

CZECH TECHNICAL UNIVERSITY IN PRAGUE

Faculty of Information Technology



Proceedings of the 13th Prague Embedded Systems Workshop

June 26 - 28, 2025

Praha-Vinoř

Czech Republic

© Czech Technical University in Prague, 2025

ISBN 978-80-01-07446-6

Editors:

prof. Ing. Hana Kubátová, CSc.

doc. Ing. Petr Fišer, Ph.D.

Ing. Jaroslav Borecký, Ph.D.

Message from the Program Chairs

The Prague Embedded Systems Workshop is a research meeting intended to present and discuss students' results and progress in all aspects of embedded systems design, testing, and applications. It is organized by members of the Department of Digital Design at the Faculty of Information Technology and supported by the Czech Technical University in Prague. The workshop is focused mainly on new technologies and methods, dependable and low-power design, embedded security, network monitoring and measurements, and algorithms and methods for anomaly detection, including artificial intelligence (AI), machine learning (ML), and neural networks (NN). The workshop aims to enhance collaboration between universities, not just within the EU. It will be based on oral presentations, mutual communication, and discussions.

There are three types of students' submissions and presentations:

- Full papers describing the students' original research. These papers were submitted to a standard reviewing process.
- Abstracts of authors' earlier published and successfully presented papers (at conferences, in journals, etc.). These contributions were not reviewed; emphasis was put on the presentation and discussion.
- Student posters submitted as abstracts of defended Bc. and MSc. theses with subsequent poster presentation. This workshop session is traditionally organized as a contest sponsored by IEEE and industry.

Ten papers were submitted for the oral presentation, from which one full paper and nine abstracts were accepted. This year, contributions from Slovenia (Jožef Stefan Institute) and Czech Republic (Masaryk University, Brno University of Technology, and Czech Technical University in Prague) will be presented.

The technical program is also highlighted by two keynote speakers:

- Toward Adaptive Embedded Systems: from Multi-Objective Design to Runtime Adaptation
Speaker: Marcello Traiola, Ph.D. (Inria centre at Rennes University, France)
- Automatic Analog IC Design Optimization
Speaker: Poki Chen (National Taiwan University of Science and Technology, Taiwan)

PESW 2025 program contains two Industrial talks as examples of current applied research in practice:

- MEMS sensors for edge AI processing
Speaker: Petr Štukjunger (STMicroelectronics, Czech Rep.)
- Optimizing and Deploying Neural Networks to Mobile Devices
Speaker: Pavel Macenauer (NXP Semiconductors, Czech Rep.)

We wish the 13th Prague Embedded Systems Workshop many heated discussions and possible establishment of mutual research cooperation.

Hana Kubátová and Petr Fišer
26th June 2025

Acknowledgment

The 13th Prague Embedded Systems Workshop was supported by SVK 61/25/F8 and SGS23/208/OHK3/3T/18.

The traditional IEEE contest - Poster session for the best diploma and bachelor's theses is, as every year, supported by sponsoring companies.

We would like to thank to our sponsors CTU in Prague, ASICentrum, NXP Semiconductors, Tropic Square, CESNET, daiteq, IMA, METIO Software, STMicroelectronics, SYSGO, UJP PRAHA.

Special thanks go to IEEE: IEEE Student Branch at Czech Technical University in Prague and IEEE Young Professionals, organizing student contest, and Czechoslovakia Section of IEEE.

CTU Poster abstracts were created with the support of the Czech Technical University (CTU) grant No. SGS23/208/OHK3/3T/18, by the Student Summer Research Program 2024 of FIT CTU in Prague, and the grant VJ02010010 of the Ministry of the Interior of the Czech Republic, "Tools for AI-enhanced Security Verification of Cryptographic Devices" in the program Impakt1 (2022-2025).

Committees

Workshop Chairs

Hana Kubátová, CTU in Prague (CZ)

Petr Fišer, CTU in Prague (CZ)

Programme Committee

A. Bernasconi, Università di Pisa (IT)

J. Bělohoubek, CTU in Prague, FIT (CZ)

A. Bosio, École Centrale de Lyon (FR)

L. Cassano, Politecnico di Milano (IT)

G. DiNatale, TIMA, Grenoble (FR)

V. Janíček, CTU in Prague, FEE (CZ)

K. Jelemenská, STU Bratislava (SK)

M. Krstić, IHP, Frankfurt (Oder) (DE)

R. Kvaček, ASICentrum, Prague (CZ)

F. Leporati, Univ. di Pavia (GR)

I. Levin, Tel-Aviv University (IL)

A. McEwan, University of Leicester (UK)

N. Mentens, KU Leuven (BE)

J. Mottok, OTH Regensburg (DE)

P. Mróz, University of Zielona Gora (PL)

M. Novotný, CTU in Prague, FIT (CZ)

A. Orailoglu, UC San Diego (USA)

Z. Plíva, TU Liberec (CZ)

M. Poupá, University of West Bohemia, FEE (CZ)

P. Puschner, Vienna University of Technology (AT)

E. Sanchez, Politecnico di Torino (IT)

J. Schmidt, CTU in Prague, FIT (CZ)

M. Skrbek, CTU in Prague, FIT (CZ)

J. Sobotka, CTU in Prague, FEE (CZ)

B. Steinbach, TU Chemnitz (DE)

A. Steininger, Vienna University of Technology (AT)

J. Strnadel, Brno University of Technology, FIT (CZ)

Z. Vašíček, Brno University of Technology, FIT (CZ)

W. Zajac, Jacob of Paradies University (PL)

Special Session on Embedded Security co-Chairs

M. Novotný, CTU in Prague (CZ)

V. Miškovský, CTU in Prague (CZ)

Student Poster Session co-Chairs

Josef Koumar, CTU in Prague (CZ)

Jaroslav Pešek, CTU in Prague (CZ)

Organizing Committee

H. Kubátová, CTU in Prague (CZ)

P. Fišer, CTU in Prague (CZ)

J. Borecký, CTU in Prague (CZ)

Secretariat

R. Kinc, AMCA (CZ)

B. Dufková, AMCA (CZ)

Contents

Keynote 1: Toward Adaptive Embedded Systems: from Multi-Objective Design to Runtime Adaptation	1
Marcello Traiola, Ph.D. (Inria centre at Rennes University, France),	
Keynote 2: Automatic Analog IC Design Optimization	2
Poki Chen (National Taiwan University of Science and Technology, Taiwan),	
Industrial Talk 1: MEMS sensors for edge AI processing	3
Petr Štukjunger (STMicroelectronics, Czech Rep.),	
Industrial Talk 2: Optimizing and Deploying Neural Networks to Mobile Devices	4
Pavel Macenauer (NXP Semiconductors, Czech Rep.),	
Explainable Anomaly Detection in Network Traffic Using LLM	5
Kamil Jeřábek, Josef Koumar, Jiří Setinský and Jaroslav Pešek	
On Evaluation of Data Fusion Methods for Network Traffic Classification	6
Richard Plný and Karel Hynek	
Multi-Agent Multi-Object Tracking with Shared Measurements and Intersecting Fields of View	7
Pavel Kopecký and Kamil Dedecius	
Innovative Methods of Signal Processing for Hydrogen-Filled Proportional Counters	16
Aleš Horna, Ondřej Kolář, Michal Kubíček, Tomáš Czako, Michal Košťál, Věra Mazánková, Jiří Petr, František Cvachovec, Aleš Jančář, Jan Král, Václav Přenosil and Zdeněk Matěj	
chisel4ml: A Tool For Generating Fast Direct Logic Implementations of Deeply Quantized Neural Networks	17
Jure Vreča and Anton Biasizzo	
The Accurate Measurement of Tilt Direction within Small Angles of Tilt	21
Veronika Junasová, Vladimír Levek and Pavel Šteffan	
Gate boosting: Novel gate driver based on an FPGA and reliability estimation	25
Pavel Tomíček and Jaroslav Boušek	
Controllability-Based Circuit Similarity Estimation	26
Michal Žáček and Petr Fišer	
Side-channel analysis of ChaCha20 stream cipher	27
Lukáš Daněk, Vojtěch Miškovský and Matúš Olekšák	
Building a Side-Channel Attack Scheme on SipHash FPGA Implementation	28
Vít Mašek, Vojtěch Miškovský and Matúš Olekšák	
Partners	29
Sponsors	30

Keynotes

Toward Adaptive Embedded Systems: from Multi-Objective Design to Runtime Adaptation

Speaker: **Marcello Traiola, Ph.D. (Inria centre at Rennes University, France),**

Embedded systems operate under tight and often conflicting constraints, such as energy, performance, accuracy, and reliability. Optimizing across these dimensions is complex, especially as systems must now operate under dynamic and unpredictable conditions.

This talk presents a vision for a two-phase approach to embedded system design that brings together design-time optimization and runtime adaptation. We begin with concrete examples of design space exploration, where multi-objective optimization techniques are used to identify Pareto-optimal configurations that span energy, reliability, and accuracy trade-offs. These configurations serve as a foundation for system flexibility.

Next, we shift focus to runtime adaptation, showcasing examples where systems can dynamically adapt to real-time conditions such as workload variations or energy constraints, enabling a new generation of adaptive, context-aware embedded architectures.

Marcello Traiola, Ph.D.

Marcello Traiola is a tenured Research Scientist with the Inria Research Institute at the IRISA laboratory in Rennes, France. He received the Laurea degree (MSc) in Computer Engineering in 2016 from the University of Naples Federico II, Italy, and the Ph.D. degree in Computer Engineering in 2019 from the University of Montpellier, France. He regularly serves as a committee member and organizing member at several international conferences. His main research topics include emerging computing paradigms, with a special interest in hardware design, testing, and reliability.

Automatic Analog IC Design Optimization

Speaker: **Poki Chen (National Taiwan University of Science and Technology, Taiwan),**

System/circuit design always takes extensive work hours with manual labor, and state-of-the-art design optimization algorithms are often need to be more practical due to their multidisciplinary expertise requirements. Taking analog IC design for a good example, software automation is one way to reduce IC design workload, particularly during fine-tuning process. One of the benefits of automating parts of the design process is progress with uptime as high as 24 hours a day, which are infeasible for design engineers. While optimization methods using genetic algorithms and machine learning can improve performance of analog ICs significantly, the operation of such complex methods requires circuit design and software development know-how. However, simple optimization methods such as parameter random spread outlined in this paper have shown to reduce the temperature coefficient of a voltage reference circuit by 25% in a typical corner and 82% in the fast-fast corner. Such performance improvement demonstrates that even a basic automated optimization procedure can help optimize circuits without further human interactions.

Poki Chen, Ph.D.

Poki Chen received his B.S., M.S., and Ph.D. degrees in Electrical Engineering from National Taiwan University. He also received a doctorate honoris causa from CTU in 2025. He is a Professor in the Electronic and Computer Engineering Department at NTUST. He has served as an Associate Editor for IEEE TVLSI since 2011 and IEEE Access since 2013. His research interests include analog and mixed-signal IC design and layout, with a special focus on time-domain signal processing circuits, such as smart temperature sensors, TDC, DTC, ADC, and high-accuracy DAC. He is also interested in developing innovative analog applications for FPGA platforms, such as FPGA smart temperature sensors, TDC and DTC.

Industrial Talks

MEMS sensors for edge AI processing

Speaker: **Petr Štukjunger (STMicroelectronics, Czech Rep.)**,

Algorithms are moving from cloud to embedded systems. In modern applications the traditional approach of gathering raw sensor data and sending the data to cloud is replaced by processing in embedded systems.

The latest MEMS sensors from STMicroelectronics have brought computing even closer to the sensor - actually, inside the sensor itself.

Processing cores of ST's MEMS sensors allow implementation of conventional as well as AI based algorithms. By leveraging on machine learning approach designers can target consumer, industrial as well as automotive applications.

This innovative approach allows to engineers reaching ultra-low power consumption of their embedded systems while offloading communications buses and the main processor.

In this talk we will see a complete overview of ST's accelerometers and IMUs with embedded processing and examine the advantages of utilizing these sensors in the design of cutting-edge embedded systems.

Petr Štukjunger

Petr Štukjunger is a highly experienced Senior Application Engineer, focusing on motion and environmental MEMS and sensors. He works for STMicroelectronics, a global semiconductor company, and is based in Prague, Czech Republic. He holds a Master of Science degree in Computer Science from the Czech Technical University in Prague and a Bachelor of Science degree in Business Computing from the Institute of Technology in Sligo, Ireland. Petr has been working for ST for over 15 years, providing technical support to numerous customers in the EMEA region.

Optimizing and Deploying Neural Networks to Mobile Devices

Speaker: **Pavel Macenauer (NXP Semiconductors, Czech Rep.)**,

Optimizing neural networks is a crucial step for their deployment to embedded and mobile devices ranging from automotive visual systems to your smart watches. We will explore the challenges in doing so, dive into specific use cases and discuss algorithms used for such tasks.

Pavel Macenauer

Pavel Macenauer is a software architect at NXP Semiconductors focused on the development of AI and Machine Learning tools & embedded runtime engines. During his life, Pavel developed software visualizing the World around us, through the lens of a camera, from the cockpit of an aircraft and eventually landing at NXP to get his hands on the silicon itself.

Explainable Anomaly Detection in Network Traffic Using LLM

Kamil Jerabek¹, Josef Koumar², Jiří Setinský¹, and Jaroslav Pesek²

¹Brno University of Technology, Czech Republic
Email: {ijerabek, isetinsky}@fit.vut.cz

²Czech Technical University in Prague, Prague, Czech Republic
Email: {josef.koumar, jaroslav.pesek}@fit.cvut.cz

Abstract

Network anomaly detection is essential for modern cybersecurity, yet existing systems often generate numerous alerts without clear explanations, leading to inefficiencies and high false-positive rates. This paper proposes a novel approach that integrates Large Language Models (LLMs) with an anomaly detection framework to enhance explainability in network traffic analysis. Instead of directly detecting anomalies, the LLM only interprets already flagged anomaly events, providing insights into their potential root causes. Our method reduces LLM over-usage while improving decision-making for security analysts. We evaluated our approach using real-world network traffic data, demonstrating its ability to enhance situational awareness, reduce false positives, and support more effective cybersecurity practices.

