1 illustrates an example of graph $\mathcal{E}$ that models an OSINT investigation of a person of interest. The root node v_1 decomposes hierarchically into the investigated subject v_2 (Person), a relative v_3 (Person), and the subject's employer v_4 (Organization), which are further decomposed into social accounts, vehicles, buildings, events, photographs, geolocations, and news articles. The resulting directed forest (V, E_h) contains $|V| = 15$ nodes and $|E_h| = 14$ hierarchical edges satisfying Equation (2). The six associative edges in the graph then capture the OSIF relations. For instance, $e_1^a = (v_5, v_7)$ links two social accounts from different nodes, revealing a mutual connection that established the family relationship $e_2^a = (v_2, v_3)$, while $e_5^a = (v_{12}, v_{13})$ connects two geolocations whose coincidence places the subject at the relative's residence. The attribute box in the figure exemplifies the mapping $\alpha(v_2)$ with the attributes of v_2.

B. Architecture and Implementation

Based on the formal definition in Section III-A, we propose an architecture and rendering system that transforms the multigraph model into an interactive visual environment. The key challenge for this application is to efficiently render potentially large instances of multigraphs with hundreds of nodes, hierarchical decompositions, and associative relations while preserving real-time interactivity. In addition, the environment is not limited only to visual exploration of the graph but enables modification of all OSIF entities, making it a fully interactive workspace for OSINT investigations.

To address this, the selection of the appropriate underlying technology was the most important factor. Standard libraries utilizing Document Object Model (DOM) or Scalable Vector Graphics (SVG) operate in a retained mode. In this paradigm, the browser maintains the entire object tree in memory. When recalculating physical models involving thousands of entities, this approach drastically reduces the frame rate (FPS) and leads to an increase in memory consumption. The solution to this bottleneck is the low-level HTML5 Canvas API, which operates in immediate mode. In these settings, visual elements are rendered directly into a raster bitmap, which means that the system does not retain complex object states, thus ensuring smooth rendering even with massive datasets.

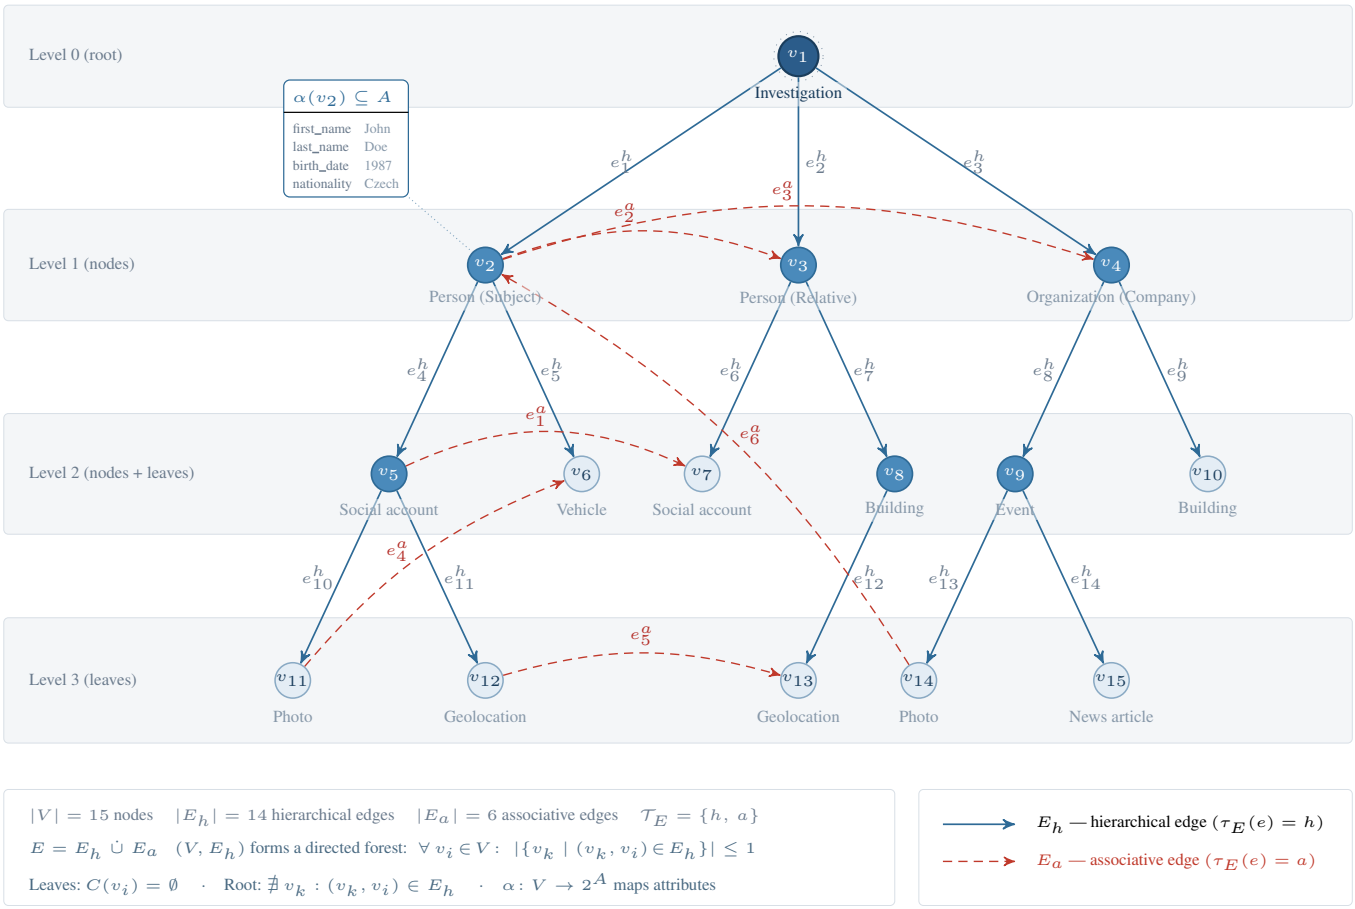

Fig. 1. Example instance of the directed attributed multigraph $\mathcal{E}$ for modeling an OSINT investigation

Based on the HTML5 Canvas API as the core of the system, we developed a custom rendering engine to fully implement the proposed graph-based environment. It strictly separates visual rendering from computational models, adhering to the principle of Separation of Concerns (SoC), and provides full control over the definition of specific shapes and rendering algorithms used in this setting.

The architecture employs a compositional approach where graph entities are not monolithic objects but are dynamically assembled from independent geometric primitives. This ensures high modularity when visualizing diverse OSINT data. Performance is further optimized through multi-canvas layering where the static grid background is isolated from the main graph, minimizing the frequency of background redraws. The interactive layer containing the graphical elements is managed by a reactive render loop. This loop is triggered exclusively upon explicit state change requests, and its CPU load is additionally mitigated by optimization techniques such as rendering frequency regulation (throttling).

An additional advantage of this solution is that the HTML5 Canvas API is natively supported by all modern web browsers. The logical workflow of the rendering system is illustrated in Figure 2 and described together with the key interactive features in Section III-C.

C. Interactive Rendering System

The spatial layout of the graph is handled by the D3.js library, which uses the d3-force simulation to calculate physical forces like collision detection (node overlap), node repulsion, and edge-length distance constraints. To maintain the optimal performance of the rendering engine, D3.js is strictly implemented as a stateless mathematical module. It operates entirely in system memory, calculating and updating the provided X and Y coordinates without any interaction with the canvas rendering context. This modular, plugin-based architecture ensures that the current layout algorithm can be easily replaced or extended with alternative approaches (e.g., hierarchical DAG models) in the future.

Although immediate-mode rendering brings significant performance benefits, it lacks native support for event listeners on individual elements. Implementing global interactive features, such as camera panning, zooming, and node manipulation, along with granular interactions such as hovering over an element or selecting a specific curved edge, requires a robust hit detection mechanism. Standard coordinate intersection algorithms have a time complexity of $\mathcal{O}(m \cdot n)$, which rapidly degrades system responsiveness when processing thousands of active elements in the graph layout.

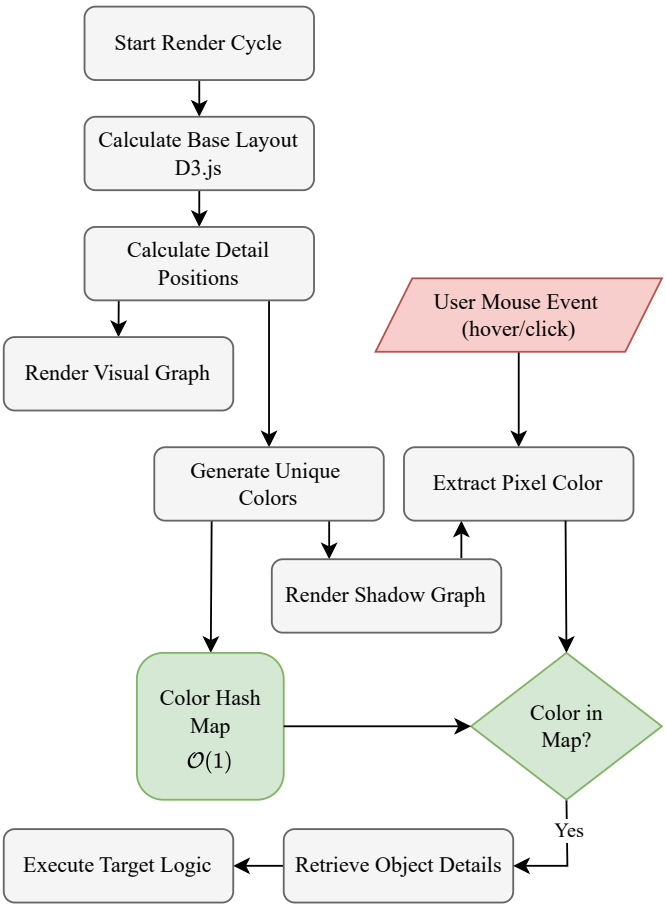

Fig. 2. Workflow of the rendering system

To overcome this limitation, an *Off-screen Hit Detection* method was implemented. Alongside the main visual output, the render loop generates a hidden shadow graph in memory. Each interactive element is assigned a unique color generated using a bitwise shift, which serves as a key in an object hash map. During interaction, the system extracts the pixel color from the hidden canvas. This approach reduces the complexity to a constant $\mathcal{O}(1)$, enabling not only immediate click detection but also real-time hover evaluation without a negative impact on the frame rate.

Rendering massive datasets also introduces unwanted visual clutter when zooming out. To maintain smooth performance, a Level of Detail (LoD) optimization was integrated. When the zoom scale drops below the defined limit, the engine automatically filters out computationally expensive details (icons, texts, backgrounds) and renders only the basic geometric primitives.

While the main purpose of the rendering engine is to maximize performance, its greatest benefit is providing rich interactivity. The system was designed to overcome the limits of traditional read-only topological views and provide users with a complete workspace for data analysis and manipulation. Thanks to fast click detection, users can navigate the canvas smoothly, dynamically edit entities, and manage their relationships. All key interactive features implemented in the system are summarized in Table I.

TABLE I
OVERVIEW OF INTERACTIVE FEATURES IN THE WORKSPACE

Category	Feature	Interaction Description
Graph	Viewport Navigation	Infinite panning, zooming, automatic fit-to-view centering, keyboard shortcuts.
	Spatial Organization	On-demand layout recalculation (D3 force), state previews, and manual position saving.
	Bulk Selection	Multi-element selection via bounding box, and integrated canvas search.
Nodes	Lifecycle (CRUD)	Interactive creation, deletion, and minimizing/restoring details directly on the canvas.
	Direct Manipulation	Drag & drop of individual nodes and groups; quick "inline" editing of text labels via double-click.
Edges	Link Management	Visual creation of new connections between nodes. Rerouting edges by dragging control points.
	Context Visualization	Automatic highlighting of associated input and output connections upon element selection.
UI, Props	Context Actions	Access to properties, preferences, and visual settings via right-click or a quick menu.

D. Practical Application

To further extend the application of the proposed solution, the interactive graph-based environment presented in this paper has been integrated into Penterep [12], a modular platform designed to support penetration testing. Penterep organizes its data in projects, each of which serves as a root node containing arbitrarily nested sub-nodes. This tree structure is persisted in a relational database using a nested-set encoding (tree traversal) that assigns left and right positional indices to each node, reducing the retrieval complexity from the quadratic cost of recursive queries $\mathcal{O}(n^2)$ to logarithmic time $\mathcal{O}(\log n)$. The platform exposes the resulting tree as a JSON object consumed by a single-page application, which then renders it and supports operations such as node addition, removal, and cloning of the whole tree or sub-tree.

From a technical perspective, the integration of the graph-based environment into the Penterep platform extends its tree structure, which was previously limited to a unidirectional data-flow architecture. To prevent direct interference with the Penterep core implementation, the visualization engine is distributed as a standalone `npm` package and operates strictly as a presentational component. It does not have direct access to the database; instead, all user interactions within the canvas trigger structured namespaced events (such as `node:add` or `edge:remove`) that are emitted upward to the host Inertia.js application. This parent layer is responsible for data validation and communication with the backend API endpoints, ensuring that the central database of the Penterep platform remains the absolute Single Source of Truth (SSoT).

To maintain seamless, real-time synchronization between traditional table views and the spatial graph, the system utilizes WebSockets. Whenever an entity is modified or selected elsewhere in the platform, a WebSocket message updates the reactive input properties (Props) passed to the graph. These changes are immediately processed by internal watchers within the engine, triggering automated layout recalculations or focusing the camera on a specific node without the need for a full page reload.

This integration brings a dual added value. First, the graph-based environment extends the existing tree structures with associative edges that capture semantic and contextual relationships between entities. Beyond OSINT, the same graph-based representation can be used to visualize penetration testing results where tested infrastructure, devices, services, and identified vulnerabilities can all be modeled as attributed nodes interconnected by hierarchical and associative edges, providing penetration testers with an interactive map of the attack surface that complements the platform's existing tree-based project structure.

IV. Experimental Testing

To evaluate the scalability of the rendering engine, automated stress tests were performed on datasets simulating dense OSINT topologies. The measurements focused on D3.js force-layout computation time, interaction FPS, and hit-detection latency. The results are presented in Table II.

TABLE II
PERFORMANCE EVALUATION OF THE CANVAS RENDERING ENGINE

No.	Nodes	Edges	Layout Calc.	FPS	Hit-Test Latency
1	100	150	81.4 ms	59	0.005 ms
2	500	500	408.1 ms	45	0.005 ms
3	1000	1500	600.2 ms	25	0.005 ms
4	2000	3000	1350.2 ms	12	0.005 ms

The force-layout calculation demonstrated predictable and stable scaling, with the largest dataset (2,000 nodes and 3,000 edges) processed in approximately 1.35 seconds. This confirms the validity of the architectural decision to decouple the mathematical physics model from the visual rendering loop, thereby preventing the user interface from freezing during position recalculation. Such dataset sizes are particularly relevant in the OSINT domain, where complex investigations can produce large interconnected graphs as analysts progressively uncover relationships across OSIF entities.

During continuous interaction with the viewport (scrolling and zooming), the frame rate gradually decreased from 59 to 12 FPS under extreme load. Profiling confirmed that this bottleneck is primarily attributable to the volume of individual Canvas API draw calls required to render edges, rather than to algorithmic inefficiency. In particular, the hit-detection latency remained constant at 0.005 ms across all dataset sizes. This proves the $\mathcal{O}(1)$ time complexity of the off-screen shadow canvas architecture, which provides instant feedback and zero perceived input latency for critical analytical interactions, even when the visual frame rate is degraded.

Based on the experimental results and user experience, we identified two optimization directions. First, rendering performance can be improved by batching draw calls for elements sharing visual properties into single Canvas API invocations, reducing the per-element overhead that limits the frame rate under extreme load. Second, the D3.js force-directed layout can be enhanced with alternative mathematical models, such as a hierarchical directed acyclic graph (DAG) mode and the constraint-based Cola.js library, to achieve faster convergence and more meaningful spatial arrangements.

V. Conclusion and Future Work

This paper presents the design, formal definition, and implementation of an interactive graph-based environment for OSINT investigations. The proposed solution models OSIF entities as nodes of a directed attributed multigraph that distinguishes hierarchical parent-child relationships from associative semantic links, enabling analysts to visualize, correlate, and manipulate OSINT data within a unified environment. The rendering engine, built on the HTML5 Canvas API, demonstrated scalability up to 2,000 nodes and 3,000 edges while maintaining constant $\mathcal{O}(1)$ hit-detection latency through an off-screen shadow canvas. The environment was successfully integrated into the Penterep platform, where it extends the existing tree-based project structure with graph-based visualization and interactive features to support OSINT investigations.

References

[1] F. Duarte. "Amount of Data Created Daily (2026)," *Exploding Topics is a Trademark of Semrush Inc*, 2026. Available: https://explodingtopics.com/blog/data-generated-per-day [Accessed: Mar. 3, 2025].

[2] Yong Shi. "Advances in big data analytics," *Adv Big Data Anal*, vol. 10, pp. 978–981, 2022.

[3] I. Böhm and S. Lolagar. "Open source intelligence: Introduction, legal, and ethical considerations," *International Cybersecurity Law Review*, vol. 2, pp. 317–337, 2021.

[4] N. A. Hassan and R. Hijazi. "Open source intelligence methods and tools," New York, USA: Apress, 2018.

[5] B. Akhgar, P. S. Bayerl, and F. Sampson. "Open source intelligence investigation: from strategy to implementation," Cham, Switzerland: Springer, 2017.

[6] G. Fenza, V. Loia, P. M. Mainardi, and C. Stanzione. "OSINT Knowledge Graph for Fact-Checking: Google Map Hacks Debunking," in *ITASEC 2023: The Italian Conference on CyberSecurity*, pp. 1–12, 2023.

[7] S. Bhandari, E. Ramos, R. Rupert, et al. "Knowledge Graph Representation of Open-Source Homicide Information," in *Proceedings of the 18th International Workshop on Value Modelling and Business Ontologies (VMBO 2025)*, pp. 1–11, 2025.

[8] Y. Su. "Open-Source Intelligence Analysis Method Based on Fine-Tuned Large Models and Knowledge Graphs," in *2025 IEEE 8th Advanced Information Technology, Electronic and Automation Control Conference (IAEAC)*, pp. 360–364, 2025.

[9] C. Bratsas, E. K. Anastasiadis, A. K. Angelidis, L. Ioannidis, R. Kotsakis, and S. Ougiaroglou. "Knowledge Graphs and Semantic Web Tools in Cyber Threat Intelligence: A Systematic Literature Review," *Journal of Cybersecurity and Privacy*, vol. 4, no. 3, pp. 518–545, 2024.

[10] F. Böhm, F. Menges, and G. Pernul. "Graph-Based Visual Analytics for Cyber Threat Intelligence," *Cybersecurity*, vol. 1, no. 16, p. 1–19, 2018.

[11] I. J. King, R. Ramirez, B. Bowman, and H. H. Huang. "Trail: A Knowledge Graph-Based Approach for Attributing Advanced Persistent Threats," in *2025 IEEE 41st International Conference on Data Engineering (ICDE)*, pp. 1207–1220, 2025.

[12] W. Lazarov, P. Seda, Z. Martinasek, and R. Kummel, "Penterep: Comprehensive penetration testing with adaptable interactive checklists," *Computers & Security*, vol. 154, pp. 1–16, 2025.

Message Recovery Attack on ML-KEM

1st Patrik Dobias
Department of Telecommunications
Brno University of Technology
Brno, Czech Republic
xdobia13@vut.cz

2nd Lukas Malina
Department of Telecommunications
Brno University of Technology
Brno, Czech Republic
malina@vut.cz

***Abstract*—While post-quantum cryptography (PQC) holds security guarantees even with presence of quantum computers, due to its higher complexity, practical implementations may contain vulnerabilities that break these security guarantees. One class of vulnerabilities present in practical implementations is side-channel leakage, which reveals information about secret data that are being processed. In this paper, we present a template-based attack against a hardware implementation of the ML-KEM scheme, compromising a secure session established between two parties. To decrease the number of traces needed and achieve a higher success rate compared to previous approaches, we identify two vulnerable regions and exploit their leakage simultaneously. Our experimental results demonstrate that highly parallelized hardware architectures are also vulnerable to attacks using side-channel analysis (SCA) and stresses the need for protected hardware implementations.**

***Index Terms*—ML-KEM, Side-Channel Analysis, Template Attack, Hardware Implementation**

I. Introduction

Security of traditional public key cryptographic algorithms is threatened by quantum computers, which enable solving their hard underlying problems (e.g., integer factorization and discrete logarithm) in polynomial time using Shor's algorithm [1]. To address this threat, researchers have shifted their focus to alternative cryptographic constructions (e.g., PQC), which are based on different underlying hardness assumptions. A major impact on this research was a standardization process organized by the National Institute of Standards and Technology (NIST), which resulted in the selection of two key encapsulation mechanisms (KEM) schemes and three digital signature algorithm (DSA) schemes for standardization. Among the selected schemes, the most successful were lattice-based CRYSTALS-Kyber and CRYSTALS-Dilithium, which were standardized as ML-KEM [2] and ML-DSA [3], respectively.

With the standards defined, an important next step in the research is the security evaluation of practical implementations, as these implementations enable different types of attacks against theoretically secure cryptographic algorithms. One class of these attacks is SCA, which exploits physical characteristics (e.g., timing behavior, power consumption, or electromagnetic (EM) emissions) of devices executing the cryptographic algorithm, that can leak information about processed secret data. To prevent these leakages, different randomization techniques (e.g., masking or shuffling) are incorporated into the implementations.

A. Related Work

In recent years, many SCA attacks against ML-KEM have been proposed, targeting pointwise multiplication (e.g., [4]–[6]), number theoretic transform (NTT) (e.g., [5], [7], [8]), message encoding or decoding (e.g., [7], [9]–[12]) or Fujisaki-Okamoto (FO) transform [13] (e.g., [14]–[16]) to recover session or long-term secret keys. Although attacks against software implementations are numerous, attacks against hardware implementations are more scarce.

Rodriguez *et al.* [4] performed a correlation electromagnetic analysis (CEMA) on the hardware implementation of Kyber [17], targeting pointwise multiplication. Zhao *et al.* [5] explored multiple vulnerable operations in their hardware implementation. They successfully recovered secret keys using correlation power analysis (CPA) and CEMA, targeting pointwise multiplication and inverse NTT. Ji *et al.* [10] performed a deep learning-based attack on hardware implementation [17], targeting message encoding. Wang *et al.* [11] presented another deep learning-based attack targeting hardware implementation of only message encoding operation.

B. Contribution

The main contributions of this paper are as follows.

- We present a template attack on hardware implementation of the ML-KEM scheme based on the deep learning-based attacks presented in [9], [10]. As our implementation employs a higher degree of parallelization compared to [10], particularly processing eight vs. two coefficients, it provides better insight into leakage exploitability in highly parallelized designs.
- We identify additional vulnerable regions to overcome additional noise, caused by the higher parallelization, and combine the leakages to reduce the required number of traces and increase the success rate of the attack.

C. Organization

The remainder of this paper is organized as follows. Section II gives a brief overview of the ML-KEM scheme and SCA. Section III presents our message recovery attack against a hardware implementation of the ML-KEM scheme and Section IV describes the results of the practical attack in our experimental setup. Section V discusses the implications of the attack along with possible countermeasures and future research, and finally Section VI concludes the paper.

II. Background

A. ML-KEM scheme

ML-KEM [2] is a lattice-based KEM that relies on the hardness of Module Learning with Errors (M-LWE) problem. It operates on a polynomial ring $\mathbb{Z}_q[X]/(X^n+1)$, where $q = 3,329$ and $n = 256$. It combines a public key encryption scheme secure under chosen plaintext attack and FO transform to offer the KEM scheme secure under chosen ciphertext attack. It consists of key generation, key encapsulation, and key decapsulation phases. Only the key decapsulation phase, outlined in Algorithm 1, is relevant for this paper. For more details and a description of the remaining phases, we refer the reader to the NIST standard [2].

1) Key Decapsulation Phase: It can be split into two public key encryption (PKE) phases extended with additional computations for shared secret computation. In the first phase, a message is decrypted from the provided ciphertext using secret key. The shared secret is derived from the decrypted message, and the message is re-encrypted. Re-encrypted ciphertext is then compared with the provided one, to prevent chosen ciphertext attacks, and an implicit rejection happens in case they do not match.

Algorithm 1 ML-KEM.Dec($sk = (\mathbf{s},\ pk,\ h,\ z),\ c = (c_1,\ c_2)$)

K-PKE.Decrypt:

1: $\mathbf{u} = \text{Decompress}_{d_u}(c_1)$
2: $v = \text{Decompress}_{d_v}(c_2)$
3: $m' = \text{Compress}_1(v - \mathbf{s}^T\mathbf{u}))$

Shared Key Derivation:

4: $(K', r') = \text{G}(m' \parallel h)$
5: $\bar{K} = \text{J}(z \parallel c)$

K-PKE.Encrypt:

6: $\mathbf{A} \leftarrow \text{SampleNTT}(\rho)$ ▷ ρ is part of pk
7: $(\mathbf{y}, \mathbf{e}_1, e_2) \leftarrow \text{SampleCBD}(r')$
8: $\mathbf{u}' = \text{Compress}_{d_u}(\mathbf{A}^T\mathbf{y} + \mathbf{e_1})$
9: $\mu = \text{Decompress}_1(m')$
10: $v' = \text{Compress}_{d_v}(\mathbf{t}^T\mathbf{y} + e_2 + \mu)$ ▷ $\mathbf{t}$ is part of pk

Implicit Rejection:

11: **if** $(c_1, c_2) == (u', v')$ **then**
12: $\quad K' = \bar{K}$
13: **end if**
14: **return** K'

B. Side-Channel Analysis

Side-channel analysis uses information leaked through the physical characteristics of a device, which is processing some sensitive data, to recover these data. Various methods can be used to exploit this leakage. These methods can be divided into non-profiling (e.g., simple or differential power analysis), which can directly exploit the leakage, or profiling (e.g., template or deep learning-based), which require full control over a clone of the device prior to the attack.

Although different physical characteristics can be used, we will consider only the power consumption of the device for the remainder of this paper.

1) Leakage Model: It defines the expected relationship between the leaked information and the processed data. In hardware implementations, the dominating factor for power consumption is the dynamic power consumption [18], due to this the Hamming distance model is selected. It can be expressed as $\text{HD}(x_i, x_{i+1}) = \text{HW}(x_i \oplus x_{i+1})$, where x_i and x_{i+1} are two consecutive values separated by a register.

2) Template Attacks [19]: They are a form of profiling attacks in which an attacker uses a clone[1] of the device under attack during the profiling phase to build templates corresponding to each possible value of the data. During the attack phase, these templates are then used to predict the value of data that were processed by the target device. An important aspect of this attack is the selection of points of interest (PoIs), which exhibit strong data-dependent leakage.

III. Message Recovery Attack

The template-based attack on hardware implementation of ML-KEM scheme aims to recover the message that is used to establish a session key during key decapsulation phase. We perform the attack against the unified hardware architecture of the ML-KEM and ML-DSA schemes that efficiently shares components to reduce the utilization of hardware resources, which was presented in [20]. The architecture uses parallelization and pipelining to offer high performance; specifically, for ML-KEM all components process eight coefficients in parallel. Moreover, it schedules operations to minimize the idle times of the individual components, and thus different components perform their operations in parallel. For a more detailed description of the architecture, we refer the reader to the original paper.

A. Adversary Model

We consider an adversary who has knowledge of the targeted implementation and can interact with the target device, i.e., choosing ciphertexts that the device will use during key decapsulation. Note that even though the attack is a profiling attack, the adversary does not require prior access to a clone device but can directly build templates on the device under attack. That is because the adversary is able to generate valid ciphertexts for any message using a public key, and thus can build templates for all byte values.

B. Vulnerable Regions

There are multiple operations that process a message during the key decapsulation phase in ML-KEM (see Algorithm 1). The most commonly exploited operation in state-of-the-art attacks is the encoding of the message coefficients into bytes (line 3). Another operation processing the message is the shared key derivation using hash function (line 4). However,

[1]It must be noted that the clone might not be always needed and the profiling might be possible directly on the target device, as is the case with this paper.

in the target hardware implementation, the message is loaded in 64-bit transactions and the keccak function processes the full state in parallel, thus making it an impractical target for the attack due to the large search space. During re-encryption, the message is first processed when the message bytes are decoded to coefficients (line 9), and then when they are added to the ciphertext part v' (line 10). In the target implementation, decoding is not performed, as it happens directly during encoding to reduce latency. But the addition is another promising region where the leakage can appear.

As was already stated, the target hardware implementation processes eight coefficients in parallel and is fully pipelined. Due to this, having the 32B message $m = (m_0, m_1, \ldots, m_{31})$, the leakage at both vulnerable regions is expected to follow equation $\text{HD}(m_i, m_{i+1}) = HW(m_i \oplus m_{i+1})$.

C. Multi-bit Error Injection Method

Predicting the exact value of the message byte from a single trace is practically impossible. To overcome this, Wang *et al.* [9] proposed a multi-bit error injection method, in which modified versions of ciphertext c are constructed in such a way that each version decrypts to a unique message. Having a message m and a corresponding ciphertext c, modified versions c_e are constructed so that $m_e = (m_0 \oplus e, m_1 \oplus e, \ldots, m_{30} \oplus e, m_{31} \oplus e)$, with $e \in [1, 255]$. The key decapsulation is then executed with the original ciphertext, and all modified versions and the corresponding power traces are used to predict the vector of probabilities $S_{i,e}$, where the probability $S_{i,e}[l \oplus e]$ is the probability that $m_i = l$. The final prediction is then made from the combination of all probabilities as follows:

$$\tilde{l} = \underset{l \in [0,255]}{\arg\max} \left(\prod_{e=0}^{255} S_{i,e}[l \oplus e] \right). \tag{1}$$

However, their attack was proposed against software implementation, where the Hamming weight is used as a leakage model. For our case with the Hamming distance model, where the consecutive message bytes are processed, the expected leakage would be $\text{HD}(m_i \oplus e, m_{i+1} \oplus e) = (m_i \oplus m_{i+1})$, effectively canceling the injected error. Because of that, we apply the error only to every second byte, i.e., $m_e = (m_0 \oplus e, m_1, \ldots, m_{30} \oplus e, m_{31})$.

Note that constructing the modified ciphertexts is straightforward, as the message is only additively hidden within v coefficients. To flip the ith bit of the message modified coefficient v'_i is computed with the following equation:

$$v'_i = \text{Compress}_{d_v}(\text{Decompress}_{d_v}(v_i + q/2 \bmod q)). \tag{2}$$

IV. Results

A. Experimental Setup

For a practical evaluation of the attack, we have used a Sakura-X[2] board as the target device and programmed it with the hardware implementation of [20]. The operating frequency of the board was set to 3 MHz and the power traces were collected using Keysight MSOS104A oscilloscope with a sampling frequency of 200 MS/s. The experimental setup is shown in Figure 1. To process the collected traces and perform the attack, we have used SCALib [21].

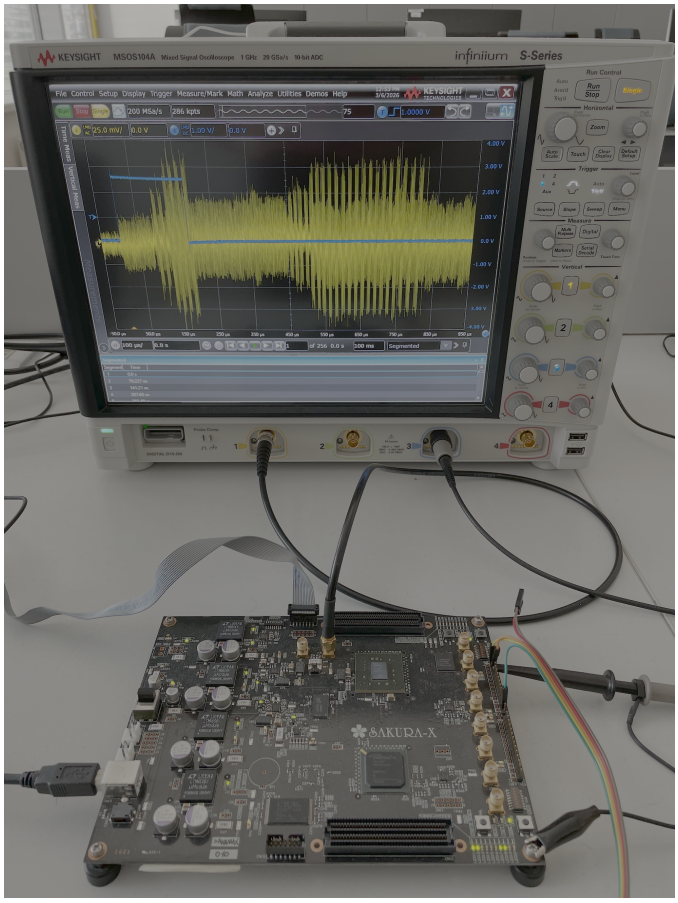

Fig. 1. Experimental setup for trace acquisition.

B. Vulnerable Regions Identification

To identify vulnerable regions, described in the previous section, in the hardware implementation, we have collected 100,000 traces of the key decapsulation phase and used the signal-to-noise ratio (SNR) metric. This is depicted in Figure 2. The first graph shows the mean power consumption of the collected traces. The second graph shows the SNR values with two clearly visible vulnerable regions. The third graph zooms into the vulnerable regions, and 31 distinct peaks corresponding to processing of individual message bytes are visible. Note that there are only 31 peaks, although the message has 32 bytes. This is because with the Hamming distance model only the differences between bytes can be recovered, and the reference value before processing the first byte was not successfully identified.

These results were then compared with the simulated run of the implementation, and the first region was identified as the encoding of message coefficients and the second region as the addition of the message during re-encryption.

[2]http://satoh.cs.uec.ac.jp/SAKURA/hardware/SAKURA-X.html

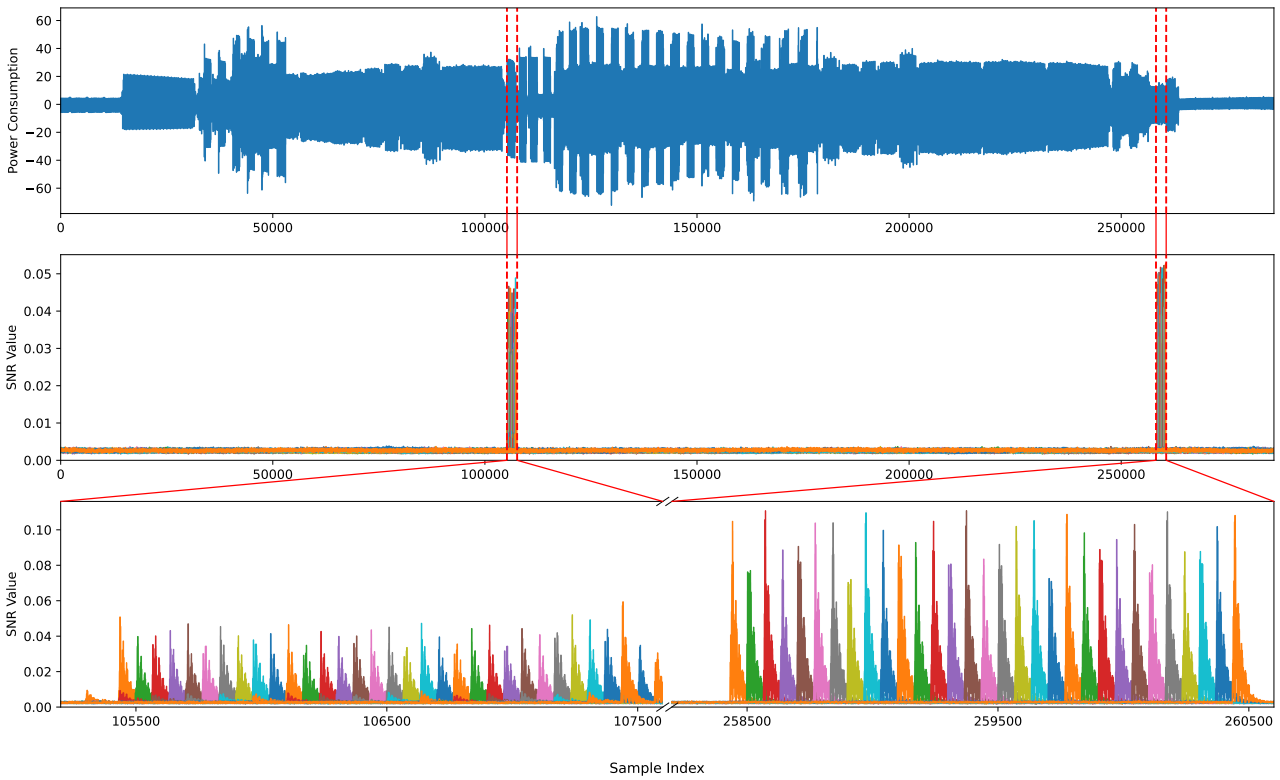

Fig. 2. Identifying vulnerable regions using SNR metric.

C. Experimental Message Recovery

Following identified vulnerable regions, we have selected PoIs corresponding to those regions and used the linear discriminant analysis (LDA) classifier from SCALib to build templates. We evaluate the success rate of the attack on 200 randomly generated messages. For each message, we first compute the ciphertext using the key encapsulation phase and then compute the modified ciphertexts by injecting errors into the v coefficients, following Equation 2. Each ciphertext is then sent to the target device and the key decapsulation is performed collecting corresponding power traces.

The accuracy of recovering each message byte and the maximum rank of the correct byte are presented in Figure 3. It shows results for different number of trace averaging. It can be observed from the graph that bytes at some indices (e.g., 1, 3, 6, etc.) were successfully recovered in all cases, while some were not. Similar bytes suffer from these problems for different averaging and can also be compared with Figure 2, where the peaks in given indices achieve lower values.

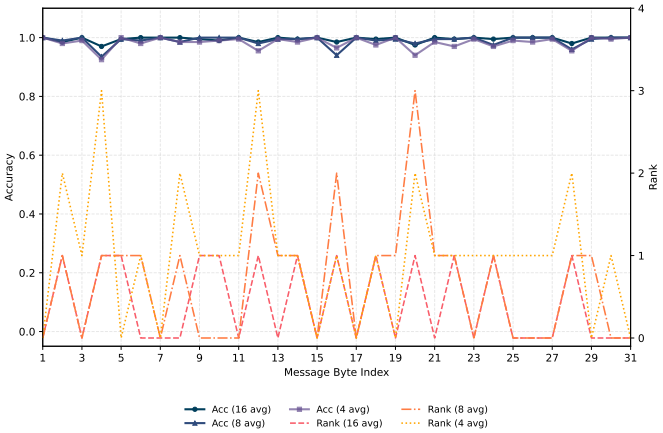

Fig. 3. Classification accuracy and maximum rank for each message byte.

V. Discussion

The presented attack recovers a secret message that is used for the session key derivation using the ML-KEM scheme, thus compromising a securely established session between two parties. For the 200 randomly chosen messages in our experimental verification of the attack, all messages were within the first two ranks, so in the worst case the attacker would have to brute-force $2^{31} \cdot 2^8 = 2^{39}$ possible messages, which is computationally feasible. Moreover, certain bytes exhibited even higher leakage, enabling 100% accuracy. Consequently, the overall search complexity can be further reduced if the attacker first iterates over those bytes with the lower SNR.

A. Secret Key Recovery

Recovering messages breaks the security guarantees of the FO transform, making the scheme vulnerable to chosen ciphertext attacks. Some authors (e.g., [7], [22]) have explored how to efficiently recover the secret key when messages for chosen ciphertexts can be recovered. For example, Xu *et al.* [7] require only eight recovered messages to reconstruct the secret key for Kyber-512. Following their method, our attack would require 16,384 traces to recover the secret key when used with 8 traces averaging.

B. Possible Countermeasures

Countermeasures (masking or shuffling the order of vulnerable operations) will prevent the attack presented. Mainly the masking countermeasure, however, comes at high resource utilization or performance costs. As the attacker is only able to recover changes in registers, a simple countermeasure against this attack is also to interleave the computations with computations on random data resulting in minimal overhead.

C. Future Work

Since we were unable to correctly identify the reference value to enable first byte recovery, further examination of this will be part of future work, together with examination of why leakage at certain bytes showed lower SNR. Another direction is to search for other vulnerable regions, providing additional information, that enable an attack with an even lower number of traces or targeting the secret key directly.

VI. Conclusion

In this paper, we have presented template-based message recovery attack on the hardware implementation of ML-KEM, demonstrating the feasibility of side-channel attacks even against highly parallelized implementations. Our attack explains how to adapt the multi-bit error injection method, proposed against software implementations to be also efficient in hardware implementations. Identifying an additional vulnerable region and combining information from both regions, the attack requires a lower number of traces compared to the attack presented by Ji *et al.* [10], while achieving a higher success rate. For reproducibility purposes, the attack scripts were released along with the datasets at https://github.com/paprikadobi/message-recovery-attack-on-ml-kem.

References

[1] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM review*, vol. 41, no. 2, pp. 303–332, 1999.

[2] NIST, "Module-lattice-based key-encapsulation mechanism standard," https://doi.org/10.6028/NIST.FIPS.203, 2024.

[3] ——, "Module-lattice-based digital signature standard," https://doi.org/10.6028/NIST.FIPS.204, 2024.

[4] R. C. Rodriguez, F. Bruguier, E. Valea, and P. Benoit, "Correlation electromagnetic analysis on an fpga implementation of crystals-kyber," in *2023 18th Conference on Ph. D Research in Microelectronics and Electronics (PRIME)*. IEEE, 2023, pp. 217–220.

[5] Y. Zhao, S. Pan, H. Ma, Y. Gao, X. Song, J. He, and Y. Jin, "Side channel security oriented evaluation and protection on hardware implementations of kyber," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 70, no. 12, pp. 5025–5035, 2023.

[6] C. Mujdei, L. Wouters, A. Karmakar, A. Beckers, J. M. Bermudo Mera, and I. Verbauwhede, "Side-channel analysis of lattice-based post-quantum cryptography: Exploiting polynomial multiplication," *ACM Transactions on Embedded Computing Systems*, vol. 23, no. 2, pp. 1–23, 2024.

[7] Z. Xu, O. Pemberton, S. S. Roy, D. Oswald, W. Yao, and Z. Zheng, "Magnifying side-channel leakage of lattice-based cryptosystems with chosen ciphertexts: The case study of kyber," *IEEE Transactions on Computers*, vol. 71, no. 9, pp. 2163–2176, 2021.

[8] M. Hamburg, J. Hermelink, R. Primas, S. Samardjiska, T. Schamberger, S. Streit, E. Strieder, and C. van Vredendaal, "Chosen ciphertext k-trace attacks on masked cca2 secure kyber," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2021, no. 4, p. 88–113, Aug. 2021. [Online]. Available: https://tches.iacr.org/index.php/TCHES/article/view/9061

[9] R. Wang, K. Ngo, and E. Dubrova, "A message recovery attack on lwe/lwr-based pke/kems using amplitude-modulated em emanations," in *Information Security and Cryptology – ICISC 2022*, S.-H. Seo and H. Seo, Eds. Cham: Springer Nature Switzerland, 2023, pp. 450–471.

[10] Y. Ji, R. Wang, K. Ngo, E. Dubrova, and L. Backlund, "A side-channel attack on a hardware implementation of crystals-kyber," in *2023 IEEE European Test Symposium (ETS)*. IEEE, 2023, pp. 1–5.

[11] J. Wang, W. Cao, H. Chen, and H. Li, "Blink: Breaking lattice-based schemes implemented in parallel with chosen-ciphertext attack," Cryptology ePrint Archive, Paper 2024/036, 2024. [Online]. Available: https://eprint.iacr.org/2024/036

[12] D. J. Bernstein, K. Bhargavan, S. Bhasin, A. Chattopadhyay, T. K. Chia, M. J. Kannwischer, F. Kiefer, T. B. Paiva, P. Ravi, and G. Tamvada, "Kyberslash: Exploiting secret-dependent division timings in kyber implementations," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2025, no. 2, p. 209–234, Mar. 2025. [Online]. Available: https://tches.iacr.org/index.php/TCHES/article/view/12046

[13] E. Fujisaki and T. Okamoto, "Secure integration of asymmetric and symmetric encryption schemes," in *Advances in Cryptology — CRYPTO' 99*, M. Wiener, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 1999, pp. 537–554.

[14] P. Ravi, S. S. Roy, A. Chattopadhyay, and S. Bhasin, "Generic side-channel attacks on cca-secure lattice-based pke and kems," *IACR transactions on cryptographic hardware and embedded systems*, pp. 307–335, 2020.

