KuVS Newsletter

2024-12 Back to Contents

Contents

KuVS Newsletter

Editor Message

Dear KuVS members,

We welcome you to the 20th edition of the KuVS newsletter.

This newsletter provides you with all the information you need to know about the KuVS community from the last half-year. This edition features reports, calls for papers and participations, as well as dissertations completed within the KuVS community.

We are excited to introduce a new member to our editorial team, Bastian Bloessl, who brings fresh energy and perspective to the newsletter. At the same time, we extend our thanks to our former editorial team members (in alphabetical order), Andreas Blenk, Mathias Fischer, Corinna Schmitt for their contributions and dedicated service to this KuVS newsletter over the past years.

At the end of the newsletter, you will again find some riddles from Rolf Windenberg, based on his mathematically oriented reform of English orthography.

More information and past editions of our newsletter are available at https://www.kuvs.de/newsletter/.

We hope you enjoy reading this edition of the KuVS newsletter. On behalf of the editorial team, we wish you a Merry Christmas and a Happy New Year!

The newsletter editors,

Oliver Hohlfeld
University of Kassel

Bastian Bloessl
TU Darmstadt

KuVS Newsletter

KuVS Calls and Announcements

In this section, you find an overview on calls for papers and participation in the german-speaking area.

- **4th GI/ITG KuVS Fachgespräch Network Softwarization**
  Submission deadline: 9.2.2025
  https://kn.inf.uni-tuebingen.de/kuvs-fg-netsoft/2025
- **NetSys 2025 – Call for Workshop Proposals**
  Proposal Submission: 24 January 2025
  https://kuvs.de/netsys/2025/calls/call-for-workshop-proposals
- **DFG SPP Resilient Worlds: Call for Proposals for 2nd Phase**
  Proposal Submission: 31 January 2025
  https://www.resilient-worlds.org/spp2378/call-for-proposals-2nd-phase/
- **SIGCOMM 2025 – Call for Papers**
  Abstract Registration: 24 January 2025
  Paper Submission: 31 January 2025
  https://conferences.sigcomm.org/sigcomm/2025/cfp/
- **SIGCOMM 2025 – Call for Workshop Proposals**
  Submission: 18 February 2025
  https://conferences.sigcomm.org/sigcomm/2025/cf-workshops/
- **NetSys 2025 – Call for Tutorials**
  Tutorial Submission: 14 March 2025
  https://kuvs.de/netsys/2025/calls/call-for-tutorials
- **NetSys 2025 – Call for Extended Abstracts and Hot Topic Papers**
  Submission: 16 May 2025
  https://kuvs.de/netsys/2025/calls
- **NetSys 2025 – Call for Demos**
  Submission: 4 July 2025
  https://kuvs.de/netsys/2025/calls

Kommunikation und Verteilte Systeme

KuVS

IEEE ICC™

IEEE ComSoc IEEE

IEEE International Conference on Communications

Montreal, Canada
08 – 12 June 2025

Communications Technologies 4Good

Call for Papers

4th Int. Workshop on Green and Sustainable Networking (GreenNet 2025)

Energy efficiency and **sustainability** have become of paramount importance in all human activities, including the information and communication technology sector. The trade-off between availability, resiliency, programmability, and energy efficiency of networks and services is a key challenge for the next decade and monitoring methods and metrics for power consumption, energy efficiency, as well as sustainability are important, as well as benchmarking of solutions based on well-defined KPIs.

For that reason, the **goal of the GreenNet Workshop** is to address emerging concepts and challenges related to energy efficiency and sustainability for networked services. An improved sustainability in all parts of the network in a time of increasing AI/ML usage, evolution of different novel network access technologies, the spreading of edge computing and micro-data centers and the demand in computational and data transport capacity across the edge-cloud continuum is essential. **Furthermore, the goal is to not only consider an energy efficient and sustainable access network but complete network, monitoring, and management solutions from data generation to data processing and further usage.**

The trade-off between availability, resiliency, programmability, and energy efficiency is a key challenge. Monitoring methods and metrics for power consumption, energy efficiency, as well as sustainability are important, as well as benchmarking of solutions based on well-defined KPIs.

TOPICS OF INTEREST
We seek original completed and unpublished work not currently under review. **Topics of interest include, but are not limited to:**

- Traffic modeling and prediction for performance and power representation
- Network and device management and control mechanisms
- Usage of digital twins to improve future networks energy consumption and efficiency
- Network measurements and simulations for sustainable and energy efficient future networks
- Benchmarking of energy efficiency and sustainability solutions
- Holistic views on networks, applications, or services from network access to data center
- Emerging networking concepts and technologies to improve energy consumption/efficiency
- AI/ML techniques in the context of energy consumption and efficiency
- Architectural solutions toward network sustainability
- Role of software in reducing network energy consumption and carbon footprint
- Coping with the end of Moore's law
- Role of standardization including network energy efficiency and sustainability metrics
- Consideration of carbon emissions or lifecycle of devices

PAPER SUBMISSION
All papers for Workshops should be submitted via EDAS.
Full instructions on how to submit papers are provided on the IEEE ICC 2025 website:
https://icc2025.ieee-icc.org/

Important Dates

Paper Submission Deadline
20 January 2025

Acceptance Notification
10 March 2025

Camera Ready
31 March 2025

Registration for Accepted Papers
31 March 2025

Workshop Date
08 or 12th June 2025

Organizing Committee
Franco Davoli
(University of Genoa, CNIT, Italy)
Hesham ElBakoury,
(Independent Consultant, Santa Clara, USA)
Timothy O'Farrell
(University of Sheffield, UK)
Tobias Hoßfeld
(University of Würzburg, Germany)
Frank Loh
(University of Würzburg, Germany)
Chiara Lombardo
(University of Genoa, CNIT, Italy)

Workshop Website
https://sites.google.com/view/greennet2025

KuVS Newsletter

Movement of persons

- **Prof. Dr. Amr Rizk** accepted a full professorship for Distributed Real-time Systems at Leibniz Universität Hannover.

