KuVS Newsletter

Contents

KuVS Newsletter

Editor Message

Dear KuVS members,

we welcome you to the 22nd edition of the KuVS newsletter.

This newsletter provides you with all information that you need to know about the KuVS community over the past six months. This edition features a summary on recent awards of KuVS members, finished PhD theses, project news, event reports, and some riddles from our dear colleague Rolf Windenberg.

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

We hope you enjoy reading this edition of the KuVS newsletter. We as editors wish you merry Christmas and a happy new year.

The newsletter editors,

Oliver Hohlfeld
University of Kassel

Bastian Bloessl
TU Darmstadt

KuVS Newsletter

Movement of persons

- Prof. Dr.-Ing. Stefan Schulte (before full professor at TU Hamburg-Harburg) accepted an offer by the Technical University of Berlin. As of November 1, 2026, he now heads the Software and Business Engineering department in Faculty IV (Electrical Engineering and Computer Science).
- Prof. Dr. Florian Klingler (before junior professor at the Technical University of Ilmenau) accepted the offer for a full professorship (W3) on computer networks at the University of Bamberg.
- Prof. Dr.-Ing. Kalman Graffi switched from the Technical University of Applied Sciences in Bingen to the RheinMain University of Applied Sciences in Wiesbaden. He is now Professor for Internet of Things. More information can be found here: https://www.hs-rm.de/personenprofil/detail/kalman-graffi.

KuVS Newsletter

Awards

WowMoM Best Paper Award. Chuyang Gao and Torsten Braun (Universität Bern) received the WowMoM 2025 best paper award for their work entitlted *Two-Stage Hybrid Edge Caching Framework for 360° VR Video.*

KuVS Newsletter

Finished PhD Theses

Dominic Laniewski

Osnabrück University (advisor Nils Aschenbruck)

Title: Rate-Adaptive LiDAR Point Cloud Streaming over Wireless Networks

Abstract:

The agricultural industry is pursuing the long-term vision of fully autonomous farming, in which fleets of autonomous robots perform tasks traditionally carried out by farmers. This cumulative dissertation addresses two key challenges towards realizing this vision: (1) The efficient streaming of Light Detection and Ranging (LiDAR) point clouds over bandwidth-constrained and instable wireless networks to enable efficient offloading of computationally expensive processing tasks from the robots to dedicated computing centers, and (2) the robust connectivity of agricultural robots operating in rural areas with limited or no access to traditional broadband networks.

This thesis proposes solutions to both challenges. Specifically, we develop the concept of on-the-fly LiDAR point cloud streaming to address the first challenge. The key idea is that points from an ongoing LiDAR scan are streamed incrementally as the sensor creates them. This differs fundamentally from existing approaches, in which point clouds are transmitted only after the LiDAR has completed the scan and the entire point cloud has been created. Our approach pipelines the LiDAR point cloud streaming process of scanning, processing, and transmission, minimizing the induced delay, and, thereby, enabling delay-sensitive applications. Additionally, we address the second challenge by assessing the ability of Starlink, as today's largest commercially usable Low Earth Orbit (LEO) satellite network, to provide reliable broadband connectivity to the agricultural robots. We conduct long-term, orchestrated, and fine-grained measurement campaigns to analyze Starlink's general performance and to dissect the impact of potential degrading factors such as weather and obstruction.

In the publications that constitute this cumulative dissertation, we present the initial concept of on-the-fly LiDAR point cloud streaming and incrementally expand and improve our approach into a practical solution compatible with LEO satellite networks. Although our work is motivated by agriculture, our solution can be flexibly applied to other domains.

KuVS Newsletter

Stefan Mehner

University of Kassel (advisor: Oliver Hohlfeld)

Title: Impact of Traffic Dynamics on Network Management

Abstract: This thesis addresses the challenges posed by traffic dynamics in network management, contrasting the perspectives of a tier-1 Internet Service Provider (ISP) and operators of Industrial Control Systems (ICS).

For ISPs, a significant portion of these dynamics arises from the behavior of large Content Providers and Content Delivery Networks (CDNs). Suboptimal mappings, such as to geographically distant servers, or unexpected traffic engineering practices, create challenges including overloaded links, suboptimal routing, and asymmetric paths.