[15] R. Ueno, K. Xagawa, Y. Tanaka, A. Ito, J. Takahashi, and N. Homma, "Curse of re-encryption: A generic power/em analysis on post-quantum kems," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2022, no. 1, p. 296–322, Nov. 2021. [Online]. Available: https://tches.iacr.org/index.php/TCHES/article/view/9298

[16] Y. Tanaka, R. Ueno, K. Xagawa, A. Ito, J. Takahashi, and N. Homma, "Multiple-valued plaintext-checking side-channel attacks on post-quantum kems," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2023, no. 3, p. 473–503, Jun. 2023. [Online]. Available: https://tches.iacr.org/index.php/TCHES/article/view/10971

[17] Y. Xing and S. Li, "A compact hardware implementation of cca-secure key exchange mechanism crystals-kyber on fpga," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2021, no. 2, p. 328–356, Feb. 2021. [Online]. Available: https://tches.iacr.org/index.php/TCHES/article/view/8797

[18] F.-X. Standaert, S. B. Örs, and B. Preneel, "Power analysis of an fpga," in *Cryptographic Hardware and Embedded Systems - CHES 2004*, M. Joye and J.-J. Quisquater, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 30–44.

[19] S. Chari, J. R. Rao, and P. Rohatgi, "Template attacks," in *Cryptographic Hardware and Embedded Systems - CHES 2002*, B. S. Kaliski, ç. K. Koç, and C. Paar, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003, pp. 13–28.

[20] P. Dobias, L. Malina, and J. Hajny, "Efficient unified architecture for post-quantum cryptography: combining dilithium and kyber," *PeerJ Computer Science*, vol. 11, p. e2746, 2025.

[21] G. Cassiers and O. Bronchain, "Scalib: A side-channel analysis library," *Journal of Open Source Software*, vol. 8, no. 86, p. 5196, 2023. [Online]. Available: https://doi.org/10.21105/joss.05196

[22] P. Ravi, S. Bhasin, S. S. Roy, and A. Chattopadhyay, "On exploiting message leakage in (few) nist pqc candidates for practical message recovery attacks," *IEEE Transactions on Information Forensics and Security*, vol. 17, pp. 684–699, 2022.

Emission Optical Spectroscopy Of Switching Arc Plasma

Ing. Karolína Fůsková
Dep. of Power Electrical and Electronic Engineering
Brno University of Technology, FEEC
Brno, Czech Republic
226708@vut.cz

***Abstract*— This paper focuses on experimental investigation of the arcing process in low-voltage circuit breakers using non-invasive optical emission spectroscopy. The circuit breaker was tested under 1.553 kA and 252 V, and the spectrometer captured lines in the range 300 to 800 nm. This work aims to analyze plasma composition during burning. Plasma contains more than the surrounding gas, specifically materials evaporated from metal components in the circuit breakers, according to the analysis. Elements such as iron, copper, carbon, sodium, and hydrogen were found in the acquired spectrum. This study demonstrates that emission optical spectroscopy is an effective tool for understanding thermodynamic processes, and it can help in development.**

Keywords*— plasma, optical emission spectroscopy, circuit breaker, switching arc plasma

I. Introduction

Reliable fault-current shutdown is one of the basic tasks of circuit breakers, which is essential for protection. This problem is complex across the fields of mechanics, physics, chemistry, and electrical engineering. This paper focuses on the electrical arc, which is formed by high-temperature plasma and is created between contacts during switching off. For proper extinguishing, it is necessary that the arc be driven into the quenching chamber, as highly concentrated plasma can wear out mechanical contacts in devices. Rapid deionization of the contact space is ensured, guaranteeing the restoration of electrical strength. Excessively high burning temperatures caused intensive evaporation of surrounding materials (contact material, circuit breaker walls, and quenching chamber) directly into the discharge space. Damaged structures of mechanical parts can lead to unreliable shutdowns and dangerous situations for workers. This article aims to examine the state and composition of plasma during switching off using non-invasive optical emission spectroscopy. Analysis of the arc itself can be done by a few methods [1,2].

For this experiment, non-invasive optical emission spectroscopy will be used, which does not interfere with the test itself. Results of optical emission spectroscopy provide information on contamination in the extinguishing environment, which can negatively affect the deionization process and result in non-compliance with the prescribed standards for dielectric testing. Optical emission spectroscopy is important for development, because it provides data on the quality of materials used in the manufacture of circuit breakers [3].

II. Standards requirements

To ensure device safety, the service must comply with the standard requirements. Two main standards relating to circuit breakers in the region are ČSN EN 60898-1 and ČSN EN IEC 60947-2. Every type of circuit breaker has to be tested according to one of these standards before being placed on the market. From this point of view, it is important to focus on dielectric tests that ensure safety. After extinguishing an electrical arc, the dielectric interval follows, during which the current in the circuit is interrupted (interaction interval). At this moment, the recovery voltage must be lower than the gas breakdown strength between the contacts at every instant. For proper extinguishing, it is necessary to ensure deionization in the space between contacts, which means rapid conversion of the electrically conductive plasma back to an insulator. If deionization were not sufficiently rapid, i.e., if the recovered voltage were higher than the dielectric strength at any given moment, the electric arc would reignite [1].

Standard ČSN EN 60898-1 includes requirements for circuit breakers for domestic and similar installations, while standard ČSN EN IEC 60947-2 focuses on industrial circuit breakers. In general, these standards are important mainly in the field of dielectric property testing, ensuring the safe and reliable shutdown of the protective element. The dielectric test is performed by using a test voltage, which can be a power frequency withstand or an impulse withstand. This voltage should be applied mainly between individual terminals, between the pole and the frame. Precise information is provided by the standard based on the type of circuit breaker [4,5].

The final part of dielectric testing is insulation resistance measurement. The purpose of dielectric testing is to test the insulation in devices and thus the safety of the entire equipment [4,5].

III. Optical emission spectroscopy

Optical emission spectroscopy is a non-invasive method that uses intense plasma radiation. From a physics point of view, there are two types of radiation (continuous and line), but only

line radiation is important for spectroscopy. Standard equipment consists of a radiation source (in this case, plasma), a radiation decomposition device, a detector, and a PC. An essential component of measurement is a monochromator, which isolates wavelengths from polychromatic radiation. It consists of an entrance slit, a dispersive element (optical grid), an exit slit, and focusing optics. Decomposed radiation is detected by sensitive detectors (CCD). These methods enable a detailed study of arc burning and help analyze. The results can be helpful in analyzing dielectric properties and thus ensuring compliance with standards [3,6].

A. Spectrum analysis

Optical spectrum provides two types of information. In the first case, it is the identification of present elements. In the spectrum, there are sharp spectral lines, which create an atomic spectrum, or a wider belt structure, which is typical of molecular transitions. Found lines are subsequently compared with measured reference values. These reference values can be used from the NIST (National Institute of Standards and Technology) database. This analysis provides information about plasma composition and pollution. In the second case, it is possible to determine the position and intensity of spectral lines [3].

IV. Electrical arc and spectroscopy

An electrical arc in protective devices occurs when electrical contacts are changing position. There are different environments that interact. The initial formation of an electric arc begins in the contact area, which is a solid substance. Industry contacts are made of alloys, which are created by atoms in an atomic lattice. The atoms interact with each other through the forces, which means they perform random motion. Atoms in solid substances oscillate around a fixed position (grid) [1,6].

As the temperature rises, the movement of oscillating atoms speeds up. If their kinetic energy exceeds the limit, the atoms cannot hold their positions and start moving randomly. This phenomenon relates to changing a solid into a liquid. After another temperature growth, the energy of the atoms also increases, so the forces are not strong enough to hold them together. That means that liquid is changing into the gaseous phase. Although the contact surface may appear smooth, from a microscopic point of view, they are uneven and only touches at a few tiny points. All the current flows through these small contact points, creating areas of increased contact resistance and extremely high current density, leading to excessive heating. If the contacts start opening, the material will start melting at a few points. Contact bridges are formed, which evaporate almost immediately due to the enormous heat [1,6].

As it was sad, the temperature increases the speed (kinetic energy) of gas molecules and evaporated metal. When the temperature exceeds the limit, a negative electron is removed from the atomic shell, forming a positive ion - this phenomenon is called thermal ionization. At the same time, the electrical field influences the charged particles. This behavior leads to the avalanche process, and space becomes conductive [1,6].

The important part during forming the arc between contacts is the cathode. At the beginning of the whole process, the electrons are released from the cathode because of the high temperature (thermal emission). In this case, the arc is not very mobile, and it stays put. However, positive ions are attracted to the cathode very fast, and they create a new strong field around. This field extracts the electrons from metal, and no high temperature is needed. Ions clump together near the cathode, thereby locally intensifying its microfield. That leads to self-emission, which relates to high arc mobility around contacts. As a result, arcs are built of plasma, which consists of conductive ions and electrons. Optical spectroscopy methods are widely used in arc analysis because they are reliable, non-invasive diagnostic tools [1, 7].

Spectroscopy is based on the spontaneous emission of photons from excited atoms or ions into the surrounding space. The theory of radiation describes these movements as a change in the energy state of a particle, whereby the frequency and the wavelength of emitted radiation are precisely determined by the difference in energy levels according to Ritz´s combination principle. Due to this property of radiation, it is possible to identify emission and absorption lines in the spectrum and subsequently determine the chemical composition of the plasma. This means that, in an electric arc in the circuit breaker, the presence of materials vaporized from contacts and walls of the extinguishing chamber can be reliably detected. Besides the spectral line position, it is possible to follow the shape and spectral line broadening, which are important for the derivation of basic state variables [1,2,6].

- Doppler broadening: It is happening due to the thermal movement of emitting particles. It is possible to determine the ion's kinetic temperature from this profile. The kinetic temperature relates to the disordered thermal motion of particles and expresses the degree of their average kinetic energy [3,8].
- Stark´s (pressure) expansion: Stark´s effect is caused by a local electric micro-region created by charged particles near the atoms. From the Stark analysis, it can be deduced that the plasma density [3,8].

This knowledge of spectral lines provides a comprehensive view of temperature and density conditions inside the plasma without mechanical intervention [3,8].

V. Experimental setup

A. Experimental setup and measurement methodology

The experiment was performed on a low-voltage circuit breaker. The spectral lines were measured in the electrical arc between contacts during switching off. For the experiment, a prototype of a circuit breaker was used, and different values of current to be interrupted were studied. This paper shows results for only one value of current for proper description.

B. Test the short-circuit circuit

The experimental measurements were carried out in the CVVOZE laboratory equipped with a short-circuit current source. This means a short-circuit three-phase synchronous generator driven by an asynchronous motor. The circuit was also created with a set of loads (resistors and reactors) used to set the required parameters, such as current and power factors.

For mounting, the tested device was mounted on a DIN rail on the testing plate. The whole circuit was controlled from the control room using the laboratory´s own control system. For this experiment, a short circuit current of 1.553 kA and a voltage of 252 V were set.

C. Tested circuit breaker

It used test samples of a low-voltage circuit breaker manufactured for testing purposes. It was a Type B circuit breaker for residential use. For experimental purposes, it was modified by cutting out a window in the circuit breaker´s contact system space. This window was covered with a plastic slide (Fig. 1), providing optical access to the arc between the contacts. The electrical arc was initialized by mechanical postponement contacts under current and voltage.



Fig. 1. Window in the circuit breaker

D. Measurement and optical equipment

The optical spectrum was captured using two high-performance spectrometers and Andor Newton CCD detectors coupled to an Andor spectrograph (Andor-500i-D2-R). To minimize thermal noise, which impairs photometric accuracy, the sensor was cooled to -30 °C. To capture the results, an optical fiber was used, connected to the spectrometer via the entrance slit. The end of the thread was attached to the copper strip holder mounting using stands. The end of the thread was 36 cm away from the mounted circuit breaker (Fig. 2). The optical thread had to be at the same level as the window in the circuit breaker. The spectrometers were set as follows.

TABLE I. SET VALUES

parameters	***values***
Acquisition Mode	*Fast kinetic*
Triggering	*External Start*
Readout Mode	*FVB*
Exposure Time	*340 µs*
Sub-area height	*10*
Number of series	*51*
Binning	*1*
Offset from the bottom now	*502*
Vertical Shift speed	*14,51 µs*
Horizontal Read-out Rate	*1 MHz at 16-bit*
Pre-Ampilifier Gain	*1x*
Output Amplifier	*High Sensitivity*
Direct input slit	10 µm

The system was calibrated with a Mercury Argon Calibration Source, Hg and Ar lines from 253 nm to 922 nm, before testing. The start of recording the spectra and electrical quantities was synchronized with a trigger from the control computer to start a short circuit.

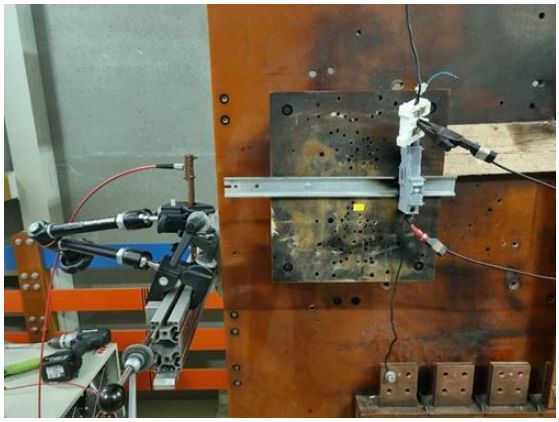

Fig. 2. Setup of optical thread

VI. RESULTS

A. Analysis of measured optical spectra

During the short-circuit test, plasma emission between the contact system was successfully recorded by using optical emission spectroscopy. The plasma of the electrical arc glows intensely as a result of the ongoing excitation. The spectra were recorded over the wavelength range 300 to 800 nm using the spectrometer settings on two spectrometers. The data was collected over a period of 6 ms. Given that the circuit current was only 1.5 kA, the duration of the arc was not very long. The largest spectral lines are therefore captured in this range, 1 ms-4 ms. Only one spectrometer with a wavelength of 300-650 nm was used for the analysis in this paper. The spectrograph, which measures wavelengths between 600 and 800 nm, did not detect any significant spectral lines that would affect the analysis. The

NIST (National Institute of Standards and Technology) database was used to identify atomic lines.

At 0.5 ms (Fig. 3), the graph shows low atomic lines of iron (Fe), copper (Cu), and sodium (Na)-relative intensity is under 70, and higher atomic lines of hydrogen (H)-relative intensity 100. Subsequently, the Na atomic line no longer appears, but higher concentrations of copper and iron appear. The graph also shows an increase in hydrogen over time. The graph in Fig.3 shows different shapes of spectral lines at different wavelengths, for example, sharp peaks and wider „scoops".

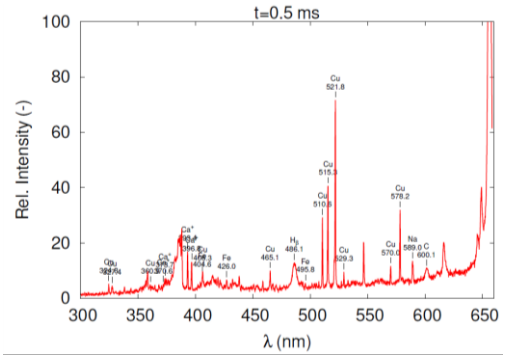

Fig. 3. Results for t = 0.5 ms

B. Identification of atomic lines

The wavelengths of emitted radiation are characteristic of a given substance. The found wavelengths were compared with the reference Ritz wavelengths from the NIST database (National Institute of Standards and Technology).

The results describe different phases of arc burning. In the first phase (Fig. 3 and Fig. 4), the plasma contains a high concentration of hydrogen, which can be caused by the evaporation of the surface layer of plastic materials, a phenomenon typical of the first phase. At this stage, hydrogen may also indicate air humidity, which is a common phenomenon. There is also an atomic line from sodium, indicating surface contamination. Pollution can cause almost anything. It can be dust on the solid components or fingerprints caused by manipulation during assembly. The arc burns on microscopic surface areas, what is the reason why there is a low concentration of metals here.

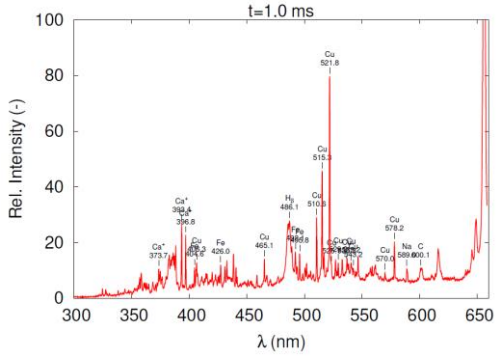

Fig. 4. Results for t = 1 ms

In the next place (time=2.5 ms - Fig. 5), there is a higher concentration of copper and iron, which indicates evaporation of the contact material and the extinguishing chamber. In the contact area, there are significant metal fumes. The graph shows the presence of carbon (C), which is visible after heating the material and at steady-state plasma radiation. It is the result of heating and evaporation of plastic materials [2,3,6].

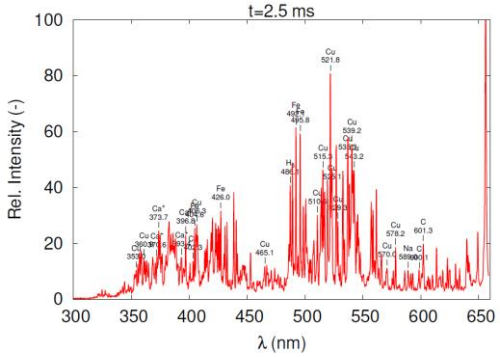

Fig. 5. Results for t = 2.5 ms

On the other hand, another figure shows the absence of Na, because plasma-cleaning surface impurities are removed right at the first stage of combustion. The hydrogen atomic line decreased over time. This is caused by the increased concentration of metal vapor, which displaced hydrogen. It is visible that the intensity of all elements descends. The largest peak occurred at 3.5 ms (Fig. 6). It corresponds to a high concentration of copper and iron.

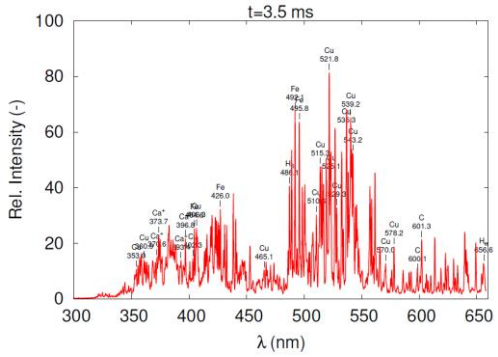

Fig. 6. Results for t = 3.5 ms

The arc gradually begins to face after exceeding the peak (fig. 7, Fig. 8). Atomic lines are lower and less concentrated. At the time 5 ms there are no atomic lines, so the process is over.

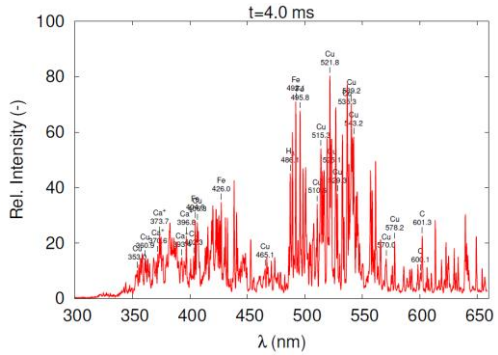

Fig. 7. Results for t = 4 ms

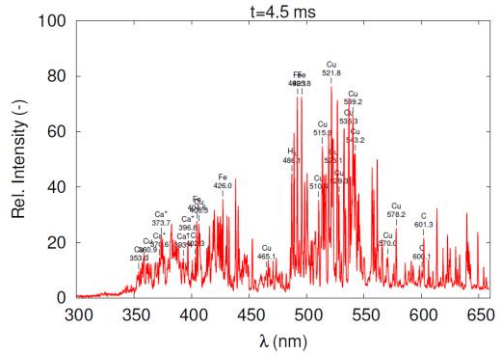

Fig. 8. Results for t = 4.5 ms

VII. Conclusion

This paper showed the results of an experimental investigation of the electrical arc in a circuit breaker using non-invasive optical emission spectroscopy. The experiment was performed in the CVVOZE lab under BUT. For the measurement, a test setup was used as described above (1.553 kA and 252 V). Two Andor Newton CCD detectors coupled with an Andor spectrograph were used for measuring. The measured data were then evaluated using the NIST database for recognition of atomic lines. Final graphs record the time evolution of spectra over time (0-6 ms). Time-spectrum recording enabled the evolution of the arc as it moved dynamically within the circuit breaker. The analysis confirmed that the plasma is not built only from surrounding gas, but it also contains iron, copper, carbon, hydrogen, and sodium.

Hydrogen was most intense at the beginning of burning and was associated with the evaporation of the surface layer of plastic materials [3]. This phenomenon can also be explained by air humidity. During this time, the intensity of the hydron decreases. In the initial phase of arc burning, there was a high concentration of hydrogen, but over time, it was replaced by metal vapor. In this phase, there was also sodium, which represents pollution apparently caused by dust or human manipulation [2,3].

Sodium immediately disappears because of the plasma cleaning effect. Increasing intensity of copper and iron confirmed the melting of the contact material and the movement of the arc to the quenching chamber. The presence of metal vapor changes the vapor properties. This is the reason why it is really important to use high-quality material for individual components. Additives can slow down the deionization process during the critical time. The same case is also carbon. It is visible after heating the material, and it means the evaporation of plastic materials. Carbon deposition on structural materials affects the subsequent dielectric strength of the intercontact space [2,3].

These experiments provided the necessary data for quality assessment. It provides detailed information on the physical processes in the plasma, which can lead to improvement. Plasma is difficult to analyze without being influenced by surrounding conditions. The importance of non-invasive optical emission spectroscopy extends beyond the optimization of new circuit breaker development. At the same time, it provides the basis and necessary information for numerical simulation in the next step of development. This experiment will be done in the future based on the obtained results [1,3].

Acknowledgment

This work was done as the first step of a complex experiment as part of doctoral studies. This paper was elaborated under the supervision of Ing. Jiří Valenta, Ph.D., at the Department of Power Electrical and Electronic Engineering and CVVOZE laboratory. Acknowledgment TAČR no. TS02020004. The author also acknowledges financial support from the Ministry of Education, Youth and Sports under institutional support and BUT specific research program (project No. FEKT-S-26-9122).

References

[1] J. Valenta, “Studium vlivu funkčních parametrů na zhášecí schopnost výkonového spínače nízkého napětí,” Ph.D. dissertation, Ústav výkonové elektrotechniky a elektroniky, Fakulta elektrotechniky a komunikačních technologií VUT v Brně, Brno, 2006.

[2] J. Barák, “Záření plazmatu elektrického oblouku,” Ph.D. dissertation, Ústav výkonové elektrotechniky a elektroniky, Fakulta elektrotechniky a komunikačních technologií VUT v Brně, Brno, 2003.

[3] Department of Macromolecular Physics MFF UK, "Diagnostika plazmatu," unpublished presentation.

[4] *Elektrická příslušenství - Jističe pro nadproudové jištění domovních a podobných instalací - Část 1: Jističe pro střídavý provoz*, ČSN EN 60898-1, Úřad pro technickou normalizaci, metrologii a státní zkušebnictví, 2020.

[5] *Spínací a řídicí přístroje nízkého napětí - Část 2: Jističe*, ČSN EN IEC 60947-2, Úřad pro technickou normalizaci, metrologii a státní zkušebnictví, 2018.

[6] F. F. Chen, *Úvod do fyziky plazmatu*. Praha: Academia, 1984.

[7] J. Bünemann, *Group Theory in Physics - An Introduction With A Focus On Solid State Physics*. Springer International Publishing, 2024.

[8] H. R. Griem, *Principles of Plasma Spectroscopy*. Cambridge: Cambridge University Press, 1997.

[9] A. Kramida, Yu. Ralchenko, J. Reader, and NIST ASD Team, "NIST Atomic Spectra Database (ver. 5.11)," [Online]. Available: https://physics.nist.gov/asd, 2026.

A Novel Self-Biased Current Reference for Ultra-Low-Voltage Operation in 22 nm CMOS

Jan Schilhab
Department of Microelectronics
FEEC, Brno University of Technology
Brno, Czech Republic
220855@vut.cz

Alexandr Matějka
Department of Microelectronics
FEEC, Brno University of Technology
Brno, Czech Republic
220848@vut.cz

Petr Mrázek
Department of Microelectronics
FEEC, Brno University of Technology
Brno, Czech Republic
220850@vut.cz

Adam Šlechta
Department of Microelectronics
FEEC, Brno University of Technology
Brno, Czech Republic
220857@vut.cz

***Abstract*—Aggressive CMOS scaling has significantly reduced the available voltage headroom for analog circuit design, making conventional self-biased current references increasingly difficult to implement in advanced digital technologies. This paper presents a novel self-biased current reference that preserves the conventional current-generation principle based on the MOS threshold voltage while enabling reliable operation at ultra-low supply voltage. The proposed topology replaces the conventional source-follower feedback path with a multi-branch current-mirror feedback structure, thereby reducing the required voltage headroom while maintaining robust self-biasing behavior. The circuit was implemented in 22 nm ultra-low-power CMOS technology and verified by post-layout simulations including parasitic extraction. The proposed reference operates from 0.72 V to 0.88 V over a temperature range from - 40 ºC to 150 ºC, providing a CTAT output current of approximately 55 µA, start-up time below 200 ns and layout area of only 216 µm². The proposed approach offers an effective and compact biasing solution for ultra-low-voltage mixed-signal systems in deeply scaled CMOS nodes.**

Keywords—Current reference, self-biasing, ultra-low-voltage circuits, CMOS analog design, threshold-voltage reference

I. Introduction

The continuous scaling of CMOS technologies has been primarily driven by the requirements of digital systems, resulting in aggressive reductions in device dimensions and supply voltages [1], [2]. While supply voltage levels have decreased significantly, the threshold voltage of MOS transistors cannot be reduced proportionally due to fundamental limitations related to leakage currents, variations of process technology, and ability to withstand stress over long period of time [1] – [3]. As a consequence, the available voltage headroom for analog circuit design continues to shrink in modern technologies. This effect is further exacerbated by pronounced short-channel phenomena, including increased channel-length modulation and reduced intrinsic gain [1], [3]. These constraints significantly complicate the design of conventional analog building blocks and motivate the development of novel circuit techniques capable of reliable operation under ultra-low supply voltages.

Current reference circuits represent essential components of integrated circuit design, as they are responsible for providing a stable and well-defined bias current for multiple functional blocks [4], [5]. To ensure reliable system operation, these circuits must maintain a constant current across variations in supply voltage, process parameters, and environmental conditions. One of the main characteristics of current references is low dependence on power supply voltage fluctuations. A common approach to achieving this property is the use of self-biasing structures, in which the operating point of the reference circuit is established internally through a feedback mechanism that stabilizes the generated current [4].

This paper proposes a novel self-biased current reference architecture designed for operation in deeply scaled CMOS technologies with limited voltage headroom. The proposed structure preserves the key characteristics of conventional threshold-voltage-based current references while enabling reliable operation at ultra-low supply voltages typical for modern digital core technologies. By modifying the self-biasing mechanism, the circuit mitigates the limitations imposed by reduced voltage margins and short-channel device effects.

The objective of this work is to develop a robust current source with a complementary-to-absolute-temperature (CTAT) current characteristic derived from the MOSFET threshold voltage. The proposed current reference is intended for use in a digitally controlled oscillator (DCO) implemented using a flip-flop-based architecture. In combination with the inherent properties of the oscillator topology, the proposed reference contributes to the realization of a temperature-stable and low-noise oscillator forming a fundamental component of a digitally controlled phase-locked loop (ADPLL).

The proposed self-biased current source was designed in a 22 nm CMOS technology for a nominal supply voltage of 0.8 V ± 10% and supports output currents of up to 55 µA, programmable through a digital-to-analog converter (DAC).

The remainder of this paper is organized as follows. Section II reviews the theoretical background of conventional self-biased current references based on the MOS transistor threshold voltage. Section III presents the proposed circuit architecture and explains how it addresses the limitations imposed by low supply voltages. Section IV discusses simulation results of the implemented design, and Section V concludes the paper and summarizes the main findings.

II. A SELF-BIASED CURRENT REFERENCE BASED ON THE MOSFET THRESHOLD VOLTAGE

Self-biasing is a technique in which the reference current is established internally by the circuit itself rather than directly from the supply voltage. As a result, the generated current is significantly less sensitive to supply variations, since it is primarily determined by the intrinsic properties of the circuit and the action of its internal feedback loop [5].

A. Topology of the Conventional Self-Biased Current Source

The topology of the conventional self-biased current reference is shown in Fig. 1. The core of the circuit consists of transistors T_1 and T_2 together with resistor R_1, forming a current-generating stage. The output current I_2 of this stage is given by

$$I_2 = \frac{V_{GS1}}{R_1}, \tag{1}$$

where V_{GS1} is the gate–source voltage of transistor T_1 [6]. Around the desired operating point, I_2 exhibits only a weak dependence on the input current I_1. This relationship can be derived from the MOS transistor equation,

$$V_{GS1} = \sqrt{\frac{2I_1L}{\mu_n C_{ox} W}} + V_{TH1}, \tag{2}$$

where V_{TH1} is the threshold voltage of transistor T_1, L and W are the channel length and width, μ_n is the carrier mobility, and C_{OX} is the gate-oxide capacitance per unit area. Substituting (2) into (1) yields

$$I_2 = \frac{1}{R_1}\left(\sqrt{\frac{2I_1L}{\mu_n C_{ox} W}} + V_{TH1}\right), \tag{3}$$

which describes the dependence of I_2 on I_1 [5].

Transistors T_3 and T_4 form a 1:1 current mirror that feeds I_2 back as I_1, thereby closing the self-biasing loop. The circuit exhibits two operating points: the desired non-zero solution $I_2 = I_1 = I_{REF}$, and the zero-current solution. The non-zero operating point defines the desired reference current.

The zero-current state is ideally unstable: a small increase in I_2 is mirrored to I_1, which increases V_{GS1} and consequently raises the voltage across R_1 through transistor T_2, acting as a source follower. This in turn further increases I_2, creating a positive feedback mechanism. The loop gain is high at very low current levels and decreases as the operating point approaches I_{REF}, where a stable equilibrium is established.

In practice, however, parasitic effects such as leakage currents and finite transistor gain may significantly reduce the loop gain near zero current [5]. As a result, the circuit may remain trapped in the zero-current state, and an auxiliary start-up circuit is typically required to guarantee proper initialization.

B. Temperature Dependence

According to (1), the temperature dependence of the reference current is determined by both resistor R_1 and transistor voltage V_{GS1}. The contribution of R_1 can be minimized by selecting a resistor with a low temperature coefficient of resistance (TCR).

The temperature dependence of V_{GS1} consists of two opposing components. The reduction of carrier mobility μ_n with temperature increases V_{GS1}, while the threshold voltage V_{TH} decreases with temperature and introduces a complementary-to-absolute-temperature (CTAT) component. In practice, the V_{TH} term usually dominates, resulting in an overall CTAT characteristic of the reference current.

Several works [7] – [9] exploit the balance of these two effects to achieve a zero-temperature-coefficient (ZTC) operating point. In this work, however, the CTAT characteristic is intentionally preserved to support subsequent temperature compensation in the oscillator core. The effect of carrier mobility is reduced by selecting a large transistor aspect ratio W/L, which also increases the self-biasing loop gain.

C. Voltage Headroom Limitation

Limited voltage headroom is a major challenge in deeply scaled CMOS technologies. In the conventional self-biased reference, the minimum supply voltage is determined by the sum of V_{GS1}, V_{GS2}, and the overdrive voltage required to keep transistor T_3 in saturation.

In modern CMOS nodes, the nominal supply voltage often does not exceed approximately twice the MOS threshold voltage, making the conventional topology difficult or impossible to operate reliably. This limitation directly motivates the modified self-biased current reference proposed in the following section.

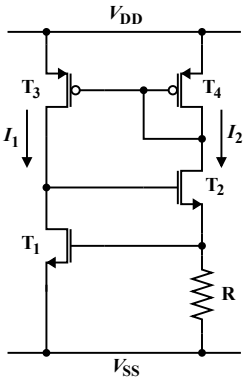

Fig. 1. Schematic of conventional self-biased current reference.

III. PROPOSED SELF-BIASED CURRENT REFERENCE FOR ULTRA-LOW-VOLTAGE OPERATION

The proposed circuit preserves the core current-generating principle of the conventional self-biased reference while fundamentally modifying its feedback mechanism to enable ultra-low-voltage operation. Instead of using a source-follower device in the direct feedback path, the feedback loop is redistributed into multiple current branches such that no branch requires more than approximately one threshold-voltage drop. This substantially reduces the voltage headroom requirement while retaining the threshold-voltage-based current generation principle.

A. Topology of the Proposed Self-Biased Current Reference

The schematic of the proposed circuit is shown in Fig. 2 and can be divided into three main functional blocks. The first block is the current-reference core, formed by transistor M_{NTH}, which provides the reference voltage, and resistor R_{TH}, which converts this voltage into the reference current according to (1). The second block is the self-biasing core, composed of NMOS and PMOS current mirrors that create a positive-feedback loop responsible for rapidly establishing the desired operating point. The third block is a start-up circuit, which assists the loop during initialization and prevents the circuit from remaining trapped in the zero-current state.

As shown in Fig. 2, the proposed structure replaces the source-follower transistor of the conventional topology with two current mirrors that together form the positive-feedback path. The NMOS mirror consists of transistors M_{NM1a} and M_{NM1b} with a unity current ratio. The PMOS mirror is formed by transistors M_{PM1a} and M_{PM1b}, and in the present design it provides a current gain of 4. Neglecting the zero-current state for the moment, this feedback arrangement amplifies the initial start-up current, causing a rapid increase of the loop current after power-up.

An additional branch of the PMOS mirror, implemented by transistor M_{PM1c}, supplies the current-reference core. As this branch current approaches I_{REF}, resistor R_{TH} develops a voltage approximately equal to the threshold voltage of transistor M_{NTH}, which then turns on. Transistor M_{NTH} is connected as a current-steering device at the input of the NMOS feedback mirror. Once it begins to conduct, part of the mirrored current is diverted away from the feedback path, reducing the effective gain of the loop below unity. The circuit then settles when the overall loop gain approaches one.

Resistor R_{FB} acts as a damping element. At very low current levels, the voltage drop across R_{FB} is negligible and therefore does not significantly affect the mirror ratio, allowing the loop gain to remain high and supporting a fast current build-up. As the current approaches the steady-state operating point, the voltage drop across R_{FB} increases and modifies the effective mirror ratio, reducing the loop gain and improving stability.

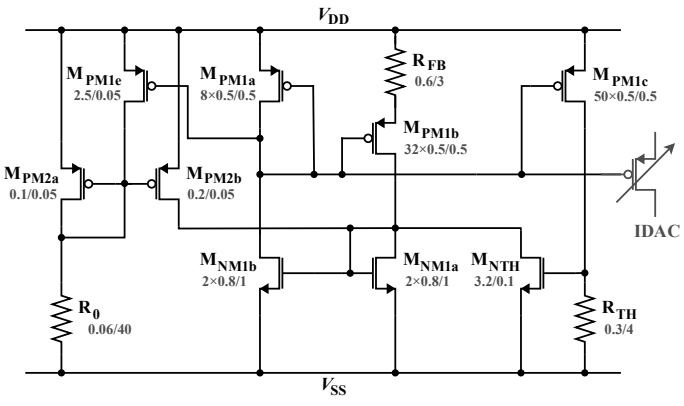

Fig. 2. Schematic of the new proposed current reference; all dimensions are in units of μm.

The current ratios of the individual branches were selected as a compromise between fast transient response and stable operation. As indicated by (3), the resulting reference current exhibits only a weak dependence on the current flowing through transistor M_{NTH}. Consequently, the final value of I_{REF} is only minimally affected by the sizing of the feedback mirror branches, particularly when a large W/L ratio is used for transistor M_{NTH}, as in the proposed design.

B. Start-Up Circuit

The final block of the proposed topology is the start-up circuit, whose purpose is to inject an initial current into the feedback loop under real operating conditions, where leakage currents may otherwise prevent self-starting. In this way, the circuit ensures reliable and sufficiently fast initialization.

The start-up circuit consists of a large resistor that sets a small current through diode-connected transistor M_{PM2a}. When the current in the main feedback loop is very small or zero, transistor M_{PM2b} mirrors this current into transistor M_{NM1a}, thereby initiating the positive-feedback action of the main loop. As the main loop current increases, auxiliary transistor M_{PM2e} reduces the gate–source voltage of M_{PM2b}, gradually suppressing the injected current. Once the main loop reaches its normal operating region, the start-up circuit effectively turns off and no longer interferes with the steady-state operation of the reference.

C. Stability analysis

Dividing the conventional self-biased current reference into multiple branches introduces an important stability challenge. In the proposed topology, the positive and negative feedback actions are no longer realized through a single compact path, but instead propagate through different branches of the circuit. The positive feedback path is formed primarily by transistors M_{NM1a}, M_{NM1b}, M_{PM1a}, and M_{PM1b}, whereas the negative feedback path includes the same loop extended by transistor M_{PM1c}, reference transistor M_{NTH}, and resistor R_{TH}.

To analyze the loop behavior, the circuit can be split at the common node between transistors M_{NM1a} and M_{NM1b}, where both feedback mechanisms interact. Since the effective branch impedance is dominated by the substantially lower resistance of the diode-connected mirror devices and resistor R_{TH}, the effect of transistor output resistances can be neglected. By further exploiting the approximately linear relationship between transconductance and transistor sizing within the mirror branches, the DC loop transfer can be approximated as

$$A_{\mathrm{FB}} = \frac{m_{\mathrm{MPM1b}}}{m_{\mathrm{MPM1a}}} \cdot \frac{1}{1+R_{\mathrm{FB}} g_{\mathrm{m,MPM1b}}} - \frac{m_{\mathrm{MPM1c}}}{m_{\mathrm{MPM1a}}} R_{\mathrm{TH}} g_{\mathrm{m,MNTH}}, \qquad (4)$$

where m denotes the number of identical unit devices in the corresponding mirror branch, and g_m is the transconductance of the relevant transistor.

From (4), it follows that stable operation at the DC operating point requires the overall loop transfer to become negative. This can be achieved either by increasing the value of R_{FB}, which was introduced specifically for stabilization, or by increasing the transconductance of the reference transistor M_{NTH}.

A second requirement for stable operation is to ensure that the positive feedback path does not respond faster than the negative feedback path. Since both mechanisms share a substantial portion of the same signal path, their dynamic behavior cannot be shaped independently. In the proposed implementation, the overall loop speed is therefore limited by the gate capacitance of the load, which in this design is formed by the IDAC mirror. The resulting AC stability is verified by simulation and discussed in the following section.

D. Implementation

The proposed circuit was implemented in a 22 nm ultra-low-power CMOS technology and designed for a supply-voltage range of 0.72 V to 0.88 V over a temperature range from - 40 ºC to 150 ºC. The final layout of the current reference is shown in Fig. 3. The layout occupies approximately 18 µm × 12 µm, corresponding to a total area of 216 µm^2, which makes the proposed reference suitable for highly area-constrained systems.

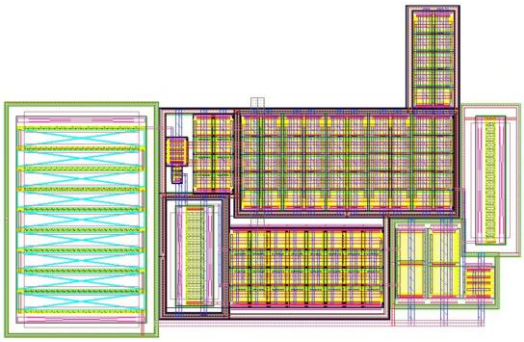

Fig. 3. The Layout of proposed structure in 22 nm CMOS Process.

IV. Simulation and Results

The proposed self-biased current reference was evaluated in the Cadence Virtuoso design environment. All results presented in this section were obtained from post-layout simulations including parasitic extraction from the final physical implementation. The circuit was verified under nominal and process-corner conditions using TT, FF, SS, SF, and FS models, representing typical, fast, and slow NMOS/PMOS device combinations. In addition, the design was evaluated over a temperature range from - 40 ºC to 150 ºC and a supply-voltage range from 0.72 V to 0.88 V. For clarity, only the worst-case corner results are presented in the following figures.

The histogram in Fig. 4 shows the Monte Carlo analysis of the output current with respect to local device mismatch. The analysis was performed for 1000 runs and indicates a mean output current of approximately 55.6 µA with standard deviation of 0.85 µA. Fig. 5 shows the DC temperature sweep over the full operating range from - 40 ºC to 150 ºC. As expected, the reference exhibits a CTAT temperature characteristic, with an estimated temperature coefficient of 2800 ppm/ ºC. The start-up transient response is shown in Fig. 6. The proposed topology achieves fast initialization, with the worst-case start-up time remaining below 200 ns. The loop stability analysis for the worst-case corner is shown in Fig. 7. The simulated loop bandwidth reaches approximately 20 MHz, while the corresponding phase margin is 59°, which confirms stable operation of the proposed structure.

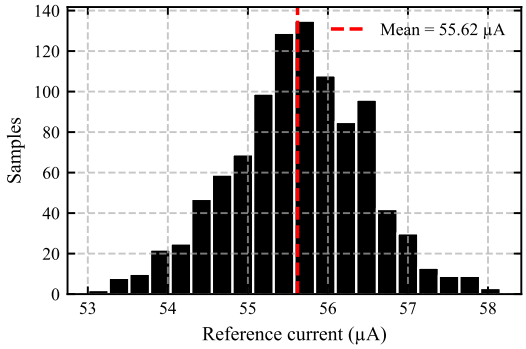

Fig. 4. Monte Carlo simulation of the reference current.

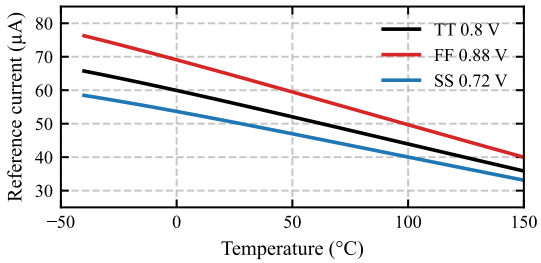

Fig. 5. Reference current vs. Temperature for nominal and worst Corners.

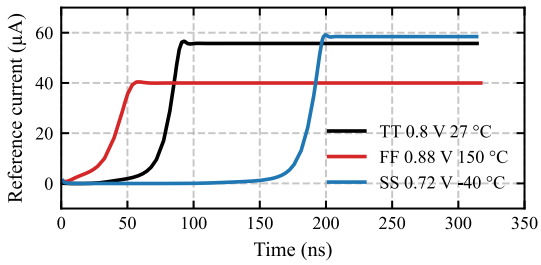

Fig. 6. Startup behavior of the reference current in transient simulation.

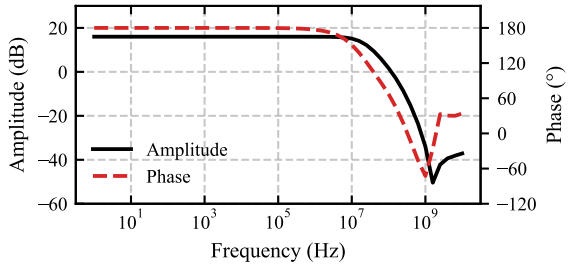

Fig. 7. Stability analysis for worst corner SS 0.88 V -40 °C.

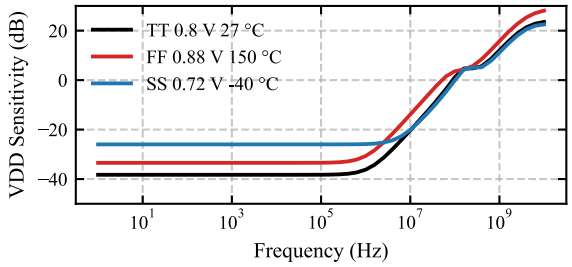

Fig. 8. Sensitivity of the reference current to supply voltage V_{DD}.

TABLE I. COMPARISON WITH RELATED CURRENT REFERENCES

Ref	Tech	VDD	Iref / Output	Power	Area	Startup time	Architecture
This Work	**22 nm CMOS**	**0.72–0.88 V**	**55 µA**	**~80 µW**	**216 µm²**	**200 ns**	**Self-biased CTAT current reference**
[10]	65 nm bulk / 22 nm FDSOI	0.7 / 0.9 V	1.1 nA / 0.9 nA	~5.4 nW	2100 µm² / 13200 µm²	~1 ms / 2 ms	Self-biased SCM current reference
[11]	0.18 µm / 0.5µm CMOS	0.5 / 2.5 V	nA range	–	22500 µm²/ –	600 µs / 20 ms	PTAT current reference with capacitive startup
[6]	130 nm CMOS	3.3 V	30 µA	–	7969 µm²	6.66 µs	Self-biased current reference
[12]	0.18 µm CMOS	0.6–1.2 V	nA range	1.8 nW	5900 µm²	0.2 ms	Capacitive startup + self-biased feedback
[13]	0.11 µm / 22 nm FD-SOI	1.2 V/ 1.5 V	2 nA / 5 nA	~16 nW	2600 µm² / 10600 µm²	~1 ms	Self-cascode current reference

The sensitivity of the reference current to supply-voltage variations is evaluated using

$$S_{V_{\mathrm{DD}}}^{I_{\mathrm{REF}}} = 20\log\left|\frac{V_{\mathrm{DD}}}{I_{\mathrm{REF}}}\frac{\partial I_{\mathrm{REF}}}{\partial V_{\mathrm{DD}}}\right| \quad (5)$$

and the corresponding frequency response is shown in Fig. 8. The worst-case sensitivity remains approximately -26 dB up to about 2 MHz, reflecting the high-speed response of the feedback loop.