Keywords— anomaly detection, network security, network traffic monitoring, time series, large language models, explainable security

ACKNOWLEDGEMENT

This research was partially funded by the Ministry of Interior of the Czech Republic in the project “Flow-Based Encrypted Traffic Analysis” (VJ02010024) and by the Grant Agency of the CTU in Prague, grant No. SGS23/207/OHK3/3T/18 funded by the MEYS of the Czech Republic.

PAPER ORIGIN

Previously published and presented at IEEE/IFIP Network Operations and Management Symposium (NOMS) 2025

On Evaluation of Data Fusion Methods for Network Traffic Classification

Richard Plný

Faculty of Information Technology, Czech Technical University in Prague, Czech Republic

Thakurova 9, 160 00 Prague 6, Czech Republic

`plnyrich@fit.cvut.cz`

Supervisor: *Karel Hynek*

Abstract

Data fusion has become a critical aspect of modern network security. As complex ML/AI models become widely adopted, the need for their interpretation and explainability of their outputs rises rapidly. One possibility to increase trust in such models is to use heterogeneous detectors and multiple data sources together with data fusion. Resulting method provides more explainable outputs together with basic reasoning. However, the impact of even common data fusion methods was rather overlooked in the network security domain. This research paper focuses on late and hard data fusion within the traffic classification and attack detection domain and evaluates seven prevalent fusion algorithms: Majority Voting, Weighted Majority Voting, Recall Combiner, Naive Bayes, Behavior-Knowledge Space, Decision Tree, and Logistic Regression. To support this study, we developed a synthetic data generator, enabling the creation of diverse experimental setups with multiple fusion scenarios and varying levels of input accuracy. Additionally, we compared the performance of these fusion methods on synthetic datasets with real-world traffic classification challenges to better understand their strengths and limitations. Our findings provide valuable insights for network security researchers and engineers, offering practical guidance in selecting the most suitable fusion method for addressing information fusion challenges in traffic analysis with an emphasis on network attack and intrusion detection.

Keywords— network traffic classification, data fusion, network security

ACKNOWLEDGMENT

This work was supported by the Grant Agency of the Czech Technical University in Prague, which is funded by the Ministry of Education, Youth and Sports of Czech Republic, grant No. SGS23/207/OHK3/3T/18.

PAPER ORIGIN

This work has been submitted as a full paper to the 21st International Conference on Network and Service Management (CNSM) 2025.

Multi-Agent Multi-Object Tracking with Shared Measurements and Intersecting Fields of View

Pavel Kopecký, Kamil Dedecius

Faculty of Information Technology, Czech Technical University in Prague

Thákurova 9, 160 00 Prague 6, Czech Republic

pavel.kopecky@fit.cvut.cz, kamil.dedecius@fit.cvut.cz

Abstract

The rapid development of smart systems such as autonomous vehicles (AVs) and unmanned aerial systems (UASs) imposes substantial requirements on the real-time monitoring of surrounding moving objects. In this paper, we propose an extension of the popular multi-object filter, namely the Joint Integrated Probabilistic Data Association (JIPDA) filter, for networks of agents (e.g., AVs and UASs) equipped with short-range communication capabilities. By sharing observations acquired within their respective fields of view, agents can significantly enhance object tracking performance in cluttered environments. The communication requirements of the proposed solution are relatively modest and theoretically suitable for the considered applications.

Keywords— Agent networks, information fusion, multi-object tracking, JIPDA filter.

I. INTRODUCTION

The relentless advancement of small devices with impressive computational performance, large memory, and communication abilities in the last decade has fuelled the deployment of many applications that were originally limited to very specific hardware and required a profound expert knowledge. This gave rise, for instance, to modern sensor networks, Internet of Things, wearable electronic devices, robotics and autonomous machinery, cyber-physical systems, autonomous self-driving vehicles and intelligent traffic systems in general, and defense and security systems (e.g., [1]–[3]). Common aspects of all these modern applications are limited size, weight, power, and cost (SWaP-C).

In this paper, we are primarily interested in *distributed multi-object tracking* (MOT) within a network of collaborating agents. Historically, *non-distributed* MOT was primarily developed for aviation and maritime domains or, more broadly, air and naval surveillance. It played a critical role in traffic safety and – inevitably – defense. Physically, sensing was based on echolocation using sonar and radiolocation using radars. The received signals were processed in several stages, ranging from signal-noise separation at the lowest level to the extraction of useful information from the resulting measurements and filtering of objects states at the highest level [4]. All this required substantial computational power, memory, and energy, which, in turn, was only possible with large and expensive hardware. However, with the advent of modern compact embedded systems, the range of object tracking applications has fundamentally shifted. They find applications, e.g., in

- Unmanned aerial systems (UAS), where they can be utilized as airborne radars in surveillance and assist in a timely manner to forfeit threats from other UASs [5].
- Urban air mobility and urban air logistics with autonomously piloted aircrafts [6], [7].
- Indoor navigation with aligned RGB camera and LiDAR sensors [8].
- Electronic warfare systems for airborne platforms [9].
- Self-driving cars – autonomous vehicle perception of obstacles, pedestrians, and other cars with various sensors like radar, lidar, and cameras, including their combinations, e.g., [10]–[13].

Also, the topic of multi-agent *collaboration* has witnessed a rapid evolution in the last two decades, mainly due to the advancements in the ad-hoc wireless networks. Modern embedded devices are capable of sharing information through a wide range of short- and long-range communication technologies. The short-range technologies comprise typically Bluetooth and its low-energy variant BLE, Wi-Fi standards IEEE 802.11a/b/g/n/ac/ax, Zigbee standard IEEE 802.15.4, IR, Ultra-Wideband, or Garmin ANT+ [14]–[20]. The long-range communication technologies are mostly based on Wi-Fi HaLow (802.11ah), LoRa, Narrowband IoT, LTE-M, and traditional 4G/5G/LTE cellular networks, see, e.g., [17], [21]. In addition, there are cross-technology communication techniques that provide the ability to transmit data between heterogeneous radios, for instance, WeBee [21], [22]. To summarize, the resulting multi-agent systems with sensing, processing, and communication capabilities give rise to a modern *joint sensing and communication* IoT paradigm [21], [23].

In this paper, we focus on networks of collaborating agent where:

- The agents monitor objects that move in the environment. Each agent has own field of view (FoV). These FoVs may be partially shared by two or more agents.
- Besides the measurements of real objects, the agents are prone to false alarms – clutter.
- Each agent performs tracking of objects in its own FoV using the Joint Integrated Probabilistic Data Association (JIPDA) filter [24].

- The agents may share raw measurements with their adjacent neighbors. The shared measurements may carry additional information about the objects in the FoVs.

Despite the omnipresence of Probabilistic Data Association (PDA) filters, the authors are not aware of any other distributed JIPDA implementation. The only exception is due to Kopfreiter et al. [25] who, however, postponed the implementation details to a future regular journal paper. Here, we present our preliminary research on this topic.

The paper is organized as follows: Section II overviews the principles of the JIPDA filter. It describes the situation, the underlying assumptions, and the filtering steps. Section III proposes a JIPDA filter with measurement sharing. The experimental validation is performed in Section IV. Finally, Section V concludes the paper.

II. JIPDA FILTER

MOT filters usually represent their object state estimates via probability densities, expressing their expected states and the variance of the estimates. The Joint Integrated Probabilistic Data Association filter [24] tackles the MOT problem by encapsulating the n_k state estimates of all present objects at time $k = 1, 2, \dots$ in the joint posterior density

$$p_k(\mathbf{x}_k^1, \dots, \mathbf{x}_k^{n_k} | \mathcal{I}_k) = \prod_{t=1}^{n_k} p_k^t(\mathbf{x}_k^t | \mathcal{I}_k), \quad (1)$$

where \mathcal{I}_k is the set of all available information about the modeled system up to time k , typically comprising the past measurements $\mathbf{Z}_1, \dots, \mathbf{Z}_k$ and the set \mathbf{Z}_0 , representing the initial knowledge, usually provided by an expert [26]. Ideally, n_k is as close to the true object count as possible. The single-object densities $p_k^t(\mathbf{x}_k^t | \mathcal{I}_k)$ represent an estimate of an object state. Each state estimate is accompanied by the posterior object existence probability $r_k^t \in [0, 1]$, describing the probability that the estimate represents a true object. The state estimate density, existence probability, as well as other information, e.g., object label, form the set of all known information about an object and are typically called a *track*.

The task of the filter is to propagate the posterior joint density in time using the available measurements, estimating the states of detectable objects as precisely as possible. The filter accomplishes this by two steps: the prediction step and the update step, described in the following sections.

A. Assumptions

First, some key assumptions of the filter need to be established.

- 1) **Target movement:** Targets in the observed area move independently of each other and their behavior is described the Markov transition density $f(\mathbf{x}_k | \mathbf{x}_{k-1})$, where \mathbf{x}_k is the object state at time step k , represented by a real finite vector.
- 2) **Target survival:** Each object present at time $k - 1$ continues to exist in the next time step with the survival probability $P_S \in [0, 1]$ and disappears with the probability $1 - P_S$.
- 3) **Target measurements:** Each object is detected with the probability $P_D \in [0, 1]$ and the measurement process obeys the measurement density $g(\mathbf{z}_k | \mathbf{x}_k)$, where \mathbf{z}_k is the object measurement, typically containing the imprecise object position relative to the agent. Each object generates *at most* one measurement per time step and each measurement originates from *at most* one object.
- 4) **Clutter measurements:** In addition to object measurements, the local sensor also generates clutter (or false) measurements, stemming from inherent signal noise or physical obstacles. It is assumed that the number of clutter measurements is Poisson with the mean $\Lambda > 0$ and that clutter is uniformly spatially distributed across the entire FoV \mathcal{F} of the agent. The density of the clutter process is therefore $\lambda = \frac{\Lambda}{|\mathcal{F}|}$, where $|\mathcal{F}|$ is the area of the agent's FoV.

These assumptions are omnipresent in the literature and significantly simplify the derivation of modern object tracking algorithms, including the JIPDA filter.

B. Prediction Step

At the start of each time step k , the filter first predicts the states of all objects based on their prior estimate densities via the Chapman-Kolmogorov equation

$$\begin{aligned} p_{k|k-1}(\mathbf{x}_k^1, \dots, \mathbf{x}_k^{n_{k-1}} | \mathcal{I}_{k-1}) &= \prod_{t=1}^{n_{k-1}} \int f(\mathbf{x}_k^t | \mathbf{x}_{k-1}^t) p_{k-1}^t(\mathbf{x}_{k-1}^t | \mathcal{I}_{k-1}) d\mathbf{x}_{k-1}^t \\ &= \prod_{t=1}^{n_{k-1}} p_{k|k-1}^t(\mathbf{x}_k^t | \mathcal{I}_{k-1}). \end{aligned} \quad (2)$$

The existence probability of each track $t = 1, \dots, n_{k-1}$ is predicted as well, resulting in

$$r_{k|k-1}^t = P_S \cdot r_{k-1}^t. \quad (3)$$

The prediction step expresses our expectation about the states of objects at time k based on the information available up to the previous time step.

C. Update Step

The update step corrects the state predictions using the available observations at the current time step $\mathbf{Z}_k = \{\mathbf{z}_k^1, \dots, \mathbf{z}_k^{m_k}\}$. First, the filter enumerates all the possible association hypotheses $\mathbf{a}_k = [a_k^1, \dots, a_k^{n_k-1}]^T$ between the tracks and the measurements

$$a_k^t = \begin{cases} j & \text{if the } j\text{-th measurement is assigned to track } t, \\ 0 & \text{if no measurement is assigned to track } t. \end{cases} \quad (4)$$

Each association hypothesis represents an assignment of measurements to local tracks. Let \mathcal{A}_k be the set of all mutually exclusive and exhaustive association hypotheses \mathbf{a}_k . Using the total probability theorem, we determine that the joint posterior density is

$$p_k(\mathbf{x}_k^1, \dots, \mathbf{x}_k^{n_k-1} | \mathcal{I}_k) = \sum_{\mathbf{a}_k \in \mathcal{A}_k} p_k(\mathbf{x}_k^1, \dots, \mathbf{x}_k^{n_k-1} | \mathbf{a}_k, \mathcal{I}_k) \Pr(\mathbf{a}_k | \mathcal{I}_k). \quad (5)$$

The association-event-conditioned densities $p_k(\mathbf{x}_k^1, \dots, \mathbf{x}_k^n | \mathbf{a}_k, \mathcal{I}_k)$ are obtained by using the standard Bayes-filter-updated estimate on the tracks with measurements assigned and the prediction on the misdetections, resulting in

$$p_k(\mathbf{x}_k^1, \dots, \mathbf{x}_k^{n_k-1} | \mathbf{a}_k, \mathcal{I}_k) \propto \underbrace{\prod_{t: a_k^t=0} p_{k|k-1}^t(\mathbf{x}_k^t | \mathcal{I}_{k-1})}_{\text{misdetections}} \underbrace{\prod_{t: a_k^t \neq 0} g(\mathbf{z}_k^{a_k^t} | \mathbf{x}_k^t) p_{k|k-1}^t(\mathbf{x}_k^t | \mathcal{I}_{k-1})}_{\text{detected tracks}}, \quad (6)$$

with $\mathbf{z}_k^{a_k^t} \in \mathbf{Z}_k$ being the measurement assigned to track t under the hypothesis \mathbf{a}_k . The association probabilities $\Pr(\mathbf{a}_k | \mathcal{I}_k)$ are

$$\Pr(\mathbf{a}_k | \mathcal{I}_k) \propto \prod_{t: a_k^t=0} (1 - P_D r_{k|k-1}^t) \prod_{t: a_k^t \neq 0} \frac{P_D r_{k|k-1}^t \int g(\mathbf{z}_k^{a_k^t} | \mathbf{x}_k^t) p_{k|k-1}^t(\mathbf{x}_k^t | \mathcal{I}_{k-1}) d\mathbf{x}_k^t}{\lambda}. \quad (7)$$