KuVS Newsletter

Awards

Best paper award at IEEE Blockchain 2024 Elmira Ebrahimi, Michael Sober, Anh-Tu Hoang, Can Umut Ileri, William Sanders, and Stefan Schulte received the Best Paper Award for the paper "Blockchain-Based Federated Learning Utilizing Zero-Knowledge Proofs for Verifiable Training and Aggregation."

https://www.tuhh.de/ide/detail/2024-12-20-best-paper-award-at-ieee-blockchain-2024

Best Paper Award at Network of the Future (NoF) Fabian Poignée, Anika Seufert, Frank Loh, Tobias Hoßfeld and Michael Seufert received the Best Paper Award for the paper "Fitting the Puzzle: Towards Source Traffic Modeling For Mobile Instant Messaging." https://www.informatik.uni-wuerzburg.de/comnet/news/single/news/best-paper-award-at-nof/

VDE ITG 2024 Award for Tobias Hoßfeld Tobias Hossfeld has been awarded the "Preis der VDE ITG 2024" for the paper: Tobias Hossfeld, Poul Heegaard, Wolfgang Kellerer: "Comparing the Scalability of Communication Networks and Systems" IEEE Access, Volume 11, 2023. DOI: 10.1109/ACCESS.2023.3314201. https://www.informatik.uni-wuerzburg.de/comnet/news/single/news/preis-der-vde-itg-2024/

KuVS Newsletter

Finished PhD Theses

Steffen Lindner

Eberhard Karls Universität Tübingen (advisor Michael Menth)

Title: Advances for Resilient Real-Time Networks using Network Softwarization

Abstract:

Traditional network devices, e.g., routers and switches, provide support for a fixed set of protocols and mechanisms. Network softwarization, and in particular software-defined networking (SDN) and programmable data planes, has emerged in recent years and breaks the previously existing vendor lock-in by decoupling the data and control plane of network devices. Thereby, the packet processing logic of network devices may be directly programmable, which enables the design of novel networking mechanisms without the manufacturer's direct support. As of today, Programming Protocol-Independent Packet Processors (P4) is one of the most common data plane programming technologies.

The research of this thesis focuses on two main objectives that aim to improve the state of the art for resilient real-time networks using network softwarization. The first objective is to develop and evaluate existing and novel resilient forwarding mechanisms using data plane programming. The second objective of this thesis is the support of real-time communication, i.e., the efficient support of BIER in P4, and concepts and algorithms for data transmission with Quality of Service (QoS) requirements. The main results of this thesis are published in five peer-reviewed publications. Seven additional peer-reviewed publications cover additional research results related to the main publications of this thesis, including a comprehensive P4 literature study.

The research presented in this thesis has been funded by different research projects by the Deutsche Forschungsgemeinschaft (DFG) under grant ME2727/1-2, ME2727/2-1, and ME2727/3-1, the German Federal Ministry of Education and Research (BMBF) under support code 16KIS1161 (Collaborative Project KITOS), and the bwNET2020+ project, which is funded by the Ministry of Science, Research and the Arts Baden-Württemberg (MWK).

http://dx.doi.org/10.15496/publikation-92488

Daniel Kopp

Universität Kassel (advisor Oliver Hohlfeld)

Title: DDoS Equilibrium: Internet Measurement Perspective on Attacks and Mitigation

Abstract:

Distributed Denial of Service (DDoS) attacks remain one of the most significant threats to Internet services and infrastructures, which leads to substantial financial and reputational risks

to their targets. Despite their existence for decades, the DDoS landscape continues to evolve as cybercriminals employ and adapt DDoS as an attack tool. Effectively understanding and countering these attacks requires multidisciplinary efforts, where security research provides insights into the latest technologies, tactics, and tools attackers employ. Persistent research enables the development of innovative detection, protection, and mitigation approaches and ensures to stay ahead of this evolving threat.

This thesis first examines one of the sources of DDoS attacks: the emergence of DDoS-as-a-Service platforms that enable anyone, without technical knowledge, to carry out DDoS attacks with the click of a button. We specifically focus on the capabilities of DDoS attacks from these so-called booter services and the effects of law enforcement interventions on the landscape of booter services and DDoS attacks. We observed a takedown leading to a temporary drop in attack traffic and found booter services quickly resuming operations. Cybercriminal activities, like booter services, often rely on bulletproof hosting infrastructures (BPH). Gaining insights into and identifying these infrastructures is essential for comprehensively understanding the DDoS phenomenon. Our study includes an analysis of a prominent BPH that was seized in 2019. We find evidence that can contribute to future identification approaches by examining domain and traffic characteristics.

Understanding attack characteristics and techniques is essential to defending against DDoS attacks and ultimately developing mitigation tools. For this objective, we observe the DDoS landscape from the vantage point of a major Internet Exchange Point (IXP). By formulating and applying detection approaches with large-scale Internet flow measurements, we are able to map the current state of employed amplification vectors and attack techniques. While we confirm the prevalence of well-known amplification attack protocols, we also identify the emergence of DDoS attacks leveraging newly discovered amplification protocols. Furthermore, we focus on attack tactics, particularly investigating the characteristics and prevalence of pulse wave DDoS attacks. Despite historical references, no study has yet quantified these attacks on the Internet. Our analysis reveals that pulse wave DDoS attacks constitute a significant proportion of the DDoS attack landscape, comprising 27% of the total attacks observed.

Leveraging distributed and large-scale Internet infrastructures offers substantial potential and benefits for mitigating DDoS attacks effectively. In this context, we evaluate the advantage of a distributed perspective on DDoS attacks across different platforms. Based on these findings, we design and evaluate an architecture that facilitates the exchange of information about ongoing attacks. Our evaluation demonstrates the feasibility and benefits of collaboratively detecting and mitigating DDoS attacks. Furthermore, we present IXP Scrubber, a Machine Learning (ML) based system for detecting and filtering DDoS traffic at the Internet's core, specifically at Internet Exchange Points (IXPs). The system continuously learns DDoS traffic properties from neighboring Autonomous Systems (ASes) and utilizes BGP signals to drop traffic for specific routes (blackholing). In this way, IXP Scrubber autonomously adapts to new attack vectors without requiring manual intervention, enabling automatic mitigation at the Internet's core.