To address these issues, ISPs require accurate ingress traffic monitoring to efficiently detect improper CDN mappings and talk to interconnected networks to optimize traffic engineering. This thesis introduces the Ingress Point Detection (IPD) algorithm, which aggregates flow-level data from border routers to assign traffic ingress points to IP prefixes with a predefined level of granularity. The algorithm's design requirements were derived from empirical observations, demonstrating an accuracy of 97.6% for the top 5 Autonomous Systems (ASes) responsible for over 56% of the total ingress traffic.

Beyond its primary purpose as an efficient network debugging tool, IPD enabled an indepth study of ingress traffic dynamics within a major tier-1 ISP. Findings from six years of deployment reveal that most prefixes ingress through the same link for only short periods, often as short as one hour. Additionally, the IPD range distribution varies significantly throughout the day, with the finest granularity observed during peak traffic hours, suggesting highly detailed traffic engineering from the CDN side. Using IPD, this work conducted the first study of path asymmetry in a major tier-1 ISP. By comparing next-hop routers from historical BGP dumps with ingress routers identified by IPD, it was observed that 9% of the ISP's traffic ingress occurs through a different link than the one it was sent out. Similar trends were identified for the top 5 ASes, while the top 20 ASes exhibited approximately 30% path asymmetry. IPD has become an invaluable tool for the ISP, supporting ingress traffic monitoring and efficient network troubleshooting.

For ICS operators, traffic dynamics pose distinct challenges due to the need for high reliability and minimal variability. Unlike the Internet, ICS networks are inherently static, employing proprietary protocols and limited security measures. This thesis presents an in-depth study of the traffic characteristics of Distributed Control Systems (DCS), alongside a comprehensive analysis of their vulnerabilities. These include a demonstrated attack on the Profinet IO protocol, capable of disrupting automation processes. Furthermore, a benchmarking framework for Time-Sensitive Networking (TSN) schedulers is introduced, enabling operators to evaluate and compare schedulers for modern ICS networks.

In conclusion, this comparative analysis highlights that while ISPs adapt to dynamic environments through fine-grained ingress monitoring and traffic engineering, ICS operators focus on strong isolation and layered security to mitigate risks effectively.

KuVS Newsletter

Regine Wendt

University of Luebeck (advisor Stefan Fischer)

Title: Application of Nanotechnologies in Precision Medicine

Abstract:

Precision medicine takes into account individual genetic, environmental, and lifestyle differences to tailor diagnoses, prevention strategies, and treatments individually. The introduction of nanotechnologies into precision medicine offers opportunities for improving personalized and targeted approaches to diagnosis, therapy, and research. Nanotechnologies have the potential to detect, monitor, and treat diseases at the molecular level within the body.

A particularly relevant nanotechnology are nanodevices. Nanodevices are miniaturized electronic, biological, or biohybrid devices with nanoscale components, which have been receiving attention in the medical context since the early 2000s. The integration of nanotechnologies and nanodevices into medical research and practice poses challenges. This work addresses these challenges and provides innovative strategies to improve the effectiveness and applicability of nanotechnologies in medicine. For example, comprehensive concepts for the construction and effective use of nanodevices are lacking. A promising and universal solution approach for this are DNA-nanonetworks based on DNA-tile-nanorobots. These are presented, and their practical application is illustrated using an example scenario of in vitro disease detection, demonstrating their potential.

Since few nanodevices can currently be tested in vivo, simulations are essential for verifying and adapting theories and concepts for nanodevice development. Realistic simulation of nanodevices in their application within the human body is crucial for evaluating research hypotheses and accelerating progress in this field. Although simulation approaches exist, a holistic architecture is lacking to realistically depict complex scenarios.

The work presents a medical holistic simulation architecture for nanonetworks in humans called MEHLISSA. This architecture enables the modeling of medical nanonetworks at various levels, from the body level to organ, capillary, and cell levels. Four relevant scenarios are modeled to demonstrate the benefits of simulation and the use of nanodevices. These scenarios include improving medical measures through individual body models as the basis for digital twins, continuous health monitoring using nanodevices in the circulatory system, metastasis prevention through nanodevices, and the application of liquid biopsy in vivo for cancer diagnosis and monitoring. The simulations show promising results, including the reliable determination of thresholds for relevant markers in the body and the detection of minimal amounts of ctDNA by nanodevices. They demonstrate the potential benefits of nanodevices for improving disease diagnostics and treatment.