In practice, the update is not performed collectively as in Eq. (5). Instead, the filter updates each track individually, forming the posterior single-object densities

$$\tilde{p}_k^t(\mathbf{x}_k^t | \mathcal{I}_k) \propto \beta_k^{t,0} p_{k|k-1}^t(\mathbf{x}_k^t | \mathcal{I}_{k-1}) + \sum_{j=1}^{m_k} \beta_k^{t,j} g(\mathbf{z}_k^j | \mathbf{x}_k^t) p_{k|k-1}^t(\mathbf{x}_k^t | \mathcal{I}_{k-1}), \quad (8)$$

where $\beta_k^{t,j}$ is the probability that the j -th measurement originated from track t and $\beta_k^{t,0}$ is the misdetection probability, both obtained by

$$\beta_k^{t,i} = \sum_{\mathbf{a}_k: a_k^t=i} \Pr(\mathbf{a}_k | \mathcal{I}_k), \quad i = 0, \dots, m_k. \quad (9)$$

The mixture in Eq. (8) is usually approximated by a density closest in the Kullback-Leibler sense

$$p_k^t(\mathbf{x}_k^t | \mathcal{I}_k) = \arg \min_{q(\mathbf{x}_k^t)} D[\tilde{p}_k^t(\mathbf{x}_k^t | \mathcal{I}_k) || q(\mathbf{x}_k^t)], \quad (10)$$

where the minimization runs over all densities in an admissible class. If the mixture in Eq. (8) and $q(\mathbf{x}_k^t)$ are Gaussian, we call this method of mixture reduction moment matching [27]. The update of the existence probability of each track utilizes the single-object association probabilities and reads

$$r_k^t = \frac{(1 - P_D) r_{k|k-1}^t}{1 - P_D r_{k|k-1}^t} \beta_k^{t,0} + \sum_{j=1}^{m_k} \beta_k^{t,j}. \quad (11)$$

D. Track Management

After the update step, the filter discards all tracks with existence probabilities below the termination threshold $\tau_T > 0$, as these tracks are unlikely to represent a real object, resulting in the modified track count $\bar{n}_{k-1} \leq n_{k-1}$. Next, tracks with existence probabilities above the confirmation threshold $\tau_C > 0$ are marked as *confirmed*. Confirmed tracks are tracks believed to represent an object and are typically displayed in the tracking system output.

During the update step, each track is typically only assigned measurements within a reasonable distance from its state prediction using a method known as *gating* [28], reducing the total number of association hypotheses by a significant margin. Measurements that are not assigned to any track give rise to $\bar{m}_k \in \mathbb{N}_0$ new *tentative* tracks with low existence probabilities $r_k^{t'}$ and densities $p_k^{t'}(\mathbf{x}_k^{t'} | \mathcal{I}_k)$, such that

$$\mathbb{E} \left[\int g(\mathbf{z}_k | \mathbf{x}_k^{t'}) p_k^{t'}(\mathbf{x}_k^{t'} | \mathcal{I}_k) d\mathbf{x}_k^{t'} \right] \quad (12)$$

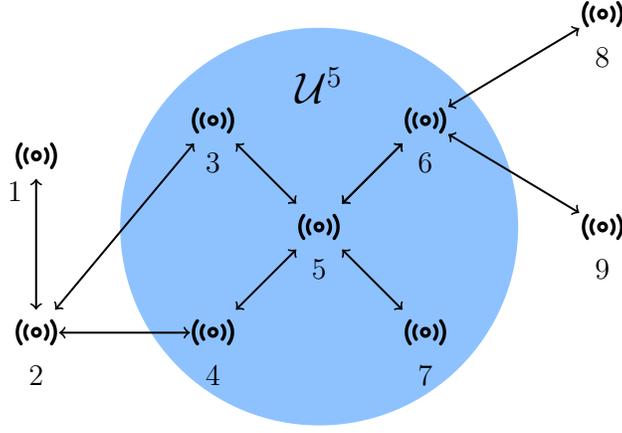


Fig. 1. Example of a network comprising eight interconnected agents. The agent set is $\mathcal{U} = \{1, \dots, 9\}$. For a selected agent $u = 5$, its closed neighborhood is $\mathcal{U}^5 = \{3, 4, 5, 6, 7\}$. u can only communicate with agents in \mathcal{U}^5 .

results in the measurement that gave rise to track t' . The initial existence probabilities of tracks are usually constant, i.e., $r_k^{t'} := r_0 \in (0, 1)$ for all $t' = 1, \dots, \bar{m}_k$. Tentative tracks are tracks suspected of representing a real object, to be deleted or confirmed later. Once a track is confirmed, it stays confirmed even when its existence probability falls below τ_C and is deleted once its existence probability falls below the threshold τ_M .

At the end of the time step, the filter replaces its posterior joint density with

$$p_k(\mathbf{x}_k^1, \dots, \mathbf{x}_k^{n_k} | \mathcal{I}_k) = \prod_{t=1}^{\bar{n}_{k-1}} p_k^t(\mathbf{x}_k^t | \mathcal{I}_k) \prod_{t'=1}^{\bar{m}_k} p_k^{t'}(\mathbf{x}_k^{t'} | \mathcal{I}_k), \quad (13)$$

where $n_k = \bar{n}_{k-1} + \bar{m}_k$. This density enters the JIPDA prediction step at the subsequent time step $k + 1$.

III. DISTRIBUTED JIPDA FILTER WITH SHARED MEASUREMENTS

The present paper extends the traditional single-agent JIPDA filter to distributed networks, enabling estimation with measurements from multiple agents. This results in an improved detection speed and tracking accuracy, which is vital in many real-time tracking systems.

In a distributed setting, instead of a single agent, the area of interest is observed by a network of agents represented by the connected graph $\mathcal{G} = (\mathcal{U}, \mathcal{E})$, where \mathcal{U} is the set of agents and \mathcal{E} is the set of communication links between agents. Each agent $u \in \mathcal{U}$ may only communicate with agents within a one network hop distance, forming its closed neighborhood \mathcal{U}^u , as demonstrated in Figure 1. By definition, $u \in \mathcal{U}^u$.

We consider that once per time step, any agent u may request measurements from its neighboring agents, which it can incorporate into its local posterior estimates. We proceed to describe how this can be achieved based on the traditional JIPDA filter.

A. Assumptions

In addition to the assumptions declared in the previous section, the distributed version of the filter presumes that the sensing capabilities of each agent vary, which can be caused by different signal acquisition techniques or hardware. Given an agent $u \in \mathcal{U}$, this is reflected in its specific clutter density λ^u , detection probability P_D^u , and likelihood function $g^u(\circ)$.

B. Integrating Neighbor Measurements

We describe the process from a viewpoint of a particular agent $u \in \mathcal{U}$, although the same applies to every agent in the network. After u executes the standard JIPDA prediction step and processes its own measurements $\mathbf{Z}_k^u = \{z_k^{u,1}, \dots, z_k^{u,m_k^u}\}$, where $m_k^u = \mathbf{Z}_k^u$, it forms the updated joint posterior

$$p_k^u(\mathbf{x}_k^{u,1}, \dots, \mathbf{x}_k^{u,n_k^u} | \mathcal{I}_k^u) = \prod_{t=1}^{n_k^u} p_k^{u,t}(\mathbf{x}_k^{u,t} | \mathcal{I}_k^u), \quad (14)$$

where \mathcal{I}_k^u is the set of all available information to agent u up to time k . Each single-object density is accompanied by the posterior existence probability $r_k^{u,t}$.

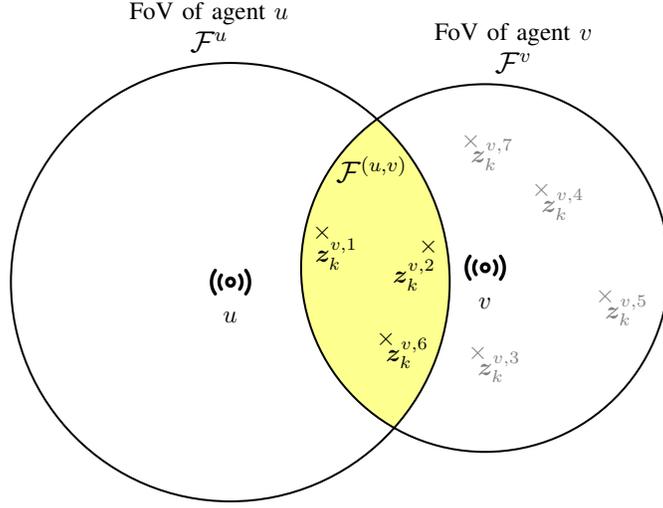


Fig. 2. FoV intersection $\mathcal{F}^{(u,v)}$ of two agents u, v . Agent u only receives the measurements from v contained within $\mathcal{F}^{(u,v)}$, i.e., $\mathbf{Z}_k^{(u,v)} = \{z_k^{v,1}, z_k^{v,2}, z_k^{v,6}\} = \{z_k^{(u,v),1}, z_k^{(u,v),2}, z_k^{(u,v),3}\}$.

Afterward, u requests measurements from each of its neighbors $v \in \mathcal{U}^u \setminus \{u\}$. As we presume u is only interested in the information specific to its FoV \mathcal{F}^u , only the measurements within \mathcal{F}^u should be transmitted. Formally, the measurements are confined to the intersection of the agents' FoVs $\mathcal{F}^{(u,v)} = \mathcal{F}^u \cap \mathcal{F}^v$, forming the measurement set $\mathbf{Z}_k^{(u,v)} = \{z_k^{(u,v),1}, \dots, z_k^{(u,v),m_k^{(u,v)}}\}$ comprising $m_k^{(u,v)} \in \mathbb{N}_0$ unique measurements. This is further illustrated in Figure 2.

Using the neighbor measurements, u then refines its local posterior estimates in a way similar to the JIPDA filter update step, forming the posterior

$$p_k^u(\mathbf{x}_k^{1,u}, \dots, \mathbf{x}_k^{u,n_k^u} | \mathcal{I}_k^u, \mathbf{Z}_k^{(u,v)}) = \sum_{\mathbf{a}_k^{(u,v)} \in \mathcal{A}_k^{(u,v)}} p_k^u(\mathbf{x}_k^{u,1}, \dots, \mathbf{x}_k^{u,n_k^u} | \mathbf{a}_k^{(u,v)}, \mathcal{I}_k^u) \Pr(\mathbf{a}_k^{(u,v)} | \mathcal{I}_k^u, \mathbf{Z}_k^{(u,v)}), \quad (15)$$

where $\mathcal{A}_k^{(u,v)}$ is the set of all association hypotheses $\mathbf{a}_k^{(u,v)} = [a_k^{(u,v),1}, \dots, a_k^{(u,v),n_k^u}]^T$ constructed the same way as in Eq. (4) and

$$p_k^u(\mathbf{x}_k^{u,1}, \dots, \mathbf{x}_k^{u,n_k^u} | \mathbf{a}_k^{(u,v)}, \mathcal{I}_k^u) \propto \prod_{t: a_k^{(u,v),t}=0} p_k^{u,t}(\mathbf{x}_k^{u,t} | \mathcal{I}_k^u) \prod_{t: a_k^{(u,v),t} \neq 0} g^v(z_k^{a_k^{(u,v),t}} | \mathbf{x}_k^{u,t}) p_k^{u,t}(\mathbf{x}_k^{u,t} | \mathcal{I}_k^u), \quad (16)$$

where $z_k^{a_k^{(u,v),t}}$ is the measurement assigned to track t under the association hypothesis $\mathbf{a}_k^{(u,v)}$.

As u may be tracking objects outside \mathcal{F}^v , the update must be modified such that tracks with position estimates (partially) outside \mathcal{F}^v are updated properly. This is achieved by the modified track-specific detection probability $P_{D,t}^{(u,v)}$, which is v 's detection probability scaled by the volume of the predictive measurement density of track $t = 1, \dots, n_k^u$ within $\mathcal{F}^{(u,v)}$

$$P_{D,t}^{(u,v)} = P_D^v \int_{\mathcal{F}^{(u,v)}} g^v(\mathbf{z} | \mathbf{x}_k^{u,t}) p_k^{u,t}(\mathbf{x}_k^{(u,v),t} | \mathcal{I}_k^u) d\mathbf{x}_k^{(u,v),t} d\mathbf{z}. \quad (17)$$

This results in the modified association hypothesis probabilities

$$\Pr(\mathbf{a}_k^{(u,v)} | \mathcal{I}_k^u, \mathbf{Z}_k^{(u,v)}) \propto \prod_{t: a_k^{(u,v),t}=0} (1 - P_{D,t}^{(u,v)} r_k^{u,t}) \prod_{t: a_k^{(u,v),t} \neq 0} \frac{P_{D,t}^{(u,v)} r_k^{u,t} \int g^v(z_k^{a_k^{(u,v),t}} | \mathbf{x}_k^{u,t}) p_k^{u,t}(\mathbf{x}_k^{u,t} | \mathcal{I}_k^u) d\mathbf{x}_k^{u,t}}{\lambda^v}. \quad (18)$$

Without the detection probability adjustment, the existence probabilities of tracks outside or near the edges of $\mathcal{F}^{(u,v)}$ would fall drastically, because, although they are not as easily detectable by v (or are undetectable at all), the standard JIPDA filter expects a measurement with the probability P_D^v .

The rest of the update proceeds identically to the single-agent JIPDA filter, i.e., the single-object association probabilities are calculated and used to update each individual track. The tracks with existence probabilities below the threshold τ_T are removed, leading to a new reduced track count \bar{n}_k^u . Tracks with existence probabilities above τ_C are marked as confirmed and

the unassigned measurements from the update are used to establish $\overline{m}_k^{(u,v)} \in \mathbb{N}_0$ new tentative tracks $t' = 1, \dots, \overline{m}_k^{(u,v)}$ with low existence probabilities r_0 and densities $p_k^{u,t'}(\mathbf{x}_k^{(u,v),t'} | \mathcal{I}_k^u, \mathbf{Z}_k^{(u,v)})$.