Finally, Table I compares the key performance parameters of the proposed reference with related works. The results indicate that the proposed topology offers a highly competitive silicon area, reliable operation at low supply voltage, and very fast start-up behavior.

V. CONCLUSION

This paper introduced a novel self-biased current reference enabling threshold-voltage-based current generation under ultra-low supply voltage in 22 nm CMOS technology. By replacing the conventional source-follower feedback path with a multi-branch mirror-based feedback structure, the proposed topology significantly reduces the required voltage headroom while preserving the key properties of conventional self-biased references. Post-layout results verified reliable operation from 0.72 V to 0.88 V over the temperature range from – 40 ºC to 150 ºC, with a CTAT output current of approximately 55 µA, start-up time below 200 ns, and active area of only 216 µm^2. The proposed approach therefore represents an effective and area-efficient biasing solution for ultra-low-voltage mixed-signal and digitally assisted systems in advanced CMOS nodes.

ACKNOWLEDGMENT

The authors gratefully acknowledge Jiří Háze (Brno University of Technology) and Paul Lo and Timothy Wu (Jmem Tek) for valuable technical discussions and their support during the development of the circuit.

REFERENCES

[1] W. Bae, "CMOS Inverter as Analog Circuit: An Overview", *Journal of Low Power Electronics and Applications*, vol. 9, no. 3, Aug. 2019, doi: 10.3390/jlpea9030026.

[2] A. Bargagli-Stoffi and et al., "Challenges of VDD scaling for analog circuits: an amplifier", *Advances in Radio Science*, vol. 3, 2005, Accessed: Mar. 15, 2026. [Online]. Available: https://ars.copernicus.org/articles/3/377/2005/ars-3-377-2005.pdf

[3] S. Pandit, C. Mandal, and A. Patra, "Nano-Scale CMOS Analog Circuits: Models and CAD Techniques for High-Level Design", Sep. 2018, doi: 10.1201/9781315216102.

[4] B. Razavi, Design of analog CMOS integrated circuits, Second edition. New York: McGraw-Hill Education, [2017]. Accessed: Mar. 15, 2026. [Online]. Available: https://picture.iczhiku.com/resource/eetop/shkFghTKAoYQtXXX.pdf

[5] P. R. Gray and et al., *Analysis And Design Of Analog Integrated Circuits*, 5. ed. Wiley India, 2009. Accessed: Mar. 15, 2026. [Online]. Available: https://picture.iczhiku.com/resource/eetop/SyiYDzqDOfoPaMVM.pdf

[6] R. M. C. Dias and et al., "Performance Evaluation of a Self-Biased Current Reference Source Using Uniformly Doped and Halo-Implanted Channel Transistors", in *Conference: 20th Microelectronics Student Forum - SFORUM 2020*, Brazil, 2020. doi: 10.6084/m9.figshare.12912998.

[7] P. Toledo, H. Klimach, D. Cordova, S. Bampi, and E. Fabris, "MOSFET ZTC Condition Analysis for a Self-biased Current Reference Design", *Journal of Integrated Circuits and Systems*, vol. 10, no. 2, Dec. 2020, doi: 10.29292/jics.v10i2.411.

[8] M. Lefebvre and D. Bol, "A Family of Current References Based on 2T Voltage References: Demonstration in 0.18-µm with 0.1-nA PTAT and 1.1-µA CWT 38-ppm/°C Designs", *IEEE Transactions on Circuits and Systems I*, vol. 69, no. 8, 2022, doi: 10.1109/TCSI.2022.3172647.

[9] L. H. Rodovalho, P. Toledo, and S. Bampi, "A 0.6 V Current Reference Based on the MOSFET Forward-Body-Biased ZTC Condition", *2019 26th IEEE International Conference on Electronics, Circuits and Systems (ICECS)*, 2019, doi: 10.1109/icecs46596.2019.8964641.

[10] M. Lefebvre and et al., "A 1.1- / 0.9-nA Temperature-Independent 213- / 565-ppm/◦C Self-Biased CMOS-Only Current Reference in 65-nm Bulk and 22-nm FDSOI", *IEEE*, 2022, Accessed: Mar. 15, 2026. [Online]. Available: https://arxiv.org/pdf/2302.04504

[11] S. Mandal, S. Arfin, and R. Sarpeshkar, "Fast startup CMOS current references", in *2006 IEEE International Symposium on Circuits and Systems*, IEEE, 2006. doi: 10.1109/iscas.2006.1693217.

[12] Y. Huang, J. Liu, F. Yan, Y. Wang, K. Sun, and W. Ge, "A 1.65-nW 11.14-ppm/°C self-biased subthreshold CMOS voltage reference with temperature compensation circuit", *International Journal of Circuit Theory and Applications*, vol. 52, no. 12, May 2024, doi: 10.1002/cta.4067.

[13] M. Lefebvre and D. Bol, "A nA-Range Area-Efficient Sub-100-ppm/°C Peaking Current Reference Using Forward Body Biasing in 0.11- µ m Bulk and 22-nm FD-SOI", *IEEE Journal of Solid-State Circuits*, vol. 60, no. 2, 2025, doi: 10.1109/jssc.2024.3406423.

Low Voltage Low Noise Amplifier Design for an Ultrasonic Receiver in Automotive Applications

Alexandr Matějka
Department of Microelectronics
FEEC, Brno University of Technology
Brno, Czech Republic
220848@vut.cz

Jan Schilhab
Department of Microelectronics
FEEC, Brno University of Technology
Brno, Czech Republic
220855@vut.cz

Adam Šlechta
Department of Microelectronics
FEEC, Brno University of Technology
Brno, Czech Republic
220857@vut.cz

Petr Mrázek
Department of Microelectronics
FEEC, Brno University of Technology
Brno, Czech Republic
220850@vut.cz

Abstract*—This paper presents the design of a low voltage low noise amplifier intended for the receiver channel of an ultrasonic sensing system for automotive applications. The amplifier was designed and verified in the Cadence Virtuoso design environment using the Spectre simulator and implemented in onsemi BCD65 technology. The work focuses on achieving low input-referred noise and sufficient gain while maintaining stable operation at low supply voltage. The proposed circuit forms the front end of an ultrasonic receiver operating around 52 kHz. Simulation results confirm that the proposed amplifier satisfies the required gain and noise specifications for the intended application. The circuit behavior was evaluated using corner analysis including process, supply voltage, and temperature variations. The presented design demonstrates a feasible approach for implementing low noise front end circuits for ultrasonic sensing systems used in automotive electronics.

***Keywords*—low noise amplifier, ultrasonic receiver, automotive applications, analog integrated circuits, low voltage design**

I. Introduction

Ultrasonic sensing systems are widely used in modern automotive applications, particularly in parking assistance and short-range obstacle detection. These systems rely on the transmission of acoustic pulses and the subsequent detection of weak reflected signals from surrounding objects. Reliable detection requires receiver circuits capable of amplifying low-amplitude signals while maintaining low noise and sufficient dynamic range [1], [2].

The analog front-end of an ultrasonic receiver typically consists of a low-noise amplifier (LNA) followed by additional gain stages or a variable gain amplifier (VGA). The purpose of the LNA is to amplify the weak signal from the ultrasonic transducer while introducing minimal additional noise, which is critical for maintaining the signal-to-noise ratio of the system [3]. In multi-stage receivers, the overall noise performance is strongly influenced by the first amplification stage according to the Friis noise formula [4].

In automotive electronics, analog circuits are often required to operate under strict constraints such as low supply voltage, wide temperature range, and process variations. These conditions significantly influence the achievable performance of analog blocks and require careful circuit design techniques [5]. Low-voltage analog design is therefore an important topic in modern integrated circuit development.

Ultrasonic receivers frequently incorporate variable gain stages in order to accommodate the wide dynamic range of received signals. As the reflected signal amplitude depends on distance and object properties, gain control mechanisms are commonly used to maintain optimal signal levels in subsequent processing stages [1]. Implementations of such amplifiers often include techniques for offset compensation and stable operation across varying conditions.

This paper focuses on the design considerations of a low-noise amplifier intended for an ultrasonic receiver channel operating in automotive applications. The circuit was designed and characterized in CMOS-based BCD technology and evaluated using circuit-level simulations. Particular attention is given to low-voltage operation and the resulting impact on noise performance and achievable gain.

The remainder of this paper is organized as follows. Section II describes the architecture of the ultrasonic receiver front-end and the role of individual analog blocks. Section III presents the design of the low-noise amplifier with emphasis on low-voltage circuit techniques. Section IV discusses the simulation environment and characterization methodology. The obtained results are summarized in Section V, and the conclusions are presented in Section VI.

II. Ultrasonic Receiver Front-End

Automotive ultrasonic sensing systems typically employ piezoelectric transducers operating in the tens of kilohertz frequency range. These sensors convert reflected acoustic waves into electrical signals that must be processed by an analog receiver chain before further digital processing. Because the received signals can exhibit very low amplitudes, the performance of the analog front-end is critical for achieving sufficient detection sensitivity and system reliability [3], [1].

A simplified architecture of the ultrasonic receiver channel considered in this work is shown in Fig. 1. The receiver consists primarily of an LNA followed by a VGA. The LNA provides the first stage of amplification directly after the sensor interface,

while the VGA adjusts the overall gain of the receiver channel according to the signal amplitude and system requirements.

The ultrasonic transducer converts reflected acoustic waves into electrical signals whose amplitude can vary significantly depending on the distance and reflectivity of the detected object. In typical automotive ultrasonic systems, the operating frequency is around 52 kHz, corresponding to the resonance frequency of the transducer used in the sensing module. The received signal levels are usually very small and can be strongly affected by thermal and electronic noise generated in the front-end circuitry. Therefore, careful design of the input amplification stage is necessary to achieve a sufficiently low input-referred noise while maintaining adequate gain and bandwidth.

In the considered receiver architecture, the LNA performs the initial amplification of the sensor signal and determines a large portion of the overall noise performance of the system. According to the well-known Friis noise equation [4], the noise contribution of the first stage dominates the total noise figure of a cascaded analog system. Consequently, the LNA must provide high gain while minimizing its intrinsic noise sources [4], [6]. Additionally, the amplifier must operate reliably under low supply voltage conditions, which introduces further design constraints for transistor biasing and signal swing [5].

Another important requirement for the receiver front-end is sufficient bandwidth around the operating frequency. The amplifier must preserve the ultrasonic signal without excessive attenuation or distortion within the required frequency range. In practice, this requires appropriate selection of the amplifier topology and careful optimization of transistor parameters and bias conditions [3].

Considering these requirements, the LNA becomes a critical block that directly influences the sensitivity and robustness of the entire ultrasonic receiver. The following section therefore focuses on the design of a low-voltage CMOS LNA optimized for low noise and sufficient gain at the operating frequency of approximately 52 kHz.

III. Low-Voltage LNA Design

The low-noise amplifier represents the first active stage of the ultrasonic receiver front-end and therefore has a dominant influence on the overall noise performance of the system. In addition to low input-referred noise, the amplifier must provide sufficient gain while operating from a limited supply voltage. The following subsections describe the selected topology and the key design considerations used in the implementation of the proposed LNA.

A. Selected Topology

The simplified schematic of the implemented LNA is shown in Fig. 2. The amplifier is designed using CMOS-based BCD technology and operates from a 1.2 V supply voltage. The selected architecture combines a transconductance input stage with a differential operational amplifier configured as a transimpedance stage.

This topology was selected in order to achieve a favorable trade-off between noise performance, gain, and circuit complexity. The first stage converts the input voltage signal from the ultrasonic transducer into a current signal, which is subsequently processed by the differential amplifier. Such an approach allows efficient amplification while maintaining relatively low input-referred noise [3].

B. Input Transconductance Stage

The input stage of the LNA is implemented as a transconductance amplifier, which converts the input voltage signal into a current. This stage plays a critical role in determining the overall noise performance of the circuit, since the dominant noise sources originate from the input transistors [3], [6]. The schematic of the implemented transconductance amplifier is shown in Fig. 3. The first stage converts an input voltage V_{in} into an output current I_{out1}:

$$I_{\text{out1}} = G_{\text{m1}} \cdot V_{\text{in}} \,. \tag{1}$$

The dominant noise source in MOS transistors operating in saturation is the channel thermal noise generated in the inversion layer [3], [6], [7], which can be expressed as:

$$i_{\text{n,M}} = \sqrt{4 \cdot k_{\text{B}} \cdot T \cdot \gamma \cdot g_{\text{m}}} \,. \tag{2}$$

Where k_{B} is the Boltzmann constant [J/K], T is absolute temperature [K], γ is thermal noise coefficient [-] and g_{m} is transconductance [A/V].

The thermal noise of a resistor, modeled as a noise current source in parallel with the resistance, is calculated using the following equation:

$$i_{\text{n},R} = \sqrt{\frac{4 \cdot k_{\text{B}} \cdot T}{R}} \,. \tag{3}$$

Where R is resistance [Ω].

Higher transconductance generally leads to improved noise performance because it reduces the input-referred noise of the amplifier [8], [9]. However, increasing the transconductance usually requires higher bias current, which must be carefully balanced with the power consumption requirements of the system.

Device sizing and bias conditions were therefore selected to achieve a compromise between low noise and low power operation. Special attention was also given to device matching properties, which are known to influence offset and noise behavior in analog CMOS circuits [10].

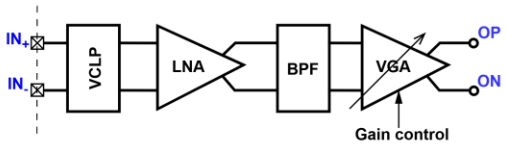

Fig. 1. Block diagram of the ultrasonic receiver front-end.

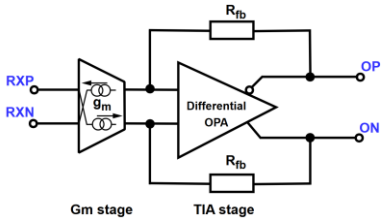

Fig. 2. Block diagram of LNA.

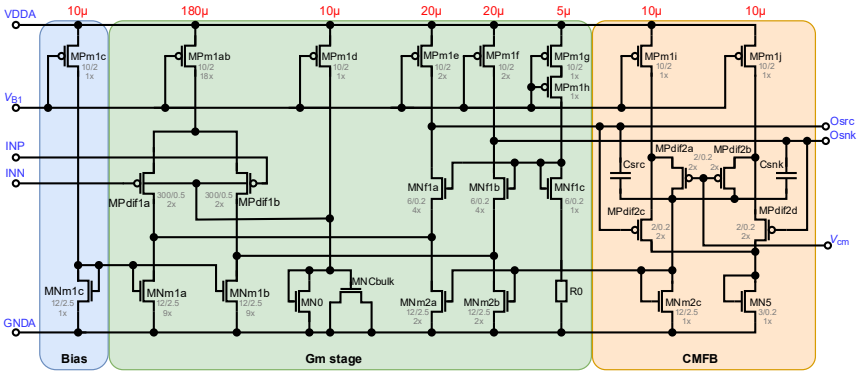

Fig. 3. Schematic of the transconductance amplifier, dimensions *W*/*L* are in µm.

C. *Differential Operational Amplifier*

The transconductance stage is followed by a differential operational amplifier that provides additional gain and signal conditioning. The differential configuration improves immunity to common-mode disturbances and allows more stable operation under varying supply and process conditions.

A simplified structure of the operational amplifier used in the design is shown in Fig. 4. The amplifier processes the current produced by the input stage and converts it back into a voltage signal suitable for further amplification or processing.

Differential architectures are commonly used in analog integrated circuits because they offer improved linearity and better suppression of common-mode noise sources [3], [8]. These properties are particularly beneficial in sensor interface circuits where small signals must be extracted from noisy environments.

D. *Transimpedance Configuration*

The LNA is implemented using a transimpedance configuration of OPA, where the current generated by the transconductance stage is converted into an output voltage V_{out} through effective feedback impedance R_{fb}.

$$V_{\text{out}} = -I_{\text{out1}} \cdot R_{\text{fb}} \tag{4}$$

This configuration allows the gain of the amplifier to be controlled by the feedback network while maintaining stable operation.

By substituting (1) into (4), we get the total voltage gain:

$$V_{\text{out}} = -(G_{\text{m1}} \cdot V_{\text{in}}) \cdot R_{\text{fb}} \tag{5}$$

$$A_{\text{v}} = \frac{V_{\text{out}}}{V_{\text{in}}} = -G_{\text{m1}} \cdot R_{\text{fb}} \tag{6}$$

Transimpedance amplifier configurations are widely used in sensor interface circuits due to their ability to efficiently convert small input currents into voltage signals while maintaining good noise characteristics [11], [12], [13].

In the presented design, the feedback network was optimized to achieve sufficient gain at the operating frequency of approximately 52 kHz while preserving adequate stability margins.

E. *Design Considerations for 1.2 V Supply*

Operating analog circuits at low supply voltages introduces several design challenges, including reduced signal swing, limited headroom for transistor biasing, and increased sensitivity to device mismatches [5], [8], [12], [14]. These issues are particularly significant in modern integrated circuit technologies where supply voltages continue to decrease [5].

To ensure reliable operation at 1.2 V, the transistor operating regions and bias currents were carefully selected so that all devices remain properly biased across expected process and temperature variations. In addition, unnecessary voltage drops within the signal path were minimized by selecting circuit topologies with reduced transistor stacking and by optimizing the bias voltages throughout the amplifier [14].

Through these design considerations, the implemented LNA achieves a compromise between low noise, sufficient gain, and low-voltage operation, making it suitable for integration in ultrasonic sensing systems used in automotive applications.

IV. Simulation and Characterization

The proposed low-noise amplifier was designed and evaluated using circuit-level simulations in the Cadence Virtuoso design environment. All simulations were performed using the Spectre analog simulator with device models provided for the selected CMOS-based BCD technology.

The objective of the simulations was to verify the functionality of the designed amplifier and evaluate its performance in terms of gain, noise, and stability under different operating conditions. The analysis included analytical estimation of noise sources, small-signal frequency response simulations, and evaluation of circuit behavior across processes, voltage, and temperature variations.

A. *Simulation Environment*

The circuit schematic of the LNA was implemented and simulated in the Cadence Virtuoso environment. Device models provided by the technology design kit (PDK) were used to ensure realistic representation of transistor behavior.

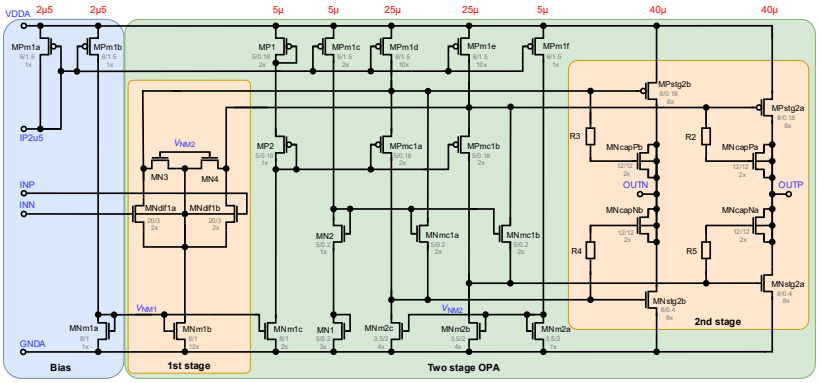

Fig. 4. Differential operational amplifier structure, dimensions *W/L* are in µm.

The simulations were performed under nominal and corner operating conditions with a supply voltage of 1.2 V ± 10%. In addition, the circuit was evaluated across a temperature range from −40 °C to 175 °C, corresponding to typical operating conditions for automotive electronic systems. The ultrasonic receiver operates at approximately 52 kHz, corresponding to the resonance frequency of the transducer used in the sensing module.

B. Analytical Noise Estimation of the Input Stage

Since the input transconductance stage represents the dominant noise contributor in the LNA, its noise performance was first analyzed using analytical expressions derived from MOS transistor noise models. Due to the relatively low supply voltage and the large channel width of the transistors forming the differential pair, the input devices are assumed to operate in the weak inversion region. In this operating regime, the transistor transconductance is primarily determined by the drain current rather than the gate overdrive voltage. The bias current of the input transistors was therefore selected as 90 µA in order to achieve sufficient transconductance while maintaining acceptable power consumption [3], [7].

The calculated noise currents of the dominant noise sources in the input stage are summarized in Table I, together with the main parameters of the individual transistors.

TABLE I. CALCULATED NOISE OF INPUT STAGE TRANSISTORS

Device	I_D [µA]	*W/L* [-]	*KP* [µAV^{-2}]	g_m [mAV^{-1}]	i_n [pA/√Hz]
M1	90	1200	30	2.163	4.23
M3	90	43	120	0.966	3.27
M5	90	45	30	0.493	2.33
M7	20	30	90	0.697	1.91
M9	20	10	30	0.251	1.14

To estimate the overall input-referred noise voltage, all dominant noise sources were referred to the gate terminal of transistor M1, as illustrated in Fig. 5.

Assuming that the individual noise contributions are uncorrelated, the resulting input-referred noise voltage can be expressed as:

$$u_{n,in} = \frac{1}{g_{m1}} \cdot \sqrt{i_{n,M5}^2 + i_{n,M1}^2 + i_{n,M3}^2 + i_{n,M7}^2 + i_{n,M9}^2 + i_{n,R}^2} \quad (7)$$

The calculated value of input-referred noise density is 2.96 nV/√Hz, which satisfies the noise requirements of the ultrasonic receiver front-end and provides a certain margin for additional noise sources that may appear in the complete circuit implementation. It should be noted that the simplified analytical model does not include flicker noise (1/f noise), which may slightly increase the overall noise level in the final circuit.

C. AC Response Analysis

The small-signal frequency response of the integrated LNA (G_m stage + OPA) was evaluated using AC analysis (Fig. 6). At the 52 kHz operating frequency, the amplifier achieves a voltage gain of 44 dB with a -3dB bandwidth of 1.6 MHz. Although the cascaded phase response shows a cumulative drop typical of multi-stage architectures, the system remains inherently stable as it operates in an open-loop configuration without global feedback.

Stability is ensured by the internal Miller compensation of the OPA stage, which maintains a phase margin of 58.0° in its local transimpedance (TIA) configuration, as confirmed by the parameters in Table II. This ensures reliable amplification of the ultrasonic signal without risk of oscillation.

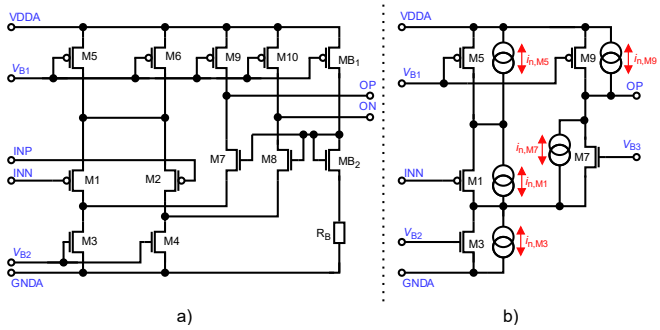

Fig. 5. a) Simplified schematic of G_m stage. b) Noise model of the G_m stage.

D. Noise Analysis

Following the analytical estimation presented in the previous subsection, noise simulations were performed to verify the expected noise performance of the designed circuit using the Spectre noise analysis tool in the Cadence Virtuoso environment.

First, the noise performance of the input transconductance stage was evaluated independently. The simulated input-referred noise spectrum of this stage is shown in Fig. 7 and shows good agreement with the analytically estimated noise density presented in Section IV-B.

The amplifier input was modeled using a PORT element, which includes the source resistance and its associated thermal noise. In order to evaluate the intrinsic noise contribution of the amplifier, the noise contribution of the source resistance was separated from the simulation result. After this correction, the obtained input-referred noise of 3.27 nV/√Hz corresponds closely to the analytically estimated value of approximately 2.96 nV/√Hz.

Subsequently, noise simulations of the complete LNA were performed. The resulting input-referred noise spectrum is shown in Fig. 8, indicating a noise density of approximately 3.58 nV/√Hz at 52 kHz after correction.

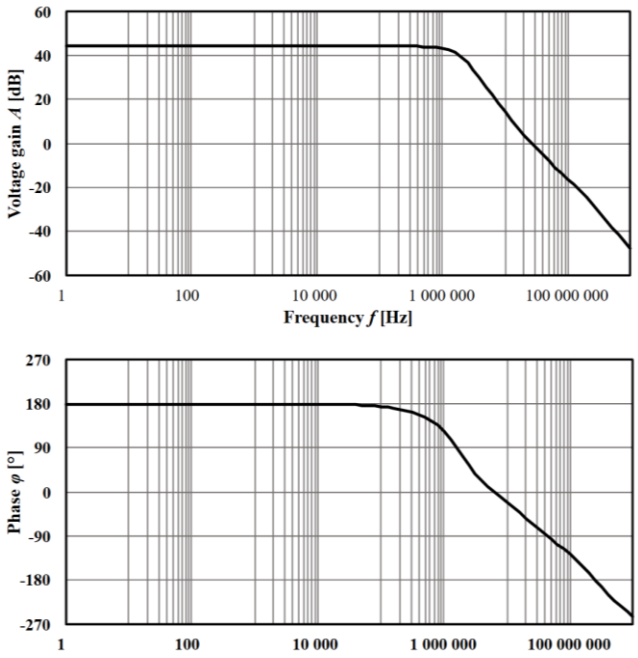

Fig. 6. Amplitude and phase response of LNA.

TABLE II. SUMMARY OF CORNER SIMULATION RESULTS OPA

Parameter OPA	**Min.**	**Typ.**	***Max.***	***Unit***
DC Gain	40.1	52.0	55.1	dB
Phase Margin	42.5	58.0	89.8	°
Gain Margin	12.9	19.5	26.0	dB
Gain Bandwidth	6.6	13.4	17.7	MHz

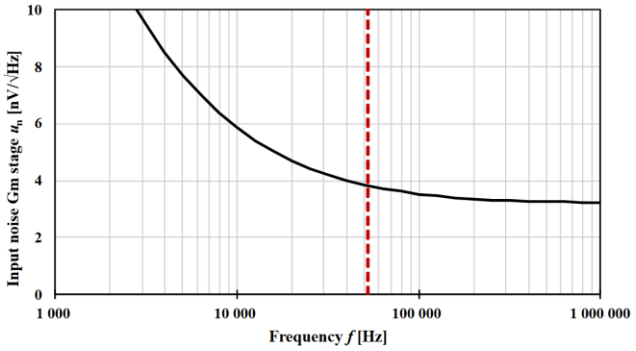

Fig. 7. Input-referred noise of G_m stage.

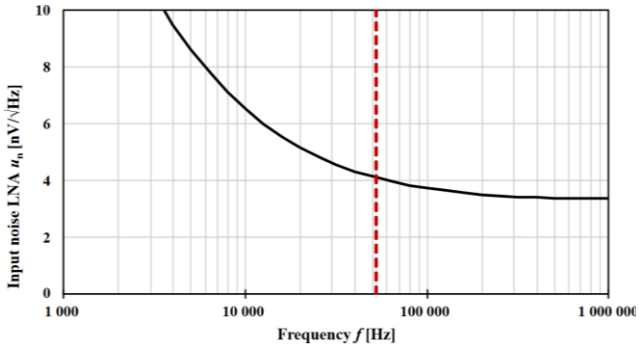

Fig. 8. Input-referred noise of LNA.

E. Corner and Operating Condition Analysis

To verify robust operation of the proposed LNA, simulations were performed across the process, supply voltage, and temperature conditions described in Section IV-A. These variations influence key transistor parameters such as threshold voltage, carrier mobility, and device matching, which may affect the gain, stability, and noise performance of the amplifier.

The simulation results confirm that the amplifier maintains stable operation and sufficient performance margins across all considered operating conditions. The most important DC and AC parameters obtained from the corner analysis are summarized in Table III.

TABLE III. SUMMARY OF CORNER SIMULATION RESULTS LNA

Parameter LNA	***Min.***	***Typ.***	***Max.***	***Unit***
Supply current	350	520	880	µA
Systematic Differential Offset		< 20		nV
Output CM Offset	-43.2	-18.7	16.3	mV
Gain Bandwidth	14.7	26.5	35.7	MHz
Bandwidth	0.25	1.60	3.10	MHz
DC Gain	38.0	44.3	44.8	dB

V. Conclusion and Future Directions

This paper presented the design and simulation of a low-voltage low-noise amplifier intended for the analog front-end of an ultrasonic receiver used in automotive sensing systems. The proposed architecture combines a transconductance input stage with a differential operational amplifier operating in a transimpedance configuration and was implemented in onsemi BCD65 technology.

Analytical noise estimation and circuit-level simulations performed in the Cadence Virtuoso environment using the Spectre simulator confirmed that the proposed design achieves a voltage gain of approximately 44 dB at the operating frequency of 52 kHz, while maintaining a low input-referred noise density of approximately 3.58 nV/√Hz. The circuit operates from a 1.2 V supply voltage and maintains stable operation across process, supply voltage, and temperature variations, demonstrating its suitability for automotive operating conditions.

The presented design provides a solid foundation for further development of analog front-end circuits. Future work will focus on extending the presented design approach to additional receiver blocks and exploring low-power implementations in more advanced technology nodes, with the goal of further improving power efficiency and system integration in automotive sensing systems.

Acknowledgment

The authors thank Vilém Kledrowetz from BUT and Pavel Horský from onsemi for valuable technical discussions and support during the development of the circuit.

References

[1] Y. Wang, M. Koen, and D. Ma, "Low-noise CMOS TGC amplifier with adaptive gain control for ultrasound imaging receivers," IEEE Trans. Circuits Syst. II: Express Briefs, vol. 58, no. 1, pp. 26–30, Jan. 2011.

[2] S.-J. Jung, S.-K. Hong, and O.-K. Kwon, "Low-power low-noise amplifier using attenuation-adaptive noise control for ultrasound imaging systems," IEEE Transactions on Biomedical Circuits and Systems, vol. 11, no. 1, pp. 108–116, Feb. 2017.

[3] P. R. Gray, P. J. Hurst, S. H. Lewis, and R. G. Meyer, Analysis and Design of Analog Integrated Circuits, 5th ed. New York: John Wiley & Sons, 2009.

[4] H. T. Friis, "Noise figures of radio receivers," Proc. IRE, vol. 32, no. 7, pp. 419–422, July 1944.

[5] S. S. Rajput and S. S. Jamuar, "Low voltage analog circuit design techniques," IEEE Circuits and Systems Magazine, vol. 2, no. 1, pp. 24–42, 2002.

[6] B. Razavi, "A study of low-noise amplifiers," IEEE Journal of Solid-State Circuits, vol. 39, no. 4, pp. 597–605, Apr. 2004.

[7] A. P. Chandrakasan, S. Sheng, and R. W. Brodersen, "Low-power CMOS digital design," IEEE Journal of Solid-State Circuits, vol. 27, no. 4, pp. 473–484, Apr. 1992.

[8] B. Razavi, Design of Analog CMOS Integrated Circuits. New York: McGraw-Hill, 2001.

[9] Y. Tsividis and C. McAndrew, Operation and Modeling of the MOS Transistor, 3rd ed. Oxford, U.K.: Oxford University Press, 2011.

[10] M. J. M. Pelgrom, A. C. J. Duinmaijer, and A. P. G. Welbers, "Matching properties of MOS transistors," IEEE J. Solid-State Circuits, vol. 24, no. 5, pp. 1433–1439, Oct. 1989.

[11] L. Luo, Y. Wu, J. Diao, F. Ye, and J. Ren, "Low power low noise amplifier with DC offset correction at 1 V supply voltage for ultrasound imaging systems," in Proc. IEEE MWSCAS, 2018, pp. 137–140.

[12] M. Kayal and M. Declercq, Analog IC Design for CMOS VLSI Systems. Boston: Kluwer Academic Publishers, 2001.

[13] H. K. Park and B. Kim, "Design of CMOS transimpedance amplifiers for sensor interfaces," IEEE Sensors Journal, vol. 11, no. 11, pp. 2938–2945, Nov. 2011.

[14] P. Harpe, "Low-power analog front-end design for sensor interfaces," IEEE Circuits and Systems Magazine, vol. 15, no. 3, pp. 22–43, 2015.



NXP Semiconductors je globální technologická společnost zaměřená na vývoj a výrobu mikroprocesorů, mikrokontrolérů a analogových součástek.

Patří ke světové špičce v oblastech automotive elektroniky, embedded systémů, bezpečnosti, konektivity, NFC technologií a bezdrátového nabíjení.

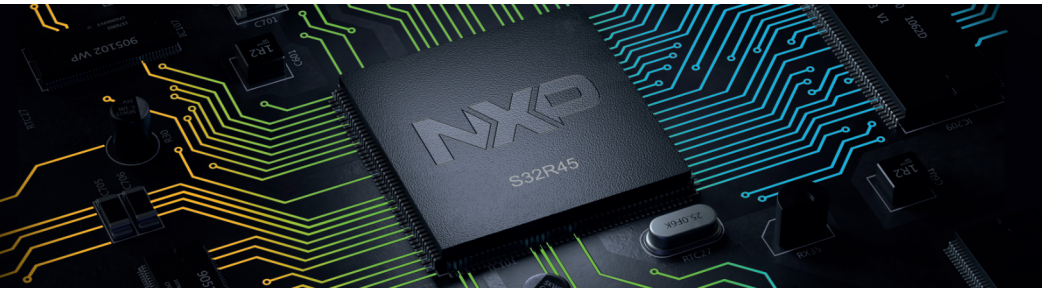

NXP Semiconductors se zabývá vývojem a výrobou mikroprocesorů a analogových součástek pro automobilový a spotřební průmysl, počítačové sítě a bezdrátové technologie. V České republice je společnost NXP Semiconductors zastoupena vývojovými centry v **Rožnově pod Radhoštěm** a v **Brně**. Nově jsme otevřeli i kanceláře na UTB Zlín a VŠB TU v Ostravě.

NXP nabízí přátelské a stabilní R&D prostředí, práci na reálných projektech a úzkou spolupráci s předními odborníky z celého světa.

Hledáme:

Každoročně nabízíme vice jak 30 příležitostí pro studenty a absolventy se zájmem o embedded technologie, především z oblastí Machine Learning, Security, Wirelesscharging, Connectivity, MotorControl, Functionalsafety, Embedded Systems, Automotive Ethernet, Signal processing či Embedded Software Validation and Verification infrastructure.

Předpoklady:

- Angličtina
- Zájem o Embedded technologie, robotiku, apod.
- Low level programování (embedded C, C++, Python)

Nabízíme:

- Příležitost zapojení do reálných projektů pod vedením světových špiček v oboru
- Flexibilní úvazky pro studenty i absolventy
- Přátelské, stabilní a inspirativní pracovní prostředí v rámci našich R&D týmů
- Motivační ohodnocení a podpora profesního růstu i osobního rozvoje
- Podpora diplomových a disertačních prací
- Workshopy, TechTalk, NXPCup a řada dalších aktivit pro studenty

Kontakt:
Careers at NXP | NXP Semiconductors
www.nxp.com/company/about-nxp/careers

NXP Semiconductors Czech Republic
1. máje 1009, 756 61 Rožnov pod Radhoštěm, Česko
Sochorova 3232, 616 00 Brno-Žabovřesky, Česko

Preparation of hard carbon derived from waste biomass for use in sodium-ion batteries

Antonín Šimek
Department of electrical and electronic Technology
Faculty of Electrical Engineering and Communication
Brno, Czech Republic
Antonin.Simek1@vut.cz

Jaromír Kotlán
Department of electrical and electronic Technology
Faculty of Electrical Engineering and Communication
Brno, Czech Republic
256652@vut.cz

Ondřej Čech
Department of electrical and electronic Technology
Faculty of Electrical Engineering and Communication
Brno, Czech Republic
cechondrej@vut.cz

David Trochta
Department of electrical and electronic Technology
Faculty of Electrical Engineering and Communication
Brno, Czech Republic
xtroch00@vut.cz

Pavel Čudek
Department of electrical and electronic Technology
Faculty of Electrical Engineering and Communication
Brno, Czech Republic
cudekp@vut.cz

Tomáš Kazda
Department of electrical and electronic Technology
Faculty of Electrical Engineering and Communication
Brno, Czech Republic
kazda@vutbr.cz

***Abstract* — This paper deals with the preparation and characterization of hard carbon derived from biomass as a promising material for anodes in post-lithium battery systems, particularly in sodium-ion batteries. Given the growing demand for sustainable and affordable energy storage, biomass, as a carbon-rich precursor, represents an environmentally friendly alternative to fossil fuels. The unique morphology of hard carbon, characterized by randomly oriented graphene domains and closed micropores, is the key to the effective intercalation and adsorption of large alkali metal ions. The paper analyses the basic composition of biomass with a focus on lignin, cellulose, and hemicellulose, and provides a theoretical review mapping the mechanisms of biomass pretreatment with a focus on the hydrothermal leaching process and subsequent high-temperature carbonization. This process was used in the experimental part of the work, where waste biomass was converted into hard carbon using a custom process and then electrochemically characterized in a Na-ion system.**

Keywords — hard carbon, biomass, carbonization, sodium-ion batteries, electrochemical energy storage

I. Introduction (*Heading 1*)

Sodium-ion batteries (Na-ion) are currently one of the most researched and promising alternatives to the dominant lithium-ion batteries (Li-ion) in materials chemistry and electrochemistry. This interest is accelerated primarily by the high global availability and low cost of mining and obtaining the necessary raw materials compared to Li-ion, which makes Na-ion ideal for large-scale stationary energy storage and renewable energy applications [1]. Although Na-ion systems share the fundamental operating principles of Li-ion, the so-called rocking chair mechanism, i.e., intercalation and deintercalation into and out of electrode materials, their commercialization faces significant physical and chemical obstacles. The main problem is the development of suitable anode materials. Commercially used graphite, which is the standard for lithium intercalation, is completely unsuitable for sodium ions. The sodium cation has a significantly larger ionic radius and different thermodynamic properties, which leads to instability of sodium-graphite intercalation compounds and the inability to achieve reversible charge storage with sufficient capacity [2].

II. Carbonaceous materials for Na-ion anodes

In the development of Na-ion batteries, attention is focused primarily on non-graphitizable carbons, collectively referred to as hard carbons (HC). This material is the first anode material for Na-ion that has successfully reached the stage of large-scale commercial implementation [3]. HC is characterized by its amorphous, macroscopically disordered structure, which at the microscopic level consists of randomly oriented graphitic microdomains surrounded by defect regions and a complex network of nanometric pores. The architecture of HC is often compared to a "falling domino" or "house of cards" model, where individual short graphene layers are spatially fixed [4].

The second amorphous material tested in the field of Na-ion anode materials is soft carbon (SC). Its structure consists of graphene layers (microcrystals) that are arranged predominantly parallel to each other. Thanks to these weaker interactions and parallel arrangement, the layers can easily rearrange when heated to high temperatures (usually above 2000 °C) and the material fully graphitizes. The fundamental difference between HC and SC therefore lies in the orientation of their graphene domains and the presence of strong bonds that determine their graphitizability [5]. The key difference between HC and SC lies in the precursors and the cross-linking mechanism. SC is typically synthesized from aromatic hydrocarbons with very low oxygen and heteroatom content, such as petroleum derivatives or coal tar. Due to weak cross-linking, SC layers are highly mobile, which allows them to shift and transform into highly

ordered graphite crystals at high temperatures [6]. In contrast, HC precursors, typically biomass or specific synthetic resins, contain large amounts of oxygen and other reactive functional groups. During thermochemical decomposition, these groups induce the formation of massive and strong cross-linking. This strong spatial interconnection immobilizes the carbon layers, incorporates crystalline domains into a rigid mass, and fundamentally prevents graphitization even at temperatures exceeding 2500 °C [7]. The resulting robust structure provides ideal interplanar distances for the diffusion of large sodium ions and closed pores for their deposition in the form of quasi-metallic clusters, which allows high capacities to be achieved at low operating voltages [1].

III. Biomass composition

For natural biomass to be effectively used as a precursor for the synthesis of HC with the required electrochemical performance, it must meet specific compositional requirements. Lignocellulosic biomass, which is the most attractive source in terms of sustainability and cost, consists of three main macromolecular components: cellulose, hemicellulose, and lignin as shown in Fig. 1. These components are integrally bound in plant cell walls and are supplemented by minor but critically important proportions of extractives and inorganic minerals (ash) [8]. Each component is characterized by a different chemical composition, thermal behaviour, and inherent ability to form a carbon matrix.

A. Lignin

It is a complex, irregular, and fully amorphous three-dimensional phenolic polymer with an enormous proportion of aromatic rings, which has the highest natural carbon content (up to 61%) of all biopolymers. It degrades over a wide temperature range based on the unique principle of in situ carbonization and close recombination – it does not depolymerize massively into a gaseous phase but gradually splits off oxygen-containing functional groups. It is structurally the most important precursor, leaving behind a mechanically resistant, strongly cross-linked carbon skeleton that graphitizes locally at high temperatures, creating a conductive hybrid architecture with ideal pathways for ion transport. [2; 9]

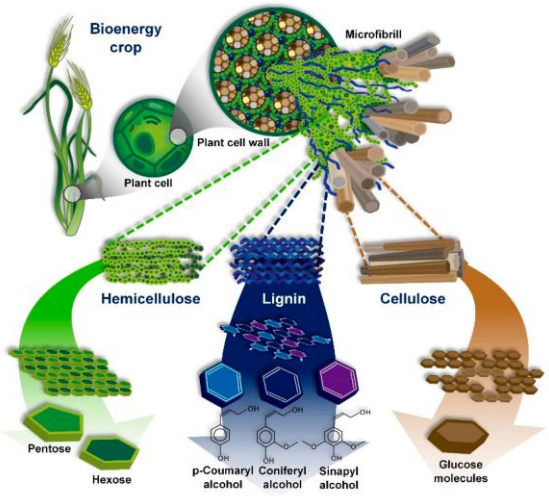

Fig. 1. Lignocellulosic biomass and its biopolymers: hemicellulose, lignin, and cellulose [10]

B. Cellulose

It is a linear homopolymer with a high degree of polymerization and considerable crystallinity, which forms the primary skeleton of biomass and contains approximately 42% carbon in the form of -D-glucopyranose. It is characterized by relative resistance to mild thermal decomposition, with its thermal stability and conversion mechanism strongly dependent on the specific polymorphic structure (monoclinic cellulose is significantly more thermally resistant than triclinic cellulose). The original rigid crystal lattice tends to limit ideal porosity; however, after its targeted disruption (e.g., by chemical modification), the material undergoes massive cross-linking during pyrolysis and excellently forms closed pores that are key for sodium storage. [8; 11]

C. Hemicellulose

It is a fully amorphous, flexible, and highly branched matrix composed of various pentose and hexose monomers (containing approximately 40% carbon), which is covalently bound into a three-dimensional network with lignin and cellulose. Due to its very low degree of polymerization and amorphous nature, it is the least thermally stable; when heated, it decomposes at significantly lower temperatures and evaporates massively in the form of gaseous products and volatile substances. Its carbonization is highly problematic in terms of battery applications, as massive evaporation leaves an extensive system of open macropores in the resulting carbon matrix, causing irreversible capacity loss due to formation of a Solid Electrolyte Interface (SEI) layer. [8; 12]

D. Extractive substances

In addition to the basic polymeric trio, raw biomass also contains minor fractions of extractive substances (sugars, tannins, resins, starches) that are highly susceptible to thermal degradation at relatively low temperatures. However, inorganic minerals, commonly referred to as ash, pose a much more serious problem for hard carbon synthesis [8]. Ash in biomass contains high concentrations of alkali metals and alkaline earth metals (potassium, calcium, and magnesium).

These inorganic impurities behave as highly aggressive gasification catalysts during high-temperature carbonization. They enable and accelerate chemical reactions between the solid carbon matrix and the released gases (CO, CO_2 and H_2O) produced during the pyrolysis of the surrounding biomass. The result is catastrophic corrosion of the carbon skeleton, in which a highly porous sponge-like structure is formed instead of the desired closed nanometric cavities. This open porosity with a huge specific surface area drastically reduces the density of the material and prevents effective sodium deposition. [1; 12]

IV. Pre-treatment of raw biomass

To achieve optimal electrochemical performance of HC in Na-ion, pretreatment of raw biomass is necessary, as its direct carbonization leads to unsuitable open porosity and low stability [13]. A crucial step is timely demineralization by acid washing, which removes catalytically active ash strictly before carbonization; if performed after carbonization, it paradoxically degrades the structure by opening blocked pores [12]. More advanced methods are used to further optimize and maximize the desired closed pores, such as hydrothermal pretreatment and

separation using eutectic solvents, which selectively remove labile hemicellulose and isolate key lignin [6; 13]. Controlled thermal pre-oxidation at low temperatures (typically up to 200 °C in the presence of air) is also suitable, as it mechanically strengthens the carbon skeleton before pyrolysis by cross-linking oxygen-containing functional groups and at the same time removes other organic impurities [14].