Another unresolved issue is the precise localization of nanodevices and disease markers, which is crucial for improved disease detection and targeted drug delivery. A new approach based on local pattern recognition and individual proteome fingerprints of important organs and tissues is developed for this purpose. By combining proteome fingerprinting and DNA-nanonetworks,

precise localization of disease markers is enabled. The simulations demonstrate the effectiveness of the method in nine major organs within minutes. The results suggest that the combination of proteome fingerprinting and DNA-nanonetworks could be crucial for more accurate and faster diagnosis and treatment.

This dissertation thus contributes to the exploration and application of nanotechnologies in medicine by providing both theoretical foundations and practical applications and simulations. It becomes clear that the use of nanotechnologies in medicine can lead to improved early detection and treatment of diseases.

https://epub.uni-luebeck.de/items/5117a759-746f-4f8b-b9b9-9cbb9f7d5490/full

KuVS Newsletter

Caspar von Lengerke

TU Dresden (advisor Frank Fitzek)

Title: Towards Practical Identification Codes

Abstract:

To prevent our networks from turning into a bottle neck for the desirable digitization of the world, their capacity has to stay ahead of the demand for connectivity. As Polar and LDPC codes can achieve the theoretical channel capacity, further evolution of channel coding techniques could eventually face diminishing returns. For one potential solution, goal-oriented communication promotes designing highly specialized codes for specific communication goals. Message identification (ID) is a communication goal that benefits from such tailored codes. In ID, the receiver wants to learn whether the sender has a copy of the exact same message as the receiver. ID codes offer better efficiency at comparable reliability than general purpose channel codes in ID tasks. The community has focused on theoretical investigation of ID codes and the computational complexity of implementing them. Specifically, tagging codes are one family of ID codes that offer strong worst-case reliability guarantees. Proposed tagging codes focus on asymptotic behavior in the infinite regime of coding parameters. Yet, trade offs between reliability and efficiency are not well understood in the practically relevant finite parameter regime. The community has investigated the computational complexity in trade offs with achieving strong worst-case reliability guarantees. However, next to the complexity, other practical performance metrics can help inform the design of practical ID codes. The community has focused on sending a single tag (codeword) and not in detail considered creating multiple tags from the same message even though this makes available another coding parameter to deliberately adjust the efficiency and reliability of tagging codes to specific use cases. This thesis complements existing ID coding metrics with practically inspired measures such as quantile and expected error probabilities. These new metrics incorporate use case characteristics such as the probability of selecting specific message pairs. This thesis proposes an uncoded ID method No-Code (NC) as an almost “free” alternative to existing ID codes in terms of the coding complexity. NC acts both as a baseline to evaluate computationally complex tagging codes and as a simple method to enable efficient ID for use cases without strict reliability requirements. Additionally, this thesis analytically characterizes how sending multiple tags reduces the number of remaining message pair candidates for the expected and the worst case. This thesis also reviews the state-of-the-art of ID coding in the practically relevant regime of finite coding parameters and thereby enables a more informed decision when selecting an ID code and its parameters to match a real world problem. Furthermore, this thesis presents an extensive tutorial part that coherently guides the reader from the formal ID definition via ID codes for noisy and noiseless channels to the construction of tagging codes and the practical optimization of the tagging code implementation. Hence, this thesis makes ID codes accessible to communication generalists without requiring expert knowledge in either information theory or ID codes.

https://nbn-resolving.org/urn:nbn:de:bsz:14-qucosa2-993912

KuVS Newsletter

Paul Schwenteck

TU Dresden (advisor Frank Fitzek)

Title: On Information Centric Content Delivery using Network Coding

Abstract:

In the age of video-on-demand, content delivery has become an integral part of everyday life. As video quality improves and the number of users increases, video streaming accounts for the largest share of total downstream traffic. Content Delivery Networks (CDNs) are continually adapting to these growing demands. In this thesis, I demonstrated how combining Information-Centric Networking (ICN) principles, peer-to-peer (P2P) networking, and Random Linear Network Coding (RLNC) can optimize content delivery systems, reduce server load, and enhance network performance while addressing energy efficiency and computational overhead challenges. I explored the potential of the InterPlanetary File System (IPFS), a protocol gaining traction due to its distributed storage capabilities and ability to enable users to act as content providers, thereby reducing server load. To enhance IPFS, I incorporated concepts from Named Data Networking (NDN), an ICN protocol. I implemented a gateway node that translates requests between the protocols and integrated IPFS's P2P storage benefits with NDN's content-based routing. This architecture enables local IPFS clusters that are connected via an NDN backbone network. The resulting hybrid network achieves comparable data retrieval times to standalone IPFS or NDN while significantly reducing server load. Using a standardized naming system for both protocols, I used content-based in-network caching, considerably reducing the download time. Furthermore, I integrated RLNC into IPFS to develop a P2P distributed network for storing and retrieving coded data. RLNC changes the client's requesting strategy: clients request a number of coded packets instead of requesting specific packets from peers. Since each encoded packet is valid, the clients can decode the data as soon as they have received a sufficient number of packets. I enhanced RLNC's requesting strategy to minimize retrieval time and network overhead while optimizing peer response times through encoding packets in-advance. I improved the hardware testbed the NET Playground and conducted energy measurements to evaluate RLNC's trade-offs. While RLNC improves network efficiency, it increases computational complexity during encoding and decoding, leading to higher energy consumption for clients and peers. These findings are critical as many video-streaming consumers rely on mobile devices with limited energy resources. Beyond RLNC's application in storage networks, I also examined its effectiveness in the transport layer as a forward error correction tool. My theoretical analysis, supported by simulations, identifies scenarios where RLNC performs optimally and where it should be avoided. Additionally, I propose improvements to RLNC in multi-hop environments to enhance throughput and packet arrival probability.

https://nbn-resolving.de/urn:nbn:de:bsz:14-qucosa2-983849

KuVS Newsletter

Jesutofunmi Ademiposi Ajayi

Universität Bern (advisor: Torsten Braun)

Title: Multi-Dimensional Network Slicing: Algorithmic and Learning-based Methods

Abstract:

Next-generation mobile networks are expected to support a wide range of demanding applications and services that have strict and varying performance requirements. To meet these requirements, Network Slicing (NS) has emerged as a powerful technique to enable cost-effective, multi-tenant communications and services over a shared physical mobile network infrastructure. However, the effective realization of the NS paradigm hinges on the ability to manage the end-to-end lifecycle of network slices in a dynamic, efficient, and automated manner. This challenge is exacerbated by the multi-dimensional slice requirements such as bandwidth, latency, and CPU and the unpredictable, online arrival of slice requests. To address the challenges of managing the end-to-end lifecycle of network slices, this thesis proposes online, data-driven solutions that are capable of adapting to dynamic conditions in mobile networks by optimizing multi-dimensional resource allocations, enabling scalable, real-time slice orchestration across heterogeneous infrastructures and proactively optimizing network slice performance to ensure service-level compliance. First, we investigate the slice admission control problem, focusing on the setting where slice requests arrive sequentially and must be admitted or rejected in real-time without prior knowledge of future resource demands. We propose an online algorithm that dynamically incorporates system resource utilization to guide admission decisions, and therefore resource allocations, with the aim of maximizing the long-term revenue of infrastructure providers. Second, building on this foundation, we address the problem of online policy selection under non-stationary network conditions. By modeling the policy selection task as a multi-armed bandit problem, we propose a data-driven solution that learns to select the most effective admission policy across time-varying network conditions. This approach balances exploration and exploitation while detecting and reacting to changes in environmental dynamics. Third, we address the problem of scalable slice provisioning in large-scale, distributed networks, and propose a hierarchical solution for the online network slice provisioning problem in which a service function chain must be effectively mapped onto the network infrastructure, while optimizing for multiple objectives. Finally, we propose a proactive optimization framework for the problem of allocating resources to heterogeneous network slices. The proposed framework aims to learn an effective resource allocation strategy in virtual radio access network environments by anticipating traffic fluctuations and proactively adjusting network slice resources for optimal performance.