This results in the joint posterior consisting of $n_k^{(u,v)} = \overline{n}_k^u + \overline{m}_k^{(u,v)}$ tracks of the form

$$p_k^u(\mathbf{x}_k^{u,1}, \dots, \mathbf{x}_k^{u,n_k^{(u,v)}} | \mathcal{I}_k^u, \mathbf{Z}_k^{(u,v)}) = \prod_{t=1}^{\overline{n}_k^u} p_k^{u,t}(\mathbf{x}_k^{u,t} | \mathcal{I}_k^u, \mathbf{Z}_k^{(u,v)}) \prod_{t'=1}^{\overline{m}_k^{(u,v)}} p_k^{u,t'}(\mathbf{x}_k^{(u,v),t'} | \mathcal{I}_k^u, \mathbf{Z}_k^{(u,v)}). \quad (19)$$

This density can be further updated with measurements from the rest of u 's neighbors in the same manner.

If we denote $S^u = \{u, v_1, \dots, v_b^u\}$, where v_b^u is the number of u 's neighbors, a single time step of the distributed JIPDA algorithm from u 's perspective can be summarized as follows:

$$p_{k-1}^u(\circ | \mathcal{I}_{k-1}^u) \xrightarrow{\text{prediction}} p_{k|k-1}^u(\circ | \mathcal{I}_{k-1}^u) \xrightarrow[\mathbf{Z}_k^u]{\text{update}} p_k^u(\circ | \mathcal{I}_k^u) \xrightarrow[\mathbf{Z}_k^{(u,v_1)}]{\text{update}} p_k^u(\circ | \mathcal{I}_k^u, \mathbf{Z}_k^{(u,v_1)}) \xrightarrow[\mathbf{Z}_k^{(u,v_2)}]{\text{update}} \dots \xrightarrow[\mathbf{Z}_k^{(u,v_b^u)}]{\text{update}} p_k^u(\circ | \mathcal{I}_k^{S^u}),$$

where

$$\mathcal{I}_k^{S^u} = \mathcal{I}_{k-1}^u \cup \bigcup_{v \in \mathcal{U}^u} \mathbf{Z}_k^{(u,v)} \quad (20)$$

replaces \mathcal{I}_k^u in the subsequent time step. The pseudocode of the algorithm is available in Algorithm 1.

C. Remarks

The computational, memory, and time requirements are driven by several factors. In principle:

- The incorporation of neighbor's measurements is equivalent to a single JIPDA update of a subset of estimated states.
- The practical selection of states that could potentially be updated by neighbor's measurements can benefit from the knowledge of the neighbor's field of view.
- The computationally challenging part is the determination of the target–measurement associations. If the spatial configuration of states and measurements rules out the case of one measurement being generated by more than one target, the JIPDA algorithm can be reduced to the basic PDA. Otherwise, the association hypotheses must be found using a convenient optimal assignment algorithm.
- It is possible to further reduce the computational cost, e.g., by using non-ellipsoidal gating [28], constant-gain Kalman filtering [29], and other techniques.

Algorithm 1 Distributed JIPDA Filter with Shared Measurements

The agent network is represented by the graph $\mathcal{G} = (\mathcal{U}, \mathcal{E})$. The track confirmation threshold τ_C , termination threshold τ_T , and initial existence probability r_0 are set. For $k = 1, 2, \dots$ and each agent $u \in \mathcal{U}$ do:

Local prediction:

- 1) Predict the states of all objects at time step k based on their estimates in the previous time step, Eq. (2).
- 2) Predict the existence probabilities of all tracks, Eq. (3).

Local update:

- 1) Construct the association hypothesis vectors, Eq. (4).
- 2) Find the association hypothesis probabilities, Eq. (7).
- 3) Calculate the single-object association probabilities, Eq. (9).
- 4) Update the estimate density of each track, Eqs. (8), (10).
- 5) Update the existence probability of each track, Eq. (11).
- 6) Remove tracks with existence probabilities below τ_T , confirm those with existence probabilities above τ_C .
- 7) Using the unused measurements, initialize new tracks with existence probabilities r_0 , Eq. (12).

Integrating neighbor measurements:

For each neighbor $v \in \mathcal{U}^u \setminus \{u\}$ do:

- 1) Get its measurements $\mathbf{Z}_k^{(u,v)}$ in u 's FoV \mathcal{F}^u .
- 2) Calculate track-specific detection probabilities $P_{D,t}^{(u,v)}$, Eq. (17).
- 3) Find the association hypothesis probabilities, Eq. (18).
- 4) Based on these probabilities, execute the rest of the JIPDA update, including track management, Eqs. (8)–(12).

Output all confirmed tracks.

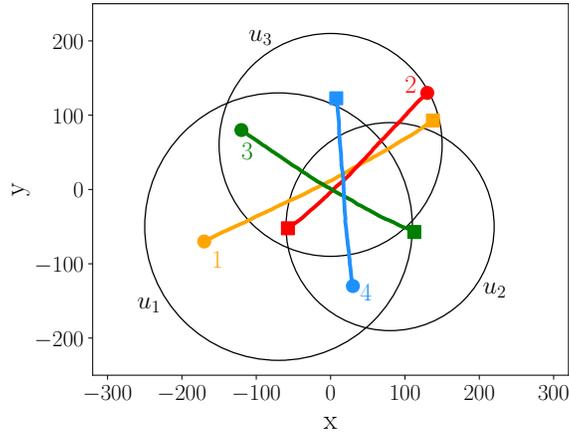


Fig. 3. Target movement in the experiments. Three agents u_1, \dots, u_3 observe the area; their FoVs are indicated by black circles. The object trajectories are visualized by the colored lines. The trajectory start is marked by a circle and the end by a square. The length of the simulation is 150 time steps.

TABLE I
TIME INTERVALS DURING WHICH THE OBJECTS ARE WITHIN THE FIELDS OF VIEW OF INDIVIDUAL AGENTS IN THE EXPERIMENTS.

	u_1	u_2	u_3
Target 1	0–120	57–142	30–150
Target 2	49–150	31–150	0–150
Target 3	0–148	54–150	0–141
Target 4	0–144	0–123	28–150

IV. EXPERIMENTS

The effectiveness of the algorithm is demonstrated in two experiments, both simulating 150 time steps of the algorithm. In both experiments, the performance is evaluated against the traditional JIPDA filter and a state-of-the-art distributed object tracking algorithm: The state-dependent weighted arithmetic average PHD (SD_WAA_PHD) filter [30]. The implementation of the JIPDA filter is based on the linear-Gaussian state-space model, with further details in, e.g., [31].

The tracking performance is assessed based on the generalized optimal sub-pattern assignment (GOSPA) metric [32] with the parameters $p = 2$, $\alpha = 2$, and $c = 10$. This metric is highly popular in contemporary object tracking literature because it captures not only the accuracy of estimates, but also whether objects are detected at all, and punishes maintaining tracks that do not correspond to true objects. The results are averaged over 100 Monte Carlo runs.

A. Experiment Setting

Both experiments utilize the same small agent network and object movements, further illustrated in Figure 3. The agents have a circular FoV of a limited range and each agent can communicate with every other agent in the network. The time intervals during which the objects are within the FoVs of individual agents are summarized in Table I.

The object survival probability is $P_S = 0.999$ and the initial existence probability is $r_0 = 0.0101$. The object confirmation and termination thresholds are $r_C = 0.9$ and $r_T = 0.01$, respectively. At every time step, each agent calculates the GOSPA error between the local confirmed tracks and the true objects within its FoV.

B. Experiment 1

In the first experiment, all agents have the same sensing capabilities. Namely, each agent has an object detection probability $P_D = 0.9$, the clutter process has the Poisson mean $\Lambda = 3$, and the measurement accuracy is the same across the whole network. The mean GOSPA error presented in Figure 4 is averaged over all agents and the objects within their FoVs.

From the results, it is clear that the distributed JIPDA filter shows a great improvement over the standalone JIPDA filter, even outperforming the SD_WAA_PHD filter. The proposed filter also establishes confirmed tracks faster than the single-agent JIPDA filter, which can be seen especially around the $k = 30$ mark. It should be noted, however, that the distributed JIPDA filter has an order higher polynomial time complexity over the SD_WAA_PHD filter.

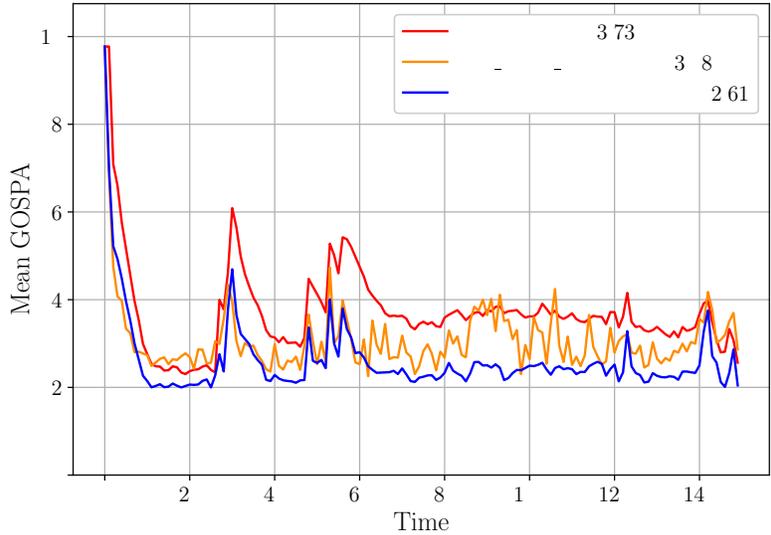


Fig. 4. Experiment 1: Mean GOSPA error over time of the noncollaborative JIPDA filter, distributed JIPDA filter, and the SD_WAA_PHD filter.

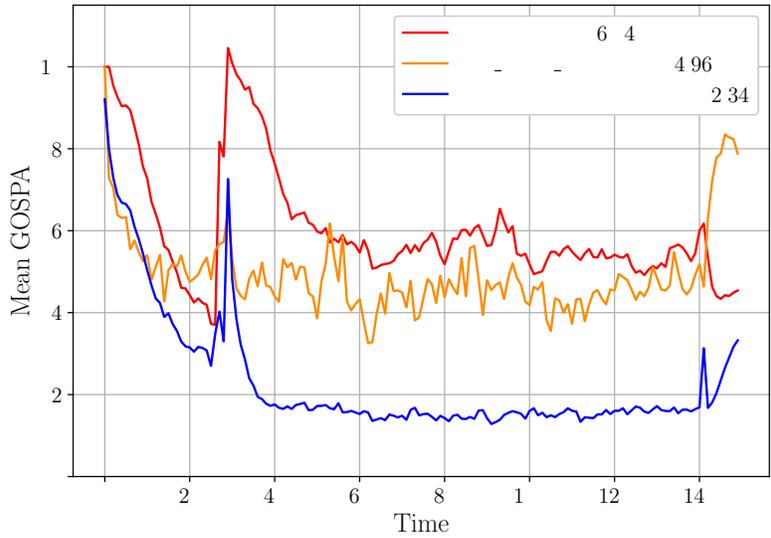


Fig. 5. Experiment 2: Mean GOSPA error over time of agent u_3 utilizing the noncollaborative JIPDA filter, distributed JIPDA filter, and the SD_WAA_PHD filter.

C. Experiment 2

In this experiment, agents u_1 and u_2 have significantly better sensing capabilities than agent u_3 . In particular, their clutter rates have a mean only $\Lambda = 1$, their detection probabilities are high – $P_D = 0.95$ – and their measurements are very precise. On the other hand, u_3 suffers from clutter with the mean $\Lambda^{u_3} = 8$, has a low detection probability $P_D^{u_3} = 0.65$, and its measurements are imprecise. The goal is to see how the measurements from the first two agents improve u_3 's tracking performance. Figure 5 shows the mean GOSPA error over time for agent u_3 specifically.

The improvement of the distributed JIPDA filter over the two other strategies is even greater than in the first experiment. When compared to the SD_WAA_PHD filter, the presented algorithm seems to enable a greater tracking improvement of worse-performing agents in a distributed setting.

V. CONCLUSION

We have proposed a multi-agent Joint Integrated Probabilistic Data Association filter that enhances object tracking performance by using an enriched set of measurements shared among agents. The algorithm imposes relatively modest communication requirements: at each time step, only the measurements of the adjacent neighboring agents with intersecting fields of view are exchanged. The approach holds significant potential in fields such as autonomous vehicles and unmanned aerial systems, where accurate monitoring of surrounding moving objects is critical. Simulation results demonstrate the superior tracking performance of the proposed algorithm compared to standard JIPDA and an alternative state-of-the-art method.