KuVS Newsletter

Eric Samikwa

Universität Bern (advisor Torsten Braun)

Title: Resource-Aware Distributed Machine Learning for Artificial Intelligence of Things

Abstract:

The integration of machine learning models within the Internet of Things (IoT) ecosystems presents significant challenges, including the expensive execution of deep neural networks (DNNs) on resource-constrained devices, ensuring privacy in data-sensitive applications, and managing the dynamic and heterogeneous nature of data and resources available in IoT devices. This thesis tackles these challenges by proposing a set of novel distributed machine learning strategies, each designed to optimize different aspects of machine learning workflows in IoT systems. These strategies include Early Exit of Computation (EEoC), Distributed Micro-Split Deep Learning (DISNET), Adaptive Resource-Aware Split-learning (ARES), and Dynamic Federated split Learning (DFL), offering solutions to manage inference and training processes across heterogeneous IoT networks efficiently. Through these approaches, the thesis contributes to the emergence of the Artificial Intelligence of Things (AIoT), where intelligent decision-making is integrated with IoT infrastructure. EEoC introduces an adaptive mechanism for optimizing DNN inference tasks by allowing early exits based on the computation intensity and resource availability, thereby reducing latency and conserving energy on IoT devices. DISNET extends this efficiency to a broader scale by implementing a micro-split deep learning approach that facilitates flexible, distributed, and parallel execution of DNN tasks, ensuring minimal inference latency and reduced energy consumption while maintaining accuracy across heterogeneous IoT devices. Building on the foundation of efficient inference, ARES introduces a dynamic, resource-aware approach for the distributed training of machine learning models specifically tailored for edge IoT environments. It significantly accelerates local training processes in resource-constrained devices and minimizes the effects of slower devices on global model performance through device-targeted split points while adjusting to time-varying training conditions. DFL complements ARES by enhancing the federated learning framework to accommodate the heterogeneity and dynamism of IoT devices and training data. It optimizes training efficiency through a resource-aware federated approach with similarity-based clustering, addressing both the heterogeneity of training data and resources. Through comprehensive evaluations in IoT testbeds with heterogeneous data and device resources, this thesis demonstrates the effectiveness of these strategies in improving the efficiency, accuracy, and adaptability of machine learning models in IoT. The contributions significantly advance the field of distributed machine learning, offering scalable and efficient methods for deploying intelligent applications in varied IoT contexts. This research paves the way for next-generation AIoT ecosystems capable of supporting complex, data-driven applications with enhanced responsiveness, energy efficiency, and privacy.

KuVS Newsletter

Julian Dreyer

Univ. Osnabrück / HS Osnabrück (advisors Nils Aschenbruck / Ralf Tönjes)

Title: Securing Communication Networks by Leveraging Hardware-Intrinsic Security Primitives

Abstract:

Creating trust in the digital world that is at least comparable to personal trust between humans is a significant challenge. Physical properties of the devices we use on a daily basis can be the starting point for the creation of such digital trust since these properties are only known by the devices themselves. Such properties and their behaviors are called Physically Unclonable Functions (PUFs) and have first been conceptually described back in 1988 by Kirk Jordan at IBM. Since then, research focussed on finding and improving new ways of exploiting various types of integrated circuits that can be used for the creation of a PUF. However, the real-world application of these implicit security improvements lags behind due to low numbers of evaluations and real-world implementation scenarios. With all of these theoretical concepts at hand, this thesis will, therefore, present such a scenario that can specifically benefit from leveraging the intrinsic hardware properties of the involved devices. In particular, the research presented in this thesis targets the trust and security on the Medium Access Control (MAC) layer in traditional local Ethernet networks due to their low security level in general. To start with, we will first discuss the concept of authentic public key exchange to guarantee an authentic execution of the involved cryptographic routines. This involves a concept called NFC-Key Exchange (NFC-KE), which creates an Air of Trust (AoT) around the involved devices. Along the way, secure random numbers are needed to properly generate cryptographic material. For this, we will derive an entire True Random Number Generator (TRNG), based on a random PUF (rPUF), that warrants the fundamental security of the involved keys. To bridge the gap between the well-researched PUF concepts and the real-world implementations, we will derive clear requirements for microcontrollers to be capable of generating Static Random Access Memory (SRAM)-based PUFs on the internal SRAM of the device. This concept is implemented and evaluated on the widely used ESP32 microcontroller and also extended for arbitrary microcontroller adaptability. By putting these concepts together we will finally be capable of securing networks against the infamous MAC address spoofing attacks that are used by attackers to impersonate other devices on the local network. The SRAM PUF is used to generate unique MAC addresses that cannot be spoofed anymore, thus creating an implicit trust among the connected devices. For this, additional network packets are crafted using the security concepts around TRNGs and the NFC-KE. Moreover, the proposed concept also allows for a distinct randomization of the PUF-derived MAC address which has previously been paradoxical to the concept of device identification. With our concept at hand, developers, researchers, and Internet of Things (IoT) customers will experience new ways of implementing well-researched security primitives in their devices. This can result in new use cases for physical hardware security, especially in the broad IoT domain.