V. Carbonization and thermochemical conversion processes

The preparation of HC from biomass is often carried out in two steps to achieve optimal properties, namely by combining hydrothermal carbonization (HTC) and subsequent high-temperature pyrolysis. The environmentally friendly HTC process can also process moist raw materials in a closed reactor at moderate temperatures (180–260 °C) and elevated pressure, whereby a cascade of reactions – hydrolysis, dehydration, decarboxylation, and aromatization—it creates a more thermally stable carbon-rich intermediate product called hydrochar [15; 16]. This precursor is then subjected to high-temperature pyrolysis in an inert atmosphere, the maximum temperature of which critically shapes the resulting porosity and structure of the carbon matrix. While annealing at lower temperatures (around 1100 °C) leads to a material with a high proportion of defects and open pores, causing undesirable formation of a massive SEI layer and low initial Coulombic efficiency, increasing the temperature to 1300 °C and above initiates a fundamental physical transformation [12]. During this transformation, structural defects are healed, aromatic nuclei grow locally, and, most importantly, open micropores are sealed (the so-called sealing effect), creating the stable morphology necessary for efficient and safe operation in Na-ion batteries [2].

A. Advanced HC structure tuning

Although biomass naturally contains foreign elements (known as auto-doping) due to its origin, its structure is further modified in a targeted manner to achieve peak performance of anodes in Na-ion batteries. Chemical activation (e.g., with phosphoric acid) is commonly used, which massively cross-links the precursor and optimizes the volume of much-needed closed pores. The key strategy is then controlled doping with heteroatoms: nitrogen (N) supplies the lattice with free electrons and rapidly improves kinetics, while larger phosphorus (P) and sulphur (S) atoms physically stretch the carbon layers due to steric tension to facilitate sodium penetration, and deficient boron (B) effectively modifies surface affinity. To maximize performance, these elements are often combined into synergistic co-doping, or natural precursors (such as lignin) are directly bonded with synthetic polymers, creating a hybrid carbon matrix with ideal spatial arrangement and high stability. [17; 18]

VI. Experiment description and results obtained

The aim of the experimental work is to prepare HC for the Na-ion anode obtained from sustainable biomass. Organic waste biomass was selected for this purpose. This type of biomass is largely landfilled (70-75%), where it gradually decomposes. A smaller portion of this biomass produced is then used in agriculture as fertilizer (15-20%). Other possible uses with minor representation include energy purposes (5%) but with limited calorific value, as well as food processing and other applications (1%) [19; 20]. Based on the above, it can be concluded that this biomass source is very suitable and sustainable for use in the development of Na-ion, as it currently serves no significant purpose and is primarily landfilled, therefore, our study presents waste valorisation and waste value addition.

A. Pre-treatment and carbonization

In the initial step, the raw biomass underwent a washing process to extract soluble compounds and free impurities. This was followed by a standard thermal drying phase. Next, the biomass was subjected to acid leaching under controlled conditions to remove further inorganic contaminants, followed by rinsing with purified water to achieve pH neutralization. After a secondary drying step, the material was mechanically ground into smaller particles. This precursor was then carbonized at elevated temperatures under an inert atmosphere. Photos of the biomass samples from selected pre-treatment steps and the subsequent carbonized material are shown in Figure 2-A. To prepare the resulting hard carbon for characterization in a battery system, a fine-milling wet grinding process in a solvent was employed to achieve the requisite particle size while preventing agglomeration. Finally, the material was dried to completely evaporate the grinding solvent. Images from Scanning Electron Microscope (SEM) of final HC are shown in Figure 2-B.

B. Electrochemical characterization

Negative electrodes were prepared for electrochemical characterization of the prepared HC. The composition of the electrode material was selected in a ratio of 85-10-5, where 85 wt.% represents the prepared HC from biomass as the active material, 10 wt.% represents SuperP as the conductive additive, and the remaining 5 wt.% represents carboxymethylcellulose (CMC) as the binder. Demineralized water was used as the solvent. The electrode slurry was mixed in a planetary micro-mill to ensure homogeneity. After mixing, the slurry was applied to a carbon-coated aluminium current collector using blade coating and an automatic coater. The layer thickness was selected as 200 µm in the wet state. After application, drying took place at 50 °C in the presence of air. Subsequently, electrodes with a diameter of 18 mm were cut out, pressed at a pressure of 1.2 $t.cm^{-2}$, and finally dried at 110 °C in a vacuum. The electrodes prepared in this way were transferred to an M-Braun glove box (argon inert atmosphere), where the electrochemical half-cells were assembled (EL-Cells). These cells contained metallic sodium as the counter electrode (CE), a Whatman GF/C separator, 1M $NaPF_6$ in EC-PC (v-v) electrolyte, and the prepared negative electrodes as working electrodes (WE). For comparison of the results achieved, a negative electrode using commercial hard carbon Kuranode Type 2 - 9 µm was prepared in the same way. As part of the electrochemical characterization, cyclic voltammetry (CV) measurements were performed in the range of 0.01-2.5 V at a scan rate of 0.1 $mV.s^{-1}$. The CV results are shown in Figure 2-C. The CV was followed by formation cycles at a load of 0.1C. The theoretical capacity for current setting was set at 200 $mA.h.g^{-1}$. A comparison of the last discharge curves from the formation cycles is shown in Figure 2-D.

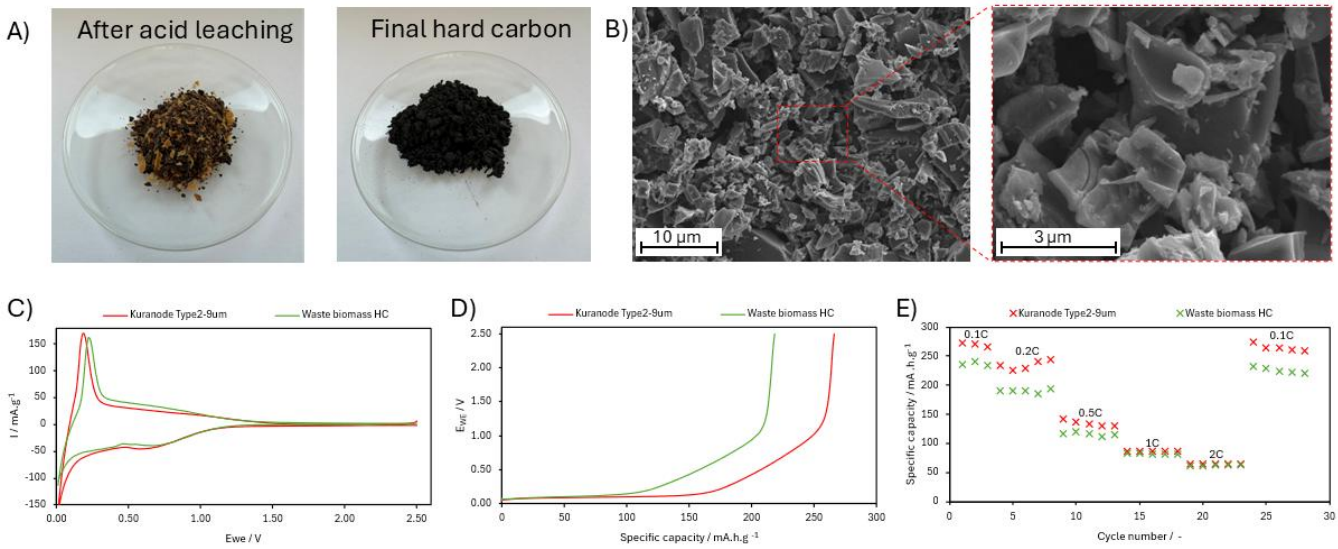

Fig. 2 A) Photos of selected biomass samples during pre-treatment and subsequent carbonization; B) SEM images of HC from organic waste biomas after carbonisation (left image HFW 50 µm, right image HFW 10 µm); comparison of electrochemical results between commercial HC and HC prepared from organic waste biomass: C) cyclic voltammetry at 0.1 $mV.s^{-1}$; D) final discharge curves from formation cycles at 0.1C and E) rate capability test.

The formation cycles were followed by a recalculation of the capacity based on the actual values obtained, and then a rate capability test (RC) was performed by gradually increasing the C-rate up to 2C, followed by a sharp decrease in load to 0.1C to determine the reversibility of the processes. The discharge capacities obtained from the formation cycles followed by RC are shown in Figure 2-E.

Cyclic voltammetry shows a shift of the main de-sodiation peak towards lower potentials and its slightly higher intensity in commercial HC compared to our HC prepared from waste biomass. This shift is due to lower polarization and better ion transport kinetics thanks to the highly optimized and homogeneous microstructure of commercial HC. The higher peak intensity indicates a higher proportion of so-called closed porosity, into which sodium is reversibly stored at low potentials (the so-called plateau region). Biomass carbon, on the other hand, usually has a more heterogeneous structure with more defects at the expense of these pores, which poses greater resistance to ion extraction (leading to higher overvoltage) and reduces the current response.

A comparison of discharge curves from the last formation cycle at a load of 0.1C shows that commercial HC achieves a higher total capacity and is characterized by a significantly longer plateau area at low potentials. Our synthesized material has a slightly lower total capacity, and its plateau section is noticeably shorter. This difference is most likely because the commercial material has a specifically optimized microstructure with a larger volume of closed nanopores, in which sodium is primarily stored in the plateau area, while our biomass carbon contains fewer of these specific pores.

The results of the RC test show that at low current loads, commercial HC achieves higher capacity (270 $mA.h.g^{-1}$) compared to our biomass HC (240 $mA.h.g^{-1}$). When the load is increased to 1C and 2C, the absolute difference in discharge capacities remains almost identical. After returning to the initial load of 0.1C, the capacity of both materials was restored, confirming their high structural stability during rapid cycling. The slightly better capacity retention of the commercial material (94.8% compared to 93.5% for biomass) is likely related to its smaller specific surface area and more stable SEI layer, which is less prone to irreversible parasitic reactions with the electrolyte.

VII. Conclusion

This work involved a brief literature review focused on the composition and processing options for biomass. As part of the experimental section, a sample of hard carbon from organic waste biomass was prepared, with the synthesis process involving precursor treatment in the form of water leaching, acid leaching, and final high-temperature carbonization. The electrochemical properties of the synthesized material were compared with commonly used, commercially available HC. Although the parameters achieved by our sample are not yet perfect, the results show that this material is competitive and comparable to the commercial standard. This opens promising opportunities for further research, in which we will focus on optimizing the internal structure of the material, for example by incorporating hydrothermal processing. For a detailed understanding of the relationship between morphology, structure, and electrochemical performance, these optimized samples will undergo comprehensive analysis in the future using scanning electron microscope, X-ray diffraction, and Raman spectroscopy.

Acknowledgment

This work was supported by the specific graduate research of the Brno University of Technology No. FEKT-S-26-8946.

AI disclaimer

During the preparation of this manuscript, the authors used DeepL for translation from Czech to English to improve grammar and readability. The authors also used Gemini 3.1 Pro

for content assistance, specifically to summarise literature and structure arguments. The authors reviewed and revised all outputs from these tools and takes full responsibility for the final content of the work.

REFERENCES

[1] T. Wang, R. Li, Q. Liu, and W. Liu, "A Review on the Upgradation of Biomass-derived Hard Carbon Materials", *Recent Patents on Nanotechnology*, vol. 19, no. 2, 2025, doi: 10.2174/0118722105287471240221094548.

[2] N. Kitchamsetti and S. Mhin, "Advances and Prospects of Lignin-Derived Hard Carbons for Next-Generation Sodium-Ion Batteries", *Polymers*, vol. 17, no. 20, Oct. 2025, doi: 10.3390/polym17202801.

[3] N. Kitchamsetti, K. -ho Kim, H. S. Han, and S. Mhin, "Biomass-Derived Hard Carbon Anodes for Sodium-Ion Batteries: Recent Advances in Synthesis Strategies", *Nanomaterials*, vol. 15, no. 20, Oct. 2025, doi: 10.3390/nano15201554.

[4] L. -F. Zhao *et al.*, "Hard Carbon Anodes: Fundamental Understanding and Commercial Perspectives for Na-Ion Batteries beyond Li-Ion and K-Ion Counterparts", *Advanced Energy Materials*, vol. 11, no. 1, Nov. 2020, doi: 10.1002/aenm.202002704.

[5] S. Ghosh, M. Zaid, J. Dutta, M. Parvin, and S. K. Martha, "Soft carbon in non-aqueous rechargeable batteries: a review of its synthesis, carbonization mechanism, characterization, and multifarious applications", *Energy Advances*, vol. 3, no. 6, 2024, doi: 10.1039/d4ya00174e.

[6] M. -S. Lim, G. H. Lim, Y. -J. Shin, J. S. Chae, J. -W. Lee, and K. C. Roh, "Partially graphitic structure-assisted hard carbon derived from lignin for sodium-ion battery anodes", *Energy Materials*, vol. 5, no. 9, May 2025, doi: 10.20517/energymater.2025.08.

[7] W. Li and J. Shi, "Lignin-derived carbon material for electrochemical energy storage applications: Insight into the process-structure-properties-performance correlations", *Frontiers in Bioengineering and Biotechnology*, vol. 11, Mar. 2023, doi: 10.3389/fbioe.2023.1121027.

[8] S. Darmawan, N. J. Wistara, G. Pari, A. Maddu, and W. Syafii, "Characterization of lignocellulosic biomass as raw material for the production of porous carbon-based materials", *BioRes*, vol. 11(2), no. 3561-3574, 2016.

[9] S. -Y. Long *et al.*, "Study on the lignin-derived sp2–sp3 hybrid hard carbon materials and the feasibility for industrial production", *Scientific Reports*, vol. 14, no. 1, Mar. 2024, doi: 10.1038/s41598-024-54190-x.

[10] J. U. Hernández-Beltrán, I. O. Hernández-De Lira, M. M. Cruz-Santos, A. Saucedo-Luevanos, F. Hernández-Terán, and N. Balagurusamy, "Insight into Pretreatment Methods of Lignocellulosic Biomass to Increase Biogas Yield: Current State, Challenges, and Opportunities", *Applied Sciences*, vol. 9, no. 18, Sep. 2019, doi: 10.3390/app9183721.

[11] J. Yi, M. Jia, H. Ai, and K. Du, "Engineering Closed-Pore Hard Carbon Anodes through Biomass Molecular Grafting to Enhance Plateau Capacity in Advanced Sodium-Ion Batteries", *ACS Applied Materials*, vol. 18, no. 8, Feb. 2026, doi: 10.1021/acsami.5c23590.

[12] R. F. Susanti, S. Alvin, and J. Kim, "Toward high-performance hard carbon as an anode for sodium-ion batteries: Demineralization of biomass as a critical step", *Journal of Industrial and Engineering Chemistry*, vol. 91, 2020, doi: 10.1016/j.jiec.2020.08.016.

[13] Y. Chen *et al.*, "Deconstruction Engineering of Lignocellulosic Biomass for the Preparation of High-performance Hard Carbon Anode Materials in Sodium-Ion Batteries", *Small*, Jan. 2026, doi: 10.1002/smll.202513464.

[14] Y. Chen, J. Cui, S. Wang, W. Xu, and R. Guo, "Biomass-Derived Hard Carbon Materials for High-Performance Sodium-Ion Battery", *Coatings*, vol. 15, no. 2, Feb. 2025, doi: 10.3390/coatings15020156.

[15] S. Sivaprasad, A. Manandhar, and A. Shah, "Hydrothermal Carbonization: Upgrading Waste Biomass to Char: Upgrading Waste Biomass to Char", *Ohioline*, 011 Jan 2021. https://ohioline.osu.edu/factsheet/fabe-6622 (accessed Mar. 10, 2026).

[16] F. Barbero *et al.*, "Hydrothermal carbonization synthesis of amorphous carbon nanoparticles (15–150 nm) with fine-tuning of the size, bulk order, and the consequent impact on antioxidant and photothermal properties", *Nanoscale Advances*, vol. 7, no. 5, 2025, doi: 10.1039/d4na00923a.

[17] A. Gopalakrishnan and S. Badhulika, "Effect of self-doped heteroatoms on the performance of biomass-derived carbon for supercapacitor applications", *Journal of Power Sources*, vol. 480, 2020, doi: 10.1016/j.jpowsour.2020.228830.

[18] A. Fahimi *et al.*, "Elemental doping and size effect-modified biomass-derived carbon: a fascinating microwave absorbing/shielding and energy saving material", *Journal of Materials Chemistry C*, vol. 12, no. 32, 2024, doi: 10.1039/d4tc01861c.

[19] L. Blinová, M. Sirotiak, A. Bartošová, and M. Soldán, "Review: Utilization of Waste From Coffee Production", *Research Papers Faculty of Materials Science and Technology Slovak University of Technology*, vol. 25, no. 40, Jun. 2017, doi: 10.1515/rput-2017-0011.

[20] B. Janissen and T. Huynh, "Chemical composition and value-adding applications of coffee industry by-products: A review", *Resources, Conservation and Recycling*, vol. 128, 2018, doi: 10.1016/j.resconrec.2017.10.001.

Design and Optimization of a Pre-treatment Line for Li-ion Battery Recycling

Jiří Báňa
Department Of Electrical and Electronic Technology
Brno University of Technology
Brno, Czech Republic
0000-0002-8119-979X

David Trochta
Department Of Electrical and Electronic Technology
Brno University of Technology
Brno, Czech Republic
0009-0000-3723-2831

Tomáš Kazda
Department Of Electrical and Electronic Technology
Brno University of Technology
Brno, Czech Republic
0000-0003-1973-02Z

The rapid transition to electromobility has triggered an exponential increase in demand for lithium-ion batteries (LIBs), necessitating their efficient recycling. The goal of mechanical and thermal pretreatment is to safely open the cells and concentrate the electrochemically active materials into a fine powder known as black mass. This article analyzes and justifies the optimal configuration of a state-of-the-art recycling line—from crushing in an inert atmosphere and air classification to pyrolysis and advanced flue gas management—with the aim of achieving maximum product purity in accordance with emerging global standards. Furthermore, the implementation of the EU Battery Regulation 2023/1542 imposes strict recovery targets, making highly efficient pre-treatment essential for the success of subsequent hydrometallurgical processing.

Keywords— Li-ion battery, recycling, pretreatment

I. Introduction

The global shift towards electric vehicles (EVs) and stationary energy storage is projected to push EV sales to approximately 38 million units annually by 2030.[1] This creates an unprecedented challenge to safely and ecologically process millions of tons of end-of-life batteries. While these batteries pose environmental risks due to toxic electrolytes, they also represent a highly concentrated strategic reserve of critical raw materials like lithium, cobalt and nickel. The legislative landscape, particularly the EU Battery Regulation 2023/1542, mandates ambitious recycling efficiencies, requiring authorized facilities to recover at least 90% of cobalt, copper, nickel, and 50% of lithium by 2027.[2]

II. Discharging of batteries

Before shredding, the residual electrical energy within the battery must be safely dissipated to prevent thermal runaway.[3] When designing a recycling line, engineers fundamentally face divergent pathways for deactivation: ohmic discharging, wet stabilization (brine discharging), or cryogenic stabilization. Ohmic discharging allows for the recuperation of residual energy back into the grid, which significantly improves the plant's overall carbon footprint. In contrast, brine discharging favors immersing batteries in aqueous salt solutions (NaCl or Na_2SO_4).[4]

The electrolytic conductivity of the solution slowly drains the voltage, while the massive heat capacity of water acts as a perfect heat sink, absolutely eliminating the risk of thermal runaway and explosions. This is particularly appealing for physically damaged or swollen cells where terminal connections are compromised. However, scaling this method up reveals severe environmental and operational flaws: water inevitably penetrates the breached cell casings, reacting violently with lithium and the conductive salt, $LiPF_6$. This hydrolysis generates highly toxic and corrosive hydrogen fluoride (HF), which dissolves into the bath.[5] Consequently, the plant produces massive volumes of heavily contaminated, toxic wastewater containing dissolved lithium salts, and fluorides.[6][7] Remediating this slurry presents a massive economic burden (OPEX). Furthermore, batteries

exit the bath saturated with water, forcing the integration of an energy-intensive drying step prior to mechanical shredding.[8] Alternatively, cryogenic stabilization uses liquid nitrogen to freeze batteries, stopping ionic flow, suppressing volatile vapor pressure to near zero, and inducing cryogenic embrittlement to ease subsequent mechanical comminution.[3]

The Superiority of Dry Inert Comminution:
To circumvent the wastewater crisis, modern high-capacity facilities have adopted dry crushing inside hermetically sealed chambers. Nitrogen or argon gas is continuously flushed into the rotary shear shredders, displacing oxygen typically under 2–4% by volume.[9] Even though the highly flammable organic carbonate solvents (EC, DEC, DMC, EMC) are released upon cell rupture, the absence of oxygen makes deflagration impossible. This dry comminution is immediately followed by a vacuum drying phase (absolute pressure < 300 mbar, temperature < 80 °C). This gentle thermal approach evaporates the solvents without reaching the thermal decomposition threshold of the $LiPF_6$ salt. The logic here is critical: by avoiding the formation of HF at this early stage, the shredder components do not require hyper-expensive alloys, and the recovered liquid electrolyte can be condensed and isolated as a high-energy fuel for subsequent processes. Additionally, bypassing high-temperature oxidation at this early stage saves nearly 1 ton of CO_2 equivalent per ton of processed batteries.[10][11]

Conversely, wet crushing systems submerge the shredder entirely in water or alkaline solutions to suppress sparks and dust. However, this method suffers from unacceptable solid mass losses. The aggressive mechanical abrasion and water turbulence wash out fine graphite and lithium compounds into the liquid.

III. Mechanical and Physical Separation

Following electrolyte extraction, the material stream is a highly heterogeneous mixture requiring precise physical separation. The separation of the black mass itself is carried out using a cascade of vibrating screens. The black mass fraction, consisting of cathode oxide powder and graphite flakes, typically has a particle size of less than 1 mm. However, screening alone is not effective enough to produce a very pure black mass.

Magnetic Separation: In modern LIB recycling, magnetic separation has evolved from a simple preliminary extraction of coarse steel casings into a sophisticated multi-stage process. Integrated into physical sorting lines, staged magnetic separation serves as a mechanism that limits subsequent contamination and enables the recovery of high-purity cathode material. By precisely calibrating the magnetic field to exploit minute differences in magnetic susceptibility, this technology can effectively isolate paramagnetic active materials—such as separating graphite from $LiFePO_4$ using exclusively physical mechanisms and without the use of chemical reagents.[12][13]

Eddy Current Separation (ECS): Non-ferrous metals (aluminium and copper current collectors) are ejected using Lorentz forces induced by high-speed magnetic rotors. The rapidly changing magnetic field induces Eddy currents within the conductive fragments, which generates a secondary magnetic field that repels the metals out of the stream.[14][15]

The Plastic Foil Bottleneck: A critical failure point in early recycling line designs was the reliance on conventional vibrating screens to remove plastic separator foils (PE/PP). These materials exhibit paradoxical rheological behaviour; they are extremely light yet possess an extreme tendency to agglomerate into concrete-like blocks that wrap around shafts and irremediably clog vibrating screens. Therefore, industrial designs must abandon physical barriers at this stage and integrate pneumatic air classification (zig-zag separators or cyclones) immediately after the mills to aerodynamically extract plastics before they can overwhelm the sieves.[11][16]

Thermal Decomposition: Pyrolysis at 500 °C: Mechanical processes alone are fundamentally limited by the presence of the polyvinylidene fluoride (PVDF) binder, which acts as a highly stable adhesive fixing the active particles to the metallic current collectors. Without addressing the binder, fine foil fragments massively contaminate the final black mass, destroying hydrometallurgical profitability. To achieve "black powder" specifications, thermal pre-treatment at temperatures approaching 500 °C in an inert or reducing atmosphere is mandatory.[17][18]

Furthermore, maintaining an inert or reducing atmosphere is critical not only for safety but for preserving metallurgical integrity. If oxygen were present during this high-temperature phase, the ultra-thin copper and aluminium current collectors would rapidly oxidize. The resulting metal oxides (CuO, Al_2O_3) are highly brittle and easily pulverize during the subsequent mechanical screening. This pulverization inextricably contaminates the fine black mass with base metals, which drastically increases the consumption of reagents and complicates the purification steps in the downstream hydrometallurgical leaching process. Figure 1 shown how oxygen atmosphere influence graphite structure. Graphite roasted at elevated temperature undergoes structural changes – cracking of graphite particles occurs caused by decomposition.

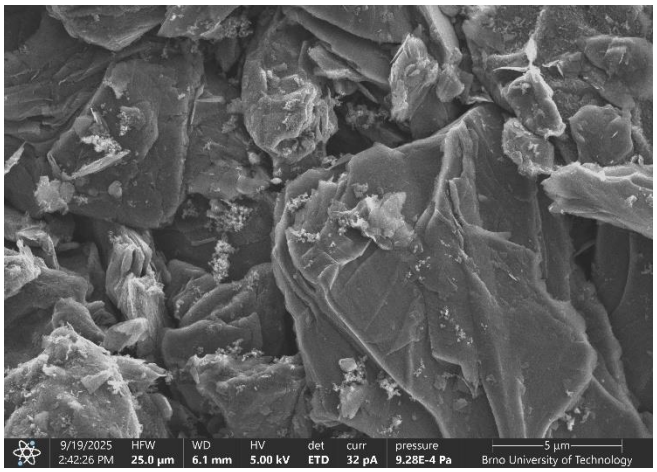

Fig. 1. Graphite sample after roasting at 500 °C in air

Binder Destruction: At temperatures between 400 and 480 °C, the macromolecular structure of PVDF collapses, cracking into short-chain alkanes and heavier tar products. The decomposition of PVDF entirely eliminates its adhesive force, leading to spontaneous delamination. Subsequent vibratory screening easily separates the pure black mass (NCM and graphite) from the clean metallic foils with a purity exceeding 84% to 96%.[19]

Roasting at 500 °C triggers a powerful phase transformation. Aluminium and carbon dust act as reducing agents, attacking the stable ternary oxide lattices (e.g., $LiNi_{1-x-y}Co_xMn_yO_2$). Lithium is selectively converted into highly soluble lithium carbonate (Li_2CO_3) or insoluble lithium fluoride (LiF), while heavy metals are reduced to their lower, insoluble oxide states. Consequently, subsequent hydrometallurgical leaching can extract over 97% of the lithium with remarkable ease using simple carbonated water.

To illustrate the underlying chemistry, carbothermic reduction of a typical cathode material (e.g., lithium cobalt oxide) by graphite can be described by the following reaction:

$$2\ LiCoO_2 + C \rightarrow Li_2CO_3 + 2\ CoO + CO$$


Fig. 2. Oxidation of LFP material during roasting in air

During this roasting phase at 500 °C, the transition metals are effectively reduced to lower oxidation states (for instance, Co^{3+} is reduced to Co^{2+}). This chemical alteration is profoundly beneficial for downstream hydrometallurgy. Instead of requiring highly aggressive inorganic acids combined with the continuous, expensive dosing of strong reducing agents (such as hydrogen peroxide) to break the stable oxide lattices, the pre-reduced black mass can be leached much faster. This significantly lowers hydrometallurgical operational costs, improves leaching kinetics, and reduces secondary chemical waste. Figure 2 shows oxidation of LFP cathode material after roasting at 500°C in air.

IV. Flue Gas Management: The Chemistry of Scrubbing and PFAS Destruction

Implementing a 500 °C pyrolysis furnace generates a massive spectrum of pollutants. The decomposition of $LiPF_6$ and PVDF produces extreme concentrations of highly toxic hydrogen fluoride (HF) and persistent per- and polyfluoroalkyl substances (PFAS).[20]

Wet Scrubbing:
At temperatures over 150 °C, the conductive salt decomposes:

$$LiPF_6 \rightarrow LiF + PF_5$$

The gaseous PF_5 is highly reactive and hydrolysis immediately upon contact with any moisture, generating phosphoryl fluoride and massive amounts of HF[21]:

$$PF_5 + H_2O \rightarrow POF_3 + 2HF$$

$$POF_3 + 3H_2O \rightarrow H_3PO_4 + 3HF$$

$$PF_5 + 4H_2O \rightarrow H_3PO_4 + 5HF$$

To neutralize these devastating acidic gases, engineers must design a liquid absorption system. A common misconception is to directly pump cheap milk of lime ($Ca(OH)_2$) into the scrubber column. However, reacting calcium directly with fluorides inside the column creates insoluble calcium fluoride (CaF_2), which rapidly forms hard scale on the column's packings, ultimately causing catastrophic clogging.

Therefore, an optimized two-step process is utilized:
Absorption in the Tower: The scrubber column utilizes a highly soluble neutralizing agent, typically a 4M solution of sodium hydroxide (NaOH). This ensures rapid, scale-free mass transfer of gases into the liquid phase (forming soluble NaF).

Precipitation in the Reactor: The saturated NaOH effluent is continuously drained from the tower into a separate chemical reactor, where it is finally mixed with the strongly alkaline

milk of lime ($Ca(OH)_2$). Here, the calcium chemically binds the fluorides, precipitating them as an inert solid sludge:

$$2F^- + Ca^{2+} \rightarrow CaF_2 \downarrow$$

The sludge is dewatered in filter presses, and the regenerated, fluoride-free water is cycled back to the scrubber. This properly dimensioned system drops HF emissions from 15 ppm down to a safe 0.1 ppm. [22][22]

VOC and PFAS Destruction: Even after wet scrubbing, the cooled gas remains saturated with volatile organic compounds (VOCs) and fluorinated radicals like CHF_3, CF_4, and COF_2 resulting from the incomplete destruction of PVDF at 500 °C.[20] While Granular Activated Carbon (GAC) physically adsorbs 85–95% of VOCs, it suffers from a fatal flaw in continuous commercial operations: rapid saturation. Once the molecular trap is full, "breakthrough" occurs, and the filter aggressively releases stored pollutants. In gigafactory settings, GAC lifespans degrade to a mere 6 months. More critically, activated carbon does not destroy PFAS; it merely stores them. For absolute legal compliance and safety, modern lines employ regenerative thermal oxidation (RTO). Operating at 1200 °C with a residence time exceeding 2 seconds, RTOs achieve near 100% destruction efficiency, completely mineralizing the most resilient carbon-fluorine bonds that survived the "mild" 500 °C pyrolysis furnace.[23]

From an economic and process engineering perspective, the strategic choice between GAC and RTO dictates the long-term financial viability of the recycling plant. While an RTO system requires a significantly higher initial capital expenditure (CAPEX) compared to simple carbon filter housings, its operational expenditure (OPEX) is drastically lower. Modern RTOs are designed to operate autothermally. Once the concentration of volatile organic compounds in the flue gas exceeds a self-sustaining threshold (typically 2–3 g/m^3), the thermal energy released by the combustion of these battery-derived VOCs is sufficient to maintain the 1200 °C operating temperature. This autothermal capability reduces auxiliary natural gas consumption to near zero, thereby permanently eliminating the exorbitant recurring costs and environmental liabilities associated with purchasing, replacing, and safely disposing of hundreds of tons of hazardous, PFAS-saturated activated carbon.[24][26][27]

V. Conclusion

LIB recycling is a topic that will become increasingly important in the future as the industry transitions towards circular economies. Recycling technologies are evolving with time and seeking efficiency, lowering recycling costs, and increasing environmental friendliness. From dry comminution under inert gas to pneumatic plastic extraction and 500 °C pyrolysis with in-situ carbothermic reduction, this advanced workflow overcomes the severe limitations of purely mechanical or wet processes. Supported by a robust multi-stage system featuring NaOH, $Ca(OH)_2$ scrubbing and RTO to completely mineralize PFAS, this optimal pre-treatment line guarantees both environmental safety and the production of highly lucrative Grade 1 black mass.

Acknowledgment

This work was supported by the specific graduate research of the Brno University of Technology No. FEKT-S-23-8286.

The authors acknowledge the use of artificial intelligence tools (specifically, large language models) to assist in the drafting, structural organization, and linguistic refinement of this manuscript. All AI-assisted content was rigorously reviewed, edited, and scientifically validated by the human authors, who assume full responsibility for the accuracy, originality, and integrity of the final published work.

BIBLIOGRAPHY

[1] O. Velázquez-Martínez, J. Valio, A. Santasalo-Aarnio, Markus A. Reuter, and R. Serna-Guerrero, “A Critical Review of Lithium-Ion Battery Recycling Processes from a Circular Economy Perspective”, *Batteries*, vol. 5, no. 4, Nov. 2019

[2] *12. European Union (EU). (2023). Regulation (EU) 2023/1542 of the European Parliament and of the Council concerning batteries and waste batteries. Official Journal of the European Union, L 191/1*. 2023.

[3] P. -A. Le, “A review of construction and sustainable recycling strategies of lithium-ion batteries across electric vehicle platforms”, *RSC Advances*, vol. 15, no. 42, Jan. 2025

[4] G. Jiang, D. Lee, D. Raimbault, Paul A. Anderson, and Gary A. Leeke, “Acetone as an agent for PVDF recovery and delamination of lithium-ion battery electrodes”, *Resources Conservation and Recycling*, vol. 209, Jun. 2024

[5] L. Zhu and M. Chen, “Development of a Two-Stage Pyrolysis Process for the End-Of-Life Nickel Cobalt Manganese Lithium Battery Recycling from Electric Vehicles”, *Sustainability*, vol. 12, no. 21, Nov. 2020

[6] M. Liu, J. Vatamanu, Han Y. H. Chen, L. Xing, K. Xu, and W. Li, “Hydrolysis of LiPFsub6/sub-Containing Electrolyte at High Voltage”, *ACS Energy Letters*, vol. 6, no. 6, May 2021

[7] M. Stich, M. Göttlinger, M. Kurniawan, U. Schmidt, and A. Bund, “Hydrolysis of LiPFsub6/sub in Carbonate-Based Electrolytes for Lithium-Ion Batteries and in Aqueous Media”, *The Journal of Physical Chemistry C*, vol. 122, no. 16, Apr. 2018

[8] Marcelo Albuquerque de Oliveira, B. Abreu, and H. Costa, “Shredding of Lithium-Ion Batteries: Overview and Industrial Perspective”, in *IntechOpen eBooks*, 2024

[9] Marcelo Albuquerque de Oliveira, B. Abreu, and H. Costa, “Shredding of Lithium-Ion Batteries: Overview and Industrial Perspective”, in *IntechOpen eBooks*, 2024

[10] N. Zachmann, M. Petraniková, and B. Ebin, “Electrolyte recovery from spent Lithium-Ion batteries using a low temperature thermal treatment process”, *Journal of Industrial and Engineering Chemistry*, vol. 118, Nov. 2022

[11] C. R. Wilke, Denis Manuel Werner, A. Kaas, and Urs A. Peuker, “Influence of the Crusher Settings and a Thermal Pre-Treatment on the Properties of the Fine Fraction (Black Mass) from Mechanical Lithium-Ion Battery Recycling”, *Batteries*, vol. 9, no. 10, Oct. 2023

[12] Z. Hu, J. Liu, T. Gan, D. Lu, Y. Wang, and X. Zheng, “High-intensity magnetic separation for recovery of LiFePO4 and graphite from spent

lithium-ion batteries", *Separation and Purification Technology*, vol. 297, Jun. 2022

[13] Daniela Romero Guillén *et al.*, "Physical Process for Li-Ion Battery Recycling from Electric Vehicles", *Industrial & Engineering Chemistry Research*, vol. 63, no. 45, Nov. 2024

[14] Y. Bai, H. Zhu, L. Zu, and H. Bi, "Eddy current separation of broken lithium battery products in consideration of the shape factor", *Journal of Material Cycles and Waste Management*, vol. 25, no. 4, Apr. 2023

[15] H. Lee, Elizabeth H. Driscoll, Kristian E. Waters, E. Kendrick, and R. Sommerville, "Pretreatment and Valorization of Critical Materials from Lithium-Ion Batteries Using Electrostatic and Magnetic Separation", *Advanced Energy and Sustainability Research*, vol. 6, no. 9, Mar. 2025

[16] J. Choi, K. Lee, H. Lee, and K. Kim, "Advancements in lithium-ion battery recycling technologies: Exploring module-scale crushing and air separation techniques", *Physicochemical Problems of Mineral Processing*, Jan. 2025.

[17] Alexander J. Bologna *et al.*, "Pyrolyzed "Black Mass" Feedstocks and Their Synthetic Proxies Relevant to Li-Ion Battery Recycling", *ACS Omega*, vol. 10, no. 24, Jun. 2025

[18] E. Feng *et al.*, "Ultrafast Material Separation by Electrochemical Delamination for Spent Lithium-Ion Batteries Recycling", *ACS Sustainable Chemistry & Engineering*, Sep. 2025

[19] T. Wang *et al.*, "Thermal Decomposition of Organic Components in Spent Lithium-Ion Batteries: Stagewise Evaluation and Kinetic Analysis", *Industrial & Engineering Chemistry Research*, vol. 64, no. 32, Aug. 2025.

[20] J. Xiao, J. Wang, J. Hong, and Z. Xu, "Unveiling and identifying the overlooked fluoride hazard derived from spent lithium-ion battery recovery", *Journal of Hazardous Materials*, vol. 496, Aug. 2025

[21] A. Rensmo, E. Savvidou, Ian T. Cousins, X. Hu, S. Schellenberger, and Jonathan P. Benskin, "Lithium-ion battery recycling: a source of per- and polyfluoroalkyl substances (PFAS) to the environment?", *Environmental Science Processes & Impacts*, vol. 25, no. 6, Jan. 2023

[22] M. Wang, Q. Tan, L. Liu, and J. Li, "A Facile, Environmentally Friendly, and Low-Temperature Approach for Decomposition of Polyvinylidene Fluoride from the Cathode Electrode of Spent Lithium-ion Batteries", *ACS Sustainable Chemistry & Engineering*, vol. 7, no. 15, Jul. 2019

[23] R. Sun *et al.*, "New Insights into Thermal Degradation Products of Long-Chain Per- and Polyfluoroalkyl Substances (PFAS) and Their Mineralization Enhancement Using Additives", *Environmental Science & Technology*, vol. 58, no. 50, Dec. 2024

[24] B. Wang and F. You, "Sustainable PFAS Removal from Electronics Wastewater through a Cost-Health Trade-Off Framework", *Environmental Science & Technology*, vol. 60, no. 6, Feb. 2026

Fabrication of electrically conductive thick films using masked stereolithography

Jaromír Jarušek
Department of Microelectronics, FEEC
Brno University of Technology
Brno, Czech Republic
jaromir.jarusek@vutbr.cz

Imrich Gablech
Department of Microelectronics, FEEC
Brno University of Technology
Brno, Czech Republic
imrich.gablech@vutbr.cz

***Abstract*—This paper presents and evaluates electrically conductive thick films, printed using a masked stereolithography 3D resin printer from prepared conductive resin mixtures. Two types of conductive fillers are tested and evaluated – silver powder and multi-walled carbon nanotubes. Sheet resistance and thickness of printed patterns are measured for different filler mass concentrations. Additionally, the effects of different exposure times on the properties of printed layers are evaluated.**

Keywords—resin, 3D printing, conductive inks, photoimageable materials, silver, multi-walled carbon nanotubes, additively manufactured electronics

I. Introduction

Thick-film technologies define a set of processes that enable the additive manufacturing of patterns for electronic devices and components. Thick films can be commonly defined by thickness in the range of (5 - 100) μm, but the key defining factor is the layer material precursors that are powders suspended in certain solvents [1]. These materials commonly take the form of inks and pastes. They can be applied by different methods such as screen printing, stencil printing, pad printing, direct ink writing, dip-coating, spray-coating, doctor blade coating, ink-jet printing, and, for flexible substrates and printed electronics, also by gravure printing, flexographic printing, and offset printing [2][3].

The most common technique for applying thick films is screen printing. The paste is applied through a mesh by a squeegee. The mesh is commonly patterned with an emulsion that forms apertures for paste. The screen is in contact with the substrate only when pushed by the squeegee and then peels off due to screen tension [4]. Printing quality is affected by parameters, such as mesh tension, mesh count, mesh material, squeegee pressure, speed, and snap-off distance [3][4]. Furthermore, the composition and properties of the inks and pastes used may differ significantly across application techniques and substrates with different material properties. [5]

In applications, such as microwave circuits and millimeter-wave waveguides, thick-film patterns with dimensions and precision commonly achieved only by photolithography are required [6][7]. This led to the development of a hybrid process called photoimageable thick-film technology. This technology utilizes low-viscosity pastes that are screen printed through a blank nylon mesh screen and then dried at 100 °C. In the next step, the dried layer is exposed to UV light in the range (320 - 360) nm through a photomask, which initiates paste photopolymerization in the paste vehicle. Exposure time varies with filler material: shortest for dielectrics, intermediate for metals, and longest for metal oxides. Unexposed areas are then developed by spraying with an aqueous solution of sodium carbonate (Na_2CO_3). Finally, the thick films are fired in a process similar to that used for common thick-film pastes.[6]

A similar technology, not commonly considered but applicable to thick films, is stereolithography (SLA), also known as vat photopolymerization. The traditional method is based on curing photopolymer resin with an ultraviolet (UV) laser, forming 3D objects layer by layer. Another exposure method is digital light processing (DLP), which uses digital micro-mirror devices (DMD) [8]. Recent advancements in liquid crystal displays (LCDs), such as high resolution, small pixel sizes, and UV resistance, have enabled their use as an exposure mask for SLA, thus enabling the MSLA method [9].

For SLA printers, some specialized electrically conductive resins are already commercially available. They are based on specially formulated resins and selected fillers, including graphene nanoplatelets, single-walled and multi-walled carbon nanotubes (MWCNTs), and silver powders [10]. Formulations of electrically conductive resins are also reported in the literature. High electrical conductivity ($\approx 1000\ S{\cdot}cm^{-1}$) can be achieved by more advanced mixtures, such as silver-coated copper flakes in combination with MWCNTs [11], while MWCNTs alone provide relatively lower conductivity ($2.7{\cdot}10^{-4}\ S{\cdot}cm^{-1}$) [12]. A key factor for an electrically conductive resin is the percolation threshold.

Another more advanced method for printed circuit board fabrication involved the use of the professional prototyping system DragonFly IV (Nano Dimension Ltd.). This instrument can print conductive and insulating inks simultaneously using piezo drop-on-demand inkjet technology and thus can fabricate multilayer printed circuit boards and similar devices [13]. Conductors are based on a silver nanoparticle ink that can be *in situ* sintered at low temperatures [14].

In this paper, the use of MSLA for printing conductive thick-film structures is proposed and evaluated as a complementary capability beyond its conventional role as a 3D printing technique. Conductive inks or pastes are formulated to remain compatible with the resin used for the structural parts of the printed device. The same resin serves as a binder for electrically conductive fillers, and the resulting printed device can thus be considered a quasi-monolithic chip. The resin and conductive fillers are mixed in a manner that requires only standard

laboratory equipment for preparation while avoiding chemical reactions or interference with the resin formulation. Potential applications include 3D-printed microfluidic devices for electrochemical detection with integrated electrode interconnections and possibly printed electrodes. In the proposed workflow, the conductive pattern corresponds to a single printing layer: the printing process is paused, the conductive resin is introduced, printing is resumed to form the conductive pattern, and then paused again for the removal of residual resin before continuing with the rest of the structure. Since the proposed process may introduce additional fabrication challenges, this study evaluates only the printing of electrically conductive patterns as a standalone process. The presented method may also enable the fabrication of conductive patterns as a digitally defined thick-film photolithographic technique on optically transparent substrates, such as glass coverslips or polymer foils, without requiring specialized or expensive equipment.

II. Resin Preparation and Printing

A. Preparation of Mixtures

The photopolymeric resin Liqcreate Bio-Med Clear was used as a base material for preparing electrically conductive mixtures. Key parameters, listed in Table I, such as the critical energy dose (E_c) and penetration depth (D_p), are coefficients of Jacobs' working curve [15], which characterizes the layer growth of photopolymerized resin as a function of the received UV energy dose via a logarithmic relationship.