REFERENCES

- [1] P. Raj and A. C. Raman, *The Internet of Things*. CRC Press, 2023.
- [2] A. Dumka, S. K. Chaurasiya, A. Biswas, and H. L. Mandoria, *A Complete Guide to Wireless Sensor Networks: From Inception to Current Trends*. CRC Press, 2022.
- [3] H. Lipson and M. Kurman, *Driverless: Intelligent Cars and the Road Ahead*. MIT Press, 2016.
- [4] M. A. Richards, J. A. Scheer, and W. A. Holm, *Principles of Modern Radar: Basic Principles*. SciTech Publishing, 2010.
- [5] X. P. Lin, B. Sul, P. Zhuang, Y. Zhang, H.-M. Chen, S. Kaniyantethu, S. Alickolli, E. Blasch, K. Pham, and G. Chen, "Low SWaP-C LFMW airborne radar for counter-UAS applications," in *NAECON 2023 - IEEE National Aerospace and Electronics Conference*, 2023, pp. 300–303.
- [6] W. A. Lies, L. Narula, P. A. Iannucci, and T. E. Humphreys, "Long range, low SWaP-C FMCW radar," *IEEE Journal of Selected Topics in Signal Processing*, vol. 15, no. 4, pp. 1030–1040, 2021.
- [7] —, "Low SWaP-C radar for urban air mobility," in *2020 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, 2020, pp. 74–80.
- [8] R. D. Hoobler, D. C. Wiberg, and M. R. Akella, "RGB-LiDAR pipeline for 3D bounding box estimation in low SWaP-C indoor navigation applications," in *2023 American Control Conference (ACC)*, 2023, pp. 3316–3323.
- [9] A. Mallika, M. U. Ravindra, T. Latha, D. S. Rao, S. Matham, C. V. Kumar, G. S. Rao, and Y. Hemalatha, "Design and development of SWaP optimized C-Ku band RF-front end for futuristic EW systems," in *2025 19th European Conference on Antennas and Propagation (EuCAP)*, 2025, pp. 1–5.
- [10] X. Song and Y. Du, "Obstacle tracking based on millimeter wave radar," in *2016 3rd International Conference on Systems and Informatics (ICSAI)*, 2016, pp. 531–535.
- [11] J. Nino and M. Campbell, "Priority tracking of pedestrians for self-driving cars," in *2022 IEEE 18th International Conference on Automation Science and Engineering (CASE)*, 2022, pp. 549–556.
- [12] L. Zhao, M. Wang, and Y. Yue, "Sem-Aug: Improving camera-LiDAR feature fusion with semantic augmentation for 3D vehicle detection," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 9358–9365, 2022.
- [13] H.-H. Jebamikyous and R. Kashaf, "Autonomous vehicles perception (AVP) using deep learning: Modeling, assessment, and challenges," *IEEE Access*, vol. 10, pp. 10 523–10 535, 2022.
- [14] S. Farahani, *ZigBee Wireless Networks and Transceivers*, 01 2008.
- [15] J. Nieminen, C. Gomez, M. Isomaki, T. Savolainen, B. Patil, Z. Shelby, M. Xi, and J. Oller, "Networking solutions for connecting bluetooth low energy enabled machines to the internet of things," *IEEE Network*, vol. 28, no. 6, pp. 83–90, 2014.
- [16] C. Gomez and J. Paradells, "Wireless home automation networks: A survey of architectures and technologies," *IEEE Communications Magazine*, vol. 48, no. 6, pp. 92–101, 2010.
- [17] J. Ding, M. Nemati, C. Ranaweera, and J. Choi, "IoT connectivity technologies and applications: A survey," *IEEE Access*, vol. 8, pp. 67 646–67 673, 2020.
- [18] C. Zhao, L. Huang, Y. Zhao, and X. Du, "Secure machine-type communications toward LTE heterogeneous networks," *IEEE Wireless Communications*, vol. 24, no. 1, pp. 82–87, 2017.
- [19] F. Restuccia, T. Melodia, and J. Ashdown, *Spectrum Challenges in the Internet of Things: State of the Art and Next Steps*, 2023, pp. 353–375.
- [20] ANT Wireless, *ANT Message Protocol and Usage*, ANT Wireless, a division of Garmin Canada Inc., 2020, revision 5.1. [Online]. Available: <https://www.thisisant.com/resources/ant-message-protocol-and-usage/>
- [21] M. Nemati, Y. H. Kim, and J. Choi, "Toward joint radar, communication, computation, localization, and sensing in IoT," *IEEE Access*, vol. 10, pp. 11 772–11 788, 2022.
- [22] Z. Li and T. He, "WEBee: Physical-layer cross-technology communication via emulation," in *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '17. ACM, Oct. 2017, p. 2–14.
- [23] S. He, K. Shi, C. Liu, B. Guo, J. Chen, and Z. Shi, "Collaborative sensing in internet of things: A comprehensive survey," *IEEE Communications Surveys & Tutorials*, vol. 24, no. 3, pp. 1435–1474, 2022.
- [24] D. Musicki and R. Evans, "Joint integrated probabilistic data association - JIPDA," in *Proceedings of the Fifth International Conference on Information Fusion. FUSION 2002. (IEEE Cat.No.02EX5997)*, vol. 2. Annapolis, MD, USA: Int. Soc. Inf. Fusion, 2002, pp. 1120–1125.
- [25] T. Kropfreiter, F. Meyer, and F. Hlawatsch, "A distributed joint integrated probabilistic data association (JIPDA) filter with soft object association," in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Seoul, Republic of Korea: IEEE, Apr. 2024, pp. 12 906–12 910.
- [26] V. Peterka, "Bayesian approach to system identification," in *Trends and Progress in System Identification*, P. Eykhoff, Ed. Oxford, U.K.: Pergamon Press, 1981, pp. 239–304.
- [27] T. Minka, "Expectation propagation for approximate Bayesian inference," in *UAI '01: Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2001, pp. 362–369.
- [28] S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*. Artech House, 1999.
- [29] D. Simon, *Optimal State Estimation: Kalman, H-infinity, and Nonlinear Approaches*. Wiley-Interscience, 2006.
- [30] W. Yi, G. Li, and G. Battistelli, "Distributed multi-sensor fusion of PHD filters with different sensor fields of view," *IEEE Transactions on Signal Processing*, vol. 68, pp. 5204–5218, 2020.
- [31] Y. Bar-Shalom, F. Daum, and J. Huang, "The probabilistic data association filter," *IEEE Control Systems*, vol. 29, no. 6, pp. 82–100, Dec. 2009.
- [32] A. S. Rahmathullah, A. F. Garcia-Fernandez, and L. Svensson, "Generalized optimal sub-pattern assignment metric," in *2017 20th International Conference on Information Fusion (Fusion)*. Xi'an, China: IEEE, Jul. 2017, pp. 1–8.

Innovative Methods of Signal Processing for Hydrogen-Filled Proportional Counters

Aleš Horna, Jiří Petr, Jan Král, Václav Přenosil, Zdeněk Matěj
*Department of Computer Systems and Communications, Faculty of Informatics, Masaryk University
Botanická 68a, 602 00 Brno*

Ondřej Kolář, Michal Kubiček
*Department of Radioelectronics, Faculty of Electrical Engineering and Communication, Brno University of Technology
Technická 10, 616 00 Brno*

Tomáš Czako, Michal Košťál
*Research Centre Řež
Hlavní 130, 250 68 Husinec-Řež*

Věra Mazánková, František Cvachovec
*Department of Mathematics and Physics, Faculty of Military Technology, University of Defence
Kounicova 65, 602 00 Brno*

Aleš Jančář
*VF, a.s.
Svitavská 588, 679 21 Černá Hora*

Abstract

Hydrogen-filled proportional counters are ionizing detectors suitable for the detection of neutrons and their spectrometry. The detection mechanism involves interaction of neutrons with hydrogen nuclei based on elastic scattering, i.e. proton recoils. Responses of these counters are proportional to the initial ionization in the gaseous filling. They respond not only to neutrons but show also significant sensitivity to gamma rays, which commonly accompany neutrons, and limit the useful energy range of neutron spectrometry. Therefore, an approach for the discrimination of neutron and gamma ray-related events was investigated. Several counters of spherical form were placed in gamma and gamma-neutron radiation fields and their response signals studied. The counters were operated in conjunction with charge-sensitive preamplifier and the output signals (pulses) captured using a custom build data acquisition unit based on FPGA, developed previously, and transferred to PC for post-processing carried out in MATLAB. Both amplitudes and rise-times of the pulses were measured and plotted into correlation diagrams showing the possibility of visual discrimination of neutron and gamma ray-related responses based on the orientation of areas in the diagrams, belonging to particular events. This approach proved efficient also for the discrimination of parasitic events, i.e. not originating due to ionizing radiation, as well. In addition to successful demonstration of ionizing radiation discrimination possibilities, the counters involved in the experiments were studied from the point of energy calibration and energy resolution which was enabled by the presence of helium-3 and irradiation by thermal neutrons, leading to characteristic and well described responses.

Keywords—FPGA, gamma-neutron discrimination, ionizing radiation, proportional counter, proton recoil

PAPER ORIGIN

The presented abstract refers to the paper Innovative Methods of Signal Processing for Hydrogen-Filled Proportional Counters, published in IEEE Sensors Journal, vol. 25, no. 11, on 1 June 2025.

chisel4ml: A Tool For Generating Fast Direct Logic Implementations of Deeply Quantized Neural Networks

Jure Vreča

*Jožef Stefan Institute & Jožef Stefan International Postgraduate School
Jamova 39, 1000 Ljubljana, Slovenia
jure.vreca@ijs.si*

Supervisor: *Dr. Anton Biasizzo*

Abstract

We present chisel4ml - a tool for generating fast direct logic implementations of deeply quantized neural networks. Such neural networks have parameters and activations quantized below 8 bits and have been shown to perform on par with their non-quantized counterparts. Combining these extreme quantization methods with other neural network compression techniques like pruning allows us to implement such neural networks as a direct logic implementation on FPGA. Such highly-parallel implementations can achieve ultra-low latencies in the order of hundreds of nanoseconds and have applications in high-energy physics and network intrusion detection among others.

Keywords— neural networks; quantization; FPGA; Chisel

I. ARTIFICIAL NEURAL NETWORKS

Artificial Neural Networks (ANN) are flexible machine learning algorithms that are useful in a plethora of use cases and can be trained both as classifiers and regression models. They are composed of layers that map the input vector to the output vector. Most layers are composed of artificial neurons. Equation 1 show how to compute a single artificial neuron:

$$y = f(b + \sum_i^N w_i \cdot x_i) = f(\vec{w} \cdot \vec{x} + b). \quad (1)$$

The input vector \vec{x} is transformed into a scalar output y . Each input has an associated weight in the weight vector \vec{w} . Additionally, there is another scalar parameter b called the bias. The function f is a non-linear activation function such as the rectified linear unit (ReLU).

Several neurons can be combined into a fully connected layer by connecting all inputs of the layer to all neurons. This yields Equation 2:

$$\vec{y} = f(W \cdot \vec{x} + \vec{b}). \quad (2)$$

Here \vec{x} represents the input vector, \vec{y} is the output vector, W is a weight matrix, \vec{b} is a vector of bias parameters, and f is a non-linear activation function.

Typically neural networks are defined for real numbers, and in practice various forms of floating-point numbers are used for this task. Because neural networks commonly have a large number of neurons and parameters, storing floating-point parameters can become expensive in terms of memory. This motivated research into quantized neural networks (QNN) that use lower bitwidth fixed-point representations of neural network parameters and activations. Methods for obtaining QNNs are divided into two general categories:

- Post-Training Quantization (PTQ) and
- Quantization-Aware Training (QAT).

PTQ methods take as input a trained floating-point neural network and optionally some sample input data for calibration. Using this, PTQ methods directly produce a QNN. PTQ methods work reasonably well to around 8 bits. Using lower bitwidths with PTQ will typically result in a significant loss in neural network performance compared to the floating-point version.

QAT methods, on the other hand, take quantization into account already during training. Neural network parameters are still stored using the floating-point representation, however, we limit them to discrete values. This is achieved by inserting special functions called fake quantization functions, which round the floating-point values in such a way that they represent some n bit fixed-point number. Fake quantization functions don't actually change the underlying datatype, they simply round the floating-point value. After training we use the fake quantization functions to produce the quantized weights. Figure 1 shows a normal computational graph of a fully connected layer on the left, and a graph with fake quantization functions inserted on the right.

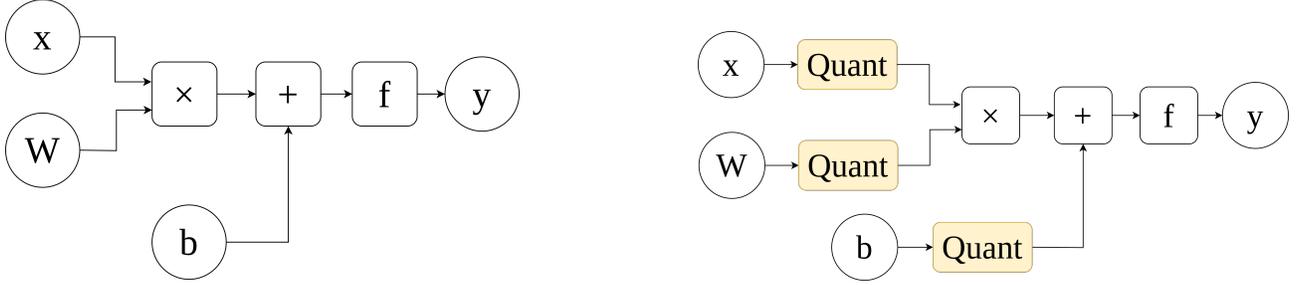


Fig. 1: A normal computational graph of a fully connected layer (a) and a computational graph of a fully connected layer for QAT (b).

QAT methods are typically more successful in retaining the accuracy of the model even when quantizing below 8 bits [10]. Neural networks with parameters and activations quantized below 8 bits are called Deeply Quantized Neural Networks (DQNN). At the extreme such training procedures are able to produce Binarized Neural Networks (BNN). Such neural networks have parameters and activations encoded with a single bit. This allows multiplications to be performed with a single XNOR operation. BNNs work surprisingly well for simple tasks, however they typically fall short for more complex tasks [9]. On the other hand using below 8 bit quantization still yields very good results even for complex problems, and offer a significant reduction in the amount of memory used per parameter [13]. A popular research direction is varying the bitwidths of the neural network layers and trying to achieve the smallest model representation while preserving the accuracy of the model [8].

An important consideration in training QNN are the scaling factors. In a typical neural network parameter range will vary significantly between layers. These factors scale the quantized values to more appropriately match the floating-point values:

$$W \approx \frac{W_q}{S}. \quad (3)$$

Scaling factors can be determined using statistical methods or they can also be learned as part of the training process [3]. The scaling factor values are usually limited to a power of two in order to simplify the implementation. Combining Equations 2 and 3, we get Equation 4:

$$y = f(W \cdot \vec{x} + \vec{b}) \approx f\left(\frac{W_q \cdot \vec{x}}{S} + \vec{b}\right). \quad (4)$$

Presuming W_q , \vec{x} and b are all integers and S is a power of two value, the neural network model consist of only integer operations that can be implemented efficiently in hardware.