KuVS Newsletter

Alexander Brundiers

Univ. Osnabrück (advisor Nils Aschenbruck)

Title: Taking the Shortcut – Enhancing Backbone Network Traffic Engineering with Segment Routing Midpoint Optimization

Abstract:

Keeping up with the continuous growth in Internet traffic volume as well as the increasing demands and expectations regarding performance and service quality is one of the major challenges for today's Internet Service Providers (ISPs) and other network operators. To achieve this, most of them complement the inevitable but time- and cost-intensive network expansion with some form of Traffic Engineering (TE). The latter enables them to make best use of their existing infrastructure or to dynamically cater to specific operational needs. Over the recent years, its expressive traffic steering capabilities combined with an exceptionally low overhead have lead to Segment Routing (SR) becoming one of if not the premier TE technology choice among many operators, thereby receiving a lot of attention from both industry and academic research. However, all of this research considers SR exclusively in an end-to-end (E2E) fashion, completely ignoring the possibility to deploy it in conjunction with other more flexible steering mechanisms (i.e. IGP Shortcut). In this thesis, we explore the applicability and potential of these so called Midpoint Optimization (MO) concepts (with a special focus on the IGP Shortcut mechanism) for different aspects of SR-based TE in backbone networks. This is done both from a theoretical perspective but also via extensive simulation experiments featuring real-world network data. Thereby, we show that, compared to conventional E2E SR, MO can provide considerable benefits in different TE use cases, including strategic long-term optimization or fast reconfiguration in the face of critical network events. Furthermore, we present corresponding TE optimization algorithms that can be used to compute virtually optimal TE configurations while simultaneously adhering to important real-world constraints and operational requirements. Thereby, we bridge the gap from pure theoretical observations towards practically deployable TE approaches that should enable (e.g., operators) to more or less immediately utilize and benefit from the findings of this thesis in practice.

KuVS Newsletter

Markus Knecht

University of Zürich (advisors Burkhard Stiller and Peter Müller)

Title: A Robust Smart Contract Approach based on Substructural Types and On-Chain Code Verification

Abstract:

The Ethereum blockchain was the first to introduce Smart Contracts, which are code executed on a blockchain guaranteed to behave as specified. After deploying a Smart Contract, its code becomes immutable, and deploying vulnerable code has led to numerous incidents of substantial financial loss. Unfortunately, not every party interested in Smart Contract development can afford the expensive security audits necessary to ensure their code is not exploitable. This financial barrier, combined with the high level of skill required, makes robust Smart Contract development a challenging prospect.

A carefully designed Smart Contract programming language can mitigate problems and challenges inherent to developing robust and secure Smart Contract systems. Such a language, along with corresponding execution environment, designed with these aspects in mind, thereby enhancing Smart Contract systems' security and robustness. This thesis explores available programming paradigms, concepts, and features for their applicability in a programming language designed to address Smart Contract-specific problems and challenges, especially regarding robust and secure code. Using such a programming language lowers the barrier to entry for secure Smart Contract development.

The outcome of this exploration is the design of two programming languages and corresponding systems: Mandala, a high-level Smart Contract programming language, and Sanskrit, a low-level byte code with a corresponding Smart Contract execution environment. Sanskrit, during the deployment of Smart Contracts, enforces many static guarantees, enabling the use of advanced concepts that are guaranteed to hold for all deployed code. Mandala and Sanskrit utilize immutable algebraic data types and pure functions grouped into modules, a departure from the traditional object-oriented approach. This design choice allows for state isolation using the functional programming paradigm only. These and other features enable the implementation of common Smart Contract concepts like tokens, assets, access control, and fee payments in a flexible, elegant, and secure manner.

KuVS Newsletter

Ahmad Khalil

TU Darmstadt (advisors Ralf Steinmetz, Tobias Meuser and Antonio Fernández Anta)

Title: Vehicular Communication for Collective Perception. Adaptive and Communication-Efficient Vehicular Perception Mechanisms

Abstract:

The global road safety report for 2023 highlights a worryingly high number of 1.19 million road traffic fatalities per year. Notably, European road accidents, predominantly caused by human errors, highlight the urgent necessity for advanced vehicle automation to supplement or replace human driver control. Effective autonomous driving systems rely on reliable vehicular perception, crucial for vehicles to accurately perceive their surroundings. Vehicular perception must deliver consistently reliable results across diverse scenarios and conditions. Various methods, such as object detection models realize vehicular perception. Presently, deep neural networks dominate object detection in vehicular contexts due to their superior performance. However, the conventional approach of centrally training object detection models involves transferring vehicle data to central servers, which raises concerns about privacy and bandwidth usage. Furthermore, this method leads to models (situation-agnostic) that perform poorly in specific conditions such as adverse weather or complex traffic scenarios. To address these challenges, the first contribution of this thesis introduces the Situation-aware Collective Perception (SCP) framework which provides two-stage perception. Environmental situations are initially detected, and then the situation-aware object detection models are employed based on the detection situation. This framework enables transitioning from conventional situation-agnostic to situation-aware object detection. When integrated with V2X communication, our framework supports collaborative and online model training directly in vehicles. Evaluations under diverse environmental conditions demonstrate that using situation-aware models notably improves detection performance. Moreover, a notable finding emerging from our first contribution highlights the negative impact of increased data heterogeneity on the performance of the object detection models. To effectively mitigate the adverse influence of data heterogeneity, the second contribution focuses on mitigating data heterogeneity by introducing two novel methods that leverage both data characteristics and model parameters to enable an optimized training process. Evaluation of our methods using camera data from the vehicular field demonstrates a remarkable enhancement in model training convergence and detection performance compared with state-of-the-art approaches. Lastly, the third contribution introduces the Adaptive Resource-aware Clustered Federated Learning (AR-CFL) framework. By integrating with the collaborative and online perception model training paradigm of our SCP framework, AR-CFL optimizes communication efficiency through the creation of efficient vehicle clusters for localized model training. Overall, the proposed approaches are promising for increasing vehicular perception by enabling adaptive perception while considering communication efficiency.

KuVS Newsletter

Christoph Gärtner

TU Darmstadt (advisors Ralf Steinmetz and Amr Rizk)

Title: Flexibility in Real-time Networks: Analytical Approaches for Adapting Time-Sensitive Networks to Dynamic Traffic Requirements

Abstract:

In the context of Industry 4.0, achieving flexible and deterministic data delivery is crucial, particularly within networks that leverage Time-Sensitive Networking (TSN) – an extension of traditional Ethernet – for real-time service guarantees. Current TSN configurations, however, are predominantly static and do not offer the adaptability required to cater to the continuously evolving demands of applications. Adaptations become increasingly complex when changes need to be incorporated during operation, and applications demand more stringent service guarantees. This complexity stems from elaborate TSN planning requirements, resulting in static switch configurations. Furthermore, these configurations lack indicators for the reconfiguration potential, thus making and predicting adaptations difficult. The static nature of TSN schedule configurations typically necessitates complete reconfigurations, resulting in service downtimes for applications in the network.