KuVS Newsletter

Lucas de Sousa Pacheco

Universität Bern (advisor: Torsten Braun)

Title: Mobility and Cloud Management with Federated and Distributed Learning

Abstract:

Modern vehicular networks face challenges supporting applications like real-time traffic prediction, collision avoidance, and adaptive signal control, which require dynamic topology management, privacy preservation, and low latency. High mobility disrupts vehicle-infrastructure connectivity, hindering data exchange, while heterogeneous sensor data from diverse onboard systems violates homogeneous data assumptions in traditional collaborative frameworks. Additional complexities include mmWave beam alignment demands and adversarial threats. Conventional approaches like FedAvg struggle due to rigid client selection, centralized aggregation bottlenecks, and one-size-fits-all compression, ignoring vehicular resource disparities. This thesis introduces four integrated frameworks to address these challenges through context-aware adaptations of FL principles. DOTFL (Chapter 3) addresses non-IID data distributions and poisoning attacks via neural similarity metrics and Wasserstein distance-based clustering, achieving 94% malicious update rejection while improving accuracy by 22% over FedAvg in urban mobility scenarios. DrivePFL (Chapter 4) optimizes bandwidth utilization through Kalman Filter-predicted contact windows and layer-wise model transmission, reducing communication overhead by 10% without compromising 83.4% inference accuracy under vehicular mobility patterns. FLIPS (Chapter 5) integrates SHapley Additive exPlanations (SHAP) for adaptive model pruning, compressing transmissions by 48% while preserving safety-critical features through layer importance scoring. eDAFL (Chapter 6) accelerates mmWave beam selection via dynamic layer clustering, achieving 84% faster sector search than exhaustive search protocols through federated multi-sensor fusion. The frameworks are validated through simulations combining realistic mobility traces, SUMO traffic models, and NS-3 network emulation. Key innovations include: first: Optimal transport-based distribution alignment for non-IID data without raw data access (DOTFL, Chapter 3); 2nd: Mobility-aware client selection using predicted link durations (DrivePFL, Chapter 4); 3rd: Selective and adaptive layer pruning for throughput optimization (FLIPS, Chapter 5); 4th: Hierarchical aggregation with Centered Kernel Alignment (CKA) for model consistency (eDAFL, Chapter 6). Experimental results demonstrate scalability to 100-vehicle networks with 200ms inference latency in DrivePFL, 91% accuracy under non-IID data distributions (CIFAR-100) via DOTFL, and 52% parameter reduction through SHAP-guided compression in FLIPS. These contributions advance distributed learning theory by formalizing trade-offs between communication efficiency, adversarial robustness, and model interpretability in vehicular systems.

KuVS Newsletter

Jan von der Assen

Universität Zürich (advisor: Burkhard Stiller)

Title: Threat-Oriented Cyber Resilience: Balancing Risk and Reward Through Prevention, Detection, and Mitigation

Abstract:

Cyber attacks have remained a key concern for organizations, nations, and individuals due to their increasing frequency and impact. Establishing a perfect posture against such attacks is non-trivial from a technical perspective. In addition, economic considerations only justify certain expenses to thwart attacks. As a result, organizations strive for optimal resilience, i.e., an ability to withstand attacks while continuing to do their business.

Practices to improve cyber resilience can be categorized as either being preventive, detective, or responsive. Prevention is a critical pillar since it shifts security efforts into a stage where changes are cheaper to implement. Detective measures are implemented, acknowledging that prevention may not be perfect. Such measures reduce the time for response, minimizing the impact of an attack. A systematically executed response then aids an organization in limiting damage and returning operations to a manageable state.

Establishing a cybersecurity program framed through preventive, detective, and responsive measures requires navigating numerous practices, models, and solutions. A risk-based perspective abandons a one-size-fits-all approach, where a specific set of security controls is considered optimal for any scenario. Instead, efforts are focused on assets that are sensitive to attacks. The conceptual frame for running a cyber risk program is well-established, and multiple methodologies, frameworks, and standards exist. Implementing such practices is nevertheless challenging, and numerous limitations exist in practice.