TABLE I. Key parameters of Liqcreate Bio-Med Clear resin [16]

Parameter	Value	Unit
Viscosity at 25°C	475	mPa·s
Density	1.18	g·cm^{-3}
E_c (405 nm)	6.10	mJ·cm^{-2}
D_p (405 nm)	0.12	mm
E_c (385 nm)	4.10	mJ·cm^{-2}
D_p (385 nm)	0.05	mm

Two types of electrically conductive fillers were used in two different types of mixtures. The first is an electrically conductive 99.99 % silver filler powder manufactured by Shenzhen Lidianrun E-commerce Co., Ltd. The second type is the MWCNTs material HQNANO-CNTs-010-0, with a length (3 - 12) μm and outer diameter (8 - 15) nm. This MWCNT length range is expected to provide a practical compromise between achievable dispersion, a low percolation threshold, and a reduced risk of electrical short circuits in printed patterns.

Silver mixtures were prepared in 80 %, 70 %, 65 %, 60 %, and 50 % mass concentrations of filler. For evaluation, 1 g samples of each concentration were prepared. For the selected concentration, additional batches were prepared as needed. Sample preparation was as follows. A target fraction of 1 g of silver filler was dispensed into a 5 mL polystyrene weighing boat using a laboratory micro spoon and weighed using a Sartorius CP255D analytical balance. The remaining resin was then dispensed into the same boat using a micropipette. Silver filler powder was then gradually mixed into the resin with a spatula.

MWCNT ink was prepared in a single mass concentration of 1.5 %. This concentration was selected to ensure that the percolation threshold is reached while keeping it low enough to preserve sufficient optical and UV transmittance and the liquid properties of the resin. To simplify handling of the ink and MWCNTs, the entire process was done inside a 25 mL amber glass bottle with a plastic cap. Direct handling of MWCNT powder was performed in a fume hood. The ink was prepared as follows. First, 300 mg of MWCNTs were gradually dispensed into a 25 mL bottle using a laboratory micro spoon, and the dispensed amount was verified using an analytical balance. Then the remaining 24.7 g of resin was transferred from a beaker into the bottle. To break down the large powder clusters of MWCNTs and roughly disperse them in the resin, the mixture was stirred using a magnetic stirrer at 1400 rpm for 30 minutes. Then, ultrasonication was performed in an ultrasonic bath with a volume of 2 L, 240 W peak ultrasonic power, 120 W effective ultrasonic power, and a frequency of 35 kHz for 1 hour.

B. Printing from Prepared Mixtures

The used printer was an Elegoo® Mars 5 Ultra, which features a monochromatic LCD with specified dimensions (153.36 × 77.76) mm, resolution of (8520 × 4320) pixels with a single pixel dimension of (18 × 18) μm^2. The printer uses a 405 nm UV source, and the average UV light intensity across the display is (3.8 ± 0.3) mW·cm^{-2}, with ≈ 4.0 mW·cm^{-2} in the center of the display, where all exposures in this study were performed. The UV light intensity values were measured using the ChituSystem® UV light meter.

Glass coverslips with dimensions (24 × 60) mm^2 and specified thickness in the range (0.13 - 0.17) mm were cut using a diamond tip to achieve a more practical size of (24 × 30) mm^2. Coverslips were rinsed using an IPA wash bottle and dried using compressed air. Additionally, to facilitate the spreading of the silver-filled resin samples with mass concentrations of 80 % and of 70 %, which exhibited a viscosity closer to a paste, a clear polypropylene foil with a thickness of 75 μm was used.

Before printing, the prepared mixtures were dispensed using two methods, depending on their viscosity. For viscous pastes, the entire 1 g of the prepared sample was dispensed onto clear plastic foil and then covered with a clean coverslip. For lower-viscosity inks, a sufficient volume was drop-cast onto a cover slip, enabling the printing of a layer > 150 μm thick. This layer thickness was estimated as an upper limit based on the optical properties of the mixtures, taking into account the penetration depth of the pure resin and the decrease in optical transmittance caused by the added fillers.

For printing, the prepared coverslips were placed at the center of the display. Coverslips with drop-cast inks were exposed under a similar regime to that used for vat cleaning. For inks confined between the coverslip and the foil, the exposure conditions were similar to those of the standard printing regime, with the exception of full-thickness layer formation. Exposure time was set by default to 60 s for all samples, to ensure proper layer formation and, for selected additional samples, decreased to 20 s and 10 s, to optimize for low lateral overcuring

After exposure, residual liquid mixtures were removed using an IPA wash bottle, a rubber brush, and compressed air. The amounts of residues on glass slides from sandwiched pastes

were smaller because the unexposed paste adhered to plastic foil. Cleaned samples were then left to dry for at least 10 minutes before further analysis.

III. Printed Thick Films Characterization and Evaluation

A. Test Structures

To evaluate electrical parameters as well as printing fidelity, patterns shown in Fig. 1a-c were designed. The first pattern in Fig. 1a is a test structure for the van der Pauw measurement method of sheet resistance. The dimensions of the internal square are (10 × 10) mm^2 and the lead width is 1 mm.

The second pattern in Fig. 1b was designed to evaluate possible limits of percolation in printed layers for thinner interconnections. It serves as a simple continuity test. Because of the silver filler structure and possible high sheet resistance of MWCNTs-filled resin, the width of the lines was in the range (0.6 - 1.5) mm, and the length was kept constant at 15 mm. This pattern was also used for 2-point sheet resistance measurements.

The third pattern in Fig. 1c consists of interdigitated electrodes (IDEs), which, compared with a simple geometry, incorporate features that can be used to evaluate printing fidelity capabilities, such as non-orthogonal angles and radii. This design can be considered as a practical test case, e.g., for thick-film sensors. The connector pads have a pitch of 5.08 mm and a width of 1.25 mm. The gap between the IDEs is 0.3 mm wide.

Because of their simplicity, all designs were drawn in Inkscape 1.4 and exported at 1411.11 dots per inch (DPI) as .png files. The exported files were then imported into GIMP bitmap editor to shift the minimum grayscale values to 128 and above. Finally, edited bitmaps were imported into proprietary ChituBox .ctb files using the open-source software UVTool, dedicated to modifying printing files for resin printers.

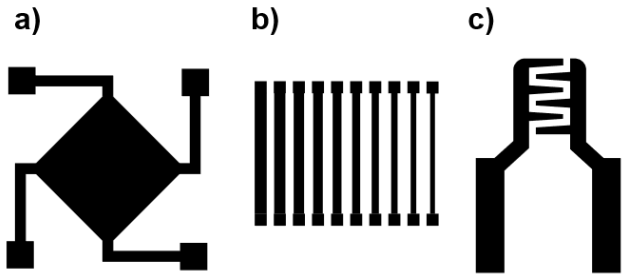

Fig. 1. Designed test structures: a) van der Pauw, b) continuity test, c) interdigitated electrode structure.

B. Optical Evaluation

Printed thick films and structures were observed using a Zeiss Axio Imager.A2m microscope with the following illumination technique combinations: reflected light, dark-field contrast, and transmitted light bright-field contrast.

The morphology of silver-filled resin is shown in Fig. 2. The silver fillers used have a coarse, granular composition, with particles resembling bulk silver fragments.

Microstructure of MWCNTs-filled resin, shown in Fig 3, was analyzed using a 50x magnification objective. It can be observed that, between small clusters of MWCNTs with a size of ≈ 20 μm, individual MWCNTs also interact with the resin and form visible submicron fibers.

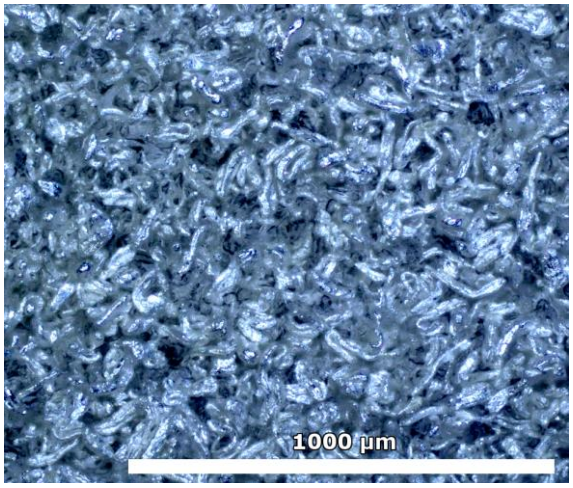

Fig. 2. Detail of resin filled with silver powder at 70 % mass, observed in dark-field reflected light.

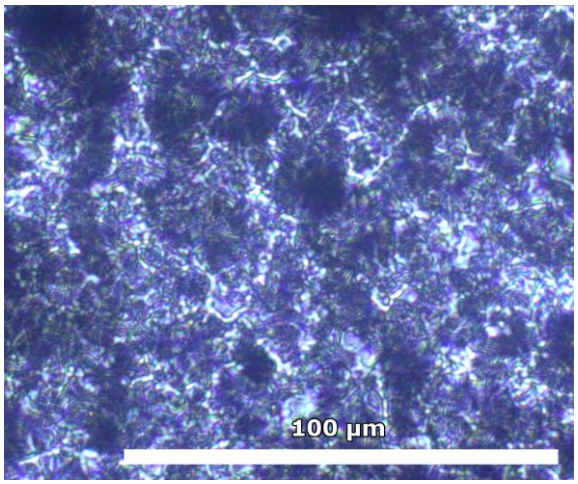

Fig. 3. MWCNTs-filled resin detail, observed in bright-field transmitted light, 50× objective magnification.

The edge definition of printed van der Pauw test structures is compared in Fig. 4a-d. An obvious lateral overcuring of resin with silver filler is apparent in Fig. 4a and Fig. 4b, in comparison with the designed shape marked by a dashed yellow line. On the other hand, the area within the marked shape in Fig. 4a is out of focus, indicating that a thicker layer formed in the intended area. Fig. 4c-d shows relatively well-defined structure edges.

Printing fidelity from silver-filler resin exposed at 10 s can be evaluated from the details of IDEs in Fig. 5a-b. By comparing transmitted and reflected light, it is evident that there is residual resin resembling the filler structure at the bottom of the glass substrate.

In Fig. 6, IDEs printed from MWCNT ink are shown, and similarly to Fig. 4d, the structure was printed relatively accurately and may require only a small adjustment of exposure time.

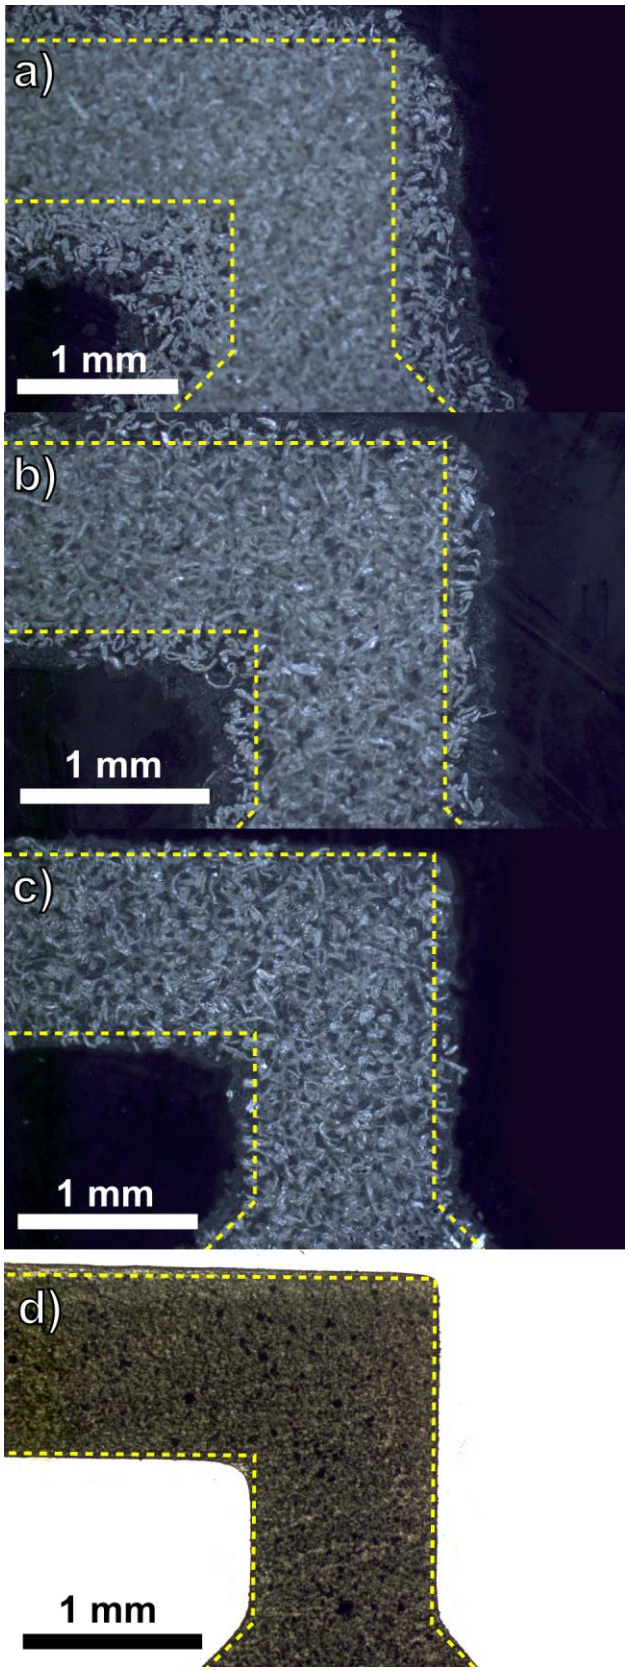

Fig. 4. Comparison of edge definition for silver-filled pastes exposed at a) 60 s, b) 20 s, c) 10 s, at reflected light dark-field and d) MWCNT ink in transmitted light brigh-field.

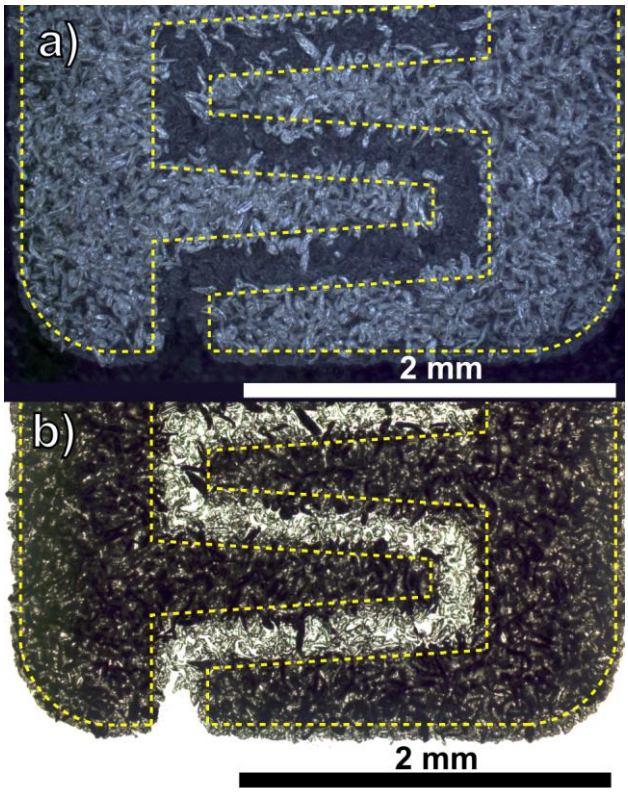

Fig. 5. Views of IDEs printed from 80 % mass concentration paste at 10 s exposure time: a) in reflected light dark-field, b) in transmitted light bright-field.

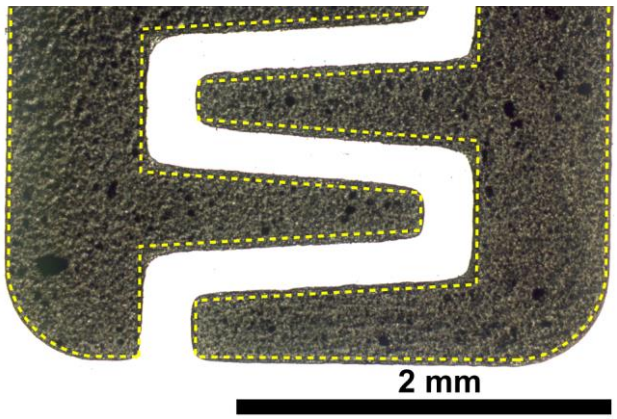

Fig. 6. Views of printed IDEs from MWCNT ink in transmitted light brigh-field.

C. Sheet Resistance and Thickness Measurement

Sheet resistance of silver-filled samples was measured on the printed van der Pauw structures using a Keithley DMM6500. The measured values were obtained by acquiring approximately 300 readings into the instrument buffer and averaging the recorded data. Results are in Table II.

The sheet resistance of the MWCNT ink was measured only on the printed continuity test pattern using a UNI-T622E LCR meter, because the high resistance of the van der Pauw test structure lead-outs caused an issue with 4-point measurement. Results are in Table III.

The thickness of layers was measured using a Bruker Dektak® XT stylus profilometer. The layer surface was measured over a trajectory of ≈ 5 mm. Profile data values were approximated by a normal distribution, with the layer thickness corresponding to the mean and the standard deviation representing surface roughness. Results are presented in Table II and Table IV. For MWCNT ink exposed for 60 s, the mean thickness was 42 µm, and the standard deviation was 14 µm.

TABLE II. PARAMETERS OF SILVER-FILLED MIXTURES FOR DIFFERENT MASS CONCENTRATIONS AT 60 S EXPOSURE TIME

	Mass concentration (%)				
	50	**60**	**65**	**70**	**80**
mean thickness (µm)	102	82	106	105	112
standard deviation (µm)	40	32	46	45	46
sheet resistance (mOhm·□$^{-1}$)	55	61	47.4	53	41
standard deviation (mOhm·□$^{-1}$)	9	15	1.5	2	3

TABLE III. RESULTS OF 2-POINT RESISTANCE MEASUREMENTS FOR MWCNT INK

	Path width (mm)									
	0.6	**0.7**	**0.8**	**0.9**	**1**	**1.1**	**1.2**	**1.3**	**1.4**	**1.5**
measured resistance (MΩ)	22	20	69	24	23	18	18	15	10	9
uncertainty (MΩ)	2	1	2	1	1	1	1	1	1	1
sheet resistance (MΩ·□$^{-1}$)	0.9	0.9	3.7	1.4	1.5	1.3	1.4	1.3	0.9	0.9

TABLE IV. THICKNESS VALUES FOR SELECTED CONCENTRATIONS OF SILVER-FILLED RESINS AND EXPOSURE TIMES

	Mass concentration (%)			
	65		**80**	
	Exposure Time (s)			
	10	**20**	**10**	**20**
mean thickness (µm)	56	72	60	88
standard deviation (µm)	24	30	26	35

Silver-filled layers prepared at lower exposure times suffered from electrical continuity issues; sheet resistance was measured only for the sample with an 80 % silver mass concentration at 20 s exposure time, using the continuity test structure. Results are listed in Table V.

TABLE V. RESULTS OF 2-POINT RESISTANCE MEASUREMENTS FOR 80 % SILVER PASTE AT 20 S

	Path width (mm)									
	0.6	**0.7**	**0.8**	**0.9**	**1**	**1.1**	**1.2**	**1.3**	**1.4**	**1.5**
measured resistance (Ω)	8	18	16	5	5	4	25	2.5	3.1	3.0
uncertainty (Ω)	2	3	4	1	1	1	3	0.5	0.5	0.3
sheet resistance (Ω·□$^{-1}$)	0.3	0.9	0.8	0.3	0.3	0.3	1.9	0.2	0.3	0.3

IV. CONCLUSION

Based on the obtained results, it can be concluded that the MSLA technique may be suitable for fabricating electrically conductive thick films, although selecting an appropriate electrically conductive filler may be difficult. The tested silver filler showed low sheet resistance at higher layer thicknesses, but the powder's large granular structure significantly limited printing fidelity and the preparation of conductive layers at lower thicknesses. It could also be concluded that MSLA was able to utilize even such a type of filler that may not be suitable for any other printing technique.

It was demonstrated that electrically conductive resin can be prepared using MWCNTs with only common laboratory equipment. MWCNTs-filled resin was found to be less suitable for electrically conductive thick films, due to its high sheet resistance. On the other hand, it showed high printing fidelity.

Therefore, it could be considered that materials such as silver nanowires could provide a combination of the benefits of both tested materials. It was demonstrated that affordable MSLA printers, in combination with well-chosen materials such as electrically conductive fillers, have the potential to enable novel rapid prototyping techniques.

REFERENCES

[1] B. Malič, D. Kuščer, M. Vrabelj, and J. Koruza, "5 - Review of methods for powder-based processing," in Magnetic, ferroelectric, and multiferroic metal oxides, in Metal oxides series. Amsterdam, Netherlands: Elsevier, [2018], pp. 95-120. Available: https://www.sciencedirect.com/science/chapter/edited-volume/abs/pii/B9780128111802000050

[2] G. M. Berube, "Thick-film – a mature technology at the cutting edge of advances in the electronics industry", IMAPSource Proceedings, vol. 2022, no. IMAPS Symposium, May 2023, doi: 10.4071/001c.74534.

[3] J. C. Rubio and M. Bolduc, "Screen Printing for Energy Storage and Functional Electronics: A Review", Electronic Materials, vol. 6, no. 2, May 2025, doi: 10.3390/electronicmat6020007.

[4] J. Pan, G. L. Tonkay, and A. Quintero, "Screen Printing Process Design of Experiments for Fine Line Printing of Thick Film Ceramic Substrates", Journal of Electronics Manufacturing, vol. 09, no. 03, 1999, doi: 10.1142/s096031319900012x.

[5] "Heraeus Electronics | Thick film technology materials", Heraeus Electronics. https://www.heraeus-electronics.com/en/products-and-solutions/thick-film-materials/thick-film-materials-overview/ (accessed Mar. 16, 2026).

[6] J. Minalgiene and S. Muckett, "Very Fine Line Photoimageable Thick Film Technology Developed at Hibridas, Lithuania", in MCM C/Mixed Technologies and Thick Film Sensors, 1995. Available: https://link.springer.com/chapter/10.1007/978-94-011-0079-3_13

[7] D. Stephens, P. R. Young, and I. D. Robertson, "Millimeter-wave substrate integrated waveguides and filters in photoimageable thick-film technology", IEEE Transactions on Microwave Theory and Techniques, vol. 53, no. 12, 2005, doi: 10.1109/tmtt.2005.859862.

[8] A. Al Rashid, W. Ahmed, M. Y. Khalid, and M. Koç, "Vat photopolymerization of polymers and polymer composites: Processes and applications", Additive Manufacturing, vol. 47, 2021, doi: 10.1016/j.addma.2021.102279.

[9] T. Siripongpreda, V. P. Hoven, B. Narupai, and N. Rodthongkum, "Emerging 3D printing based on polymers and nanomaterial additives: Enhancement of properties and potential applications", European Polymer Journal, vol. 184, 2023, doi: 10.1016/j.eurpolymj.2022.111806.

[10] "Conductive 3D resins", 3Dresyns. https://www.3dresyns.com/collections/3d-resins-for-electronics (accessed Mar. 16, 2026).

[11] S. -C. Tsai, L. -H. Chen, C. -P. Chu, W. -C. Chao, and Y. -C. Liao, "Photo curable resin for 3D printed conductive structures", Additive Manufacturing, vol. 51, 2022, doi: 10.1016/j.addma.2021.102590.

[12] Q. Mu et al., "Digital light processing 3D printing of conductive complex structures", Additive Manufacturing, vol. 18, 2017, doi: 10.1016/j.addma.2017.08.011.

[13] "DragonFly IV | A new era in 3D electronics printing", Nano Dimensions. https://www.nano-di.com/dragonfly-iv (accessed Mar. 15, 2026).

[14] "AME Materials | Innovation in 3D printing of electronics." https://www.nano-di.com/ame-materials (accessed Mar. 15, 2026).

[15] P. F. Jacobs, "Fundamentals of Stereolithography", Accessed: Mar. 16, 2025. [Online]. Available: http://hdl.handle.net/2152/64400

[16] "Biocompatible 3D printing resin - Liqcreate Bio-Med Clear", Liqcreate. https://www.liqcreate.com/product/bio-med-clear-biocompatible-resin/ (accessed Mar. 15, 2026)

Development of a Reduced Order Thermal Model for Real-Time Lithium-ion Battery Storage Simulation

Marek Sedlařík
Department of Electrical and Electronic Technology
Faculty of Electrical Engineering and Communication, Brno University of Technology
Brno, Czech Republic
Marek.Sedlarik@vut.cz

Petr Vyroubal
Department of Electrical and Electronic Technology
Faculty of Electrical Engineering and Communication, Brno University of Technology
Brno, Czech Republic
vyroubal@vut.cz

Martin Mačák
Electric Vehicle Technologies
Transport Technologies
AIT Austrian Institute of Technology GmbH
Vienna, Austria
Martin.macak@ait.ac.at

***Abstract*— This study presents the formulation and preliminary assessment of a reduced-order thermal modelling approach for lithium-ion battery cells. Accurate prediction of battery temperature is critical for reliable operation and thermal management of battery systems, while detailed physics-based models often impose significant computational demands. The proposed approach derives a computationally efficient thermal representation from the response of a detailed thermal model of a lithium-ion pouch cell. Material properties of the cell were obtained from the geometrical structure of the electrode stack and subsequently homogenized to determine effective thermophysical parameters suitable for thermal modelling. Based on these parameters, a thermal model of the cell was established and excited by heat generation within the active material. The resulting temperature response was compared with experimentally measured temperatures obtained at several locations on the cell surface. Preliminary simulations performed in ANSYS Fluent indicate a reasonable agreement between the simulated and measured temperature evolution. The presented methodology establishes a framework for the development of computationally efficient thermal representations of lithium-ion batteries suitable for system-level simulations.**

Keywords — Li-ion battery, Reduced Order Model, Thermal Model, Electrochemical Model

I. Introduction

Lithium-ion batteries (LIBs) are widely used electrochemical energy storage devices due to their high specific energy, high power capability, good charge-discharge efficiency, and long cycle life [1]. These properties have led to their extensive deployment in a wide range of applications, including portable electronics, electric vehicles (EVs), hybrid electric vehicles (HEVs), and stationary energy storage systems [2-6]. In stationary storage networks, LIBs support grid stabilization and renewable energy integration, where increasing demands for higher energy density and charge-discharge rates lead to more demanding operating conditions [7,8]. Under demanding conditions, internal electrochemical and thermal processes strongly influence battery performance, degradation, and safety. For this reason, accurate modelling of LIBs is important for the analysis, design, and control of battery systems operating under demanding conditions [9, 10].

Despite their favourable electrochemical properties, the performance, degradation behaviour, and operational safety of LIBs are strongly influenced by their internal thermal state [10]. The recommended operating temperature range for most lithium-ion chemistries is relatively narrow, typically between 15 °C and 35 °C [11,12]. Operation outside this temperature window can trigger several degradation mechanisms that negatively affect battery performance and lifetime [13,14]. At low temperatures, the electrolyte viscosity increases and ionic conductivity decreases, which leads to higher charge-transfer resistance and increased cell polarization, significantly limiting the available power capability [12, 15]. In addition, charging at low temperatures may result in lithium plating on the graphite anode surface. This process irreversibly consumes active lithium and may lead to dendrite formation capable of penetrating the separator and causing internal short circuits [2, 12, 25]. At elevated temperatures, parasitic side reactions are accelerated, resulting in faster degradation, continuous modification of the solid electrolyte interphase (SEI) layer, and dissolution of active materials. In extreme cases of thermal abuse or uncontrolled heat accumulation, exothermic reactions may initiate thermal runaway, potentially leading to venting, ignition, or explosion [3, 2, 12].

Beyond the mean cell temperature, the spatial distribution of temperature within the battery, both within individual cells and between cells in a module or pack, represents an important challenge in large-scale battery systems used in EVs and stationary energy storage [16, 17, 18]. High-power discharge and fast-charging conditions generate significant internal heat, primarily due to irreversible ohmic (Joule) heating, polarization losses, and reversible entropic heat generation during electrochemical reactions [16, 17, 19]. Due to the anisotropic thermal conductivity of the tightly packed jelly-roll or stacked electrode structures, considerable temperature gradients can develop between the core and the surface of the cell [20]. Experimental studies have shown that surface-to-core temperature differences can exceed 20 °C under dynamic load conditions typical for automotive applications [21]. Non-uniform temperature distributions lead to uneven electrochemical behaviour, where regions operating at higher temperatures experience accelerated ageing, reduced

impedance, intensified heat generation, and reduced cycle life [16, 22, 23]. For this reason, an effective Battery Thermal Management System (BTMS) is required to control heat dissipation, maintain temperature uniformity, and ensure safe and reliable operation of battery systems [24].

The design of effective BTMS and the development of safe operating strategies require a detailed understanding and accurate prediction of the transient thermal behaviour of LIBs [24]. Direct measurement of internal temperature distribution within commercial sealed cells is technically challenging and impractical for large-scale applications. For this reason, numerical thermal modeling has become an essential tool for battery analysis and development [25]. High-fidelity modeling approaches, such as 3-dimensional Computational Fluid Dynamics (CFD) and finite element models (FEM) coupled with electrochemical pseudo-two-dimensional (P2D) frameworks, enable detailed investigation of the coupled electrochemical and thermal processes occurring within the battery [16, 26]. These high-resolution models are commonly used during the offline design phase to analyse temperature distributions, evaluate cooling concepts, and investigate battery behaviour under demanding conditions [27, 28, 29].

The high spatial and temporal resolution of full-order multi-dimensional models comes at the cost of significant computational complexity, limiting their real-time implementation in onboard Battery Management Systems (BMS) that must continuously estimate internal temperatures for predictive thermal control and safe operation [27, 30, 31]. However, the computational capabilities of typical automotive microcontrollers are insufficient to run high-fidelity multi-physics models, such as detailed finite element or three-dimensional CFD simulations, at the sampling rates required for real-time control [17]. Consequently, the direct implementation of full-order physics-based models in embedded BMS is generally infeasible.

Reduced Order Models (ROMs) are commonly employed to address the trade-off between modelling accuracy and computational efficiency in battery thermal simulations. High-fidelity models provide detailed spatial resolution but require significant computational resources, limiting real-time implementation in BMS [11]. In contrast, ROMs, particularly lumped equivalent circuits, use simplified thermal networks with substantially lower computational demands, making them suitable for embedded applications with limited processing capabilities [35]. In addition to physics-based and equivalent circuit approaches, data-driven thermal models have also been increasingly investigated. These models rely on experimental datasets to learn the relationship between operating conditions and battery temperature without explicitly describing the underlying physical processes. Techniques such as statistical regression or machine learning methods can capture nonlinear thermal behaviour and estimate temperature dynamics from measured electrical variables. Although such approaches may achieve high prediction accuracy, their applicability strongly depends on the availability of representative training data and they generally provide lower physical interpretability compared to physics-based models [33].

This study focuses on the development of a reduced order thermal model for real-time simulation of LIB systems based on a Linear Time-Invariant (LTI) formulation. The objective is to bridge the gap between the high accuracy of physics based thermal models and the limited computational resources available in onboard BMS. In the proposed approach, the battery thermal behaviour is represented as an LTI system in which heat generation acts as the input and the resulting temperature evolution represents the system response. The thermal model is coupled with an Equivalent Circuit Model (ECM) describing the electrical behaviour of the battery, enabling consistent electro-thermal interaction between current, heat generation, and temperature dynamics. Due to its low computational demand, the model is suitable for real-time implementation in BMS, where it can support temperature estimation, thermal control strategies, and safety monitoring in LIB systems used stationary energy storage applications.

II. Experiment

The experiment was conducted on LIBs employing lithium–nickel–manganese–cobalt oxide (NMC) as the cathode material and graphite as the anode material. nce the proposed ROM thermal model is derived from the response of a high-fidelity thermal model, it is necessary to first define the material properties of the cell. Based on these properties, a full thermal model of the cell is established and excited by a defined heat generation within the active material. The resulting thermal response is then used to derive the ROM, where the spatial temperature distribution is represented by the average cell temperature. The electrical ECM is subsequently coupled with the ROM to provide the heat generation input based on the electrical operating conditions of the cell.

A. ECM definition

A second-order (2RC) ECM was used to represent the dynamic electrical response of the investigated lithium-ion cell. The model consists of a series resistance and two RC branches that capture the transient voltage behaviour during current excitation. The model parameters were identified from pulse-based measurements designed to excite the dynamic response of the cell under controlled laboratory conditions

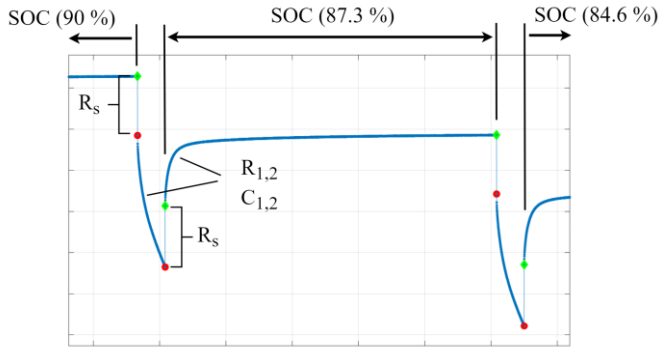

Fig. 1: Characteristic voltage response to a current pulse during GITT measurement used for ECM parameter identification.

For parameterization, a protocol derived from the galvanostatic intermittent titration technique (GITT) was employed. The procedure began with a constant-current (CC) charge at 0.2 C until the cell voltage reached 4.2 V, followed by a constant-voltage (CV) step at 4.2 V until the current decreased to 0.05 C. After a 1 h rest period, the discharge phase was initiated using CC pulses at 0.3 C with a duration of 5 min, each followed by a 1 h relaxation period to allow the voltage to approach equilibrium conditions. The pulse sequence was repeated until the lower cut-off voltage of 2.5 V was reached. The principle of parameter extraction from an individual pulse is illustrated in Fig. 1, where the instantaneous voltage drop is used to determine the series resistance, while the subsequent voltage relaxation is used to identify the RC parameters.

The ECM parameters were identified from the voltage response to current pulses using numerical fitting implemented in MATLAB. The instantaneous voltage drop at the beginning of each pulse was used to determine the series resistance R_0. The subsequent voltage relaxation during the rest period was approximated using a sum of exponential functions corresponding to the RC branches. The time constants were defined as $\tau_i = R_i C_i$. Parameter identification was performed for individual pulses under controlled temperature conditions.

B. Battery material properties

The material properties of the investigated pouch cell were derived from the geometrical structure of the electrode stack and available material data for the individual cell components. The thickness of the positive electrode, separator, and negative electrode layers was measured using a micrometer. Measurements were performed both on complete electrodes and on electrodes after removing the active material layer, which allowed the thickness of the current collectors and the active layers to be determined separately. Based on these measurements, the layer structure of the cell stack was reconstructed and used to estimate the effective material properties of the cell.

TABLE I: MATERIAL PROPERTIES [38-40]

Zone	Pc	Pe	S	Ne	Nc	Results	Model
Thickness [µm]	10	60	20	100	20	-	-
Density [kg/m³]	2700	4530	492	2660	8900	3380	2000
Heat capacity [J/kg/K]	903	800	1978	1437	385	1205	1000
Thermal conductivity [W/m/K]	238	4.3	0.334	1.04	398	28/1.0*	25/1.0*

** In-plane/cross-plane (results cross-plane value from [38])*

The thermophysical properties of the layered cell structure were homogenized to obtain effective parameters suitable for use in the thermal model. The effective density, heat capacity, and thermal conductivity were calculated using thickness-weighted averaging of the individual layers, including the current collectors, electrode materials, and separator. In this approach, the contribution of each component is weighted according to its thickness fraction within the total stack thickness. The formulation used for the homogenization is given in Eq. (1). [37]

$$x_{eff} = \frac{0.5x_c^p\delta_c^p + x_e^p\delta_e^p + x_s\delta_s + x_e^n\delta_e^n + 0.5x_c^n\delta_c^n}{\delta_{total}} \quad (1)$$

$$\delta_{total} = 0.5\delta_c^p + \delta_e^p + \delta_s + \delta_e^n + 0.5\delta_c^n \quad (2)$$

In Eq. (1), the subscripts *c*, *e*, and *s* denote the current collector, electrode, and separator layers, respectively, while the superscripts *p* and *n* refer to the positive and negative sides of the cell. The total stack thickness δ_{total} represents the combined thickness of all layers within one repeating unit of the electrode stack. Since the current collectors are shared between adjacent electrodes, their contribution is considered with a factor of one half.

Due to the laminated structure of pouch cells, heat transfer is strongly anisotropic. Therefore, an orthotropic thermal conductivity tensor was adopted in the model. The in-plane directions (x, y) correspond to heat conduction along the electrode layers, while the through-plane direction (z) represents heat transfer across the layered structure. Based on the homogenization procedure, the effective in-plane thermal conductivity was approximately 25 $W \cdot m^{-1} \cdot K^{-1}$, while the through-plane conductivity was significantly lower, approximately 1 $W \cdot m^{-1} \cdot K^{-1}$. The resulting effective properties used in the model are summarized in Table I.

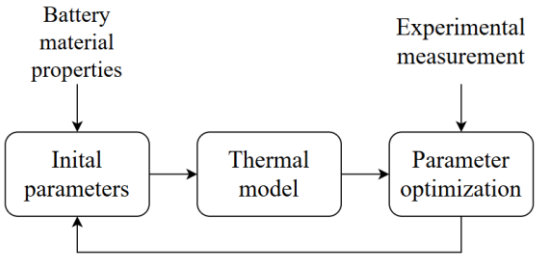

Fig. 2: Model parameter optimization.

The obtained material properties served as initial parameters for the thermal model. To improve the agreement between the numerical model and experimental measurements, selected parameters were further refined using an optimization procedure. As illustrated in Fig. 2, the optimization iteratively compares the simulated temperature response with experimental measurements and updates the model parameters until satisfactory agreement is achieved. This approach allows the model to account for uncertainties associated with material heterogeneity and simplifications introduced during the homogenization process.

III. Results

Figure 3 presents preliminary results obtained from the material parameter optimization performed in Ansys Fluent. The plot shows the simulated temperature response of the thermal model subjected to heat generation within the active material and its comparison with experimentally measured temperatures at three locations of the cell (middle, positive tab side, and negative tab side). The applied power corresponds to the heat generation in the active material during the experimental test. The heat generation rate used in the simulation was derived from the electrical measurements obtained during the experimental test and applied as a volumetric heat source in the active material region. The results demonstrate that the optimized material parameters provide a good agreement between the simulated and measured temperature evolution, particularly for the temperature measured in the center of the cell.

The observed agreement indicates that the selected thermophysical properties and boundary conditions adequately represent the dominant heat transfer mechanisms within the cell. Slight deviations between simulated and measured temperatures at the tab locations can be attributed to localized heat generation and contact resistance effects, which are difficult to fully capture using homogenized material properties. These results confirm that the applied parameter optimization procedure improves the predictive capability of the thermal model.

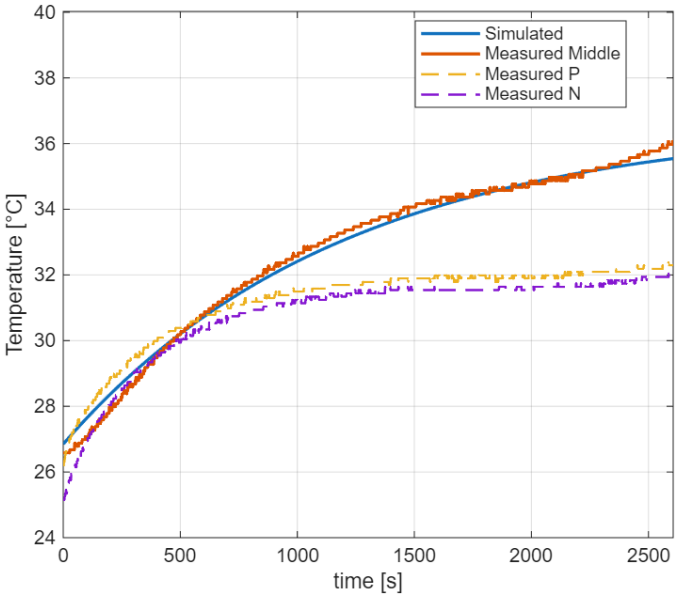

Fig. 3: Model parameter optimization.

IV. Conclusion

This study presented the definition of a thermal ROM for LIB cells based on the response of a detailed thermal representation. Material properties of the pouch cell were derived from the electrode stack structure and homogenized to obtain effective thermophysical parameters suitable for thermal modelling. Preliminary simulations performed in ANSYS Fluent were compared with experimentally measured temperature data at several locations of the cell. The results indicate a reasonable agreement between the simulated and measured temperature evolution, particularly in the central region of the cell. These findings demonstrate the feasibility of the proposed modelling approach and provide a basis for further development and validation of computationally efficient thermal models for battery system simulations. The obtained results demonstrate that the proposed reduced-order modeling approach provides sufficiently accurate temperature prediction while maintaining low computational complexity, supporting its applicability for real-time thermal management and digital twin applications.

Acknowledgment

This work was supported by the BUT specific research program (project No. FEKT-S-23-8286).

The authors would like to acknowledge Antonín Šimek, Martin Šedina and David Trochta for their support during the experimental measurements and their valuable collaboration.

References

[1] C. X. He *et al.*, "A transient multi-path decentralized resistance-capacity network model for prismatic lithium-ion batteries based on genetic algorithm optimization", *Energy Conversion and Management*, vol. 300, 2024, doi: 10.1016/j.enconman.2023.117894.

[2] S. S. Madani, C. Ziebert, and M. Marzband, "Thermal Behavior Modeling of Lithium-Ion Batteries: A Comprehensive Review", *Symmetry*, vol. 15, no. 8, Aug. 2023, doi: 10.3390/sym15081597.

[3] X. Feng, M. Ouyang, X. Liu, L. Lu, Y. Xia, and X. He, "Thermal runaway mechanism of lithium ion battery for electric vehicles: A review", *Energy Storage Materials*, vol. 10, 2018, doi: 10.1016/j.ensm.2017.05.013.

[4] K. Li *et al.*, "Assessment method of the integrated thermal management system for electric vehicles with related experimental validation", *Energy Conversion and Management*, vol. 276, 2023, doi: 10.1016/j.enconman.2022.116571.

[5] Q. L. Yue, C. X. He, M. C. Wu, J. B. Xu, and T. S. Zhao, "Pack-level modeling of a liquid cooling system for power batteries in electric vehicles", *International Journal of Heat and Mass Transfer*, vol. 192, 2022, doi: 10.1016/j.ijheatmasstransfer.2022.122946.

[6] J. Kim, J. Oh, and H. Lee, "Review on battery thermal management system for electric vehicles", *Applied Thermal Engineering*, vol. 149, 2019, doi: 10.1016/j.applthermaleng.2018.12.020.

[7] A. A. Kebede, T. Kalogiannis, J. Van Mierlo, and M. Berecibar, "A comprehensive review of stationary energy storage devices for large scale renewable energy sources grid integration", *Renewable and Sustainable Energy Reviews*, vol. 159, 2022, doi: 10.1016/j.rser.2022.112213.

[8] M. M. Hasan *et al.*, "Advancing energy storage: The future trajectory of lithium-ion battery technologies", *Journal of Energy Storage*, vol. 120, 2025, doi: 10.1016/j.est.2025.116511.

[9] A. Okan, N. Sözbir, and M. Bulut, "Reduced Thermal Mathematical Model Generation of a Li-ion Battery Block", *Energies*, vol. 18, no. 24, Dec. 2025, doi: 10.3390/en18246374.

[10] S. Ma *et al.*, "Temperature effect and thermal impact in lithium-ion batteries: A review", *Progress in Natural Science: Materials International*, vol. 28, no. 6, 2018, doi: 10.1016/j.pnsc.2018.11.002.

[11] B. E. Lebrouhi, B. Lamrani, M. Ouassaid, M. Abd-Lefdil, M. Maaroufi, and T. Kousksou, "Low-cost numerical lumped modelling of lithium-ion battery pack with phase change material and liquid cooling thermal management system", *Journal of Energy Storage*, vol. 54, 2022, doi: 10.1016/j.est.2022.105293.

[12] M. Alipour, C. Ziebert, F. V. Conte, and R. Kizilel, "A Review on Temperature-Dependent Electrochemical Properties, Aging, and Performance of Lithium-Ion Cells", *Batteries*, vol. 6, no. 3, Jun. 2020, doi: 10.3390/batteries6030035.

[13] H. Liu, Z. Wei, W. He, and J. Zhao, "Thermal issues about Li-ion batteries and recent progress in battery thermal management systems: A

review", *Energy Conversion and Management*, vol. 150, 2017, doi: 10.1016/j.enconman.2017.08.016.

[14] J. Cen, Z. Li, and F. Jiang, "Experimental investigation on using the electric vehicle air conditioning system for lithium-ion battery thermal management", *Energy for Sustainable Development*, vol. 45, 2018, doi: 10.1016/j.esd.2018.05.005.

[15] M. Petzl, M. Kasper, and M. A. Danzer, "Lithium plating in a commercial lithium-ion battery – A low-temperature aging study", *Journal of Power Sources*, vol. 275, 2015, doi: 10.1016/j.jpowsour.2014.11.065.