To address these challenges, this thesis introduces a novel flexibility metric called ”flexcurve” designed to enhance the management of dynamic TSN networks. The introduction of the flexcurve metric aims to quantify the flexibility at network bottlenecks by considering frame size and other traffic requirements, thereby measuring possibilities for the integration of new traffic without extensive rescheduling. Consequently, it enables a central TSN controller to perform more informed traffic admissibility checks, path selection, and scheduling adjustments during the configuration planning phase. Utilizing a path-based approach, the flexcurve metric captures the essence of end-to-end scheduled traffic. Furthermore, the formulation allows for both disaggregation and aggregation, enhancing efficiency. Disaggregations enable traffic admissibility checks of different sources simultaneously, while aggregations facilitate the rapid construction of the metric from existing configurations.

The thesis builds on the flexcurve metric and introduces a search heuristic algorithm that employs the flexcurve for dynamic scheduling. This algorithm is designed to generate eligible candidates and incorporates a secondary heuristic for pruning to select the most promising candidates. Additionally, methods are included for choosing paths that optimize the value of the flexcurve.

As final contributions, the deployment capabilities of TSN to non-TSN hardware have been augmented by leveraging programmable Push-In-First-Out (PIFO) queues and by developing the Residence Delay Aggregation (RDA) method, which ensures per-flow delay guarantees on programmable switches. Additionally, the flexcurve has been further extended by integrating multi-mechanism support, significantly improving network flexibility and broadening TSN deployment opportunities.

KuVS Newsletter

Leonie Reichert

TU Darmstadt (advisor Björn Scheuermann)

Title: Privacy-Preserving Data Analysis and Distributed Processing in Pandemic Settings and Beyond

Abstract:

Privacy is acknowledged as a fundamental human right and essential for the functioning of modern democracies, particularly as research and the economy become increasingly data driven. The Covid-19 pandemic has given rise to many new applications necessitating the processing of sensitive information such as health, location, and proximity data. Notable examples include discovering new infections by retracing the contacts of diagnosed individuals and identifying super-spreader events through presence tracing. To fight a pandemic, gaining meaningful statistical insight on the current epidemiological situation based on health data - or other sensitive information - is important. For all these applications, the processed data can reveal private information regarding the data providers. Therefore, data providers require concrete privacy guarantees at every step.

This thesis focuses on solutions for processing and analyzing sensitive data in a privacy-preserving way without requiring trust being place in a central authority. Multiple approaches are proposed to ensure privacy during distributed data processing for Digital Contact Tracing (DCT). To establish a general understanding of the topic, an introduction to the problems and solutions for DCT is presented. The literature is systematized and common challenges with regard to privacy, security, and functionality are identified. Based on the shortcomings of existing contact tracing applications, novel designs for privacy-preserving DCT are presented, along with their respective advantages and drawbacks. The focus is on distributing the tracing process and risk-scoring tasks to users while mitigating the leakage of private data through metadata. Strong privacy guarantees are also provided by using cryptographic primitives such as blind signatures, Oblivious Random Access Memory (ORAM), and Private Set Intersection (PSI). Such techniques allow the design of protocols that only reveal the minimal required amount of information to all parties involved. Systems for super-spreader detection through presence tracing are also presented that can be integrated with DCT systems in a privacy-preserving manner.

While decentralized processing provides better privacy than the centralized alternative, it limits the ability to observe the epidemic situation through statistical analysis. By reviewing common approaches for collecting and analyzing health data for research purposes, we identify various threats to the privacy of people who are willing to share their data. Both in the pandemic and post-pandemic settings, privacy guarantees are a tool to ensure to data providers that their data can not be misused. To this end, a platform is presented that leverages Trusted Execution Environments (TEEs) in combination with oblivious algorithms that safeguard sensitive data during data collection and analysis. To combat the drawbacks of TEEs, new methods are introduced to hide the access patterns and volume patterns of database queries. All contributions presented in this thesis aim to improve the privacy of individuals through solutions that follow the concept of privacy by design.

KuVS Newsletter

Pegah Golchin

TU Darmstadt (advisors Ralf Steinmetz and Andreas Mauthe)

Title: Machine Learning Models in Network Intrusion Detection Systems: Self-Supervised Detection of Malicious Flows and Traffic Patterns Recognition in Programmable Networks

Abstract:

The recent increase in cyber-attacks highlights the critical need for reliable Network Intrusion Detection Systems capable of detecting anomalies before they inflict substantial damage. Conventional intrusion detection methods often fail to classify previously unseen intrusion patterns accurately. This shortfall is exacerbated by the emergence of new network intrusion types and the evolving nature of network structures. Machine Learning (ML) models address this need by learning representations of network traffic flows. Nonetheless, challenges persist, particularly in ensuring their adaptability and ability to generalize in detecting various network traffic patterns and integrating them into programmable networks. The first contribution of this thesis highlights the presence of diverse flow feature patterns in existing network traffic patterns. To mitigate the impact of these disparities on the final detection performance and minimize noise in flow features, thereby reducing the complexity of ML models, an Ensemble Feature Selection approach is devised. This method integrates statistical and ML-based feature selectors, taking into account the imbalance of benign and attack traffic to avoid biased feature extraction. Evaluation results demonstrate the potential to attain high detection performance with a reduced flow feature dimension. Additionally, a data-driven approach is incorporated into the proposed feature selection method to improve the transferability of selected flow features across different network traffic patterns. The second contribution aimed at tackling two main challenges: the limited availability of annotated network traffic flow data required for training ML models and the limited ability of ML models to generalize across various network traffic patterns. To overcome these challenges, a Self-Supervised Contrastive Learning approach is introduced, which is specifically trained on benign flows to learn the abstract representation of benign flow patterns. The results illustrate improvements in the generalization of detection performance across diverse network traffic patterns. These improvements surpass the performance of both supervised and unsupervised ML models used as baselines. The last contribution explores integrating ML models into programmable networks, particularly following the Software-Defined Networking paradigm, which separates the data plane from the control plane. However, deploying complex ML models in the control plane can increase the risk of overwhelming it, given the necessity to forward flows through it. Conversely, employing lightweight models with few trainable parameters in the data plane may compromise detection performance. To tackle these challenges, we propose a collaborative ML-based intrusion detection approach. This approach facilitates cooperation between ML models deployed in the data plane and the control plane based on the confidence level of the deployed ML model in the data plane. Using this approach, a balance is achieved between attaining high detection performance and speed while reducing network load.