This PhD thesis considers those challenges within a risk-based security program as research opportunities. Through a set of design studies, novel solutions are proposed and evaluated in the context of specific challenges in prevention, detection, and response. To prevent risks, the proposed threat modeling solutions leverage visual architectural modeling to address the role of collaboration, the significance of insider attacks, and the emergence of AI-related attacks. In addition, the problem of quantification is addressed through solutions exploiting continuity measures. The results indicate improved effectiveness and efficiency in those problem domains while acknowledging and raising attention to underlying limitations, such as the lack of data available or the infeasibility of achieving full automation. Detection and response are considered closely related practices, and novel solutions are proposed to gather attack data and use it in the scope of proactive and reactive defense against sets of threats. Testing the developed solutions attests to the viability of both defense paradigms while being closely coupled to the specific attack and defense scenarios.

KuVS Newsletter

Markus Dahlmanns

RWTH Aachen University (advisor: Klaus Wehrle)

Title: Identifying Security Issues in the Industrial Internet of Things

Abstract: The Internet of Things (IoT) allows the collection and communication of data from physical things as well as sending and receiving commands for interaction and actuation. Thus, it is one of the key enablers for digitalization and indispensable to seizing promised potentials like increasing the efficiency of processes. As a result, the IoT grows continuously, expands in dedicated strains like the Industrial IoT (IIoT) for industry or the Consumer IoT for everyday applications, and comes in touch with more and more sensitive data and commands, making it a sweet spot for attackers. Therefore, a secure and safe operation is essential to protect, e.g., workers laboring in the operating space of machines and the environment, which were shown to suffer from the misoperation of critical IIoT deployments. As a foundation for such a secure and safe operation IIoT protocols nowadays include security features. However, it is unclear whether these developments of protocol specifications lead to a secure IIoT in practice or whether and why insecure deployments remain in operation.

In this dissertation, we address the open research gap of the current uncertainty on IIoT security and encourage secure-by-default deployments. To this end, we first set out to assess the security of today's IIoT. Second, we analyze the pitfalls that hinder operators from operating securely despite the existence of secure protocols. Last, we propose novel approaches to help operators secure their future deployments.

Our first three contributions show from various angles that the majority of Internet-exposed IIoT deployments are insecurely configured, independently of their potential deployment date and the protocol used being either secure-by-design or retrofitted. For example, 92% of 1,114 Internet-exposed OPC UA deployments suffer from deficient security configurations. Similarly, only a minority of 6.5% of 967,551 deployments protect their communication via Transport Layer Security (TLS). Even if deployments implement secure communication, they are often only seemingly secure: Most prominently, various deployments reuse compromised secrets, neglecting any security benefits. In our fourth contribution, we trace this problem back to modern technologies like containerization that ease deployment processes but also disguise security issues. Specifically, we analyze 337,171 public container images and show that 8.5% of them include secrets. Even more alarming, 275,269 TLS and SSH hosts on the Internet rely their authentication on these secrets. To overcome such issues, in our last two contributions, we (i) propose ColPSA that—without requiring security expertise—gives deployment-agnostic security advice to users after crowdsourcing configuration possibilities and (ii) present LUA-IoT, our user-friendly authentication scheme for the whole IoT, that reduces the configuration barrier for secure communication and enables automatic certificate enrollment to IoT deployments.

Our contributions underpin the vast amount of issues in IIoT security despite the existence of strong security features in today's protocol specifications and sketch countermeasures to tackle the identified pitfalls that operators tap into when configuring their deployments. Overall, this dissertation encourages shifting from secure-by-design to secure-by-default protocols and establishes two approaches as a foundation to support operators in configuring today's protocols.

KuVS Newsletter

Jan Schlichter

TU Braunschweig (advisor Lars Wolf)

Title: Leveraging Industrial Wireless Sensor Networks for Energy-Efficient HVAC System Operations

Abstract:

In industry, many processes require specific environmental conditions to guarantee product quality. Most of the time, these conditions, as well as the well-being of human workers, are ensured by Heating, Ventilation and Air Conditioning (HVAC) systems. In Germany, such systems account for at least 6.4 % of all energy used by the industry. In order to reduce this amount while maintaining or improving the conditions for the human workers, we propose a Cyber Physical System (CPS) to control the systems based on the actual demand. To realize the CPS, we focus on three main points: (i) data acquisition through the sensor nodes of an Industrial Wireless Sensor Networks (IWSNs), (ii) reliable data transmission in the IWSN and (iii) closing the control loop through a decision framework.