[16] E. Hosseinzadeh, R. Genieser, D. Worwood, A. Barai, J. Marco, and P. Jennings, "A systematic approach for electrochemical-thermal modelling of a large format lithium-ion battery for electric vehicle application", *Journal of Power Sources*, vol. 382, 2018, doi: 10.1016/j.jpowsour.2018.02.027.

[17] L. Mattia, H. Beiranvand, W. Zamboni, and M. Liserre, "Lithium-ion battery thermal modelling and characterisation: A comprehensive review", *Journal of Energy Storage*, vol. 129, 2025, doi: 10.1016/j.est.2025.117114.

[18] M. R. Khan and S. K. Kaer, "Multiphysics based thermal modeling of a pouch lithium-ion battery cell for the development of pack level thermal management system", *2016 Eleventh International Conference on Ecological Vehicles and Renewable Energies (EVER)*, 2016, doi: 10.1109/ever.2016.7476442.

[19] S. Du *et al.*, "An investigation of irreversible heat generation in lithium ion batteries based on a thermo-electrochemical coupling method", *Applied Thermal Engineering*, vol. 121, 2017, doi: 10.1016/j.applthermaleng.2017.04.077.

[20] O. S. Burheim, M. A. Onsrud, J. G. Pharoah, F. Vullum-Bruer, and P. J. S. Vie, "Thermal Conductivity, Heat Sources and Temperature Profiles of Li-Ion Batteries", *ECS Transactions*, vol. 58, no. 48, Apr. 2014, doi: 10.1149/05848.0145ecst.

[21] U. S. Kim, C. B. Shin, and C. -S. Kim, "Modeling for the scale-up of a lithium-ion polymer battery", *Journal of Power Sources*, vol. 189, no. 1, 2009, doi: 10.1016/j.jpowsour.2008.10.019.

[22] D. Werner, A. Loges, D. J. Becker, and T. Wetzel, "Thermal conductivity of Li-ion batteries and their electrode configurations – A novel combination of modelling and experimental approach", *Journal of Power Sources*, vol. 364, 2017, doi: 10.1016/j.jpowsour.2017.07.105.

[23] Z. An, L. Jia, L. Wei, and C. Yang, "Numerical modeling and analysis of thermal behavior and Li transport characteristic in lithium-ion battery", *International Journal of Heat and Mass Transfer*, vol. 127, 2018, doi: 10.1016/j.ijheatmasstransfer.2018.07.147.

[24] W. Zichen and D. Changqing, "A comprehensive review on thermal management systems for power lithium-ion batteries", *Renewable and Sustainable Energy Reviews*, vol. 139, 2021, doi: 10.1016/j.rser.2020.110685.

[25] M. R. Khan, M. J. Swierczynski, and S. K. Kaer, "Determination of the behavior and performance of commercial Li-Ion pouch cells by means of isothermal calorimeter", *2016 Eleventh International Conference on Ecological Vehicles and Renewable Energies (EVER)*, 2016, doi: 10.1109/ever.2016.7476441.

[26] M. Doyle, T. F. Fuller, and J. Newman, "Modeling of Galvanostatic Charge and Discharge of the Lithium/Polymer/Insertion Cell", *Journal of The Electrochemical*

[27] A. De Vita, A. Maheshwari, M. Destro, M. Santarelli, and M. Carello, "Transient thermal analysis of a lithium-ion battery pack comparing different cooling solutions for automotive applications", *Applied Energy*, vol. 206, 2017, doi: 10.1016/j.apenergy.2017.08.184.

[28] Y. -M. Dai, Y. -C. Ting, C. M. Chang, and C. -T. Huang, "Thermal management analysis of fast-charging lithium-ion battery packs: Effects of cooling strategies", *Next Energy*, vol. 9, 2025, doi: 10.1016/j.nxener.2025.100465.

[29] H. Ahmad, R. Prabakaran, and S. C. Kim, "Electrochemical–thermal coupled modeling of lithium-ion batteries: a comprehensive review of mechanisms, methods, and applications", *Journal of Thermal Analysis and Calorimetry*, Mar. 2026, doi: 10.1007/s10973-026-15302-z.

[30] M. J. Hunt, F. Brosa Planella, F. Theil, and W. D. Widanage, "Derivation of an effective thermal electrochemical model for porous electrode batteries using asymptotic homogenisation", *Journal of Engineering Mathematics*, vol. 122, no. 1, Apr. 2020, doi: 10.1007/s10665-020-10045-8.

[31] A. Schmidt, D. Oehler, A. Weber, T. Wetzel, and E. Ivers-Tiffée, "A multi scale multi domain model for large format lithium-ion batteries", *Electrochimica Acta*, vol. 393, 2021, doi: 10.1016/j.electacta.2021.139046.

[32] B. E. Lebrouhi, B. Lamrani, M. Ouassaid, M. Abd-Lefdil, M. Maaroufi, and T. Kousksou, "Low-cost numerical lumped modelling of lithium-ion battery pack with phase change material and liquid cooling thermal management system", *Journal of Energy Storage*, vol. 54, 2022, doi: 10.1016/j.est.2022.105293.

[33] Y. Wang *et al.*, "Temperature prediction of lithium-ion battery based on artificial neural network model", *Applied Thermal Engineering*, vol. 228, 2023, doi: 10.1016/j.applthermaleng.2023.120482

[34] A. Okan, N. Sözbir, and M. Bulut, "Reduced Thermal Mathematical Model Generation of a Li-ion Battery Block", *Energies*, vol. 18, no. 24, Dec. 2025, doi: 10.3390/en18246374.

[35] K. Saqli, H. Bouchareb, N. K. M'sirdi, and M. Oudghiri Bentaie, "Lithium-ion battery electro-thermal modelling and internal states co-estimation for electric vehicles", *Journal of Energy Storage*, vol. 63, 2023, doi: 10.1016/j.est.2023.107072.

[36] C. Chen, F. Diwoky, Z. Pavlovic, and J. Wurzenberger, "An Application of the Linear and Time-Invariant Method for the System-Level Thermal Simulation of an EV Battery", *SAE Technical Paper Series*, vol. 1, Apr. 2015, doi: 10.4271/2015-01-1197.

[37] P. Vyroubal and T. Kazda, "Equivalent circuit model parameters extraction for lithium ion batteries using electrochemical impedance spectroscopy", *Journal of Energy Storage*, vol. 15, 2018, doi: 10.1016/j.est.2017.10.019.

[38] P. Peng and F. Jiang, "Thermal safety of lithium-ion batteries with various cathode materials: A numerical study", International Journal of Heat and Mass Transfer, vol. 103, 2016, doi: 10.1016/j.ijheatmasstransfer.2016.07.088.

[39] S. Landini, J. Panter, A. Roy, G. D. Valasai, and M. F. Ismail, "Numerical Investigation of Hybrid Immersion Cooling Strategies for Battery Packs in Light Electric Vehicles", *Journal of Fluid Flow, Heat and Mass Transfer*, vol. 12, 2025, doi: 10.11159/jffhmt.2025.024.

[40] E. J. Cheng, K. Hong, N. J. Taylor, H. Choe, J. Wolfenstine, and J. Sakamoto, "Mechanical and physical properties of LiNi0.33Mn0.33Co0.33O2 (NMC)", *Journal of the European Ceramic Society*, vol. 37, no. 9, 2017, doi: 10.1016/j.jeurceramsoc.2017.03.048.

Semi-Automated Parameter Identification and Rapid Control Prototyping of a DC Motor Drive Using an STM32 Microcontroller

Miroslav Kozumplík
Department of Control and Instrumentation
FEEC, Brno University of Technology
Brno, Czech Republic
Miroslav.Kozumplik@vut.cz

***Abstract*—This paper presents a cost-effective laboratory platform designed for semi-automated DC motor parameter identification and rapid control prototyping (RCP). The hardware setup comprises an STM32 Nucleo G474RE microcontroller, an STM32 Nucleo IHM07M1 expansion power board, a DC power supply, and a tested DC motor equipped with an incremental encoder. The paper describes a semi-automated parameter identification process, defined here as the combination of automatically generated microcontroller software with a concise procedural guide for acquiring motor parameters and subsequently tuning the current and speed controllers. Developed within MATLAB/Simulink, the control software leverages *Embedded Coder* for automatic C-code generation, supporting real-time tuning and data logging. Due to its affordability and modularity, the STM32 platform is an ideal educational tool for modelling and identification courses. Experimental results confirm the device's functionality.**

***Index Terms*—Rapid Control Prototyping, Parameter Identification, Motor Drive, MATLAB/Simulink, Embedded Coder, Automatic Code Generation, STM32 Microcontroller**

I. Introduction

This work aims to design and realize a functional and cost-effective experimental platform for electric motor parameter identification and subsequent rapid control prototyping. The proposed solution enables the semi-automated determination of both electrical and mechanical motor parameters, utilized for designing and experimental verifying control structures. The platform is conceived as a universal development tool suitable for academic and educational purposes.

Accurate parameter knowledge is essential for effective drive control. Unlike ideal simulation models, real-world systems are affected by manufacturing tolerances, thermal dependencies, and magnetic nonlinearities. Factors such as aging, load variations, or faults further alter motor parameters. Consequently, parameter identification serves not only the primary role of refining inaccurate or unknown parameters but also facilitates diagnostics, as deviations can indicate incipient faults or motor degradation. [1]

Implementing advanced real-time identification and control algorithms imposes specific requirements on the control hardware. While traditional FPGA-based RCP systems like dSPACE or Speedgoat provide high performance, their cost limits accessibility in lower-budget applications. This study demonstrates the viability of using affordable modern microcontrollers paired with MATLAB/Simulink as an RCP environment. [2], [3]

A DC motor was selected to verify the semi-automated identification methods. Although permanent magnet synchronous motors (PMSM) dominate modern application, the DC motor retains relevance in low-cost servos, compact battery-operated applications, automotive actuators, and specific industrial or railway traction applications. The presented platform aims to demonstrate the identification process and the design of closed-loop regulators. A DC motor was chosen mainly because of its simple model and ease of control, but the platform's capabilities can be readily extended to PMSMs as well. [3], [4]

II. Rapid Control Prototyping

Rapid Control Prototyping represents a modern approach to developing electric motor control algorithms. RCP enables developers to test and validate control structures in real-time without the constraints of final embedded systems or the necessity of manual low-level coding. In practice, RCP radically accelerates the development of control algorithms and reduces costs. The iterative RCP development cycle can be divided into five consecutive steps executed in a loop: Design and Simulation, Code Generation, Target Deployment, Performance Evaluation, and finally, Tuning and Optimization. [5]

Within RCP, the control algorithm can be deployed on various types of platforms that differ in computational performance, availability, simplicity, and cost:

- **Dedicated real-time systems:** The most high-end platforms are dedicated real-time systems such as dSPACE or Speedgoat. These systems are primarily used in industrial research and development, especially in the automotive industry. They offer immense computational power but come with high costs. [6]
- **Field-Programmable Gate Arrays (FPGA)**: FPGA enable parallel algorithm execution with deterministic nanosecond latency. However, development is costly due to synthesis times and required HDL expertise. However,

some modern approaches already support automatic HDL code generation. [5]

- **Microcontroller-based platforms:** The most accessible RCP platforms are modern 32-bit microcontrollers combined with automatic code generation tools. Although their computational performance is limited compared to FPGAs, it remains sufficient for control algorithms running at 20 kHz sampling rate. These MCUs often include fixed-point and floating-point arithmetic units, bringing them close to the capabilities of specialized DSP processors. [7], [8]

A. STM32 NUCLEO G474RE MCU

The experimental setup utilizes an STM32 NUCLEO G474RE microcontroller paired with a NUCLEO-IHM07M1 power expansion board. The STM32 G474RE features a 170 MHz ARM Cortex-M4F core, optimized for digital power applications. The MCU contains accelerators for motor control, such as CORDIC for trigonometric calculations, the FMAC unit for signal filtering and five 12-bit ADCs. The expansion board is facilitated via ST Morpho headers. [9]

The NUCLEO-IHM07M1 power stage is based on the L6230 three-phase driver, operating at 8–48 V with a 2.8 A peak current limit. It includes configurable 1-shunt or 3-shunt topology for phase current sensing. While the L6230 is a 3-phase driver, it is configured here as a full H-bridge for DC motor control. For higher current requirements, the NUCLEO-IHM08M1 board can be used as an alternative, supporting peak currents up to 15 A. [10]

B. Automatic Code Generation from MATLAB/Simulink

Communication with the STM32 NUCLEO-G474RE is supported by official software tools, specifically STM32CubeMX and STM32CubeProgrammer. The board connects to a PC via USB, enumerating as a virtual serial COM port (ST-Link VCP). Automatic code generation tools facilitate effective RCP deployment. The MATLAB/Simulink environment provides the *Embedded Coder* and *Simulink Coder* add-ons, which automatically translate the Simulink model into C code and compile it for the target MCU. The official *Embedded Coder Support Package for STM32 Processors* ensures seamless integration of Simulink with STM32CubeMX tools and Arm Cortex-M compilers. Real-time interaction is achieved via Simulink's External Mode with Monitor and Tune feature, allowing parameter adjustment and data logging during operation. A limitation of this approach is the bandwidth of the serial link used by the XCP protocol, which may cause data loss at high frequencies. Solutions include using a higher-bandwidth interface or data decimation.

III. Mathematical Model and Identification of DC Motor Parameters

A. Mathematical Model of the DC Motor

The permanent-magnet DC motor can be mathematically modelled by separate equations for its electrical and mechanical subsystems. Applying Kirchhoff's voltage law to the armature circuit yields the voltage equation:

$$u(t) = R \cdot i(t) + L \cdot \frac{di(t)}{dt} + K_e \cdot \omega(t) \tag{1}$$

where $u(t)$ and $i(t)$ denote terminal voltage and current, and K_e is the back-EMF constant, which, together with $\omega(t)$, represents the induced back-EMF. Mechanical dynamics follow Newton's law:

$$J\frac{d\omega(t)}{dt} = K_t i(t) - b\omega(t) - T_L \tag{2}$$

where J is the system inertia, b is the viscous friction coefficient, T_L is the load torque, and K_t is the torque constant. For an ideal motor, the energetic equivalence $K_t = K_e$ holds in SI units.

The right side of Fig. 1 displays the block diagram of the DC motor model, explicitly showing the coupling between the electrical part, represented by the transfer function $1/(L \cdot s + R)$, and the mechanical part with the transfer function $1/(J \cdot s + b)$. The control architecture on the left side of Fig. 1 employs a cascade topology: an inner PI current loop regulates torque, while an outer PI speed loop generates the current reference. This hierarchy decouples fast electrical dynamics from slower mechanical ones, simplifying controller design. [11], [12]

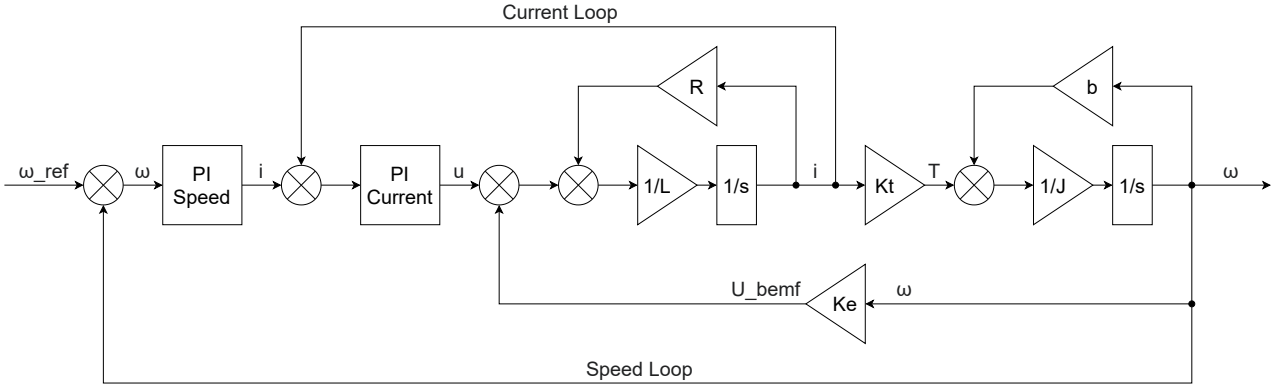

Fig. 1. Block diagram of DC motor model and cascade control

TABLE I
PARAMETERS OF THE MATADOR PMDC MOTOR

Rated speed (RPM)	Rated current (A)	Stall torque (Nm)	No. of poles (-)
3000	4.45	0.32	2

Voltage const. (V/kRPM)	Torque const. (Nm/A)	Rotor resist. (Ω, 25°C)	Rotor induct. (mH)
7.3	0.07	0.82	1.34

B. Identification of DC Motor Parameters

Precise parameter identification is essential for analytical tuning and diagnostics.

Electrical parameters (R, L) are determined via the Locked-Rotor test. Since zero speed eliminates back-EMF, the model simplifies to a first-order RL circuit, allowing identification from the current step response. Mechanical parameters (J, b) are identified experimentally using a PRBS as an input in the current loop. Parameters are computed offline from the recorded speed response using Ordinary Least Squares method.

IV. EXPERIMENTAL PLATFORM IMPLEMENTATION

The experimental platform integrates two permanent magnet direct current (PMDC) motors, an incremental encoder, a power supply, and an STM32 NUCLEO-G474RE MCU paired with the NUCLEO-IHM07M1 expansion power board.

A. Hardware Setup

The setup utilizes a PMDC Matador DCM2B motor. The detailed parameters of this motor are listed in Tab. I. The test bench features a dual-motor configuration (indicated by red marker 1 in Fig. 3); where one motor operates as the actuator, while the other acts as a load.

Power is supplied by a 12 V/7 A BKE source (marker 2 in Fig. 3). Control hardware consists of the STM32 Nucleo G474RE MCU (marker 3 in Fig. 3) and the IHM07M1 power board (marker 4 in Fig. 3). Feedback is provided by a 1024-pulse incremental TTL encoder (marker 5 in Fig. 3). A mechanical fixture (marker 6 in Fig. 3) allows rotor locking for electrical parameter identification.

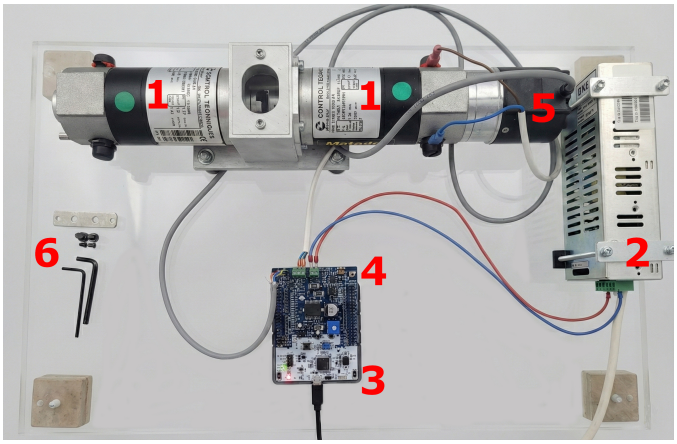

Fig. 3. The experimental laboratory test bench

B. Software Architecture

MCU peripherals including ADCs, GPIO pins and communication baud rates were configured using STM32CubeMX v6.4.0. Control algorithms for identification and RCP were developed in MATLAB/Simulink R2024b using the *Embedded Coder Support Package for STM32* add-on. Real-time communication is facilitated by STM32CubeProgrammer v2.6.0 via Simulink External Mode. Simulink Data Inspector tool handles real-time signal logging and monitoring.

The control loop runs at 20 kHz. To minimize noise, center-aligned Timer1 generates PWM and triggers ADC1 synchronously with the PWM cycle. An ADC end-of-conversion interrupt triggers the main control step, minimizing measurement-to-actuation latency. Timer2 and GPIOB process encoder signals to derive position and angular velocity. The complete Simulink model is illustrated in Fig. 2.

The application, shown in Fig. 2, offers extensive configurability, including a 500 Hz IIR current filter, speed unit switching or selection of measured or constant DC voltage. It also features four operation modes:

- **Mode 1 (Open-Loop):** Sets PWM duty cycle directly, bypassing the current controller. Used for the Locked-Rotor test.

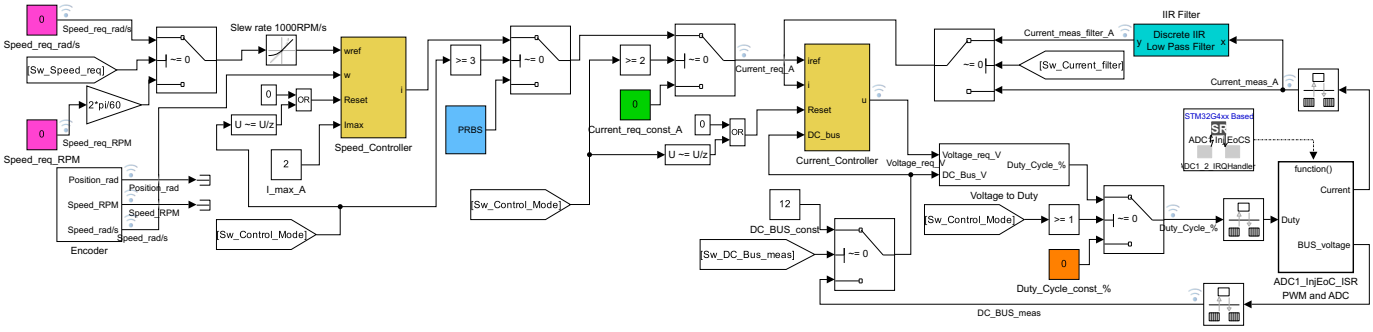

Fig. 2. Cascade control architecture implemented in MATLAB/Simulink.

- **Mode 2 - (Current Control - Constant):** Uses a constant reference for current loop tuning.
- **Mode 3 (Current Control - PRBS):** Applies a PRBS reference to acquire excitation data for mechanical identification.
- **Mode 4 - (Cascade Control):** Engages the full cascade structure for RCP verification.

The discrete PI controllers for speed and current are implemented in a parallel form with clamping anti-windup logic, described by the transfer function $F_r(z) = K_p + K_i \frac{T_s}{z-1}$. Controllers operate at 20 kHz using Forward Euler integration method and support external reset. A 1000 RPM/s slew rate limiter on the speed request ensures safety.

V. Parameter Identification Experiment

This chapter describes the experimental identification of the DC motor parameters and the subsequent design of the control loops. The process comprised four stages: electrical parameter identification via a locked-rotor test, current controller synthesis, mechanical parameter identification using the Least Squares method, and speed controller synthesis.

A. Electrical Parameter Identification

A locked-rotor test was performed to decouple electrical dynamics from mechanical motion. By mechanically fixing the shaft to eliminate back-EMF, the system response to a 50 % PWM voltage step was strictly governed by the RL circuit dynamics. To mitigate inaccuracies from voltage drops across the L6230 driver ($2 \times R_{DS(on)}$), the shunt resistor (R_{Shunt}), dead-time effects, and cabling, the actual voltage was measured at the motor terminals directly using a voltmeter, yielding 2.72 V. The resulting current response was recorded and analyzed, as shown in Fig. 4. The armature resistance R was calculated from steady-state terminal voltage and the measured current, while the inductance L was derived from the current rise time constant. The identified motor armature resistance R corresponds $1.69\,\Omega$, and the inductance L was determined to be 2.08 mH. The deviation from the datasheet nominal resistance ($0.82\,\Omega$) is attributed to significant age of the motor, contact resistances and temperature variations.

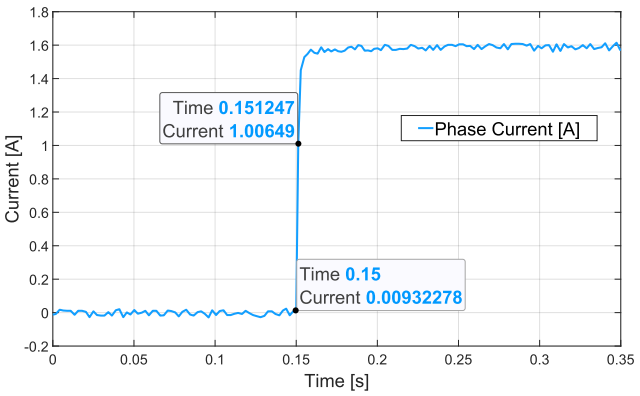

Fig. 4. Locked-rotor current response

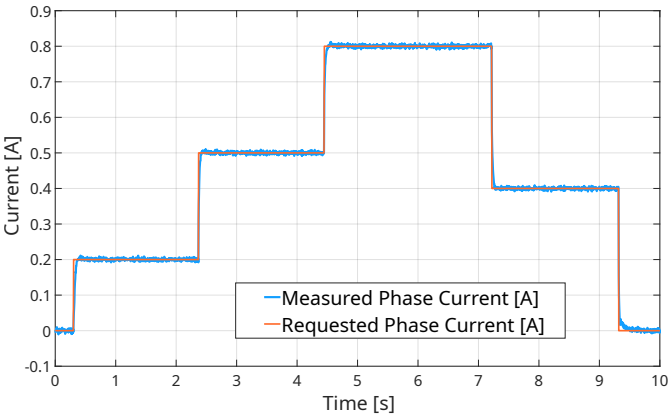

Fig. 5. Closed-loop current response

B. Current PI Control Design

Using the identified R and L, a PI current regulator was tuned to stabilize the inner control loop. The controller design uses pole placement, targeting a phase margin of 60° and a crossover frequency of 500 rad/s corresponding to a response time 2 ms. The current PI loop was then tested in program Mode 2 with a constant current reference; Fig. 5 shows the measured closed-loop current response.

C. Mechanical Parameter Identification

Following the current loop tuning, the mechanical parameters—moment of inertia (J) and viscous friction coefficient (b)—were identified. Together, these parameters constitute a first-order system with a static gain of $1/b$ and a time constant of J/b. In control Mode 3, the current PI loop is active, and a PRBS current reference was applied. The resulting angular velocity response is recorded, as illustrated in Fig. 6.

Subsequently, the Ordinary Least Squares (OLS) method was applied to the measured data based on the discretized form of (2). The parameters are estimated using the equation:

$$\hat{\theta} = (\mathbf{\Phi}^T \mathbf{\Phi})^{-1} \mathbf{\Phi}^T \mathbf{Y} \tag{3}$$

where $\mathbf{Y}$ is the output speed column vector, $\mathbf{\Phi}$ represents the regression matrix of delayed current inputs and speed outputs,

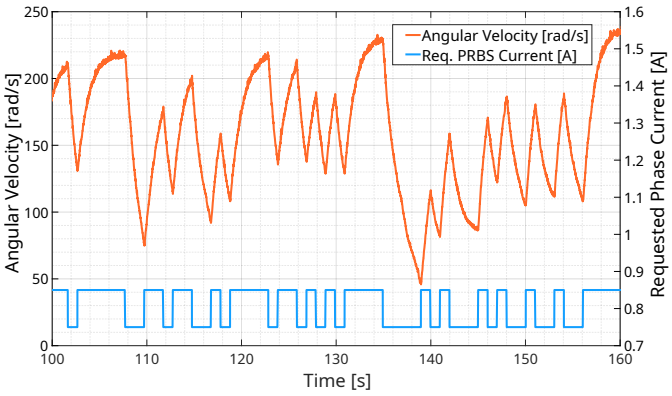

Fig. 6. Angular velocity response to PRBS current excitation

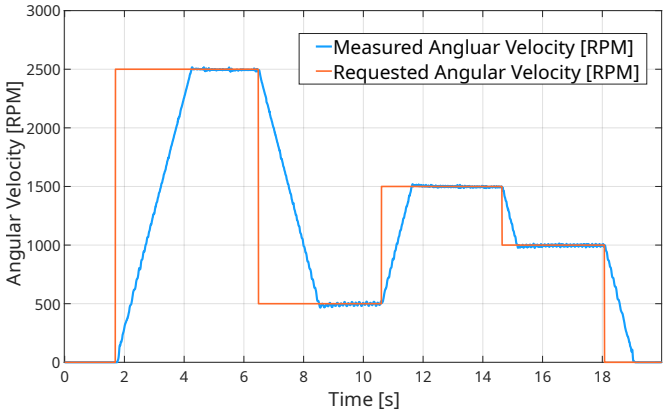

Fig. 7. Closed-loop speed response

and $\hat{\theta}$ is the vector of unknown parameters to be estimated. The identified values were $b = 4.64 \cdot 10^{-5}\,\mathrm{N \cdot m \cdot s/rad}$ and $J = 3.62 \cdot 10^{-5}\,\mathrm{kg \cdot m^2}$.

D. Speed PI Control Design

Finally, the speed PI controller was synthesized using the estimated mechanical parameters J and b. The speed PI controller was tuned for a phase margin of 60° and a bandwidth of 10 rad/s ensured a response 50 times slower than the inner loop. The speed control was evaluated in program Mode 4 using a constant speed reference with a slew-rate limit of 1000 rpm/s; the motor speed response is shown in Fig. 7.

Conclusion

A low-cost platform for DC motor identification and RCP was developed using the STM32 Nucleo-64 G474RE MCU with the STM32 NUCLEO-IHM07M1 power expansion board. MATLAB/Simulink facilitated automatic code generation and real-time tuning.

Armature resistance $R = 1.69\,\Omega$ and inductance $L = 2.08$ mH were experimentally identified using a locked-rotor test. The values differ slightly from the nominal values $R_{\text{tab}} = 0.82\,\Omega$ and $L_{\text{tab}} = 1.34$ mH as shown in Tab. I. These differences are attributed to the relatively high on-resistance of the L6230 driver ($R_{DS(on)} = 0.73\,\Omega$), uneven temperature distribution, dead-time or motor wear. After identifying the transfer function of the motor's electrical subsystem, a current PI controller was designed and demonstrated in Fig. 5.

Using a PRBS current as the input signal, the motor's rotational speed was measured as the output signal. The least squares identification method was applied to these input-output data, yielding the inertia parameter $J = 3.62 \cdot 10^{-5}\,\mathrm{kg \cdot m^2}$. and the viscous friction coefficient $b = 4.64 \cdot 10^{-5}\,\mathrm{N \cdot m \cdot s/rad}$. Inaccuracies in the identification arise from neglecting constant Coulomb static friction, which leads to overestimation of the coefficient b, as well as from the motor's low load during the test.

Experimental validation confirmed the platform met development objectives. Future work may involve testing advanced identification methods for DC motors and PMSMs. With the addition of the IHM08M1 power board, the platform can accommodate higher-power motors.

Acknowledgment

The completion of this paper was made possible by the grant No. FEKT-S-26-8988 - "Advanced Methods in Cybernetics, Robotics, and Artificial Intelligence" financially supported by the Internal science fund of Brno University of Technology.

References

[1] LI, Meidong and MA, Yanjiao, 2020. Parameter Identification of DC Motor based on Compound Least Square Method. Online. 2020 IEEE 5th Information Technology and Mechatronics Engineering Conference (ITOEC). Pp. 1107-1111. Available at: https://doi.org/10.1109/itoec49072.2020.9141652. [accessed 2026-01-16].

[2] MEAH, Kala; HIETPAS, Steven and ULA, S., 2007. Rapid Control Prototyping of a Permanent Magnet DC Motor Drive System using dSPACE and Mathworks Simulink. Online. APEC 07 - Twenty-Second Annual IEEE Applied Power Electronics Conference and Exposition. Pp. 856-861. ISSN 1048-2334. Available at: https://doi.org/10.1109/apex.2007.357615. [accessed 2026-01-16].

[3] KIM, Pyung Soo; KIM, Su Min and KIM, Su Yeol, 2024. Model Parameter Identification for DC Motor Using Rapid Control Prototyping System. Online. 2024 International Conference on Artificial Intelligence in Information and Communication (ICAIIC). 2024-2-19, pp. 816-819. ISBN 9798350344349. Available at: https://doi.org/10.1109/icaiic60209.2024.10463479. [accessed 2026-01-16].

[4] HUGHES, Austin and DRURY, Bill, 2019. Electric motors and drives: fundamentals, types and applications. Fifth Edition. Oxford: Newnes. ISBN 978-0-08-102615-1.

[5] KOZUMPLÍK, M, 2025. Rapid Prototyping Platform for the Control and Testing of Electric Drives. Online. Proceedings I of the 31st Conference STUDENT EEICT 2025. 1st Edition. Brno: BUT FEKT, 2025. p. 296-300. ISBN: 978-80-214-6321-9. Available at https://www.eeict.cz/eeict_download/archiv/sborniky/EEICT_2025_sbornik_1.pdf. [accessed 2026-01-16].

[6] KHAN, Jihas, 2016. Rapid Control Prototyping (RCP) solutions for the validation of motor control applications. Online. 2016 International Conference on Emerging Technological Trends (ICETT). Pp. 1-6. Available at: https://doi.org/10.1109/icett.2016.7873699. [accessed 2026-01-16].

[7] ARUNI, J; TITUS, Jose and HATUA, Kamalesh, 2021. Design of a DSP-based Control Platform for Real-Time Simulation of Power Electronic Systems. Online. 2021 IEEE 2nd International Conference on Smart Technologies for Power, Energy and Control (STPEC). 2021-12-19, pp. 1-6. Available at: https://doi.org/10.1109/stpec52385.2021.9718761. [accessed 2026-01-16].

[8] SHAMASHERGY, Abdullah El; TOKHY, Hazem El; ABDO, Mahmoud; FARAG, Mohamed; SAMIEH, Mohannad et al., 2025. A Cost-Effective Hardware-in-the-Loop (HIL) Platform for Power Electronics Control Using DSP. Online. 2025 IEEE Conference on Power Electronics and Renewable Energy (CPERE). 2025-9-28, pp. 1-8. Available at: https://doi.org/10.1109/cpere65146.2025.11240030. [accessed 2026-01-16].

[9] STM32G4 Nucleo-64 boards - User Manual. Online. Rev 6, December 2025. Available at: https://www.st.com/resource/en/user_manual/um2505-stm32g4-nucleo64-boards-mb1367-stmicroelectronics.pdf. [accessed 2026-01-16].

[10] Getting started with the X-NUCLEO-IHM07M1 motor driver expansion board based on the L6230 for STM32 Nucleo - User Manual. Online. Rev 4, December 2023. Available at: https://www.st.com/resource/en/user_manual/um1943-getting-started-with-the-xnucleoihm07m1-motor-driver-expansion-board-based-on-the-l6230-for-stm32-nucleo-stmicroelectronics.pdf. [accessed 2026-01-16].

[11] SUL, Seung-Ki, c2011. Control of electric machine drive system. IEEE Press series on power engineering. Hoboken: Wiley-IEEE Press. ISBN 978-0-470-59079-9.

[12] KIM, Sang-Hoon, 2017. Electric motor control: DC, AC, and BLDC motors: DC, AC, and BLDC motors. Amsterdam, Netherlands: Elsevier. ISBN 978-0-12-812138-2.

Fatigue Life Analysis Methods for Vibration Shakers

1st Vít Šimeček
Department of Control and Instrumentation
Brno University of Technology
Brno, Czech Republic
vit.simecek@vut.cz

***Abstract*—This paper addresses fatigue life estimation of vibration shakers used in accelerated vibration testing for product lifetime verification. Accelerated tests subject the shaker to significant operational loading, and unexpected failures may lead to costly production delays. The study presents the principles and methodological variants of fatigue life calculation and provides experimental verification of some proposed approaches.**

***Index Terms*—Fatigue life, vibration shaker, accelerated testing, Dirlik, Miner's rule, spectral analysis.**

I. Introduction

Nowadays, lifetime management is important in almost every industrial section. Vibration testing using vibration shakers is a tool for verifying the quality of the product in terms of stress during its planned lifetime and to find defective components before their completion into a product. Vibration testing using electrodynamic shakers is commonly employed to verify product quality under simulated lifetime loading and to identify defects. However, accelerated vibration tests reproduce the operational stress of a component within hours or days, which subjects the shaker itself to significant loading. Moreover, shaker downtime is costly and may delay production due to interruption of vibration testing procedures. For this reason, accurate fatigue life calculation is required. This article presents its principles, variants and some experimental verification of fatigue life calculation to reduce shaker failures and improve operational reliability.

II. Vibration Shakers

An electrodynamic vibration shaker is a laboratory actuator used to reproduce prescribed acceleration or force histories at the mounting table to evaluate the dynamic response and durability of a device under test (DUT). It operates on the Lorentz-force principle, analogous to a loudspeaker, and is typically driven in closed loop using an accelerometer feedback signal. The typical operational frequency range is up to thousands Hz, and the DUT can be mounted on the shaker table. [1] The mathematical description of the vibration shaker (model) is described in [2], [3]. The vibration shaker is capable of various test profiles [4]:

- **Sine sweep** test (chirp signal) that can detect oscillations at specific frequencies. The fundamental frequency is increased according to the specified linear/logarithmic rate. Using sine sweep is recommended for shaker diagnostics.
- **Random** is defined by frequency range (for instance 5-2000 Hz) and energy levels in specified bands of the range (predefined PSD curve), duration time
- **Sine on Random (SoR)** as a linear combination of sine and random excitation with the possibility of multiple sine frequencies
- **Shock** excitation, one event with defined amplitude, acceleration, and other parameters. Time-series data or SRS (Shock Response Spectrum) is obtained.

Since the shaker is used primarily in industrial testing, high priority is given to its availability and operational efficiency, as failure of the device may postpone batch sample testing and, in extreme cases, interrupt the manufacturing process. Consequently, continuous assessment of fatigue life after each test and proactive failure prediction are desirable.

III. Fatigue analysis

The aim of fatigue life estimation is to estimate the life cycle/time to failure under specific loading of the vibration shaker. For this purpose, there are methods both in time and frequency domain.

A. Cumulative fatigue

Fatigue damage is calculated with following principle in general: using the response of SDOF systems; each of them has a specific natural resonance frequency dependent on the covered bandwidth. This results in a particular damage at the frequency band. Afterwards, a Miner's rule of fatigue damage accumulation (total damage D) is applied [5]:

$$D = \sum_i d_i = \sum_i \frac{n_i}{N_i} \tag{1}$$

where d_i is damage at specified stress level,
n_i number of cycles at stress level and
N_i number of cycles to failure at stress level (defined by the SN curve).
Thus, the D should be always lower than 1. [5]

SN curve is a (usually experimentally) defined curve of the maximal number of cycles under a certain stress that the component should withstand without failure/damage. In methods listed in [6] are used simplified models of SN curves with the parameters of the inverse slope k and intercept C

[MPa^k]. Example of a SN curve is in the Figure 1, parameters in Table I.

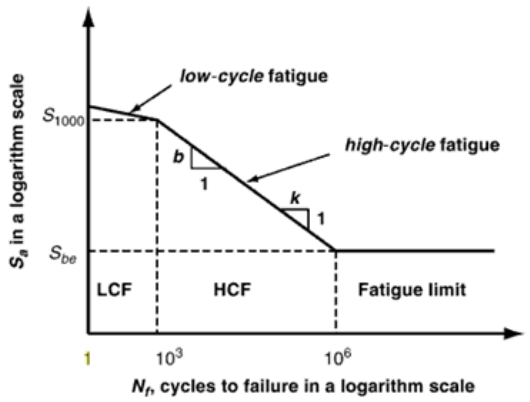

Fig. 1: SN curve example, source: [7].

TABLE I: SN curve parameters for various materials. Source:[6].

Material	C [MPa^k]	k
Steel	$1.9 \cdot 10^{12}$	3.3
Aluminium	$6.8 \cdot 10^{19}$	7.3
spring steel	$1.4 \cdot 10^{37}$	11.7

B. Time domain fatigue analysis

Time domain methods might be used when we acquire a time signal from the vibration test and calculate the mechanical stress for following workflow. The most common method is *Rainflow counting*. This method decomposes a load time history into an equivalent set of closed cycles characterized by the stress (or strain) range and mean value, which are then mapped to N_i via the S–N curve and the Miner's rule, and so determine the fatigue life. As a general signal does not only consist of repeated periods but the general cycles vary in amplitude, period, components etc., Rainflow counting evaluates those cycles and provides numerical result for fatigue analysis. [8]

There is also a possibility to evaluate the test by the exact number of cycles when running the sine sweep profile by counting the cycles derived from test time and sine sweep rate and combine it with frequency domain analysis.

C. Frequency domain fatigue analysis

The typical vibration test produces a large amount of time-domain data. Thus, it is useful to analyze the test directly in the vibration shaker software and store only the frequency domain data in many applications. Furthermore, we obtain information for every single frequency that would be more complicated to detect in a time domain.

Methods are not generally applicable on every test profile due to the statistical properties (e.g. stationarity) of specific profile which are factored into the calculation. Thus, this article is mainly handling with sine sweep and random profiles. There are, on the other hand, further studies about other profiles, for example Sine on Random in [5], [9].

For sine sweep, we use usually the Short-time Fourier transform spectrum (STFT, Figure 2). For random we use the form of Power Spectral Density (PSD, Figure 3). Even if this implies a reduction in the information contained in the stored data, there are several methods for fatigue life calculation from PSD. Usually, acceleration data are measured, which implies that data have to be converted by appropriate transfer function into equivalent stress [MPa] (for Rainflow, Dirlik and Tovo-Benasciutti).

Overview of the methods is given in [6], there is also an open-source python library for computation of them in [10]. Selection of the methods is also mentioned below:

1) Dirlik method: (random) Dirlik introduced one of the first methods for fatigue life calculation. Dirlik, like many other frequency domain methods, demands a stationary Gaussian process [6]. This implies that the method is suitable for dataset obtained from a random vibration test.

This method uses approximation of the cycle-amplitude distribution by three probability densities, namely exponential and two Rayleigh densities [11].

2) Tovo–Benasciutti: (random) Next to Dirlik, Tovo–Benasciutti method is another well performing and well-established method. It calculates the fatigue life as a linear combination of the upper and lower fatigue damage intensity limits. The method has two variants as were proposed by Benasciutti and Tovo. [6], [11].

Note that both Dirlik and Tovo–Benasciutti methods are dependent on the SN curve slope. These methods perform better on lower slopes k. For higher values, the fatigue error increases, as described in [6] with the proposal of other better-performing methods.

3) THD: (sine sweep) is used for the shaker analysis using the sine sweep. The THD is then calculated for every fundamental frequency and generally increases over time as the device suffers from degradation. A common rule-of-thumb threshold used in practice is THD $> 10\%$. When THD reaches higher values, the shaker is expected to be defective. [12].

4) Resonance frequency: (sine sweep) Resonance frequency shift over time may contain information about the wear of mechanical components, etc [13]. The example resonance can be found in the x and y axis in the Figure 2 around 780 Hz. This metrics is further evaluated on the experimental data as it is not a classical approach as decribed in [14].

D. Fatigue analysis using AI

Neural networks have been increasingly adopted in many scientific disciplines, incl. vibrodiagnostics, especially for tasks such as fault classification and remaining useful life estimation, particularly because their ability for feature learning from high-dimensional measurements or detection of nonlinear dependencies.

An example of neural network (NN) utilization in time–frequency analysis is reported in [15], in which conventional signal processing techniques, including the Short-Time

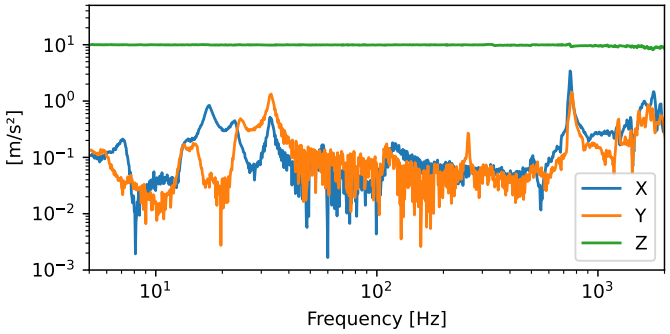

Fig. 2: Sine sweep spectrum, z-acceleration 10 g, 5-2000 Hz.

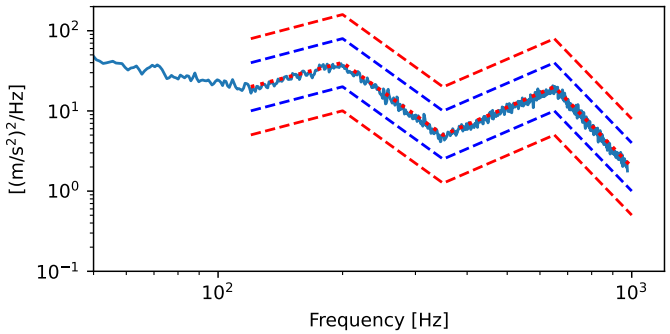

Fig. 3: PSD example, random. Red dotted line is target PSD, dashed blue are warnings, dashed red aborts. Limits are defined only in the specified interval.

Fourier Transform (STFT), are employed for feature extraction, followed by interpretation of the generated spectrogram by a pre-trained convolutional neural network.

Another option would be implementation of Physics-Informed Neural Network (PINN). The advantage of this model is that the physics is expressed in the loss function of the PINN. PINNs provide result interpretation based on physical laws modeled by a (partial) differential equation, which is their advantage and also a system identification u sing PINNs is possible [16].