KuVS Newsletter

Pratyush Agnihotri

TU Darmstadt (advisor Ralf Steinmetz)

Title: Accurate Performance Modeling for Distributed Stream Processing - Methods for Performance Benchmarking and Zero-shot Parallelism Tuning in Distributed and Heterogeneous Environments

Abstract:

Distributed Stream Processing (DSP) systems have emerged as a pivotal paradigm, enabling real-time data analysis using distributed cloud resources. Major Internet companies like Amazon and Google, build on DSP systems for their real-time data workloads. For instance, Amazon provides Apache Flink as a service for implementing DSP workloads. Parallelism is often a desired property of DSP workloads to meet the timeliness and scaling requirements of current applications, necessitating the use of distributed and multi-core cloud resources. However, cloud resources are heterogeneous in nature, which makes understanding the performance of DSP workloads very difficult, as it depends on highly varying resources, i.e., compute, storage, and network. Therefore, (i) understanding the performance and (ii) predicting it for distinct DSP workloads on such heterogeneous cloud environments are both very challenging problems. This thesis solves these two fundamental research challenges by contributing methods for accurate performance modeling of DSP workloads in heterogeneous cloud environments.

First, this thesis contributes to methods for performance understanding by proposing PDSP-BENCH, a novel benchmarking system. It tackles three primary challenges of existing work: lack of expressiveness in benchmarking parallel DSP workloads, the need for heterogeneous hardware support, and the need for integration of learned DSP models. Unlike existing systems, PDSP-BENCH enables the evaluation of parallel DSP applications and workloads using both synthetic and real-world applications, offering an expressive and scalable solution. Further, it facilitates the systematic training and evaluation of learned DSP models on diverse streaming workloads, which is crucial for optimizing performance. The extensive evaluation of PDSP-BENCH demonstrates its benchmarking capabilities and highlights the impact of varying query complexities, hardware configurations, and workload parameters on system performance. The key observations of our experiments show the non-linearity and paradoxical effects of parallelism on performance.

Second, this thesis contributes to methods on performance prediction and optimization by proposing ZEROTUNE, a novel learned cost model for DSP workloads and an optimizer for parallelism tuning. It provides highly accurate cost predictions while generalizing to (unseen) heterogeneous hardware resources of the cloud. The generalizability of the model is based on transfer learning, the same technique that is used in Large Language Models like ChatGPT. The main idea is to learn from so-called transferable features and parallel graph representation that together enable the model to generalize to unseen DSP workloads and hardware. Our extensive evaluation demonstrates ZEROTUNE's robustness and accuracy across workloads, various parallelism degrees, unseen operator parameters, and training data efficiency. The evaluations show significant speed-ups with parallelism tuning compared to existing methods. Most notably, our approach has been adopted by Amazon Redshift for query execution time prediction.

KuVS Newsletter

Yassin Alkhalili

TU Darmstadt (advisor Ralf Steinmetz)

Title: Content-aware Adaptation of Point Cloud Streams – A Model-based Perspective on Processing of Point Cloud Streams

Abstract:

The utilization of point clouds, a three-dimensional (3D) data representation, has recently experienced a significant rise in popularity. Recent advancements and affordability in 3D sensor hardware have played an essential role in driving the widespread adoption of 3D point clouds across diverse domains. These domains include virtual reality, augmented reality, volumetric video, 3D sensing for robotics, smart cities, telepresence, and automated driving applications. Consequently, the availability of point clouds, characterized by millions of data points per frame, has been increasing steadily since. However, the substantial size of point cloud data presents significant challenges regarding efficient transmission, compression, and processing. Existing methods for point cloud compression tend to prioritize data quality preservation, often overlooking the practical utilization of the data. For instance, the primary concern in autonomous vehicles is machine perception tasks, such as vehicle positioning and object detection. Thus, the focus should be more on relevant objects and less, or not on other irrelevant surrounding objects. Dedicating resources to these task-specific objects conserves valuable transmission bandwidth and enhances the overall utility of the point cloud data at the recipient's end.

In this thesis, we consider the unique characteristics of point clouds, particularly the capability to extract individual objects from the original point cloud, which are then processed and streamed independently. Our first contribution improves the understanding of how various configurations, including compression-related parameters, distance, and reduced frame rate, influence the quality of point cloud objects. For our second contribution, we investigate how these configurations affect both output quality and resource demands. Examining these relationships aims to identify configurations that maximize quality while minimizing resource consumption. Based on the collected data, we generalize our contributions by building a machine learning-based model to predict the quality of a given point cloud.

To enable the adaptability of the point cloud content to changing conditions, our third contribution explores how incorporating object-related information, such as object semantics, into point cloud content streaming impacts adaptability and delivery efficiency compared to conventional methods. Through experiments, we extensively evaluate our contributions and show the significant benefits of content-aware streaming. This approach has the potential to enhance point cloud streaming by enabling dynamic content delivery that adapts to changing scenarios.