II. CHISEL HARDWARE CONSTRUCTION LANGUAGE

The state of the art approaches to implementing DQNNs in hardware are based on High-Level Synthesis (HLS) technology. Two most prominent examples are the hls4ml tool [1], and the FINN tool [4]. The chisel4ml tool, on the other hand, uses the Chisel Hardware Construction Language (HCL) to implement the DQNN in hardware.

Chisel HCL is a language embedded in Scala, a general-purpose functional programming language. It uses the Register-Transfer Level (RTL) abstraction, which means that the designer must provide a structural description of a circuit. In Chisel you are writing a program that constructs an interconnected graph representation of a hardware structure when executed. The interconnected graph representation can be exported as a Verilog code. This is in contrast with HLS methods that rely on algorithmic descriptions, from which hardware structure is inferred. HLS methods have difficulty scaling to larger circuit sizes, as they rely on polynomial time algorithms [6].

To illustrate the Chisel design flow an example of a FIR filter implementation is given. A structure of a simple FIR filter is depicted in Figure 2 and the Chisel code that generates its hardware implementation is given in Listing 1. The FIR structure can be generalized to an arbitrary N -order FIR filter shown in Figure 3. The Scala language enables us to generalize the Chisel description in a concise and precise manner. A generic FIR filter implementation in Chisel is given in Listing 2. This generic implementation represents an arbitrary N -order FIR filter, where the filter is determined by the provided coefficient sequence c . Note that this generic Chisel implementation generates an equivalent hardware design to the previous specific implementation presented in Figure 1 when appropriate coefficients are given. The generic implementation does not introduce any additional hardware structures.

```

1 class FirFilterSimple(bitWidth: Int) extends Module {
2   val io = IO(new Bundle {
3     val in = Input(UInt(bitWidth.W))
4     val out = Output(UInt(bitWidth.W))
5   })
6
7   val z1 = RegNext(io.in)
8   val z2 = RegNext(z1)
9
10  io.out := (io.in * 1.U) + (z1 * 2.U) + (z2 * 3.U)
11 }

```

Listing 1: A simple FIR filter [2].

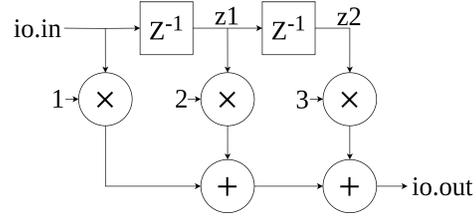


Fig. 2: A simple FIR filter circuit diagram.

```

1 class FirFilter(bitWidth: Int, c: Seq[UInt]) extends Module {
2   val io = IO(new Bundle {
3     val in = Input(UInt(bitWidth.W))
4     val out = Output(UInt(bitWidth.W))
5   })
6   // Create the serial-in, parallel-out shift register
7   val zs = Reg(Vec(c.length, UInt(bitWidth.W)))
8   zs(0) := io.in
9   for (i <- 1 until c.length) {
10    zs(i) := zs(i-1)
11  }
12
13  // Do the multiplies
14  val products = VecInit.tabulate(c.length)(i => zs(i) * c(i))
15
16  // Sum up the products
17  io.out := products.reduce(_ + _)
18 }

```

Listing 2: Generic FIR filter design [2].

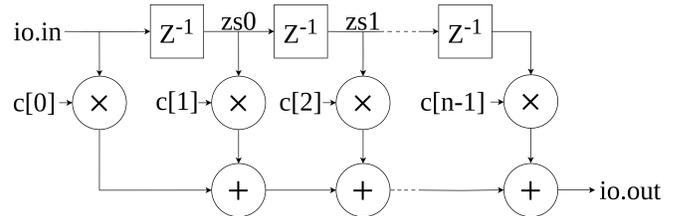


Fig. 3: A generic FIR filter circuit diagram.

III. THE CHISEL4ML TOOL

The chisel4ml tool is composed of two main parts: a Python frontend, which provides the user interface, and a Chisel backend, which generates and simulates the circuits. The high level software architecture is given in Figure 4.

The input to the user interface is a QONNX model [11]. This is a format that introduces fake quantization function to the standard graph based neural network description format ONNX. The Python frontend translates the QONNX format into an internal model representation called Low-Bitwidth Intermediate Representation (LBIR). LBIR is able to represent quantized neural networks of arbitrary bitwidths. It supports different granularity of quantization, including per-tensor, per-kernel and per-channel granularity. The LBIR description is sent to the Chisel backend which then generates the Verilog implementation.

The generated circuit can be simulated directly from Python by sending properly quantized neural network inputs to the backend in the form of a numpy array. This array is fed into an RTL simulation and the results of neural network inference are returned to the Python frontend as a numpy array.

Listing 3 demonstrates the functionality described above with a short Python code snippet.

```

1 from chisel4ml import transform, generate
2 from qonnx.core.modelwrapper import ModelWrapper
3
4 qonnx_model = ModelWrapper('/path/to/model.onnx')
5 accelerators, lbir_model = generate.accelerators(qonnx_model, input_shape)
6 circuit = generate.circuit(accelerators, lbir_model)
7 res = circuit.predict(test_data)
8 circuit.package(directory='~/my/pkg/dir/')

```

Listing 3: Demonstration of chisel4ml Python frontend.

Because chisel4ml uses highly-quantized and possibly pruned neural networks, it is able to generate highly-parallel direct logic implementations, i.e., a combinational circuit with added pipeline registers. Such implementations have ultra-low latencies and are useful in specific applications like high-energy physics [5], network intrusion detection [12], and elsewhere [7].

The chisel4ml tool uses structural descriptions of DQNN, as opposed to algorithmic descriptions like hls4ml or FINN. Consequently, it is able to generate direct logic implementations in orders of magnitude faster, and the implementations use less resources than comparable designs made with hls4ml.

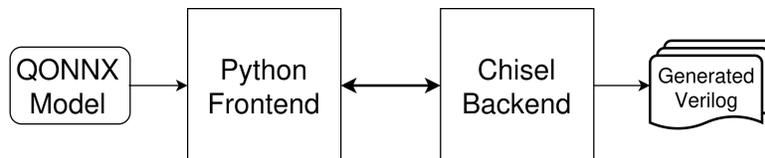


Fig. 4: Software architecture of chisel4ml.

ACKNOWLEDGMENT

The authors acknowledge the financial support from the Slovenian Research and Innovation Agency under grant: No. P2-0098. This work is also part of projects that are funded by the ECSEL Joint Undertaking under grant agreement No 101007273 (DAIS) and by the Chips Joint Undertaking under grant agreement No 101139892 (EdgeAI-Trust).

PAPER ORIGIN

This paper has been published in the journal MDPI Electronics on the 21 February 2025 in the special issue Embedded AI. The DOI number: <https://doi.org/10.3390/electronics14050849>.

REFERENCES

- [1] Thea Aarrestad et al. “Fast convolutional neural networks on FPGAs with hls4ml”. In: *Machine Learning: Science and Technology* 2.4 (July 2021), p. 045015. DOI: 10.1088/2632-2153/ac0ea1. URL: <https://dx.doi.org/10.1088/2632-2153/ac0ea1>.
- [2] Steve Bailey et al. *Chisel Bootcamp*. <https://github.com/freechipsproject/chisel-bootcamp>. Accessed: 2024-11-13.
- [3] Yash Bhalgat et al. “LSQ+: Improving low-bit quantization through learnable offsets and better initialization”. In: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. 2020, pp. 2978–2985. DOI: 10.1109/CVPRW50498.2020.00356.
- [4] Michaela Blott et al. “FINN-R: An end-to-end deep-learning framework for fast exploration of quantized neural networks”. In: *arXiv* 1 (1 2018). ISSN: 23318422.
- [5] Claudionor N. Coelho et al. “Automatic heterogeneous quantization of deep neural networks for low-latency inference on the edge for particle detectors”. In: *Nature Machine Intelligence* 3.8 (Aug. 2021), pp. 675–686. ISSN: 2522-5839. DOI: 10.1038/s42256-021-00356-5. URL: <https://doi.org/10.1038/s42256-021-00356-5>.
- [6] Jason Cong and Zhiru Zhang. “An efficient and versatile scheduling algorithm based on SDC formulation”. In: *Proceedings of the 43rd Annual Design Automation Conference. DAC '06*. San Francisco, CA, USA: Association for Computing Machinery, 2006, pp. 433–438. ISBN: 1595933816. DOI: 10.1145/1146909.1147025. URL: <https://doi.org/10.1145/1146909.1147025>.
- [7] Allison McCarn Deiana et al. “Applications and Techniques for Fast Machine Learning in Science”. In: *Frontiers in Big Data* 5 (2022). ISSN: 2624-909X. DOI: 10.3389/fdata.2022.787421. URL: <https://www.frontiersin.org/journals/big-data/articles/10.3389/fdata.2022.787421>.
- [8] Zhen Dong et al. “HAWQ-V2: Hessian Aware trace-Weighted Quantization of Neural Networks”. In: *Advances in Neural Information Processing Systems*. Ed. by H. Larochelle et al. Vol. 33. Curran Associates, Inc., 2020, pp. 18518–18529. URL: <https://proceedings.neurips.cc/paper/2020/file/d77c703536718b95308130ff2e5cf9ee-Paper.pdf>.
- [9] Itay Hubara et al. “Binarized neural networks”. In: *Advances in Neural Information Processing Systems (NIPS 2016)*, pp. 4114–4122. ISSN: 10495258.
- [10] Itay Hubara et al. “Quantized Neural Networks: Training Neural Networks with Low Precision Weights and Activations”. In: *J. Mach. Learn. Res.* 18.1 (Jan. 2017), pp. 6869–6898. ISSN: 1532-4435.
- [11] Alessandro Pappalardo et al. “QONNX: Representing Arbitrary-Precision Quantized Neural Networks”. In: *4th Workshop on Accelerated Machine Learning (AccML) at HiPEAC 2022 Conference*. June 2022. arXiv: 2206.07527 [cs.LG]. URL: [https://accml.dcs.gla.ac.uk/papers/2022/4thAccML_paper_1\(12\).pdf](https://accml.dcs.gla.ac.uk/papers/2022/4thAccML_paper_1(12).pdf).
- [12] Jure Vreča et al. “Detecting Network Intrusion Using Binarized Neural Networks”. In: *2021 IEEE 7th World Forum on Internet of Things (WF-IoT)*. 2021, pp. 622–627. DOI: 10.1109/WF-IoT51360.2021.9595961.
- [13] Shuchang Zhou et al. “DoReFa-Net: Training Low Bitwidth Convolutional Neural Networks with Low Bitwidth Gradients”. In: 1 (1 2016), pp. 1–13. URL: <http://arxiv.org/abs/1606.06160>.

The Accurate Measurement of Tilt Direction within Small Angles of Tilt

Veronika Junasová, Vladimír Levek, Pavel Šteffan
Department of Microelectronics, Faculty of Electrical Engineering and Communication, Brno University of Technology
Technická 3058/10, Brno, Czech Republic
xjunas00@vutbr.cz

Abstract

This work focuses on the measurement of tilt and the tilt direction specifically in small angles. Particularly problematic is the measurement of the tilt direction within tilt angles up to 10° . In such small angles the direction is excessively impacted by noise and it is impossible to measure it precisely. However, in application where the accurate information about it is needed one accelerometer might not be sufficient. The method proposed in this paper introduces two accelerometers shifted from each other by 45° counterclockwise in the x and z axes at the same time. The axes of the shifted accelerometer are not affected by the noise as considerably and the measurement of tilt direction is more accurate. This claim is supported with mathematically derived equations and measured data. The proposed optimisation allows to measure tilt from vertical line and its direction relatively easily with enlarged precision. This method can be an alternative to the time-consuming utility programs and complex digital filters which would be otherwise needed in applications concerning precise measurement of the tilt direction. Overall, this is a simple solution that does not bring unwanted additional computational or financial costs.

Keywords— accelerometer, tilt azimuth, direction, inclination, multi-accelerometer application

I. INTRODUCTION

The information about tilt is used in applications such as aviation, space applications, medicine, geodesy, mobile phones, and many other areas. The tilt direction is typically needed in mobile phones to determine the orientation of the screen when the mobile phone is tilted [1]. The various applications require different accuracy of measurement. The special instance is when the tilt angle is small. In such case, the error of the determined value of the tilt direction angle can be significant. The tilt can be defined as the inclination from the vertical line. The tilt direction can be also referred to as the tilt heading or tilt azimuth. It contains information about the orientation of the device in the space. It is an angle defined in the plane formed by the axes perpendicular to the vertical line. One of those axes has to be the initial axis, to which the tilt direction is defined.

The three-axis accelerometers are the best choice for the accurate measurement due to better distribution of the applied acceleration over its axes [2]. In the typical configuration of the device, there is one accelerometer, possibly accompanied by another inertial sensor, such as gyroscope or magnetometer. The measurement of tilt using only one accelerometer is relatively accurate over the whole range of motion. The tilt direction can be too inaccurate within smaller angles of tilt. For tilt smaller than 10° the measurement of tilt direction is unreliable. For these circumstances, the tilt direction has to be measured differently or the measurement has to be altered by software. This would add processing time and complexity. Moreover, the system would most likely require digital filtering in addition, which would rise the computational costs and the design of such application could be time consuming.