For the data acquisition we develop and evaluate a versatile wireless sensor node platform. The platform captures data with a high spatial resolution in factories through its low-cost approach. Together with a custom 3D airflow sensor, which has been specifically developed for this use case, the node provides valuable insight into the actual conditions inside a factory.

To enable reliable data transmissions through the IWSN, we evaluate the usage of Multi-Connectivity (MC), which enables simultaneous data transmission through multiple independent network links. The benefit of using MC depends on several factors, such as the OSI layer on which it is implemented. Especially in the multi-hop IWSNs considered here, MC can either be implemented below or above the network layer. This decides if the MC approach is applied on a per-hop basis or end-to-end basis. To evaluate both approaches, we consider them analytically, through simulations and in a real-world deployment. Through this, general influencing factors on reliability, such as fragmentation, are extracted and discussed.

The collected data is used as input for a simple control algorithm, which closes the control loop to the HVAC system. The complete CPS is exemplarily implemented in real industrial environments, showing the demand driven control approach. Finally, the monetary and environmental benefit of the CPS is calculated for projected energy savings of 5 % and 10 %.

KuVS Newsletter

Erik Daniel

TU Dresden (advisor Florian Tschorsch)

Title: Private Content Discovery for P2P Data Networks

Abstract:

The volume of data that needs to be stored and transferred is increasing. Centralized solutions, i.e., cloud providers offer scalable storage solutions, making data availability and control dependent on a single logical entity. A self-scalable, distributed, and content-centric alternative for data exchange are peer-to-peer (P2P) data networks. P2P data networks can allow users to regain control over their data. The content-centric approach enables users to directly request the content from the network via mostly pseudonymous identifier. Unfortunately, a query reveals interest in specific content and the source of a query can be linked to a user. The combination of source and content of a query reveals the interest of a user, enabling profiling. Furthermore, performance improvements and fail-safe mechanisms spread interests widely through the network. In this thesis, we design and explore different protocols using privacy-enhancing mechanisms. The goal of the privacy enhancement is to improve content discovery privacy with an adjustable overhead in P2P data networks. Specifically, we focus on source and query privacy mechanisms, providing different privacy guarantees. Source privacy aims to obfuscate the source of a query. With network-level methods, we increase the number of possible sources of a query, obfuscating the source, improving source privacy. The modifications to the content discovery process aim to provide plausible deniability up to relationship anonymity for messages. The second focus, query privacy, aims to hide the item of a request. By adding probabilistic data structures and cryptographic protocols to the content discovery process, we obfuscate the item of interest in a query, increasing query privacy. Specifically, we increase the anonymity set of a query from one item to a range of items, increasing the anonymity set ideally to all existing items in the network. Private set intersection reduces the privacy impediments of the content providers. We apply, explain, and evaluate our protocols in the context of a prime example of the current P2P data networks: the InterPlanetary File System (IPFS). The protocols are evaluated using code experiments with a simulated network, analyzing the overhead's impact on performance. We analyze the protocols in the context of the public IPFS network by using passive measurements to evaluate the protocols' real-world feasibility. Each proposed protocol improves the privacy of content discovery, by reducing the overall information leak, effectively limiting large-scale attacks. A combination of the underlying mechanisms with additional mechanisms like cover traffic can provide strong protection.

https://nbn-resolving.org/urn:nbn:de:bsz:14-qucosa2-994295

KuVS Newsletter

Jannis Weil

TU Darmstadt (advisor Ralf Steinmetz)

Title: Deep Reinforcement Learning in Communication Systems – Addressing Partial Observability in Decentralized Multi-Agent Systems with Learned Communication

Abstract:

Communication systems are the backbone of the modern digital world. To provide services that support our everyday lives, countless interconnected components around the globe process and exchange data. With rising complexity, it becomes increasingly challenging to handcraft algorithms that efficiently control these systems.