In conclusion, this thesis introduces the concept of content-aware point cloud adaptation and compares it with alternative state-of-the-art approaches. The contributions of this thesis represent an initial step towards demonstrating the feasibility of content aware adaptation for enhancing point cloud delivery efficiency.

KuVS Newsletter

Project News

State of Hesse to Continue Funding emergenCITY Until the End of 2026

The LOEWE boards approved the continued funding of around nine million euros.

The researchers at the LOEWE Center "emergenCITY" are working on resilient infrastructures of digital cities that can withstand crises and disasters. Photo: Gerd Keim

"Once again, the LOEWE center emergenCITY was able to convince the LOEWE boards and the Hessian Ministry of Science and Research, Arts and Culture with a comprehensive concept paper. On December 4, they decided to release a further nine million euros in state funding from the LOEWE program, which promotes excellent research in Hesse, for the years 2025 and 2026.

This will enable the interdisciplinary research network, in which the Technical University of Darmstadt plays a leading role as well as the University of Kassel and the Philipps University of Marburg, to continue its scientific work in phase 2. Since 2020, the center has been researching resilient information and communication technology to better protect cities from disasters. emergenCITY has already received around 22 million euros from the LOEWE funding program."

Further Information: www.emergencity.de

KuVS Newsletter

Ilmenau School of Green Electronics – Advancing Sustainable IT

The Ilmenau School of Green Electronics (ISGE), funded by the Carl-Zeiss-Stiftung, is a newly established graduate school at Technische Universität Ilmenau. Focused on cutting-edge research in sustainable information technology, ISGE unites 13 PhD projects, each targeting transformative approaches to energy- and resource-efficient computing.

Key topics include low-energy data processing, neuromorphic architectures, and bio-inspired superconducting circuits. PhD projects under participation of the Distributed and Operating Systems group cover optical in-network processing for data-intensive applications, and superconducting circuits inspired by dendritic architectures as an ultra-efficient computing paradigm. A defining feature of ISGE is its strong emphasis on interdisciplinary collaboration across computer science, electrical engineering, physics, material sciences, and media technology. This fosters holistic solutions and enables groundbreaking research outcomes.

Further information: https://www.tu-ilmenau.de/imn/isge/projects

KuVS Newsletter

Event Reports

3rd Würzburg Workshop on Next-Generation Communication Networks (WueWoWAS)

Stefan Geißler, Amr Rizk, and Tobias Hoßfeld (Organization Committee)

The third edition of the Würzburg Workshop on Next-Generation Communication Networks (WueWoWAS) took place on September 30th and October 1st, 2024, in the beautiful city of Würzburg. More than 50 participants from industry and academia discussed recent advances in 6G Networks.

Participants of WueWoWAS'24 at the Infosim building in Würzburg.

The scope of the WueWoWAS Workshop is to provide a platform for connecting young and senior researchers from industry and academia who are working on next-generation communication networks and their performance evaluation. The focus this year's iteration was 6G networks, and invited speakers from industry presented their vision on 6G. Maciej Muehleisen (Ericsson) presented his view on the (radio) network concept for verticals. Christian Henke (emnify) showed advances in converged cellular and non-terrestrial networks (NTN) with a particular focus on global coverage for IoT. Andra Lutu (Telefonica Research, Spain) presented key challenges to be addressed in 6G, like local breakouts for international mobile roaming in the presence of trust issues or viable global operator models for next-generation networks.

The program included 20 presentations across five technical sessions. Session 1, chaired by Maciej Muehleisen, explored advancements in P4 and hardware, featuring dynamic eBPF integration, universal control plane development, and frameworks for multi-tenant P4 network devices. Other highlights included leveraging hybrid P4 solutions for 5G programmability and innovative approaches to energy measurement and traffic generation. Session 2, chaired by Christian Henke, focused on modeling advancements, including synchronization techniques for network digital twins and data-driven approaches for digital twins in microservice architectures. Additionally, the session introduced LRU-BottomUp, a caching strategy combining LRU-level performance with FIFO-like efficiency. Session 3, chaired by Prof. Stanislav Lange, highlighted critical issues in security and reliability. Topics included the impact of thermal effects on timestamping in high-speed 100G networking devices, an exploration of attack vectors in 5G and future networks, and real-time monitoring solutions in SDN to counter denial-of-service attacks. Session 4, chaired by Andra Lutu, explored innovations in 5G networks and beyond. Presentations covered frameworks for non-terrestrial networks and V2X, the impact of 5G New Radio on dataplane traffic, and energy consumption studies in converged networks like LTE-M and NB-IoT. Additional highlights included a generalizable fine-granular 5G resource model and adaptive traffic steering techniques using predictive QoS metrics for optimized application performance. Session 5, chaired by Prof. Andreas Kassler, discussed advancements in Wireless Time-Sensitive Networking (TSN). Key topics included P5G-TSN, a private 5G TSN simulation framework, and a review of deterministic communication demands in hybrid 5G-TSN networks. Furthermore, the integration of 5G with TSN was discussed for over-the-air communication and robust TSN scheduling models designed to handle wireless performance uncertainties.

For the first time this year, we organized a speedmentoring session. Participants had the opportunity to book a mentoring slot with experts from industry and academia to discuss their own work, recent developments, or other professional matters. The workshop's mentoring panel featured esteemed experts Andra Lutu (Telefonica Research), Wolfgang Kellerer (Technical University of Munich), Maciej Muehleisen (Ericsson), and Christian Mannweiler (Nokia), providing participants with valuable insights and guidance. The post-workshop survey indicated that the speed mentoring was well received by mentees and mentors. Due to the very positive feedback, we will integrate the speedmentoring again in the next WueWoWAS edition.