Similar approaches as described above might be used in vibration shaker analysis and will be the aim of future research.

IV. Shaker validation

The shaker must be periodically tested to ensure that it is in sufficient o perating c onditions f or v ibration t esting. A s tests must be executed periodically, there is an option to watch the shaker parameters evolve over time. Furthermore, it is recommended to run the test after any maintenance.

The recommended periodic tests (from those introduced above) or parameters to follow are the following [14]:

- Sine sweep - resonance frequency changes, visible in x, y axes (results from this metrics are in the section Experiments). Moreover, THD parameter characteristics.
- Random - fatigue calculation (Dirlik, Tovo-Benasciutti)
- Drive signal - electricity supply changes, when unexpectedly increases, a failure is to happen.

The interval of the testing depends on the recommendation of the specific shaker or the laboratory guidelines. The resonance frequency test was selected for further analysis, as it is an additional test to the others proposed in [14].

V. Experiments

An example of the periodic test executed regularly in the Test lab CVVOZE at the Department of Control and Instrumentation, Faculty of Electrical Engineering, Brno University of Technology, is described and illustrated below.

Approximately once per two months or after any repairs or changes in the shaker, a sine sweep test of following parameters is executed:

- Acceleration: 1 g
- Frequency span: 5–2000 Hz
- Sensor output: 3 axial, z matches the accelerations axis as in Figure 2.
- Test cube (small variant, fixed on shaker).

There are 16 spectra from various acquisition dates with the same parameters in total. For testing purposes, three synthetic test were added to simulate upcoming failure (freq decrease of resonance around 750 Hz) marked test number 17 to 19 in Figure 4 and 5.

A. Analysis

Spectra are preprocessed by the logarithm in order to enhance also the less significant resonances and afterwards filtered by a simple averaging filter. An odd length and filter delay compensation are used, followed by a peak index correction according to the original unfiltered data.

Spectra measured by the acceleration sensor perpendicular (y axis in following examples, x axis might be analyzed equivalently) to the direction of movement (z axis) are analyzed using a Scipy *find_peaks* method. As the method has peak parameters *height, threshold, distance, prominence, width, etc.* (see [17] for explanation), it has to be fine-tuned (adapted for every single dataset) for a general use case. In this study, only parameters distance, prominence and wlen were used, furthermore, *find_peaks* parameters were not adapted between spectral datasets; however, significant peaks were always detected sufficiently.

Other options for future algorithm tuning is using the polynomial or spline interpolation in the peak detection to obtain the peak frequency even with the frequency-domain sub-bin precision.

B. Results discussion

Resonances (peak frequencies) were tracked across datasets, as is illustrated in Figure 4. The most important resonances are the following:

- dark green – around 780 Hz
- light orange – 260 Hz
- blue – 30 Hz
- orange and brown – 110 Hz, at dataset 12 not detected, thus, the trend is interrupted

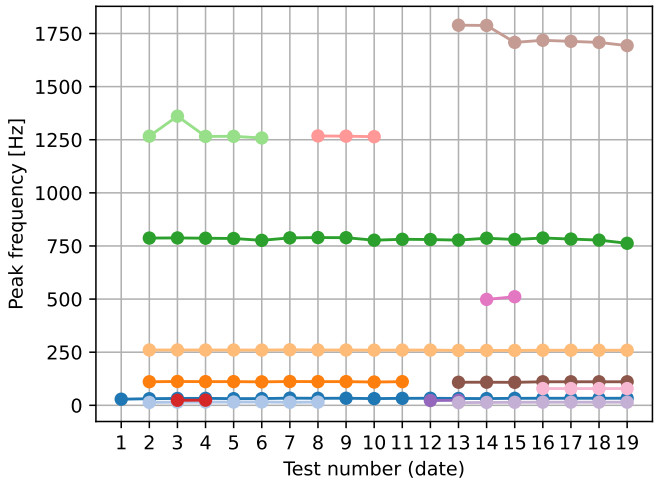

(a) Frequency trends.

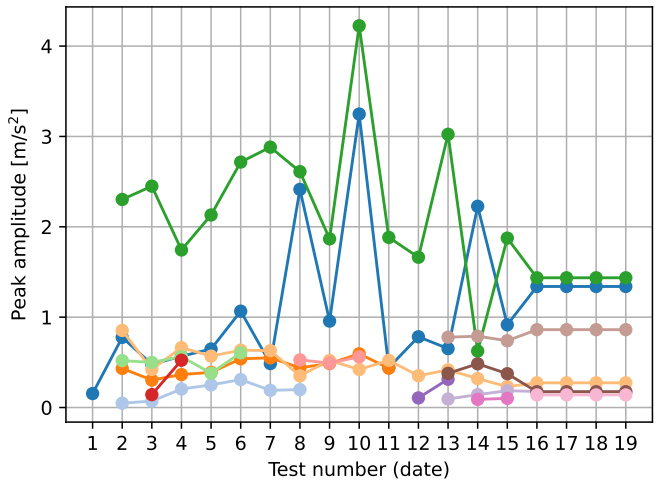

(b) Amplitude trends.

Fig. 4: Detected peak trends across the datasets, y axis.

The upper graph in Figure 4 shows the frequency trends between datasets. Usually, an initiating crack may lower the resonant frequency [13] which is visible in the trend. For algorithm testing, datasets 17 to 19 show the (synthetically) decreasing resonant frequency as is more visible in Figure 5. Every resonance change over certain threshold (5 % for example), which has to be adapted to the specific device, should lead to the shaker inspection and component replacement in purpose to satisfy the principles of the predictive maintenance.

The lower graph shows the amplitudes - intensities - of the resonances. Two most significant resonances are 780 and 30 Hz, peak amplitudes are in [m/s^2]. There is a significant deviation at the 10th date, however, it returns in the 11th. Amplitude analysis provides further possibilities for the shaker analysis and will be the aim of the next research as well as mentioned AI methods.

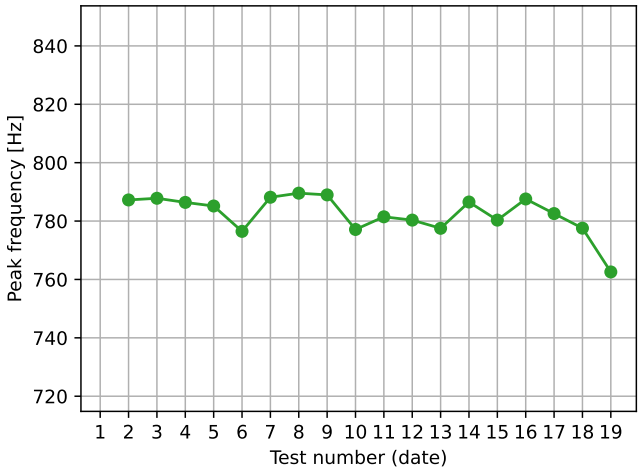

Fig. 5: Detail of the frequency increase around 800 Hz in synthetic datasets, y axis.

VI. Conclusion

This article presents methods for fatigue life estimation of vibration shakers used in accelerated testing. The vibration shaker and commonly used test profiles were described, from which the sine sweep profile is suitable for shaker parameter validation due to its ability to detect resonances. Since most tests produce spectral data, spectral methods for random loading, particularly the Dirlik and Tovo–Benasciutti approaches, were reviewed as reliable techniques for fatigue evaluation.

A peak tracking method based on the sine sweep test profile was further proposed and applied to shaker analysis in the test laboratory CVVOZE at the Brno University of Technology with possible future modifications to mitigate the effects of noise and improve algorithmic robustness. Synthetic data demonstrate the capability of the algorithm to detect shifts in resonant frequencies over time, which may indicate impending failure.

Acknowledgment

The completion of this paper was made possible by the grant No. FEKT-S-26-8988 - "Advanced Methods in Cybernetics, Robotics, and Artificial Intelligence" financially supported by the Internal science fund of Brno University of Technology.

References

[1] J. Martino and K. Harri, "Virtual shaker modeling and simulation, parameters estimation of a high damped electrodynamic shaker," *International Journal of Me-chanical Sciences*, vol. 151, pp. 375–384, 2019. DOI: 10.1016/j.ijmecsci.2018.11.025. Accessed: Jan. 22, 2026.

[2] G. F. Lang and D. Snyder, "Understanding the Physics of Electrodynamic Shaker Performance," *Sound and Vibration*, vol. 35, no. 10, pp. 24–33, 2001.

[3] C. C. Smith and F. L. Staffanson, "Basic Mechanical Interactions in Shaker Testing," *Shock and Vibration*, vol. 4, no. 4, pp. 269–280, 1997. DOI: 10.1155/1997/ 190274. Accessed: Jan. 29, 2026.

[4] Vibration Research Corporation. "Products," Vibration Research, Accessed: Feb. 2, 2026. [Online]. Available: https://vibrationresearch.com/products/.

[5] A. Angeli, B. Cornelis, and M. Troncossi, "Synthesis of Sine-on-Random vibration profiles for accelerated life tests based on fatigue damage spectrum equivalence," *Mechanical Systems and Signal Processing*, vol. 103, pp. 340–351, 2018. DOI: 10.1016/j.ymssp.2017.10.022. Accessed: Feb. 4, 2026.

[6] A. Zorman, J. Slavič, and M. Boltežar, "Vibration fa-tigue by spectral methods—A review with open-source support," *Mechanical Systems and Signal Processing*, vol. 190, p. 110 149, 2023. DOI: 10.1016/j.ymssp.2023. 110149. Accessed: Jan. 22, 2026.

[7] Altair OptiStruct. "Fatigue using E-N (Strain - Life) Method," Accessed: Feb. 4, 2026. [Online]. Available: https://2021.help.altair.com/2021.2/hwsolvers/os/topics/solvers/os/fatigue en strain life method r.htm.

[8] C. Amzallag, J. Gerey, J. Robert, and J. Bahuaud, "Stan-dardization of the rainflow counting method for fatigue analysis," *International Journal of Fatigue*, vol. 16, no. 4, pp. 287–293, 1994. DOI: 10.1016/0142-1123(94) 90343-3. Accessed: Feb. 2, 2026.

[9] F. Kihm, A. Halfpenny, and N. S. Ferguson, "Fatigue Life from Sine-on-random Excitations," *Procedia Engi-neering*, vol. 101, pp. 235–242, 2015. DOI: 10.1016/j. proeng.2015.02.033. Accessed: Feb. 4, 2026.

[10] Laboratory for Dynamics of Machines and Structures at the University of Ljubljana, Faculty of Mechanical Engineering. "FLife/setup.py at main · ladisk/FLife," GitHub, Accessed: Jan. 23, 2026. [Online]. Available: https://github.com/ladisk/FLife/blob/main/setup.py.

[11] M. Mršnik, J. Slavič, and M. Boltežar, "Frequency-domain methods for a vibration-fatigue-life estimation – Application to real data," *International Journal of Fatigue*, vol. 47, pp. 8–17, 2013. DOI: 10.1016/j.ijfatigue.2012.07.005. Accessed: Jan. 22, 2026.

[12] Vibration Research Corporation. "Total Harmonic Dis-tortion (THD) Graph - Vibration Testing," VRU, Ac-cessed: Feb. 2, 2026. [Online]. Available: https://vru.vibrationresearch.com/lesson/total-harmonic-distortion/.

[13] S. Doebling, C. Farrar, M. Prime, and D. Shevitz, "Damage identification and health monitoring of struc-tural and mechanical systems from changes in their vi-bration characteristics: A literature review," LA–13070-MS, 249299, 1, 1996, LA–13070-MS, 249 299. DOI: 10.2172/249299. Accessed: Mar. 30, 2026.

[14] Vibration Research Corporation. "Shaker Validation," Vibration Research, Accessed: Mar. 30, 2026. [Online]. Available: https://vibrationresearch.com/webinar/shaker-validation/.

[15] Y. Liao, I. Ragai, Z. Huang, and S. Kerner, "Manu-facturing process monitoring using time-frequency rep-resentation and transfer learning of deep neural net-works," *Journal of Manufacturing Processes*, vol. 68, pp. 231–248, 2021. DOI: 10.1016/j.jmapro.2021.05.046. Accessed: Feb. 3, 2026.

[16] W. Zhai, D. Tao, and Y. Bao, "Parameter estimation and modeling of nonlinear dynamical systems based on Runge–Kutta physics-informed neural network," *Non-linear Dynamics*, vol. 111, no. 22, pp. 21 117–21 130, 2023. DOI: 10.1007/s11071-023-08933-6. Accessed: Feb. 3, 2026.

[17] The SciPy community. "Find peaks — SciPy v1.17.0 Manual," Accessed: Feb. 17, 2026. [Online]. Available: https://docs.scipy.org/doc/scipy/reference/generated/scipy.signal.findpeaks.html.

RENESAS

SHAPE THE FUTURE



PROGRESS WITH US

EMBEDDED EVERYWHERE

GLOBAL & FLEXIBLE

RENESAS – Embedded Everywhere

Renesas Design Czech s.r.o. as part of Renesas Electronics Corporation delivers trusted embedded design innovation with complete semiconductor solutions that enable billions of connected, intelligent devices to enhance the way people work and live. A global leader in microcontrollers, analog, power, and SoC products, Renesas provides comprehensive solutions for a broad range of automotive, industrial, infrastructure, and IoT applications that help shape a limitless future. Learn more at renesas.com.

WHO IS THIS FOR?

We are looking for pro-active and results-oriented engineers who would like to work on top notch technology and make a difference. Essential would be experience in one of the following disciplines: **Analog IC Design**, **Digital IC Design**, **Verification**, **Project Management**, S**W Development (embedded systems, C, ASM)**, **HW Testing**, **Microelectronics**.

WE OFFER

At Renesas we treat each other with respect and fairness. We see our differences as strengths and passionately believe that working at Renesas, you will be joining the brightest, most diverse and ambitious talent in the Semiconductor Industry. Talented people with strong motivation and a passion for winning can enjoy excellent career opportunities at Renesas. Renesas gives you opportunities for growth in personality and knowledge at any level of responsibility.

Part-time or Full-time employment /
Internships / Bachelor's or Master's thesis

jobs.renesas.com

www.renesas.com /company/renesas/

Comparative Analysis of Fitness Functions in Coverage Path Planning

1st Jan Čtvrtníček
Department of Theoretical and Experimental Electrical Engineering
Brno University of Technology
Brno, Czech Republic
xctvrt06@vutbr.cz

***Abstract*—Coverage Path Planning (CPP) is an important field in autonomous robotics, applicable in various areas from agriculture to indoor robot navigation. With the rise of Unmanned Aerial Vehicles (UAVs), applications using autonomous flight have become more prevalent. While numerous articles focus on developing novel path planning methods,few studies compare the influence of different fitness functions on the resulting paths. This paper aims to analyse fitness functions, ranging from generic shortest-path optimization to UAV-specific algorithms, and compare their advantages and disadvantages. The results are used to determine the best fitness function for UAV CPP and to specify the optimal use cases for each function.**

***Index Terms*—UAV, metaheuristic, path planning, optimization, fitness function**

I. Introduction

In modern robotics, Coverage Path Planning (CPP) is a common problem applicable to all types of mobile robots (Unmanned Ground Vehicles - UGVs, Unmanned Aerial Vehicles - UAVs, etc.). CPP is a specific type of planning that distinguishes itself from typical point-to-point planning by its main goal, which is to cover every point in Area of Interest (AoI), while still maintaining the typical requirement of obstacle avoidance. Although still actively researched, CPP applications have long been used in various fields, for example vacuum robots [1], de-mining robots [2], robot lawn mowers [3], and many more.

This paper specifically focuses on UAV coverage planning, which comes with its own set of advantages and disadvantages. Applications requiring a fast exploration of complex areas commonly use UAVs because of their bird's-eye view and free 3D movement compared to conventional ground mobile robots. Some use cases include search and rescue missions, crop monitoring, or photogrammetry.

The goal of this work is to analyse different fitness functions used in path planning and determine which one is best suited for the field of UAVs. The main challenge that comes with using drones is their energy consumption. Since each drone has a limited battery life, planning an optimal path is crucial for creating a reliable system that can cover large areas and operate for the maximum amount of time.

Using optimal planning algorithm is one part of creating optimal paths. The second part is choosing the correct fitness function, which determines path quality based on given criteria. Path quality can mean a number of different things; e.g., the shortest path is the best, the smoothest path is the best, etc.

To test different fitness functions, the metaheuristic Ant Colony Optimization (ACO) algorithm was used. Metaheuristics were chosen because their framework provides a general structure, making them great for path planning and testing different algorithm constraints.

II. Energy consumption

UAVs' energy consumption is generally made up of two components. The first is the energy needed to sustain on-board circuitry (on-board computer, communication interface, etc.), which is considered constant to simplify this task. The second component is energy needed for propulsion. This energy can be divided into three phases: acceleration, deceleration and constant-speed cruise [4].

Before calculating movement, the power to keep the drone in flight needs to be defined. Assuming the thrust needed to keep UAV in the air is equal to its weight, we get (1), where m is the mass and g is the gravitational acceleration [5].

$$T = mg \tag{1}$$

Next from (2) we can calculate the power necessary to keep the UAV hovering. This value is calculated using thrust T from (1), air density ρ, rotor disc area A and efficiency η [5].

$$P_h = \frac{T^{3/2}}{\eta\sqrt{2\rho A}} \tag{2}$$

To calculate the power necessary to move the UAV, we need to determine the drag force. This is a resistive force that opposes the motion of an object moving through a fluid. From (3), we can see that it is calculated from air density ρ, drone velocity V, reference area S and drag coefficient C_D. In this equation, velocity is function of time, which results in the drag force changing its value during acceleration and deceleration. The reference area is usually the frontal projected area (in m^2)—the cross-section of the drone that the wind "sees" as it flies forward. When the drone tilts to accelerate or decelerate, this area changes. The drag coefficient is a dimensionless number that accounts for all the complex flow physics (skin friction, shape drag, flow separation) that cannot be modelled

by simple geometry. Multiplying the drag force with velocity results in parasitic power (4), that the drone must overcome when accelerating [4] [6] [7].

$$F_{drag} = \frac{1}{2} C_D S \rho V^2 \tag{3}$$

$$P_{parasite} = F_{drag} V = \frac{1}{2} C_D S \rho V^3 \tag{4}$$

The total power consumed during a UAV flight can be calculated using the Standard Aerodynamic Power Model (5). Inducted power is the power used to keep the drone in the air, profile power is determined from friction of spinning the propellers, and parasite power is the energy spent pushing the drone through the air [4] [8].

$$P_{total}(V) = P_{induced}(V) + P_{profile}(V) + P_{parasite}(V) \tag{5}$$

This equation can be simplified by substituting inducted and profile power with our hover power from (2). From this, our total power can be calculated with (6) [4] [8].

$$P(V) = \frac{T^{3/2}}{\eta \sqrt{2 \rho A}} + \frac{1}{2} C_D S \rho V^3 \tag{6}$$

$$E(t) = \int_{t_1}^{t_2} P(t) dt \tag{7}$$

The total energy consumption of a quadcopter during a flight segment is the integral of power over time (7). This total energy is the sum of the energy required for acceleration, constant-speed cruise and deceleration. Throughout all phases, the drone must expend a baseline hover power to remain aloft (2). During the acceleration phase to a target cruise velocity V_c at a constant rate a, the battery must supply power to maintain lift, overcome increasing aerodynamic drag, and add kinetic energy, formulated as (8) [9] [10].

$$E_{acc} = \int_0^{V_c/a} \left(P_h + \frac{1}{2} C_D S \rho V(t)^3 + maV(t) \right) dt \tag{8}$$

This can be directly solved as (9).

$$E_{acc} = P_h \left(\frac{V_c}{a} \right) + \frac{C_D S \rho V_c^4}{8a} + \frac{1}{2} m V_c^2 \tag{9}$$

Once at the constant cruise velocity V_c for a duration t_{cruise}, the kinetic power requirement drops to zero, yielding the cruise energy integral (10).

$$E_{cruise} = \int_0^{t_{cruise}} \left(P_h + \frac{1}{2} C_D S \rho V_c^3 \right) dt \tag{10}$$

Solving this integral, we get (11).

$$E_{cruise} = \left(P_h + \frac{1}{2} C_D S \rho V_c^3 \right) t_{cruise} \tag{11}$$

Finally, depending on the braking method, the energy consumption is either (12) or (13), which can be solved as (14) or (15). The first scenario occurs when the UAV uses aerodynamic drag to slow down, requiring only the energy to stay aloft. The second scenario involves active braking, where motors actively brake [10].

$$E_{dec} = \int_0^{V_c/d} P_h dt \tag{12}$$

$$E_{dec} = \int_0^{V_c/d} \left(P_h - \frac{1}{2} C_D S \rho V(t)^3 + mdV(t) \right) dt \tag{13}$$

$$E_{dec} = P_h \left(\frac{V_c}{d} \right) \tag{14}$$

$$E_{dec} = P_h \left(\frac{V_c}{d} \right) - \frac{C_D S \rho V_c^4}{8d} + \frac{1}{2} m V_c^2 \tag{15}$$

These equations give us insight into UAV energy consumption depending on the type of movement. We can clearly see that acceleration requires the most energy, while passive braking is most energy-efficient. The influence of this is tested later in this paper.

III. Planner

This section describes the structure of the implemented planner. Planning is generally done in two to three steps, depending on the application. First, we need to setup the Area of Interest (AoI), which is then divided into points based on the mission parameters (flight altitude, etc.). The second step is dividing the created mission waypoints into sub-polygons, depending on the number of UAVs used. Lastly, the algorithm creates suitable flight paths.

In this paper, the focus is on the fitness function for generating paths, so the second step, which focuses on swarm planning, can be skipped.

A. Area of Interest

The first step for coverage mission planning is to set an area of interest. In real-world applications, a user typically creates a polygon on a map using latitude and longitude. With this format, it is essential to convert coordinates from WGS84 to a better-suited coordinates system, like East, North, Up (ENU) system. The converted polygon can now be used with desired units (meters).

Next is the coverage waypoints creation. Using flight altitude, sensor field of view (FOV), and aspect ratio, the correct waypoint distance (footprint) can be determined. This is done using (16), (17) and (18). For complete coverage, the AoI is divided into regular cells using footprint distance. Implemented methods for waypoint creation are Poisson disk division, square division and hexagonal division, and their respective usage is specified for each given testing sample.

$$D = 2x \tan \left(\frac{\theta}{2} \right) \tag{16}$$

$$A = \frac{D}{\sqrt{1+r^2}} = \frac{2x\tan\left(\frac{\theta}{2}\right)}{\sqrt{1+r^2}} \tag{17}$$

$$B = rA = \frac{2rx\tan\left(\frac{\theta}{2}\right)}{\sqrt{1+r^2}} \tag{18}$$

Here, *D* represents footprint diagonal, *A* and *B* are footprint width and height, θ is the field of view, *x* is flight altitude, and *r* is aspect ratio [11] [12].

B. Ant Colony Optimization

Metaheuristic algorithms can be categorised as a branch of artificial intelligence, and their structure provides a general framework that can be applied regardless of the specific task. This makes them highly suitable for solving optimization problems like CPP. Ant Colony Optimization falls into the category of nature-inspired swarm metaheuristics. It utilizes the behaviour of ant when foraging for food. At the beginning, ants move in a random fashion, leaving a trail of chemical pheromones. When an ant finds a food source, it evaluates the quality and quantity of the food and returns with some of it back to the nest. On this trip back, ant deposits pheromones based on the overall quality of the food source. During other explorations, ants follow these trails, and over time, they converge on the best path to the given food source [13], [14].

ACO implements this behaviour by initially conducting a random search and then leaving a pheromone trail based on the quality of the solution. Used in path planning, this quality is determined by a fitness function [14] [15].

$$p(c_{ij} \mid s^p) = \frac{\tau_{ij}^{\alpha}\eta_{ij}^{\beta}}{\sum_{c_{il}\in N(s^p)} \tau_{il}^{\alpha}\eta_{il}^{\beta}} \tag{19}$$

Equation (19) determines the probability of selecting a given waypoint. The value α changes the influence of pheromones, and β the influence of the heuristic information (fitness function) on the resulting probability. The variable τ represents the level of pheromones between waypoints and the variable η is the heuristic value determined by the fitness function [15].

IV. Fitness functions

Determining path quality is an essential part of planning algorithms. Let a path consist of N discrete waypoints where each point is defined as $p_i = (x_i, y_i)$. UAVs are unique in the sense that they require precise battery management, since fully draining the battery results in a crash and most likely the destruction of the drone. To avoid this outcome, planning algorithms need to create paths that have the best value-to-cost ration (value being area covered and cost being battery drainage). Planning algorithms such as Ant Colony Optimization work iteratively and the algorithm "learns" from previous iterations to find a better solution. ACO does this by updating pheromones based on path fitness. Suboptimal paths do not accumulate as much pheromone, while optimal paths are prioritized. In this paper, three approaches for determining path quality are tested.

Other factor that needs to be specified is the type of UAV used. For this paper we assume that UAV is a quadcopter, which can stop and hover in place, making sharper turns possible. Planning for fixed-wing aircraft or VTOLs (Vertical Take-Off and Landing) must account for turning radii, which is a separate problem outside the scope of this paper.

A. Path length

The first fitness function is based on path length. This approach works with the assumption that the shortest path is the best. This approach is widely used, but it can be suboptimal when the shortest path also requires many sharp turns. Path length is calculated using the Euclidean distance between each point of the path (20). Total length is then calculated using (21) [16].

$$d(p_i, p_{i+1}) = \sqrt{(x_i - x_{i+1})^2 + (y_i - y_{i+1})^2} \tag{20}$$

$$D_{path} = \sum_{i=1}^{N-1} d(p_i, p_{i+1}) \tag{21}$$

B. Path smoothness

The second tested approach is based on overall path smoothness. This is generally an uncommon approach, as relying solely on smoothness tends to create paths that, while smooth, are otherwise impractical. For the sake of testing, this fitness function is included to provide a comprehensive comparison.

To evaluate the smoothness of a planned trajectory, we calculate the cumulative turning angle across all sequential waypoints. For any three consecutive waypoints, the heading angle of each segment is first determined using the two-argument arctangent function, as defined in (22). The raw angular difference δ_i between consecutive segments is then computed in (23). To ensure the angle reflects the shortest physical rotation and prevents the turn from exceeding π radians, the angle is normalized according to the piecewise function in (24). Finally, the total path smoothness cost S_{path} is formulated in (25) as the sum of all normalized turning angles, where a lower total indicates a smoother trajectory with fewer sharp, energy-draining turns [17].

$$\phi_i = \text{atan2}(y_{i+1} - y_i, x_{i+1} - x_i) \tag{22}$$

$$\delta_i = |\phi_{i+1} - \phi_i| \tag{23}$$

$$\Delta\theta_i = \begin{cases} 2\pi - \delta_i & \text{if } \delta_i > \pi \\ \delta_i & \text{otherwise} \end{cases} \tag{24}$$

$$S_{path} = \sum_{i=1}^{N-2} \Delta\theta_i \tag{25}$$

C. Hybrid approach

The hybrid approach combines both previous functions to ensure the path is efficient in both length and smoothness. The main challenge with combining smoothness and length is that they do not return normalized values. Generally, the length is much larger value, which can dominate the smoothness influence. To ensure the correct functionality of this fitness function, various weight values were tested. With the implemented planner, the smoothness value was around 10 to 100 times smaller than the length. For this reason, smoothness was assigned a weight of 50, while length had a weight of 1.

V. Testing setup

To test the different polygon partition methods and fitness functions, a testing setup was created in SITL (Software in the Loop). It is an ArduPilot simulator allowing the user to simulate different types of UAVs. When running in SITL, the sensor data comes from a flight dynamics model in a flight simulator. ArduPilot has a wide range of built-in vehicle simulators and can interface with several external simulators [18].

The planner was provided with a concave polygon, which was preprocessed for each of the three grid types: Poisson disk, square, and hexagonal division. The algorithm was tuned separately for each fitness function to ensure optimal results. To provide a robust evaluation, each grid method was tested across all fitness functions, generating 10 distinct paths for every combination. The resulting paths were then simulated in SITL, and the results are analysed in the next section.

VI. Results

To test all methods and fitness functions thoroughly, each was planned and simulated ten times. The first measured value was mean energy efficiency, measured in Joules per meter. In Table I, we can see that the smoothness fitness function was the most efficient for all grid methods, closely followed by hybrid approach, while the least efficient was the length fitness function.

TABLE I
Mean energy efficiency

Grid method	Fitness function	Mean efficiency [J/m]
Square grid	Length	59.7
	Smoothness	50.7
	Hybrid	51.7
Hexagonal grid	Length	59.5
	Smoothness	50.6
	Hybrid	51.7
Poisson grid	Length	59.4
	Smoothness	49.3
	Hybrid	53.2

Next the mean total energy consumption per path was evaluated. As shown in Table II, the results indicate that even though smoothness function was the most efficient per meter, it ultimately required the most overall power to complete the entire path. The same problem persisted with hybrid approach, and the lowest total energy consumption resulted form optimizing the path length. The relationship between path length and energy consumption is displayed in Fig. 1, where the clear winners per grid method were length-optimized paths.

TABLE II
Mean total energy consumption

Grid method	Fitness function	Mean energy [J]
Square grid	Length	404028.9
	Smoothness	468577.8
	Hybrid	434738.0
Hexagonal grid	Length	457922.4
	Smoothness	512682.6
	Hybrid	483690.9
Poisson grid	Length	313550.6
	Smoothness	402127.9
	Hybrid	339291.9

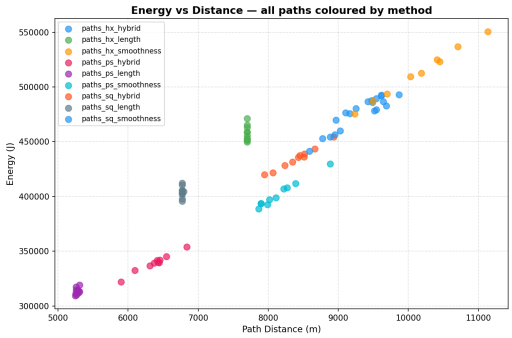

Fig. 1. Energy consumption based on path length

The last significant value measured was the mean path flight time. Table III clearly shows that the total mission time did not vary significantly between fitness functions, with the fastest being length-optimized paths and the slowest smoothness-optimized paths.

TABLE III
Mean flight time

Grid method	Fitness function	Mean flight time [s]
Square grid	Length	108.0
	Smoothness	123.8
	Hybrid	115.6
Hexagonal grid	Length	122.5
	Smoothness	135.7
	Hybrid	128.3
Poisson grid	Length	83.9
	Smoothness	105.5
	Hybrid	89.9

Overall, the fitness function analysis solidified the fact, that the largest quadcopter energy consumption comes from simply keeping the UAV aloft. Optimizing path smoothness or using the hybrid approach become viable when planning paths for fixed-wing planes or VTOLs, as they generate an upward aerodynamic lift force, but formally quantifying that benefit is outside the scope of this paper.

In Fig. 2 and Fig. 3, we can see flight speed comparison for the length-optimized path and the smoothness-optimized path respectively. The length-optimized path does not allow the UAV to reach top speed because it contains numerous turns and sharp corners. The opposite can be seen in the smoothness-optimized path, where the UAV often travels at top speed due to the long straight sections.

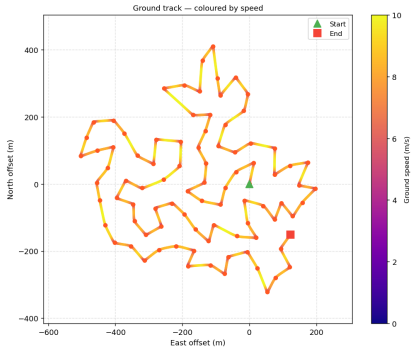

Fig. 2. Flight speed in length optimized path

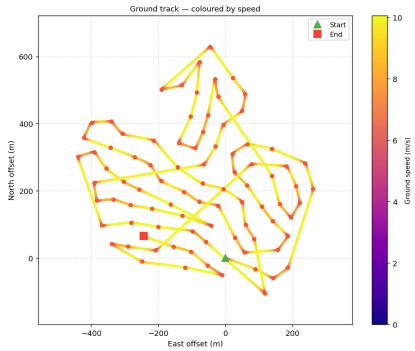

Fig. 3. Flight speed in smoothness optimized path

VII. Conclusion

This paper focuses on a comparative analysis of different path quality determination methods. Achieving the best value-to-cost ratio is a crucial aspect of planning missions for UAVs. The goal of this comparison was to find the most effective way to plan paths for real-world applications. Unsurprisingly, the planner optimized for path length provided the shortest paths. However, a somewhat unexpected result is that even though planning for smoothness or using a hybrid fitness function yields the best energy efficiency per meter, the resulting longer paths lead to overall much higher total energy consumption. While these results indicate that smoothness optimization is less ideal for quadcopters, these alternative fitness functions could prove highly beneficial for fixed-wing planes or VTOLs, which utilize the upward lift force generated by airflow over their wings.

Acknowledgment

This work was supported by the Ministry of the Interior of the Czech Republic under the programme Strategic Support for the Development of Security Research 2019–2025 (IMPAKT 1), project "An Artificial Intelligence-Controlled Robotic System for Intelligence and Reconnaissance Operations" (VJ02010036), and by the Czech Ministry of Defence under the PRODEF programme, project "Multirobotic System for Autonomous Field Operations with Multisource Navigation — ROJ II" (OZ01020024).

References

[1] F. Yasutomi, M. Yamada, and K. Tsukamoto, "Cleaning robot control", in Proceedings. 1988 IEEE International Conference on Robotics and Automation, IEEE Comput. Soc. Press, pp. 1839-1841. doi: 10.1109/robot.1988.12333.

[2] D. W. Gage, "Randomized search strategies with imperfect sensors", in SPIE Proceedings, SPIE, Feb. 1994, pp. 270-279. doi: 10.1117/12.167503.

[3] Z. L. Cao, Y. Huang, and E. L. Hall, "Region filling operations with random obstacle avoidance for mobile robots", Journal of Robotic Systems, vol. 5, no. 2, 1988, doi: 10.1002/rob.4620050202.

[4] Y. Zeng, J. Xu, and R. Zhang, "Energy Minimization for Wireless Communication With Rotary-Wing UAV", IEEE Transactions on Wireless Communications, vol. 18, no. 4, 2019, doi: 10.1109/twc.2019.2902559.

[5] "Momentum theory", Wikipedia: the free encyclopedia. Accessed: Mar. 10, 2026. [Online]. Available: https://en.wikipedia.org/wiki/Momentum_theory

[6] J. Anderson, Fundamentals of aerodynamics. Dubuque: McGraw-Hill Education, 2010.

[7] J. L. Marins, T. M. Cabreira, K. S. Kappel, and P. R. Ferreira, "A Closed-Form Energy Model for Multi-rotors Based on the Dynamic of the Movement", in 2018 VIII Brazilian Symposium on Computing Systems Engineering (SBESC), IEEE, 2018, pp. 256-261. doi: 10.1109/sbesc.2018.00047.

[8] João L. Marins, Tauã M. Cabreira, Kristofer S. Kappel, and Paulo R. Ferreira, "A Closed-Form Energy Model for Multi-rotors Based on the Dynamic of the Movement", Nov. 2018, doi: 10.1109/sbesc.2018.00047.

[9] R. Mahony, V. Kumar, and P. Corke, "Multirotor Aerial Vehicles: Modeling, Estimation, and Control of Quadrotor", IEEE Robotics & Automation Magazine, vol. 19, no. 3, 2012, doi: 10.1109/mra.2012.2206474.

[10] C. Di Franco and G. Buttazzo, "Energy-Aware Coverage Path Planning of UAVs", in 2015 IEEE International Conference on Autonomous Robot Systems and Competitions, IEEE, 2015, pp. 111-117. doi: 10.1109/icarsc.2015.17.

[11] T. Cabreira, L. Brisolara, and P. R. Ferreira Jr., "Survey on Coverage Path Planning with Unmanned Aerial Vehicles", Drones, vol. 3, no. 1, 2019, doi: 10.3390/drones3010004.

[12] S. Srinivasan, "Calculating a drone camera's image footprint OR How I learned to love tan again", Technology for Wildlife Foundation. Accessed: Mar. 11, 2025. [Online]. Available: https://www.techforwildlife.com/blog/2019/1/29/calculating-a-drone-cameras-image-footprint

[13] K. Sörensen and F. W. Glover, "Metaheuristics", in Encyclopedia of Operations Research and Management Science, Boston, MA: Springer US, 2013, pp. 960-970. doi: 10.1007/978-1-4419-1153-7_1167.

[14] C. Blum, "Ant colony optimization: Introduction and recent trends", Physics of Life Reviews, vol. 2, no. 4, 2005, doi: 10.1016/j.plrev.2005.10.001.

[15] M. Dorigo and C. Blum, "Ant colony optimization theory: A survey", Theoretical Computer Science, vol. 344, no. 2-3, 2005, doi: 10.1016/j.tcs.2005.05.020.

[16] "Euclidean distance", Wikipedia: the free encyclopedia. [Online]. Available: https://en.wikipedia.org/wiki/Euclidean_distance

[17] S. Guillén Ruiz, L. V. Calderita, A. Hidalgo-Paniagua, and J. P. Bandera Rubio, "Measuring Smoothness as a Factor for Efficient and Socially Accepted Robot Motion", Sensors, vol. 20, no. 23, Nov. 2020, doi: 10.3390/s20236822.

[18] "SITL Simulator (Software in the Loop)", ArduPilot Development Site. Accessed: Mar. 10, 2026. [Online]. Available: https://ardupilot.org/dev/docs/sitl-simulator-software-in-the-loop.html

Zero-Shot Traffic Accident Detection in Urban Surveillance Video Using Large Vision-Language Models

1st Stepan Jezek
Dept. of telecommunications, FEEC
Brno University of Technology
Brno, Czech Republic
Stepan.Jezek1@vut.cz

2nd Daniel Duraj
Dept. of telecommunications, FEEC
Brno University of Technology
Brno, Czech Republic
257095@vutbr.cz

3rd Radim Burget
Dept. of telecommunications, FEEC
Brno University of Technology
Brno, Czech Republic
burgetrm@vutbr.cz

***Abstract*—Traffic accident detection in urban surveillance video is a critical component of smart city infrastructure, yet remains a challenging open problem. Conventional deep learning approaches depend on large annotated datasets and fixed definitions of normal behavior, making them prone to failure across diverse road environments and camera viewpoints. The rapid advancement of large vision-language models (LVLMs) offers a compelling alternative: zero-shot, text-guided video understanding that requires no task-specific training. In this paper, we systematically evaluate three LVLM-based inference strategies for traffic accident detection – a context-aware visual question answering pipeline based on AnyAnomaly [17], direct video-input inference with Qwen3-VL-8B, and frame-level image inference at 1 and 2 fps with the same model. We evaluate on two datasets: the Traffic Accident Detection (TAD) benchmark for broad crash recognition, and a purpose-built dataset of real-world footage from Seattle city cameras, comprising 91 accident and 91 normal-traffic clips that more faithfully represent fixed-camera urban traffic monitoring conditions. Our results show that while Qwen3-VL-based approaches outperform the context-aware visual question answering (VQA) pipeline across both datasets, all methods struggle with low crash recall on fixed-camera surveillance footage, where accidents frequently occupy only a small region of the frame. Our results suggest that zero-shot LVLMs are a promising direction for training-free traffic accident detection, although further improvements in accuracy may be needed depending on the requirements of the target application.**

***Index Terms*—video anomaly detection, traffic surveillance, large vision-language models (LVLM), zero-shot learning**

I. Introduction

Traffic surveillance cameras are now widely deployed across urban road networks, forming a core part of modern smart city infrastructure. However, the large volume of video data generated by these cameras makes manual monitoring impractical, creating a need for automated systems that can detect hazardous events in real time. Traffic accident detection is one of the most important such tasks, as it has direct implications for emergency response, urban planning, and road safety [9].

Traditional approaches to video anomaly detection (VAD) rely on learning the distribution of normal behavior and flagging deviations from it [3]. While effective in controlled settings, these methods struggle with the inherent variability of traffic scenes and require large annotated datasets that are expensive to collect. More fundamentally, traffic accidents are rare and highly diverse. As described in [9], they exhibit long-tailed, unbalanced, highly dynamic, and uncertain characteristics, which make them essentially out-of-distribution events. Due to the diversity of possible collision types, weather conditions, road layouts, and camera perspectives, it is practically difficult to train a supervised model that generalizes well across different deployments. For instance, a model trained on highway dashcam footage may perform poorly on a fixed intersection camera, and the reverse situation also applies.



Fig. 1. Example frames from urban traffic surveillance cameras showing car crash events – the type of scenario this work aims to detect automatically.

The recent rise of Large Vision-Language Models (LVLMs) offers a good alternative. By aligning visual features with natural language, these models can be queried in a zero-shot way – instead of training on thousands of accident clips, we can simply ask the model whether a given video segment contains a traffic accident. Early experiments have shown that models such as GPT-4V have strong cognitive and reasoning abilities on typical traffic events, while also revealing some limits in more complex or ambiguous scenarios [4]. More broadly, LVLMs have demonstrated promise for semantic understanding of traffic scenes [1], for map and road-scene comprehension [7], and for safety-critical event recognition [6], suggesting that the wide world knowledge encoded in these models can be transferred meaningfully to the traffic do-

main [2].

Despite the growing literature on large-vision language models (LVLMs), their application to fixed-camera urban traffic surveillance remains largely unexplored, as most existing work focuses on dash-cam or autonomous-driving scenarios. Fixed-infrastructure cameras bring a distinct set of challenges such as wide-angle perspectives, strong vehicle-scale variation, and the need to detect collisions that may occupy only a small region of the frame. It is also not clear which inference strategy best fits this setting – whether a context-aware frame-level analysis, a direct video-understanding approach, or a lightweight per-frame querying method. Fig. 1 shows representative examples of the crash events targeted in this work, illustrating the fixed-camera viewpoint.

In this paper we systematically evaluate three LVLM-based approaches for zero-shot traffic accident detection: the AnyAnomaly [17] context-aware visual question answering pipeline, direct video-input inference with Qwen3-VL-8B, and frame-level image inference at 1 and 2 fps using the same model. We evaluate them on the TAD benchmark [3] as a general crash-recognition baseline and on a purpose-built dataset collected from Seattle city traffic cameras, which more faithfully represents the fixed-camera conditions of real urban monitoring. Our primary contributions are as follows:

- We provide the first systematic comparison of LVLM-based inference strategies for traffic accident detection from surveillance cameras, covering context-aware VQA, video-level, and frame-level approaches.
- We introduce a purpose-built traffic surveillance dataset collected from real Seattle city cameras, offering a more representative evaluation environment than existing dashcam-oriented benchmarks.
- We identify input resolution and the small spatial scale of accidents in fixed-camera footage as key factors that drive low crash recall, highlighting a systematic challenge that is specific to the traffic-surveillance domain.

II. Related Work

A. Traditional Video Anomaly Detection

VAD has historically been framed as an unsupervised one-class classification task, where models learn the distribution of normal behavior and flag deviations from it. Early methods exploited reconstruction errors from autoencoders or GANs, while later frame prediction approaches flagged large deviations from predicted future frames as anomalies [13]. This approach was extended to traffic by predicting future object locations in dashcam video and using prediction inconsistency as an anomaly signal [11]. More recent supervised approaches combine deep spatio-temporal features with optical flow for accident detection [3], while others apply trajectory conflict analysis to fixed-camera intersections [14]. However, a fundamental limitation shared by most of these approaches is the need for large annotated training datasets, which are difficult to collect and do not generalize well across environments, viewpoints, and rare event categories [9], [15].

B. Vision-Language Models for Traffic Understanding

The emergence of LVLMs has enabled a zero-shot scene-understanding capability without task-specific training data. Models such as GPT-4V have shown strong reasoning ability on common traffic events, although they still struggle with visually subtle or ambiguous scenarios [4]. Broader evaluations on traffic benchmarks [1], [7] and safety-critical event recognition [5], [6] confirm that large-scale pretraining transfers well to the traffic domain [2]. In [8], LVLMs were employed for multi-view reasoning of pedestrian-vehicle incidents. However, most of this work is evaluated on dashcam or autonomous driving footage, where collisions are close-range and visually prominent, which is a very different setting from fixed infrastructure cameras.

C. Zero-Shot Anomaly Detection and the Surveillance Gap

AnyAnomaly [17] introduced customizable VAD (C-VAD), which uses context-aware VQA with position and temporal context to detect user-defined anomalies without any training. Although it achieves strong results on benchmarks such as UBnormal [12] and UCF-Crime [18], these datasets are mostly pedestrian-centric or synthetic and do not reflect the vehicle-dominated, wide-angle conditions of real urban traffic cameras. Looking at both the traditional VAD and LVLM literature, a consistent gap emerges: fixed-camera traffic surveillance – where vehicles are small and distant, viewpoints are oblique, and there are no ego-motion cues – remains largely unaddressed. This paper directly targets that gap.