II. METHODS OF MEASUREMENT

A. Tilt Measurement

The tilt of the device is the inclination from the vertical line. In another words, it is the angle between the accelerometers z axis and the reference axis labeled as z_{ref} . This tilt is conventionally marked ψ ; labels for the inclination from x axis is φ and from the y axis it is θ [3]. The equation for the tilt angle can be derived by applying basic trigonometry and the Pythagorean theorem. Let the output acceleration on axes of accelerometer be marked as A_x , A_y and A_z . When the accelerometer is steady the vector sum of all three outputs A_x , A_y and A_z gives $1g$. The equations for angle of tilt ψ and φ or θ using only information of acceleration from the axis A_z , respectively A_x or A_y of the accelerometer and the known value of $1g$ is [2]:

$$\cos \psi = \frac{A_z}{1g}, \quad \sin \varphi = \frac{A_x}{1g}. \quad (1)$$

When the accelerometer is tilted and still, the gravitational acceleration, the only present acceleration, splits among all accessible axes of the accelerometer in accordance with the tilt angle. To obtain the most accurate measurement, three axes are necessary. In case of an accelerometer tilting in the direction of the missing axis, the part of the gravitational acceleration would not be captured, and the resultant quantification would not be as accurate as when three axes are present [2; 4]. This is part of the reason why three axis accelerometers are the most suitable for tilt quantification. The tilt angle ψ using three-axis accelerometer is given by the quotient of plane, formed by outputs from the axes x and y, and the output of the z axis [2]:

$$\psi = \tan^{-1} \left(\frac{\sqrt{A_x^2 + A_y^2}}{A_z} \right). \quad (2)$$

The tilt from (2) is in degrees and obtains angle in the range from 0° to 180° . Its graphical interpretation is depicted in Fig. 1. There are shown reference axes, and the output accelerations of all axes of accelerometer at certain tilt. It is evident that the output accelerations of the axes x and y are small compared to output acceleration of the z axis. That means that the tilt ψ is not very large because the main portion of gravitational acceleration is divided to the z axis. The resultant of the A_x and A_y acceleration vectors is shown with green arrow. The vector sum of this resultant and the acceleration A_z would be equal to $1g$. With the use of trigonometric tangential function, the (2) can be easily derived.

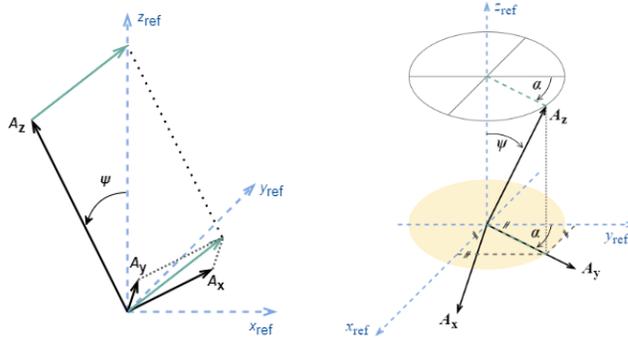


Fig. 1. The graphical interpretation of tangential function of tilt ψ (left) and tilt direction angle α [5].

B. Tilt Direction Measurement

The tilt angle itself does not bring information about its orientation in the space and additional information about the heading of tilt is needed. For this purpose, the tilt direction needs to be defined. The tilt direction can be also called the tilt heading or tilt azimuth. It is the projection of the tilt to the plane formed with the other two axes perpendicular to the vertical line. The tilt angle is the same for the whole circle around the reference axis z_{ref} . That means that for one value of tilt angle the tilt direction can reach angles from 0° to 360° . The tilt direction can be determined using appropriate quotient of the output accelerations A_x and A_y . The denominator determined the axis to which the angle of tilt direction will be related. In other words, it assigns where the 0° angle is placed. When the tilt direction angle α is related to axis y then its equation is given as [6]:

$$\alpha = \tan^{-1} \left(\frac{A_x}{A_y} \right). \quad (3)$$

The tangential function can obtain angles from 0° do 180° . However, the tilt direction can have values up to 360° . Therefore, some modifications need to be provided. There can be used function atan2 , which gives output in radians and broadens the range up to 360° . In that case the output needs to be transformed from radians to degrees. Following up the definition stated in previous part, the z axis of the accelerometer is antiparallel to the gravitational acceleration and its output value is $+1g$ when tilt is 0° . Then tilt direction is the angle between the axis x or the axis y and projection of the z axis to the plane formed by axes x and y . In Fig. 1 the angle of the tilt is defined in relation to the reference axis y_{ref} .

III. MATHEMATICAL HYPOTHESIS AND PROPOSED OPTIMISATION

A. Noise and Distribution of Gravitational Acceleration

Noise is present in any measurement and cannot be eliminated without the use of complex filters. However, it is possible to define the theoretical impact of noise on measurement and then propose the solution for optimization. For this, (1) will be used. The acceleration on each axis has superimposed noise. Let the axis z measure $A_z + \Delta_{Az}$ where A_z is the actual acceleration and Δ_{Az} is the superimposed noise. Similarly for axis x and y , let the output be $A_x + \Delta_{Ax}$, respectively $A_y + \Delta_{Ay}$, where A_x , A_y is acceleration on given axis, and Δ_{Ax} , Δ_{Ay} is superimposed noise on them. Then the equations can be written as follows:

$$\cos \psi = \frac{A_z + \Delta_{Az}}{1g}, \quad \sin \varphi = \frac{A_x + \Delta_{Ax}}{1g}, \quad \sin \theta = \frac{A_y + \Delta_{Ay}}{1g}. \quad (4)$$

When the tilt angle is relatively small, the gravitational acceleration is then distributed in such a way that dominant part of it is measured on axis z and minor portion is measured on axes x and y . Let the tilt angle be so small that the acceleration on the axis z of the accelerometer is $A_z \approx 1g$, the accelerations on the axes x and y are approximately $A_x \approx A_y \approx 0g$. From this presumption, the angles ψ , φ and θ can be written as:

$$\cos \psi \approx \frac{1g + \Delta_{Az}}{1g} \approx 1 + \frac{\Delta_{Az}}{1g}, \quad \sin \varphi \approx \frac{0g + \Delta_{Ax}}{1g} \approx \frac{\Delta_{Ax}}{1g}, \quad \sin \theta \approx \frac{0g + \Delta_{Ay}}{1g} \approx \frac{\Delta_{Ay}}{1g}. \quad (5)$$

From equation (5) it is theoretically possible to deduce that the noise on the axis z should have a lower impact on the calculation of angle ψ than the noises on the axes x and y on angles φ and θ . This is given by the mathematical behaviour of the used functions and derived equations. Equation (3) to determine azimuth utilises only accelerations from axes x and y . When the angle of tilt ψ is small, most of the gravitational acceleration is measured on A_z . Therefore, the accelerations $A_z \approx 1g$ and $A_x \approx A_y \approx 0g$ apply. From there, the equation (3) for azimuth obtaining can be altered as:

$$\tan \alpha = \frac{A_x}{A_y} \approx \frac{\frac{\Delta A_x}{1g}}{\frac{\Delta A_y}{1g}} \approx \frac{\Delta A_x}{\Delta A_y}. \quad (6)$$

In the quantification of tilt azimuth only noise is involved and minimal portion of actual gravitational acceleration. The measurement of the tilt azimuth is inaccurate and of no significance. To improve it, the impact of noise needs to be eliminated.

B. The Proposed Optimisation

The proposed optimization incorporates two three-axis accelerometers; the second accelerometer is shifted by 45° counterclockwise in the axes x and z from the first accelerometer. When there is applied tilt close to the 0° on the first accelerometer, the second one observes tilt close to 45°. The first accelerometer in this instance measures tilt direction with significant error. The tilted accelerometer is able to measure it more precisely compared to the first one. This way the tilt direction can be acquired from the second accelerometer. From it, a 45° angle can be subtracted to get the actual value. The designed device comprises of two boards with two three-axis accelerometers from inertial measurement unit IMU IIM-42652 [7]. The first primary accelerometer is located on the motherboard which has a cutout prepared for insertion of daughterboard. The daughterboard has the second accelerometer mounted on itself. The whole device can be seen in the Fig. 2.

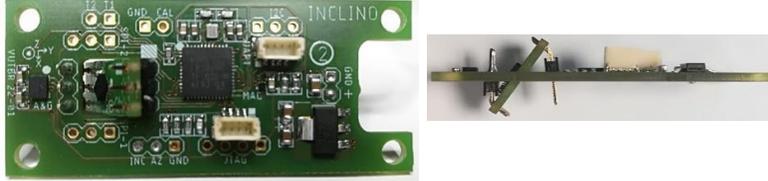


Fig. 2. The designed device with two three-axis accelerometers shifted by 45° [5].

IV. EXPERIMENTAL RESULTS AND DISCUSSION

A. Noise Measurement

To prove assumption and the correctness of the (5) and (6), the experiment was carried out. In Fig. 3 the noise on all axes after calibration of the accelerometer is depicted. It is visible that the noise taken on 200 samples is stochastic on all axes.

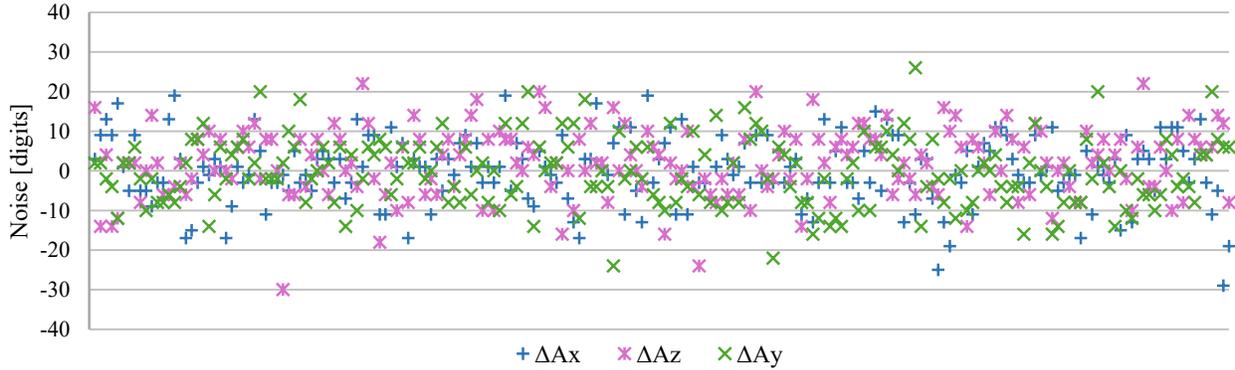


Fig. 3. The measured noise on axis x (blue colour), y (green colour) and z (purple colour) of the accelerometer without tilt.

1000 samples were taken in total from which the angles ψ , ϕ and θ were quantified using (1) and statistically processed. The average value and standard deviation of each angle were calculated, and the results are summarised in Table I. From there it is proved that the noise on the axes x and y has a bigger impact on the accuracy of measurement compared to the noise on the axis z . Therefore, the measurement confirms the theoretical hypothesis.

TABLE I
THE AVERAGE VALUE AND STANDARD DEVIATION OF MEASURED ANGLES

	ϕ	θ	ψ
Average value [°]	-0,0005	-0,0008	0,03740
Standard deviation [°]	0,031	0,029	0,020

B. Tilt Direction Measurement

The tilt direction was measured using both classical and optimized configuration. During the process, the y axis of the primary accelerometer was oriented antiparallel to the gravitational acceleration. Therefore, the tilt was near 0°. The perfect 0° tilt can not be achieved. The obtained tilt angle values were measured close to 1°. On the secondary accelerometer located on the daughterboard was measured tilt angle approximately 44°. The values of tilt on primary accelerometer ψ_{prim} and secondary accelerometer ψ_{sec} are shown in the Fig. 4. There can be also seen the measured values for tilt direction for the primary accelerometer and tilted secondary accelerometer. It is evident that tilt direction of the primary accelerometer a_{prim} (blue colour)

is unstable and noisy. The measurement from secondary accelerometer α_{sec} (orange colour) is more stable compared to the first accelerometer. The difference between tilt directions is caused by small sensitivity on the axes x and z in the classical configuration and substantial noise that is present there.

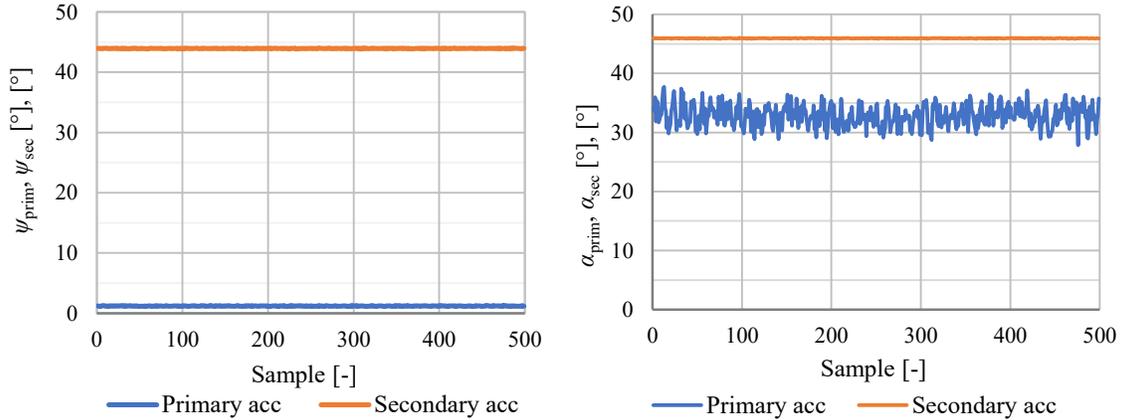


Fig. 4. The measured values of tilt angle (left) and tilt direction (right) from primary accelerometer and tilted secondary accelerometer [5].

The average values of tilt and tilt direction for both configurations are summarized in the Table II. There is also stated the achieved accuracy for tilt direction measurement.

TABLE II
THE AVERAGE VALUES OF TILT, TILT DIRECTION AND ACCURACY OF TILT DIRECTION

	ψ	α	$\Delta\alpha$
Primary accelerometer	1,21	32,77	9,87
Secondary accelerometer	43,94	45,93	0,19

V. CONCLUSION

This article introduces new method for precise measurement of tilt direction angle within the angles of tilt up to 10° . The classical configuration of devices for such measurement uses only one accelerometer. The optimized method adds second accelerometer that is shifted by 45° counterclockwise in the axes x and z from the first accelerometer. The measurement with one accelerometer has several problems. They are small sensitivity of axes used to determine the tilt direction angle and significant noise. The added second accelerometer solves the stated problems. It is shifted in such a way that the sensitivity of the measuring axes is larger, and noise is omissible compared to the output value of these axes. The proposed method shows improved accuracy in the tilt direction measurement. The accuracy of classical configuration is $9,87^\circ$. The newly proposed configuration reaches better resolution, and that is $0,19^\circ$. The future research will be focused on experiments with different angles of shift between accelerometers and their impact on the measurement accuracy.