In the workshop, we also integrated a session on the VDE ITG expert group KT 2 "Communication networks and systems" (https://lists.uni-wuerzburg.de/mailman/listinfo/itg-kt2). Thereby, Damian Dudek (Managing Director VDE ITG) gave an overview on the research topics, but also the role of VDE ITG in Germany. The discussion session showed that there is a strong interest in industry collaboration and industry perspectives for young but also senior researchers. Conducting research, validating ideas, and comparing market needs with the current state of the art are crucial for aligning industry insights with real-world implications, use cases, and ongoing research activities. Effective technology collaboration involves leveraging shared platforms, accessing advanced hardware, and utilizing test beds to bridge the gap between innovation and market readiness. As key topics and interest for the next workshop, participants mentioned the evolution of mobile networks (5G and 6G), wireless connectivity advancements, and emerging technologies like neuromorphic computing and large language

models, since they are driving innovation in network programmability, energy efficiency, and AI integration. Key areas also include satellite communication, network testing tools, and digital twins to enhance network operations and performance.

Razvan-Mihai Ursu has received the Best Paper Award for his contribution “Towards Data-Driven Approaches for Network Digital Twins of Microservice-Based Architectures”, authored by Razvan-Mihai Ursu, Navidreza Asadi, Leon Wong, Wolfgang Kellerer.

We would like to thank all the authors and attendees for their contributions to the success of the workshop! We are looking forward to welcoming you in Würzburg for the **WueWoWAS’25, 6–8 October 2025**.

KuVS Newsletter

Open Positions

Faculty Positions - Professorships:

- **Embedded Systems and Networks (W2)**, Frankfurt University of Applied Sciences https://www.frankfurt-university.de/fileadmin/standard/Aktuelles/Stellenangebote/P50-2024_Ausschreibung_VES_BK_en.pdf

Open Postdoctoral and PhD positions:

- Open Postdoctoral and PhD positions at HPI, Potsdam in the area of Internet Measurements: https://hpi.de/bajpai/jobs.html
- Open Postdoctoral and PhD positions at TU Berlin in the area of telecommunication networks: https://www.tkn.tu-berlin.de/jobs/

KuVS Newsletter

COMSYS @ RWTH Aachen is Hiring

We currently have open positions for PhD students and PostDocs available that are to be filled within the next 12 months. In particular, we are looking for (young) talents who are dedicated and passionate about their interests and would like to work in and contribute to our collaborative team atmosphere.

Research Topics

The available positions are funded by different projects spanning our entire research portfolio, ranging from Industry 4.0 and large-scale Internet measurements to developing security and privacy mechanisms for health data or improving the reliability of critical infrastructure. Moreover, we conduct research in the area of coordination, networked control, and security/privacy in the (Infrastructural) Internet of Things, including (Smart) Energy Grids. You can find a general overview of our research at our website (https://www.comsys.rwth-aachen.de/research/research-vision) and more details by looking through the individual profiles of our staff.

Application Process

We are accepting applications on a rolling basis and strive for timely notifications.

Please submit your meaningful application, at least consisting of a short CV, a transcript of records, and a motivation letter that describes your research interests and your motivation to pursue your PhD with us, as pdf to **phd[[at]]comsys.rwth-aachen.de**. For PostDoc applications, please also sketch a research outline covering what you would like to work on in the next 12 months.

More information:
https://www.comsys.rwth-aachen.de/research/phd-and-postdoc-positions

Details on the PhD process:
https://www.comsys.rwth-aachen.de/teaching/phd-process

Fun

How 2 Shor10 English Texts

Riddles Based on a "Mathematically Oriented Reform" of English Orthography

Rolf Windenberg (alias: Nigel Fred Brown)

The Rules:

1. Usage of mathematical symbols and of numbers
2. Capital letters are pronounced as in the alphabet

Examples:

(Trafalgar)^2 [meaning: Trafalgar Square]

$\sqrt{66}$ [meaning: Route 66]

Y R U so Z 2dA ? [meaning: why are you so sad today ?]

The Riddles *(Solutions, see on next page)***:**

- *Beginners:* **with U, I Njoy T 4 2**
- *Playing with Capital Letters:* **Nd – ly**
- *Advanced Persons:* **V ∀wAs used a str8 4ward wA 2 communic8**
- *Experts:* **I will stRt 2 plA gol50 will not B served soon**
- *Geniuses:* **the 5er 5d his 5**

Fig. 1: Illustration to assist the reader in solving the fifth riddle (source: [1])

[1] Windenberg, R., Hasselfang, R.W.: How 2 Shor10 English Texts. Shaker Media Verlag, Düren, ISBN 978-3-95631-590-9, 2017

Solutions of the riddles (by Rolf Windenberg):

- with you, I enjoy tea for two [because: with-*U, I-N*-joy-*T-four-two*]
- endlessly [because: *N*-d-*less*-ly]
- we always used a straightforward way to communicate [because: *V-all*-w-*A*-s-used-a-str-*eight-four*-ward-w-*A-two*-communic-*eight*]
- I will start to play golf if tea will not be served soon [because: *I*-will-st-*R*-t-*two*-pl-A-gol-*f.if.ty*-will-not-*B*-served-soon]
- the fifer fifed his fife [because: the-*five*-er-*five*-d-his-*five*].

KuVS Newsletter

Next Newsletter - Deadline 15 May

Next newsletter: June 2025

Deadline for submissions and contributions: 15 May 2025

We ask you for submissions in English. Topics can be from the following time frame: December 2024 - May 2025

- Fachgruppe KuVS
  - Geschäftsberichte der GI – KuVS – Fachgruppe
  - ...
- News from the working groups
  - Dissertations
  - Awards
  - News form persons
  - Open positions
  - ...
- New projects (DFG, BMBF, KMU, etc.)
  - Initiatives
  - Larger projects
  - ...
- Calls and news from events, conferences, etc.
  - Reports (Conferences, workshops, Fachgespräche, Dagstuhl, doctoral summer/winter schools, ...)
  - Call for papers and participation
    (conferences (supported by or hosted in Germany/Austria/Switzerland), Fachgespräche, Summer Schools, ...)
  - ...
- Announcements and important dates

The preferred submission format is text, e.g., using markdown language. Call for papers can also be submitted as PDFs.

Submissions should be done by sending emails to the editors:

mailto:oliver.hohlfeld@uni-kassel.de mailto:mail@bastibl.net