III. Methodology

A. Problem Definition

Given a video clip V comprising T frames recorded by a dashboard or a fixed traffic surveillance camera, the goal is to produce a binary prediction $\hat{y} \in \{0, 1\}$ indicating whether the clip contains a traffic accident ($\hat{y} = 1$) or depicts normal traffic ($\hat{y} = 0$). All three inference strategies operate in a fully zero-shot manner, requiring no task-specific training data. To ensure a fair comparison of inference strategies rather than model capacity, all three approaches use Qwen3-VL-8B-Instruct [16] as the shared backbone. Fig. 2 provides an overview of the three zero-shot inference strategies evaluated in this work.

B. Method 1: AnyAnomaly

AnyAnomaly [17] implements customizable video anomaly detection (C-VAD) by decomposing a video into fixed-length segments of 24 frames and scoring each segment against a user-defined anomaly keyword using an LVLM. We use Qwen3-VL-8B-Instruct as the backbone and retain all original preprocessing steps, including a fixed resize to 224×224 pixels. For each segment, a representative key frame is selected via image–text similarity and queried with a prompt that incorporates spatial and temporal context, asking the model to rate the presence of the anomaly keyword on a continuous

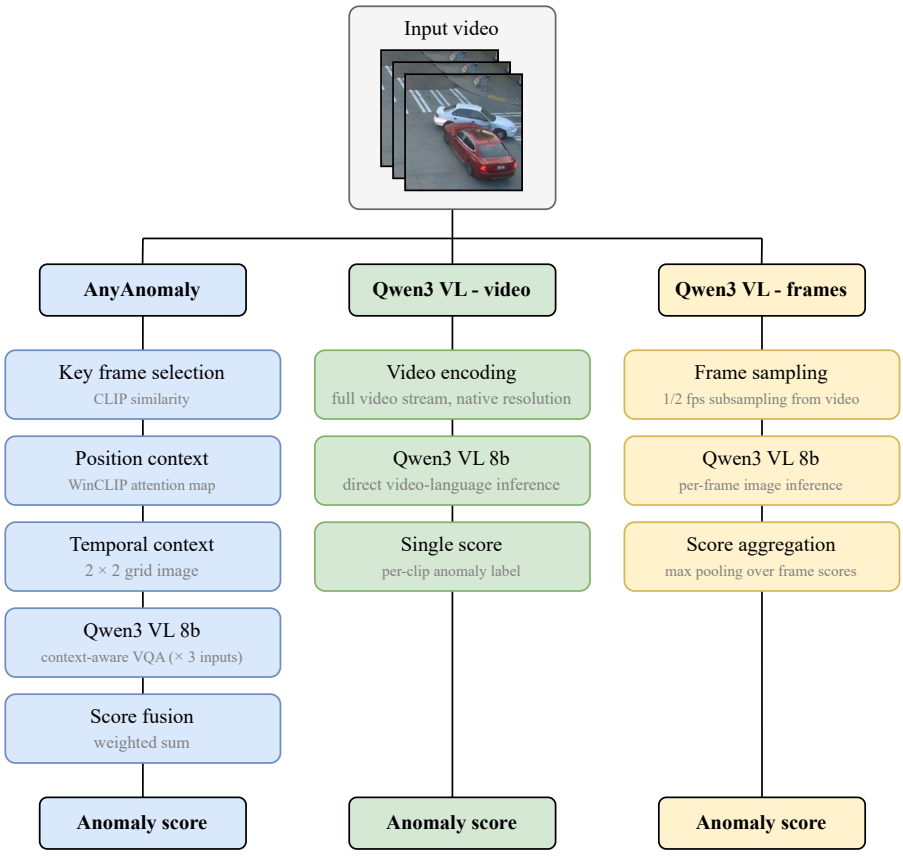

Fig. 2. Overview of the three zero-shot LVLM inference strategies evaluated in this paper.

scale from 0 to 1. We calculate the clip-level prediction by thresholding the maximum segment score:

$$\hat{y} = \begin{cases} 1 & \text{if } \max_i s_i > 0.5 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where s_i is the predicted anomaly score of the i-th segment, representing the model's confidence that the segment contains anomalous activity.

C. Method 2: Qwen3-VL Direct Video Inference

The full video clip is passed directly to Qwen3-VL-8B-Instruct as a single video input, without any manual frame selection or segmentation, and processed through the default preprocessing pipeline provided by the *Qwen3VLProcessor* in the Hugging Face *transformers* library, which samples frames at 2 FPS and preserves near-native resolution. The model is prompted with a minimal binary instruction: *Does this video show a car crash? Respond with yes or no.* A *yes* response yields $\hat{y} = 1$ and *no* yields $\hat{y} = 0$.

D. Method 3: Qwen3-VL Frame-Level Inference

Individual frames are sampled from the clip at a fixed rate and processed independently using the same *Qwen3VLProcessor* preprocessor as in the previous video inference approach. We evaluate two sampling rates: 1 fps and 2 fps. Each frame is queried with: *Does this image show a car crash? Respond with yes or no.* An early-exit strategy is applied – as soon as any frame returns *yes*, the clip is classified as a crash and remaining frames are skipped. If all frames return *no*, the clip is classified as normal.

IV. Experiment

A. Datasets

TAD: The Traffic Accident Detection benchmark [9] comprises crash and normal-traffic clips sourced predominantly from dashcam recordings. We use it as a general-purpose baseline representative of conditions under which many prior work has been evaluated. For our evaluation, we created a subset of 349 normal and 355 crash clips.

Seattle Traffic Surveillance Dataset: We constructed a purpose-built dataset from publicly available footage recorded by Seattle city traffic cameras, comprising 91 crash clips and 91 normal-traffic clips. Unlike TAD, all footage originates from fixed overhead or roadside cameras. Accidents may often appear as small regions within high-resolution frames, making the dataset substantially more challenging than TAD. All clips in the dataset are captured in the Full HD resolution.

TABLE I
CLASSIFICATION RESULTS ON THE TAD DATASET.

	Normal			Crash			
Method	**Precision**	**Recall**	**F1-Score**	**Precision**	**Recall**	**F1-Score**	**Accuracy**
AnyAnomaly	0.58	0.93	0.71	0.83	0.33	0.47	0.63
Qwen3-VL Video	0.73	0.89	0.80	0.86	0.68	0.76	0.78
Qwen3-VL Frames (1 fps)	0.73	0.87	0.79	0.84	0.68	0.75	0.77
Qwen3-VL Frames (2 fps)	0.79	0.83	0.81	0.82	0.78	0.80	0.80

TABLE II
CLASSIFICATION RESULTS ON THE SEATTLE TRAFFIC SURVEILLANCE DATASET.

	Normal			Crash			
Method	**Precision**	**Recall**	**F1-Score**	**Precision**	**Recall**	**F1-Score**	**Accuracy**
AnyAnomaly	0.54	0.95	0.69	0.77	0.19	0.30	0.57
Qwen3-VL Video	0.64	1.00	0.78	1.00	0.44	0.61	0.72
Qwen3-VL Frames (1 fps)	0.66	0.93	0.77	0.89	0.52	0.65	0.73
Qwen3-VL Frames (2 fps)	0.67	0.91	0.77	0.86	0.55	0.67	0.73

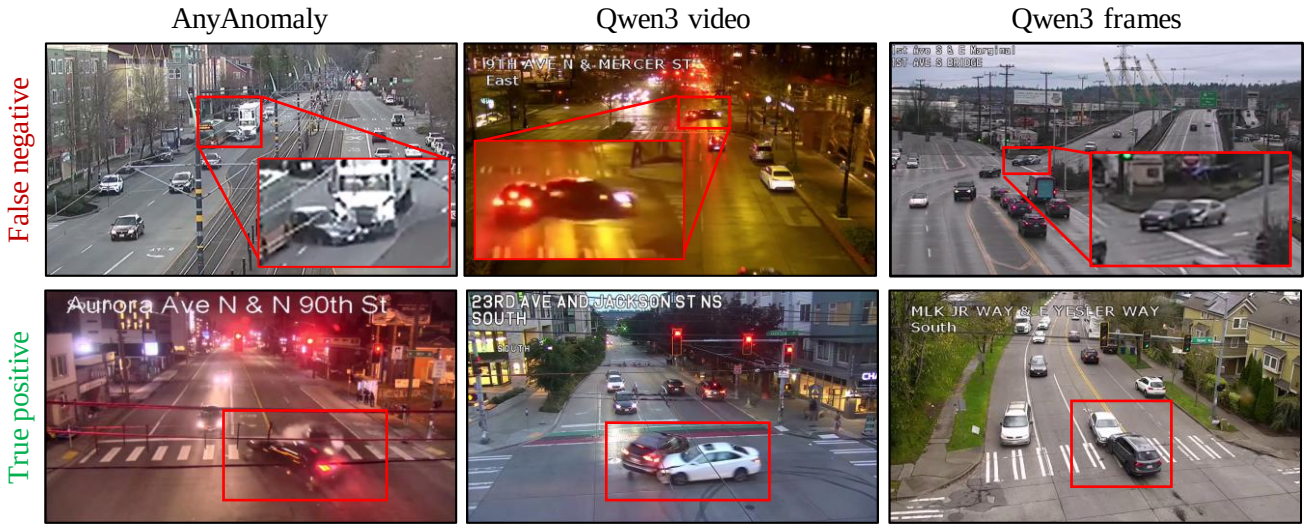

Fig. 3. Qualitative results for each evaluated method. The top row shows false negative examples – crash clips incorrectly classified as normal traffic. The bottom row shows true positive examples – crashes correctly identified by each method.

B. Evaluation Protocol

All methods produce a binary per-clip prediction. We report precision, recall and F1-score for both classes at the clip level. Accuracy is also reported to evaluate the overall performance across both classes.

V. RESULTS AND DISCUSSION

Overall performance: On TAD (see Table I), Qwen3-VL Frames at 2 fps achieves the highest accuracy (0.80), followed closely by Qwen3-VL Video (0.78) and Frames at 1 fps (0.77), with AnyAnomaly trailing at 0.63. On the Seattle dataset (see Table II), all methods show a marked performance drop, with accuracies ranging from 0.57 to 0.73. Fig. 3 illustrates representative true positive and false negative predictions for each method, highlighting the difficulty of detecting crashes that occupy a small spatial region in fixed-camera surveillance footage.

Precision–recall asymmetry on Seattle dataset: Both Qwen3-VL approaches exhibit high normal recall (0.91–1.00) but low crash recall (0.44–0.55) on Seattle, indicating a systematic bias toward predicting normal traffic. We attribute this to the scale difference between datasets: while dashcam crashes typically occupy a large portion of the frame, surveillance footage may show accidents on small pixel area, at which scale visual crash indicators become imperceptible and the model defaults to a normal prediction. AnyAnomaly exhibits the most extreme version of this pattern, achieving near-perfect normal recall (0.95) while detecting only 19% of crashes, resulting in the lowest overall accuracy (0.57).

Frame-level vs. video-level inference: On TAD, frame-level inference at 2 fps outperforms direct video inference (0.80 vs. 0.78), suggesting that processing individual frames independently can be more effective than holistic video understanding for dashcam crash detection. On Seattle however, both sampling rates yield identical accuracy (0.73), and the

difference between 1 fps and 2 fps is negligible across both datasets, indicating diminishing returns from higher temporal sampling.

AnyAnomaly and the resolution gap: AnyAnomaly substantially underperforms the Qwen3-VL approaches on both datasets, and the gap is most pronounced on Seattle (0.57 vs. 0.73). A key confounding factor is input resolution: AnyAnomaly resizes all frames to 224×224 pixels, whereas the Qwen3-VL approaches operate at near-native resolution via the default *Qwen3VLProcessor* preprocessing. This compression disproportionately affects fixed-camera surveillance footage where accidents already occupy a small spatial extent, effectively destroying the visual evidence needed for detection. We emphasize that this reflects a practical limitation of the current AnyAnomaly implementation rather than an inherent weakness of context-aware VQA as a detection strategy

VI. Conclusion

We evaluated three zero-shot LVLM-based inference strategies for traffic accident detection on both a dashcam benchmark (TAD) and a purpose-built fixed-camera surveillance dataset collected from Seattle city traffic cameras. Frame-level inference with Qwen3-VL-8B-Instruct at 2 fps achieved the best overall performance on TAD, while all methods showed substantially reduced recall on the Seattle dataset, which we attribute to the small spatial scale of accidents in overhead surveillance footage. AnyAnomaly's considerably weaker results on both datasets, particularly on Seattle, highlight input resolution as a critical practical factor.

Overall, our results confirm that zero-shot LVLMs are a viable approach to training-free accident detection. Future work should explore higher-resolution preprocessing for context-aware VQA pipelines, surveillance-specific prompting strategies, and larger fixed-camera benchmarks to better characterize and close this gap.

Acknowledgment

This work was supported by the following project: CZ.01.01.01/01/24_063/0006315

References

[1] S. Jain, S. Thapa, K. -T. Chen, A. L. Abbott and A. Sarkar, "Semantic Understanding of Traffic Scenes with Large Vision Language Models," 2024 IEEE Intelligent Vehicles Symposium (IV), Jeju Island, Korea, Republic of, 2024, pp. 1580-1587, doi: 10.1109/IV55156.2024.10588373.

[2] X. Zhang, "Advancing ITS Applications with LLMs: A Survey on Traffic Management, Transportation Safety, and Autonomous Driving," in Rough Sets, 2024, pp. 295–309.

[3] V. A. Adewopo and N. Elsayed, "Smart City Transportation: Deep Learning Ensemble Approach for Traffic Accident Detection," in IEEE Access, vol. 12, pp. 59134-59147, 2024, doi: 10.1109/ACCESS.2024.3387972.

[4] X. Zhou and A. C. Knoll, "GPT-4V as Traffic Assistant: An In-depth Look at Vision Language Model on Complex Traffic Events," arXiv.org, 2024. https://arxiv.org/abs/2402.02205 (accessed Mar. 11, 2026).

[5] HAN, X., et al, "From regional to general: A vision-language model-based framework for corner cases comprehension in autonomous driving," in ECCV 2024 Workshop on Multimodal Perception and Comprehension of Corner Cases in Autonomous Driving, 2024.

[6] Shi, L., et al, "Scvlm: Enhancing vision-language model for safety-critical event understanding," in Proceedings of the Winter Conference on Applications of Computer Vision, 2025, pp. 1061–1071.

[7] Cao, X., et al, "Maplm: A real-world large-scale vision-language benchmark for map and traffic scene understanding," in Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2024, pp. 21819–21830.

[8] H. Zhen, Y. Yang, J. Yang. "Multi-view Phase-aware Pedestrian-Vehicle Incident Reasoning Framework with Vision-Language Models," in arXiv preprint arXiv:2511.14120, 2025.

[9] J. Fang, J. Qiao, J. Xue and Z. Li, "Vision-Based Traffic Accident Detection and Anticipation: A Survey," in IEEE Transactions on Circuits and Systems for Video Technology, vol. 34, no. 4, pp. 1983-1999, April 2024, doi: 10.1109/TCSVT.2023.3307655.

[10] N. Dogru and A. Subasi, "Traffic accident detection using random forest classifier," 2018 15th Learning and Technology Conference (L&T), Jeddah, Saudi Arabia, 2018, pp. 40-45, doi: 10.1109/LT.2018.8368509.

[11] Y. Yao, M. Xu, Y. Wang, D. J. Crandall and E. M. Atkins, "Unsupervised Traffic Accident Detection in First-Person Videos," 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Macau, China, 2019, pp. 273-280, doi: 10.1109/IROS40897.2019.8967556.

[12] Acsintoae, A., et al, "Ubnormal: New benchmark for supervised open-set video anomaly detection," in Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2022, pp. 20143–20153.

[13] Liu, W., W. Luo, D. & Gao, S. Future Frame Prediction for Anomaly Detection – A New Baseline. *2018 IEEE Conference On Computer Vision And Pattern Recognition (CVPR)*. (2018)

[14] H. Ghahremannezhad, H. Shi and C. Liu, "Real-Time Accident Detection in Traffic Surveillance Using Deep Learning," 2022 IEEE International Conference on Imaging Systems and Techniques (IST), Kaohsiung, Taiwan, 2022, pp. 1-6, doi: 10.1109/IST55454.2022.9827736.

[15] U. Khalil, T. Javid and A. Nasir, "Automatic road accident detection techniques: A brief survey," 2017 International Symposium on Wireless Systems and Networks (ISWSN), Lahore, Pakistan, 2017, pp. 1-6, doi: 10.1109/ISWSN.2017.8250025.

[16] Bai, S., et al. "Qwen3-vl technical report," in arXiv preprint arXiv:2511.21631, 2025.

[17] Ahn, S., et al, "Anyanomaly: Zero-shot customizable video anomaly detection with lvlm," in Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, 2026, pp. 3026–3035.

[18] W. Sultani, C. Chen, M. Shah, "Real-world anomaly detection in surveillance videos," in Proceedings of the IEEE conference on computer vision and pattern recognition, 2018, pp. 6479–6488.

Orpheus XAI: Autonomous Navigation and Immersive VR Telepresence for a Military-Grade Unmanned Ground Vehicle

1st Miloš Cihlář
FEEC, Dept. of Control Instrumentation
CEITEC – Central European Institute of Technology
Brno University of Technology
Brno, Czech Republic
Milos.Cihlar@vut.cz

2nd Petr Šopák
FEEC, Dept. of Control Instrumentation
CEITEC – Central European Institute of Technology
Brno University of Technology
Brno, Czech Republic
Petr.Sopak@vut.cz

3rd Stanislav Svědiroh
FEEC, Dept. of Control Instrumentation
CEITEC – Central European Institute of Technology
Brno University of Technology
Brno, Czech Republic
Stanislav.Svediroh@vut.cz

4th Tomáš Horeličan
FEEC, Dept. of Control Instrumentation
Brno University of Technology
Brno, Czech Republic
Tomas.Horelican@vut.cz

***Abstract*—This paper presents Orpheus XAI, an extension of the military-grade unmanned ground vehicle Orpheus AC3 designed to support both autonomous navigation and immersive telepresence. The system combines a ROS 2-based autonomy layer with a virtual reality interface developed using the OpenXR standard. The navigation subsystem provides GNSS/RTK-based localization, terrain-aware planning based on a 2.5D elevation map, and several navigation modes defining different levels of trajectory compliance. Missions are defined as waypoint sequences with optional task actions executed at selected locations.**

The telepresence subsystem enables the operator to perceive the robot environment through a VR headset with head-tracked camera control and multi-spectral visualization combining RGB and thermal imagery. Because consumer VR hardware is not designed for outdoor operation, several fail-safe mechanisms were implemented, including tracking loss detection, headset presence monitoring, and automatic fallback to a conventional control interface.

***Index Terms*—UGV, telepresence, virtual reality, autonomous navigation**



Fig. 1: Orpheus XAI is an enhanced version of the Orpheus AC3 UGV equipped with sensors enabling autonomous mis-sion execution.

I. Introduction

Unmanned ground vehicles (UGVs) are increasingly deployed in military reconnaissance, search-and-rescue, and security patrol scenarios [1], [2]. Two interrelated challenges stand at the centre of current UGV research: (i) increasing the level of autonomy while preserving predictable, fail-safe behaviour, and (ii) improving the quality of telepresence so that remote operators can perceive the environment as naturally as possible. Although both topics have been studied extensively in isolation, their joint integration into a single fielded system remains an open challenge.

On the navigation side, the ROS 2 Navigation2 framework [3] provides a widely adopted foundation, yet outdoor operation demands robust multi-sensor localization (e.g. GNSS/LiDAR fusion) [4] and flexible mission execution, for which behavior trees have recently gained traction [5]. On the telepresence side, virtual reality head-mounted displays (HMDs) offer intuitive head-tracked camera control and improved situational awareness [7], [8]. Several systems have demonstrated VR-based robot teleoperation, from military UGVs [1] to mobile platforms with shared obstacle avoidance [6] and immersive 3D scene reconstruction [10], but these operate predominantly indoors and do not address the fail-safe concerns specific to outdoor deployment, where consumer VR headsets suffer from inside-out tracking degradation in direct sunlight and featureless terrain [12], end-to-end latency must stay below approximately 100 ms for acceptable operator experience [11], and multi-spectral visualization (RGB + thermal fusion) within VR has received little attention.

This paper addresses two questions: (1) Can consumer-grade VR headsets serve as a reliable outdoor telepresence

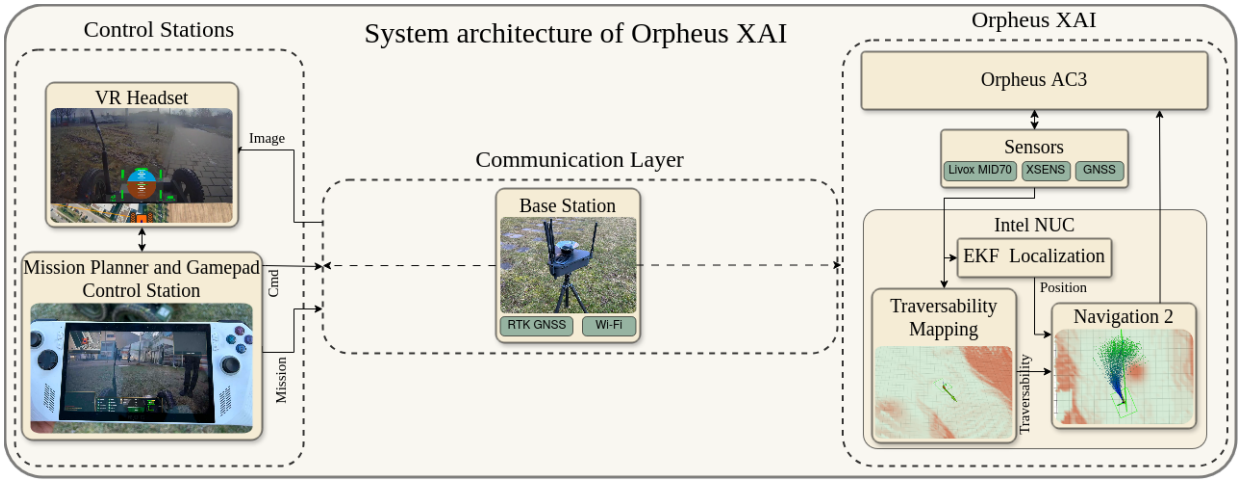

Fig. 2: System architecture of Orpheus XAI. The low-level platform (Orpheus AC3) is extended by a ROS 2 autonomy layer and a VR telepresence module. The VR headset communicates with the Gamepad Control Station, which transmits commands and mission data to the robot over WiFi.

interface for UGVs, and what fail-safe mechanisms are required? (2) How can autonomous navigation and immersive telepresence be integrated so that the operator transitions seamlessly between modes? We present **Orpheus XAI**, a demonstrator developed within the IMMAA project, extending the military-grade Orpheus AC3 platform (LTR s.r.o.) with a ROS 2-based autonomous navigation layer (GNSS/RTK localization, configurable navigation modes, AI-powered waypoint actions, web-based mission planner) and an OpenXR-based VR telepresence system (multi-spectral visualization, PID-controlled head tracking, gaze-based interaction, Wi-Fi-adaptive streaming, and fail-safe fallback to non-immersive control). The integrated system was validated outdoors at TRL 6.

II. System Overview

The Orpheus XAI system extends the capabilities of the production robotic platform Orpheus AC3. While the original platform is designed primarily for teleoperation, the proposed extension introduces two additional capabilities: autonomous mission execution and immersive virtual telepresence.

The motivation for this combination is practical. Autonomous navigation allows the robot to perform routine tasks without continuous operator attention. However, fully autonomous operation cannot guarantee reliable behaviour in all environments. For this reason, the system must allow the human operator to intuitively take control at any moment. The operator may also remain only an observer and supervise the execution of the autonomous mission.

The overall concept is illustrated in Fig. 2. The system provides three main operational capabilities. First, it enables autonomous navigation supported by a human-friendly mission planner. The robot can execute missions defined as a sequence of waypoints while performing obstacle avoidance and terrain-aware motion planning. Depending on the selected operating mode, the robot may either actively avoid obstacles or stop safely when the planned path becomes blocked.

Second, the system supports immersive teleoperation. In this mode, the operator controls the robot through a virtual telepresence interface using a VR headset. The surrounding environment is presented through live video streams from the robot sensors, providing situational awareness and enabling intuitive intervention when autonomous navigation encounters situations it cannot reliably resolve.

Third, the system provides an operator control interface. The operator can monitor the robot state through a graphical interface displaying system status, telemetry, and mission progress.

From a system perspective, the architecture consists of four main physical components:

1) **Mobile robot platform** – the Orpheus AC3 chassis equipped with a telepresence head consisting of a pan–tilt camera unit with RGB and LWIR sensors. The perception stack is further extended by a 3D LiDAR (Livox MID-70), an inertial measurement unit (Xsens MTi-680G), and GNSS receivers integrated with the IMU providing RTK correction support. The onboard computer runs ROS 2 nodes responsible for localization, terrain representation, navigation, and action execution.
2) **Operator control station** – a ruggedized Winmate G101AD-A (Fig. 3) tablet equipped with dual joysticks. The device runs the operator graphical interface and a web-based mission planner. Communication with the robot is established via Wi-Fi or through another existing network infrastructure.
3) **VR headset** – an OpenXR-compatible headset (currently Meta Quest 3) running a Unity-based immersive telepresence application. The headset receives RTSP video streams via GStreamer and exchanges JSON-based telemetry and control messages with the Gamepad



Fig. 3: The image shows a ruggedized Winmate G101AD-A tablet displaying the planning interface GUI.

Control Station (GCS) over UDP. Both the headset type and the video sources are configurable through an external XML configuration file, which allows deployment with different hardware without recompilation.

4) **Base station** – a communication and positioning infrastructure providing RTK correction data and long-range connectivity between the operator and the robotic platform.

The system architecture intentionally separates the certified low-level control layer of the Orpheus AC3 platform from the ROS 2-based autonomy layer. If the autonomy layer fails, the operator can immediately fall back to direct manual control. This design satisfies the fail-safe requirement and preserves the operational reliability of the base platform.

III. Mission Planning

To enable practical deployment of the robot in real-world environments, a mission planning system was developed. The main objective is to allow mission creation and execution by non-expert operators. The system therefore focuses on simplicity of interaction while preserving the ability to define structured autonomous tasks. The mission planner is integrated into the Gamepad Control Station (see Fig. 3). The interface is implemented as a web application, which allows the system to be executed both on desktop computers and on mobile devices. The mission planning system was developed as part of the IMMAA project.

The planner allows the operator to create, edit, store, and execute missions. At the same time, the system provides a visualization of the mission progress. During execution, the current robot position and orientation are displayed directly in the map interface. Mission definition is based on a sequence of waypoints. For each waypoint it is possible to define additional actions that the robot should perform. These actions extend the mission from pure navigation to task execution. The actions allow the operator to specify robot behavior at each waypoint. A variety of actions can be implemented on the robot. For the IMMAA project, the implemented actions include *takePhoto* for image capture, *start* and *stop* video recording, and *wait*. Additionally, AI-based actions were implemented, including object detection (typically people, vehicles, or other prominent objects in the camera view) and text recognition (e.g., reading vehicle plates). The system also provides an interface to an LLM server via the Sonar-pro API from Perplexity, enabling the operator to define custom actions that process input from the robot's camera.

The output of the planner is a mission specification stored in the YAML format. This representation provides a human-readable description of the mission. Before execution, the mission is transformed into a custom ROS 2 message, which is used by the robot navigation stack. The message consists of a header followed by a list of waypoints. The header contains parameters defining the mission execution. It includes the number of mission repetitions, allowing the robot to execute the same mission multiple times. A typical use case is patrolling or guarding a specified area. The header further contains the selected navigation mode, which is described in more detail in the following section IV.

The resulting ROS 2 message is then passed to the robot navigation system, which is responsible for executing the mission.

IV. Autonomous Navigation

When a mission execution request is received, the robot takes control over the drive system and starts to autonomously execute the mission defined by the user. The navigation pipeline is divided into three main categories: state estimation, terrain estimation and representation, and motion planning and control.

The mission contains a sequence of waypoints defined in the WGS84 coordinate system. The robotic platform Orpheus XAI is equipped with a single GNSS receiver. From a single GNSS measurement it is not possible to determine the robot heading. Therefore the robot must be calibrated after each system start. The calibration is performed by moving the robot forward for several meters. The displacement vector obtained from this motion provides an initial estimate of the robot heading. After this initial calibration the estimation of position and orientation is handled by an Extended Kalman Filter (EKF). The EKF fuses high-frequency measurements from the inertial measurement unit, wheel odometry, and GNSS measurements. Waypoints defined by GNSS coordinates are transformed into a local coordinate frame defined at the robot start position.

The advantage of this approach is robustness to temporary GNSS signal degradation. In outdoor environments the GNSS signal is often partially blocked by trees, buildings, or terrain structures. In such cases the robot continues operation using fused inertial and odometry measurements, effectively performing short-term dead-reckoning. In practice the robot can travel tens of meters without significant loss of localization accuracy. The localization drift during GNSS absence depends strongly on terrain properties and wheel slip conditions. Detailed analysis of this problem is outside the scope of this paper.

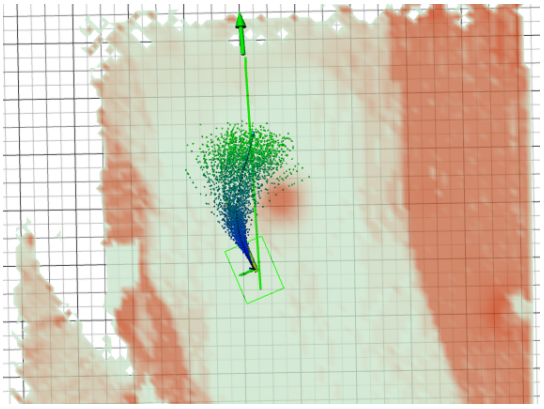

Fig. 4: The figure shows a visualization of the elevation map, representing terrain traversability. It also illustrates the operation of the Nav2 controller using Model Predictive Path Integral (MPPI), which dynamically avoids obstacles detected along the global path plan.

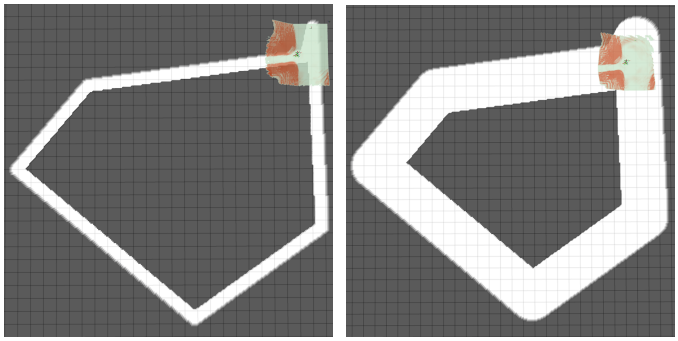

Fig. 5: The figures illustrate two modes of the global planner. The strict mode, shown on the left, creates a corridor of 1m width around the defined waypoints. The constrained mode, shown on the right, generates a corridor of 5m width. All three modes, strict, constrained, and free differ in the corridor width, which can be adjusted as needed.

The system extending the certified robotic platform Orpheus AC3 is based on the Ubuntu 22.04 operating system and uses the ROS 2 middleware. The use of these technologies allows integration of established open-source projects widely used in the robotics community. One of these projects is the Navigation2 [3] framework which implements navigation for mobile robots in a two-dimensional environment. Our system extends this concept to outdoor operation by introducing a 2.5D elevation grid map [16]. This representation stores terrain elevation and information relevant for terrain traversability.

From the elevation map a traversability estimate is computed. The primary metric is derived from analysis of terrain surface normals which indicate the local terrain slope. These traversability values are subsequently transformed into a two-dimensional occupancy cost map used by the Navigation2 planning stack. The resulting cost map is then used for trajectory planning and motion control.

The navigation system supports three operation modes (Fig. 5): strict, constrained, and free. These modes define the allowed freedom of robot motion with respect to the planned trajectory.

1) In the **strict mode** the robot has almost no freedom of motion. The robot must follow the planned trajectory and it is not allowed to leave it. The only available reaction to an obstacle is to stop and wait until the mission timeout is reached.
2) The **constrained mode** allows the robot to avoid obstacles within a limited distance from the planned trajectory. In this case the robot can temporarily deviate from the path while still maintaining the overall mission structure.
3) The **free mode** removes this limitation. The robot can avoid obstacles without restriction and can deviate significantly from the original trajectory if necessary.

The advantage of this approach is that the user can define how strictly the robot must follow the planned trajectory. In autonomous operation this allows the designer to enforce predefined behavior. In difficult environments such restriction may increase operational safety, because the robot is not forced to search for a feasible path in situations where the local planner could generate undesirable motion.

Another safety mechanism is the possibility for the operator to take control over the robot. In such case the robot can be operated remotely using a standard screen interface or a virtual reality headset. This allows the operator to directly resolve situations where autonomous navigation is insufficient or ambiguous. The telepresence system is described in the following section.

V. Immersive VR Telepresence

The VR telepresence system was developed in Unity using the OpenXR API [13], making it compatible with any OpenXR-compliant headset (validated with Meta Quest 3). Its goal is to provide the operator with a natural, spatially consistent perception of the robot's environment while maintaining immediate fallback to conventional (non-immersive) control. The application communicates with the GCS via UDP, receiving JSON-encoded status messages (telemetry, sensor data, control authority state) and transmitting operator commands back over the same channel.

A. Spatial GUI Layout and Interaction

The VR interface supports two viewing modes. In the *world-locked* mode, UI panels are anchored in the virtual space around the operator: the live camera feed is placed in the direction corresponding to the actual orientation of the telepresence head relative to the robot chassis (Fig. 6), control elements are accessible by looking left or right, and a real-time map is visible when looking down. In the *head-locked* mode, only the camera feed and essential robot status indicators follow the operator's head, which is useful during rapid camera repositioning when the full set of surrounding panels is not

needed. The operator switches between these modes via a double-click gesture on the operator station button.

All in-VR interaction is based on gaze-directed raycasting: a crosshair follows the operator's gaze across the UI canvases, and selection is triggered by a physical button on the operator station. A state machine distinguishes single click, double click, and long press gestures to map different actions to the single available button.



Fig. 6: The VR telepresence interface. The left and right panels display robot parameters, system information, and available control commands.

B. Head-Tracked Camera Control

When activated, the operator's head rotation is mapped to the pan and tilt of the robot's telepresence head via a PID controller running independently for both the yaw and pitch axes. The controller computes the angular error between the desired head orientation (derived from the HMD rotation) and the current sensor-reported joint angles, and outputs a speed command in the range $[-1, 1]$ transmitted to the robot at each frame. The PID gains are configurable via an external XML file, allowing field-tuning without recompilation.

The robot chassis is always controlled via the physical joysticks on the operator station - only the camera orientation follows the head. This ensures that robot motion control remains available even if the VR headset experiences tracking issues in outdoor conditions. A long-press gesture temporarily disengages head tracking, allowing the operator to reorient without moving the robot's camera.

C. Multi-Spectral Video Streaming

The video feed is received via RTSP using a GStreamer pipeline [14] integrated into Unity. The system supports three visualization modes (Fig. 7): standard RGB, thermal (LWIR), and a fused overlay computed by a custom GPU shader. The fusion shader applies a calibrated UV transformation (scale, offset, and radial distortion correction) to align the thermal image with the RGB feed before blending them with a configurable transparency ratio.

To adapt to varying network conditions, the system implements Wi-Fi signal strength-based adaptive streaming. A background monitor polls the headset's Wi-Fi RSSI and link speed, and maps the result to one of five quality profiles (Ultra through Minimal), each associated with a different RTSP URL providing a different resolution/bitrate encode of the same camera feed. Profile switching uses a minimum hold time to prevent oscillation.

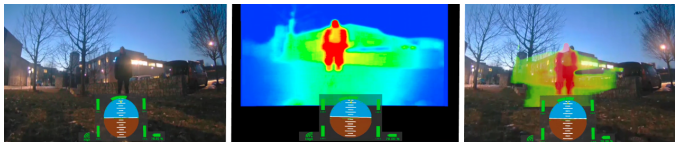

Fig. 7: Multi-spectral visualization modes: RGB (left), thermal LWIR (center), and fused RGB + thermal (right).

D. Fail-Safe Mechanisms

Commercial VR headsets, designed primarily for indoor use, can experience tracking failures in outdoor environments. The system addresses this at multiple levels:

1) **Tracking loss detection** – the application continuously monitors the XR tracking state. When inside-out tracking is lost, head-tracked camera control is automatically released (the robot retains its last commanded orientation) and a notification overlay is displayed.
2) **Headset presence detection** – when the operator removes the headset, camera and UI control authority is released, and the operator seamlessly continues on the station's display.
3) **Adaptive streaming quality** – as described above, degrading Wi-Fi conditions trigger automatic bitrate reduction, maintaining a smooth frame rate at the expense of image quality.

This approach reflects a design philosophy favoring affordable, expendable equipment with graceful degradation over expensive military-grade VR systems.

E. Mission Planner Integration

The web-based mission planner (Section III) is also accessible from within the VR headset via an embedded WebView [15] component (Fig. 8). The planner runs as a standard web page; the only difference is that mouse input is replaced by the gaze-based cursor, allowing the operator to interact with the map, define waypoints, and launch missions without removing the headset. Navigation commands are confirmed via a double-click gesture and transmitted to the robot's ROS 2 navigation module.

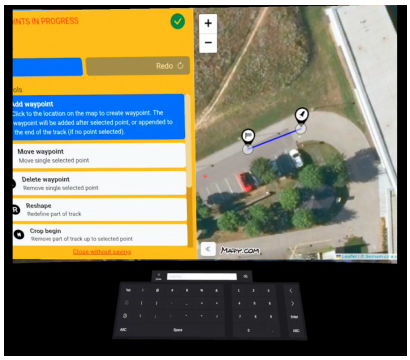

Fig. 8: The mission planner as viewed from the VR headset.

VI. Results

The Orpheus XAI system was evaluated during a series of outdoor field experiments conducted under varying environmental conditions, including winter operation. The objective of the experiments was to verify the functional capabilities of the integrated system at Technology Readiness Level 6.

All primary system functions were successfully demonstrated. The robot was able to execute autonomous missions defined as sequences of GNSS waypoints while operating in all three navigation modes: strict, constrained, and free. The navigation stack reliably performed terrain-aware planning and obstacle avoidance based on the elevation map representation.

The VR telepresence interface enabled the operator to observe the robot environment and control the telepresence head through head-tracked motion. The system supported RGB, thermal, and fused multi-spectral visualization. During testing, the implemented fail-safe mechanisms behaved as intended. When tracking loss or headset removal occurred, control of the telepresence head was released and the operator continued operation through the control station without interruption.

The only requirement that was not fully achieved was the end-to-end video latency, which in a telepresence system effectively constitutes the operator's motion-to-photon delay. The observed latency exceeded the ideal 60 ms threshold [17] but remained within the range reported as acceptable for telepresence in [11], where 121 ms did not significantly degrade operator performance. This is consistent with the operator experience observed during our field trials.

VII. Conclusion

This paper presented Orpheus XAI, an extension of the military-grade Orpheus AC3 unmanned ground vehicle integrating autonomous navigation with immersive VR telepresence. The system combines a ROS 2-based autonomy layer with a VR interface built on the OpenXR standard. Field trials confirmed that a consumer VR headset can serve as a practical outdoor telepresence interface when combined with tracking loss detection, headset presence monitoring, and adaptive streaming. The decision to exclude VR controllers and retain physical joysticks for robot steering proved essential for reliable outdoor operation. The shared mission planner, accessible both from the control station and from within VR, enables seamless transitions between autonomous and manual operation. Future work will focus on reducing the video latency through hardware-accelerated encoding and on conducting a formal user study to evaluate operator experience.

Acknowledgment

The completion of this paper was made possible by the grant No. FEKT-S-26-8988 - "Advanced Methods in Cybernetics, Robotics, and Artificial Intelligence" financially supported by the Internal science fund of Brno University of Technology. The work has been performed in the project IMMAA: Imerzivita a automatizace pozemních bezpilotních prostředků (TQ03000946), which is co-funded with state support from the Technology Agency of the Czech Republic under the SIGMA Program. The work was supported by the infrastructure of RICAIP that has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 857306 and from Ministry of Education, Youth and Sports under OP RDE grant agreement No CZ.02.1.01/0.0/0.0/17_043/0010085.

References

[1] T. Kot and P. Novák, "Application of virtual reality in teleoperation of the military mobile robotic system TAROS", *International Journal of Advanced Robotic Systems*, vol. 15, no. 1, Jan. 2018, doi: 10.1177/1729881417751545.

[2] S. Opiyo, J. Zhou, E. Mwangi, W. Kai, and I. Sunusi, "A Review on Teleoperation of Mobile Ground Robots: Architecture and Situation Awareness", *International Journal of Control, Automation and Systems*, vol. 19, no. 3, Oct. 2020, doi: 10.1007/s12555-019-0999-z.

[3] S. Macenski, F. Martin, R. White, and J. G. Clavero, "The Marathon 2: A Navigation System", in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, Oct. 2020, pp. 2718-2725. doi: 10.1109/iros45743.2020.9341207.

[4] J. Lee et al., "ODS-Bot: Mobile Robot Navigation for Outdoor Delivery Services", *IEEE Access*, vol. 10, 2022, doi: 10.1109/access.2022.3212768.

[5] W. Yoon, J. Choi, K. Park, J. Shin, S. Sung, and J. Kim, "ATBT: Adaptive Topological Map-Based Behavior Tree for Quadruped Robots", in *2025 25th International Conference on Control, Automation and Systems (ICCAS)*, IEEE, Nov. 2025, pp. 733-738. doi: 10.23919/iccas66577.2025.11301131.

[6] J. E. Solanes, A. Muñoz, L. Gracia, and J. Tornero, "Virtual Reality-Based Interface for Advanced Assisted Mobile Robot Teleoperation", *Applied Sciences*, 2022, doi: 10.3390/app12126071.

[7] M. Wonsick and T. Padir, "A Systematic Review of Virtual Reality Interfaces for Controlling and Interacting with Robots", *Applied Sciences*, 2020, doi: 10.3390/app10249051.

[8] K. Darvish et al., "Teleoperation of Humanoid Robots: A Survey", *IEEE Transactions on Robotics*, 2023, doi: 10.1109/tro.2023.3236952.

[9] R. Alonso, A. Bonini, D. Reforgiato Recupero, and L. D. Spano, "Exploiting virtual reality and the robot operating system to remote-control a humanoid robot", *Multimedia Tools and Applications*, vol. 81, no. 11, Feb. 2022, doi: 10.1007/s11042-022-12021-z.

[10] K. Li, R. Bacher, S. Schmidt, W. Leemans, and F. Steinicke, "Reality Fusion: Robust Real-time Immersive Mobile Robot Teleoperation with Volumetric Visual Data Fusion", in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, Oct. 2024, pp. 8982-8989. doi: 10.1109/iros58592.2024.10802431.

[11] M. Thavendra and K. Jayarajah, "GazeLight: Understanding the Feasibility of Gaze-based Adaptive Streaming for AR-Mediated Human-Robot Systems", in *2025 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops)*, IEEE, Mar. 2025, pp. 195-200. doi: 10.1109/percomworkshops65533.2025.00064.

[12] A. K. Chesang and D. D. Uyeh, "VR for Situational Awareness in Real-Time Orchard Architecture Assessment", *Sensors*, vol. 25, no. 21, Nov. 2025, doi: 10.3390/s25216788.

[13] Khronos Group, "OpenXR Specification", 2026. [Online]. Available: https://www.khronos.org/openxr/ (Accessed: Mar. 11, 2026)

[14] Mrayy, "mrayGStreamerUnity: GStreamer integration with Unity using a native plugin," *GitHub repository*, 2022. [Online]. Available: https://github.com/mrayy/mrayGStreamerUnity (Accessed: Mar. 11, 2026)

[15] t-34400, "SimpleUnity3DWebView: Simple Unity 3D WebView library for Android platform," *GitHub repository*, 2025. [Online]. Available: https://github.com/t-34400/SimpleUnity3DWebView (Accessed: Mar. 11, 2026)

[16] ANYbotics, "GitHub - ANYbotics/grid_map: Universal grid map library for mobile robotic mapping," *GitHub*, Jan. 06, 2021. https://github.com/ANYbotics/grid_map?tab=readme-ov-file (accessed Mar. 12, 2026).

[17] S. M. LaValle, A. Yershova, M. Katsev, and M. Antonov, "Head tracking for the Oculus Rift", in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2014. doi: 10.1109/icra.2014.6906608.



BRNO UNIVERSITY OF TECHNOLOGY
FACULTY OF ELECTRICAL ENGINEERING AND COMMUNICATION
TECHNICKÁ 3058/10, 616 00 BRNO, CZECH REPUBLIC