ACKNOWLEDGMENT

The article was supported by project no. FEKT-S-23-8162, Modern micro- and nanoelectronics for future.

PAPER ORIGIN

This abstract comprises parts from two papers, that have been previously published and presented at 46th International Spring Seminar on Electronics Technology ISSE [5] and 31th Conference Student EEICT 2025, which will be published soon.

REFERENCES

- [1] J. FRADEN, Handbook of Modern Sensors: Physics, Designs, and Applications, 5th ed, San Diego: Springer International Publishing, 2016. ISBN 978-3-319-19302-1.
- [2] C. Fisher, "AN-1057: Using An Accelerometer for Inclination Sensing", Analog Devices, rev. 0, 2010.
- [3] T. Young, "Performance of the Jet Transport Airplane: Analysis Methods, Flight Operations, and Regulations", 1st ed. Wiley, 2018, doi: 10.1002/9781118534786.ch9.
- [4] S. Luczak, "Single-axis tilt measurement realized by means of MEMS accelerometers", Int. J. Mech. Eng. Educ., vol. 18, pp. 341–351, 2011.
- [5] V. Junasová, V. Levek, and P. Šteffán, "The proposal for optimalization of direction of tilt measurement in small angles", in 2023 46th International Spring Seminar on Electronics Technology (ISSE), IEEE 2023. ISBN 979-8-3503-3484-5.
- [6] Ch. KONVALIN, "Compensating for Tilt, Hard-Iron, and Soft-Iron Effects," Questex, 2009.
- [7] IIM-42652 Datasheet: High-performance 6-Axis SmartIndustrial™ MotionTracking MEMS Device for Industrial Applications, TDK InvenSense, 2021.

Gate boosting: Novel gate driver based on an FPGA and reliability estimation

Pavel Tomiček, Jaroslav Boušek
Faculty of Electrical Engineering and Communication, Brno University of Technology
Technická 3580/10, 616 00 Brno
{xtomic05, bousek}@vut.cz

Abstract

Electronic switches are commonly used in all areas of electronics. They are the basis for all digital electronics and are a key part of all switched power supplies. Another important area is instrumentation for experimental physics where devices such as Pockels cells and particle accelerators require high voltage pulse power supplies with sub-nanosecond rise and fall times. These applications usually use a power MOSFET with a gate driver.

In the first part a novel gate driving method called gate boosting is presented, and a gate driver using this method is designed. Gate boosting utilizes three-level driving instead of conventional two-level driving. Adding a short high voltage pulse allows for significant improvements in switching speed. Amplitude of the pulse is higher than the maximum V_{GS} specified by the manufacturer and thus it is necessary that the pulse duration is short enough to ensure the reliability of gate oxide. The timing circuit of the driver consists of a four-channel digital delay generator with 5 ns resolution implemented in an FPGA. This generator creates all the waveforms necessary to turn the driver on and off given the user input. The driver is designed in such a way that allows a direct comparison of switching speed of conventional gate driving and gate boosting.

Second part of this paper consists of an investigation on the effects of gate boosting on the gate oxide reliability of commercially available MOSFETs. Based on parasitic components it is demonstrated that properly designed gate boosting driver will not stress the gate-oxide more than conventional two-level gate drive. However, it is difficult to properly design a gate boosting gate driver as it often relies on the precise value of several parasitic components. As such a theoretical model of gate oxide reliability capable of working on limited data regarding gate-oxide reliability supplied by MOSFET manufacturers is proposed. This model shows that even an improperly designed driver can be safely used in an application requiring long lifetime.

Keywords— gate boosting, gate driver, switching speed, FPGA, gate oxide reliability, TDDB

ACKNOWLEDGMENT

The article was supported by project no. FEKT-S-23-8162, Modern micro- and nanoelectronics for the future.

PAPER ORIGIN

This paper is a combination of two papers previously published and presented at conferences ISSE 2023 and EEICT 2025.

Controllability-Based Circuit Similarity Estimation

Michal Žáček, Petr Fišer

Faculty of Information Technology, CTU in Prague

Thákurova 9, 160 00 Prague 6, Czech Republic

{zacekmi2, fiserp}@cvut.cz

Abstract

Assessing the structural similarity of different implementations of logic functions is of importance in many areas of digital design, such as iterative resynthesis, engineering change order (ECO) based design, design of reliable redundant systems (duplex, TMR), etc. In general, numerous metrics exist that describe such similarity, mostly based on its intended application. In this paper, we introduce a novel metric based on a calculation of the functional equivalence of subcircuits. As this approach requires repeated calls of time-consuming functional equivalence checking, we propose a linear-time approximation of this method based on signal controllability calculation. These two approaches are compared to the state-of-the-art fault detection-based design diversity estimation technique and applied to assess the fault-security capabilities of duplex systems.

Keywords— circuit similarity, diversity, reliability, COP, SAT.

PAPER ORIGIN

Published in: 2025 IEEE 28th International Symposium on Design and Diagnostics of Electronic Circuits and Systems (DDECS)

DOI: 10.1109/DDECS63720.2025.11006818

Side-channel analysis of ChaCha20 stream cipher

Lukáš Daněk, Vojtěch Miškovský, Matúš Olekšák
Faculty of Information Technology, Czech Technical University in Prague
Prague, Czech Republic
{daneklu2, miskovoj, oleksmat}@fit.cvut.cz

Abstract

This paper investigates the side-channel resistance of the ChaCha20 stream cipher on 8-bit and 32-bit microcontrollers. Using the low-cost ChipWhisperer platform, we demonstrate full keystream recovery under nonce reuse via Correlation Power Analysis (CPA). For proper nonce usage, we propose a novel two-phase CPA strategy that successfully recovers key bytes in the first column of the cipher's internal state by exploiting second-round leakage. While full key recovery remains out of reach, the results confirm exploitable leakage and highlight practical risks in ChaCha20 implementations on embedded systems.

ACKNOWLEDGMENT

This work was supported by the Czech Technical University (CTU) grant No. SGS23/208/OHK3/3T/18, by the Student Summer Research Program 2024 of FIT CTU in Prague, and the grant VJ02010010 of the Ministry of the Interior of the Czech Republic, "Tools for AI-enhanced Security Verification of Cryptographic Devices" in the program Impakt1 (2022-2025).

Building a Side-Channel Attack Scheme on SipHash FPGA Implementation

Vít Mašek, Vojtěch Miškovský, Matúš Olekšák
Faculty of Information Technology
Czech Technical University in Prague
Prague, Czech Republic
{masekvit,miskovoj,oleksmat}@fit.cvut.cz

This abstract addresses the vulnerability of SipHash, an ARX-based cryptographic algorithm, against side-channel attacks when implemented on FPGA platforms. Although ARX cryptographic functions—comprising modular addition, rotation, and XOR operations—are widely believed to offer inherent resistance against side-channel analysis, recent evidence has called this assumption into question. The present work systematically investigates these vulnerabilities, demonstrating practical feasibility of side-channel attacks against FPGA-based implementations.

Initially, an FPGA-based SipHash implementation is created using the ChipWhisperer CW308 platform. The design enables flexible communication via UART and AHB bridges, with parameterizable SipHash configurations, and is optimized for performance and resource efficiency. Subsequently, rigorous side-channel leakage assessments employing Welch’s t-test and Pearson’s X^2 test reveal measurable information leakage.

Building upon these insights, a Differential Power Analysis (DPA) attack methodology is developed and refined. Initially leveraging a simplified success rate model, the attack evolves to employ a realistic Hamming-distance-based leakage model. This iterative approach strategically expands the keyspace under consideration, gradually identifying and eliminating incorrect subkeys through targeted analysis of specific internal state bits, specifically beginning from bit $i = 160$ of the SipRound function.

Although the practical attack demonstrates promising results with simulated power models, preliminary experiments with actual FPGA power measurements highlight the complexities and challenges in identifying “bad bits”—state bits whose correlated leakage complicates the analysis. Despite these challenges, the attack provides substantial evidence that ARX-based algorithms like SipHash require careful consideration of side-channel vulnerabilities during hardware implementation.

The findings underscore the critical need for thorough side-channel analysis and dedicated countermeasure designs to ensure secure ARX-based cryptographic implementations. This work thereby contributes important insights and methodologies for future research aimed at enhancing side-channel resilience in cryptographic hardware.

Acknowledgment: This work was supported by the Czech Technical University (CTU) grant No. SGS23/208/OHK3/3T/18, by the Student Summer Research Program 2024 of FIT CTU in Prague, and the grant VJ02010010 of the Ministry of the Interior of the Czech Republic, “Tools for AI-enhanced Security Verification of Cryptographic Devices” in the program Impakt1 (2022-2025).

Partners



ASICentrum, established in 1992 in Prague is a design center of EM Microelectronic and a competence center of ETA, belonging to the Swatch Group. EM Microelectronic is one of the most innovative IC providers. It developed and manufactured the smallest and the lowest power consuming Bluetooth chip on the market, the top performing optical sensors for optical office as well as gaming mice and it was the first to release the award-winning world-first dual-frequency NFC + RAIN RFID emlecho.



NXP Semiconductors: Driving the Future of Technology

NXP Semiconductors is at the forefront of technological innovation, pushing boundaries in machine learning, vision, voice recognition, security, wired and wireless connectivity, wireless charging, embedded operating systems, multicore software, configuration tools, and functional safety.

With a global presence and a team of over 33,100 innovators, NXP develops solutions that power smart cities, autonomous vehicles, and intelligent IoT applications. Our advances in AI-driven processing, secure communications, and high-performance computing are shaping the future of embedded systems.

In the Czech Republic, NXP Semiconductors Czech Republic operates two research centers, strategically located in Rožnov pod Radhoštěm and Brno. Additionally, we foster collaboration through two university meeting points in Ostrava and Brno, supporting the next generation of engineers and researchers.

Our strong, collaborative culture keeps teams stable, engaged, and thriving, with employees enjoying long-term growth, meaningful connections, and competitive rewards.



Tropic Square is a fabless chip and IP company based in Prague that is at the forefront of open and transparent secure semiconductor innovation. The company stands for TRuly Open Integrated Circuits where the focus is on advancing hardware security through developing open source secure cryptographic coprocessors and IPs. The designs are tested by experts in the open source community and made available for community use. This approach enables embedded system designers to verify the security implementations while making informed decisions about their products' security frameworks and threat management strategies.

Sponsors



CTU

**CZECH TECHNICAL
UNIVERSITY
IN PRAGUE**

cesnet
■■■■■■

CESNET is an association of universities of the Czech Republic and the Czech Academy of Sciences. It operates and develops the national e-infrastructure for science, research and education which encompasses a computer network, computational grids, data storage and collaborative environment. It offers a rich set of services to connected organizations.

daiteq

daiteq, established in 2013, provides advanced arithmetic solutions for space-grade processors and FPGAs. For floating-point processing we offer a highly configurable floating-point units compliant with IEEE 754 (2019) that support complex arithmetic and user-defined floating-point number formats, beneficial e.g. for deep learning. For fixed-point processing we offer an implementation of SIMD operations targeted at GNSS processing, low-precision deep learning inference and video compression. Our arithmetic units are complemented by our customised version of the LLVM compiler that supports user-defined floating-point and integer data types.



IMA is a Czech ICT company specializing in the areas of identification, location detection, evidence and Internet of Things. IMA has a long-standing profile as an independent centre focused on the development and application of microcomputer electronics.

In 2017, IMA started working with the German company WITTE Automotive on innovative gesture recognition systems. WITTE Automotive Group, of which IMA became a full partner on January 1, 2021, is a leader in the field of mechatronic locking systems and a major business group with a global presence.

Our systems are used daily by hundreds of thousands of people worldwide by customers such as Škoda Auto, ČVUT, ČEZ, mBank, LEGO ...

We develop smart and innovative identification solutions and always strive to stay a few steps ahead of the competition. Participation in international grant projects aimed at finding new useful solutions for the future helps us to do this.

METIO SOFTWARE

Metio Software is a software development company that develops various kinds of software projects.



STMicroelectronics is a world leader in providing the semiconductor solutions that make a positive contribution to people's lives, today and into the future.

ST is a global semiconductor company with net revenues of \$17.3 billion in 2023.

Offering one of the industry's broadest product portfolios, ST serves customers across the spectrum of electronics applications with innovative semiconductor solutions for Smart Driving and the Internet of Things. By getting more from technology to get more from life, ST stands for life.augmented.



SYSGO is the leading European provider of real-time operating systems for critical embedded applications. Our products have been designed to meet the highest requirements when it comes to Safety and Security. Our customers are leading players in the Avionics & Defense, Space, Railway, Automotive and Industrial Automation and Medical industries, who use our PikeOS product as a platform for critical systems that need to be certified against industry-specific Safety and Security standards.



UJP PRAHA is a Czech company focused on research, development and manufacturing of globally distributed radiotherapy medical systems. Its research and development activities in the field of microelectronics focus on radiation-hardened microelectronics and the development of custom ASICs designed for extreme radiation operating conditions. It is currently the project leader of the Novel Radiation Tolerant Microelectronics for Medical, Scientific, Industrial, Space and Environmental applications which is supported by EU in IPCEI (Important Projects of Common European Interest).



**Student Branch
CTU in Prague**



IEEE Czechoslovakia Section



Proceedings of the 13th Prague Embedded Systems Workshop
Editors: prof. Ing. Hana Kubátová, CSc.; doc. Ing. Petr Fišer, Ph.D.; Ing. Jaroslav Borecký, Ph.D.
First electronic edition
Published by the Czech Technical University in Prague, 2025
Prepared by the Faculty of Information Technology

ISBN 978-80-01-07446-6