With Reinforcement Learning (RL), agents autonomously learn to solve a sequential decision-making problem via trial and error. In recent years, the use of neural networks in RL has led to multiple breakthroughs, from surpassing human world champions in games to real-world applications in robotics. Also in the domain of communication systems, RL outperforms many handcrafted heuristics and is expected to play a major role in the future Internet. Orthogonally, recent works explore communication in RL, where in addition to selecting actions, agents learn to encode and interpret messages. The goal of this message exchange is to reduce uncertainty and improve decision-making, for example, by sharing intentions and partial observations. However, existing methods often neglect essential properties of real communication systems, rendering the transfer to the real world challenging.

This thesis investigates the intersection of communication in RL and decentralized communication systems. We show that agents can not only learn to communicate but also make efficient use of limited communication resources. Given an unreliable communication channel with limited capacity, we propose a method that allows agents to adapt the size of their messages and improve the efficiency of the communication. Restricting communication to a message exchange between agents, as common in related works, limits its applicability in decentralized systems. Our main contribution is a novel framework that conceptually separates decision-making agents and components in a communication system. Components learn to exchange messages and iteratively refine local representations, which, in turn, are leveraged by agents to solve tasks. We show that this framework enables learning of a decentralized and generalizing routing protocol, given only a sparse reward signal. Overall, we propose and extend three multi-agent environments that can be leveraged by future work to investigate subsequent challenges.

This thesis provides evidence that RL methods with learned communication allow tackling the ever-increasing complexity of future communication systems. When components in a system have dependent goals, they may benefit from an autonomous message exchange. Future work could investigate methods for decentralized training, study remaining challenges such as asynchronicity and message delays, and evaluate learned communication in real communication systems.

https://doi.org/10.26083/tuprints-00030200

KuVS Newsletter

Youming Tao

TU Berlin (advisor Falko Dressler)

Title: Trustworthy Collaborative Machine Learning for Edge AI: Privacy, Unlearning, and Robustness

Abstract:

Collaborative Machine Learning (CML) enables edge devices to train models jointly without sharing raw data, offering a promising solution for privacy-sensitive and latency-critical applications in areas such as mobile computing, healthcare, and industrial automation. Among the various CML frameworks, Federated Learning (FL) has become the most widely adopted paradigm. However, deploying FL in real-world edge environments raises fundamental trust-related challenges, each linked to one of the three core components of CML: data, model, and collaboration mechanism. Specifically, at the data level, strong privacy guarantees are difficult to achieve due to the risk of gradient leakage. At the model level, machine unlearning becomes essential to remove outdated, erroneous, or sensitive data from trained models. At the collaboration level, robustness against adversarial or faulty clients is critical to ensure reliable training in fully open environments. This dissertation systematically addresses these challenges by developing trustworthy FL algorithms with formal theoretical guarantees on convergence, differential privacy, machine unlearning, and Byzantine robustness. First, to mitigate privacy risks in wireless edge learning, the dissertation explores FL over analog over-the-air aggregation channels, where both inherent wireless noise and participant-injected perturbations jointly contribute to privacy amplification. A privacy-aware OTA-FL algorithm is proposed that jointly leverages wireless noise and structured perturbation to achieve user-level local differential privacy, even under malicious server-side manipulation of channel state information (CSI). A bandwidth-adaptive gradient compression scheme is integrated to reduce communication overhead and accommodate limited bandwidth. Theoretical guarantees on privacy and convergence are established, and experimental results demonstrate robust protection against CSI attacks, achieving improved trade-offs between privacy, accuracy, and communication efficiency compared to state-of-the-art OTA-FL baselines. Second, to address the need for data forgetting, the dissertation introduces a provably exact federated unlearning framework. Exact unlearning is defined as the requirement that the distribution of the output model and all intermediate states of the unlearning algorithm be statistically indistinguishable from those of a retraining process conducted from scratch without the deleted data. By establishing a novel connection between federated unlearning and an optimal transport problem, the work identifies a sufficient condition for exact unlearning: the underlying FL algorithm must satisfy Total Variation (TV) stability. Guided by this insight, a TV-stable FL algorithm called FATS is proposed, based on local SGD with periodic averaging and structured sub-sampling. Efficient unlearning protocols are further developed for both sample-level and client-level deletion. Theoretically, FATS is proven to guarantee exact unlearning with favorable convergence rates. Empirical evaluations across multiple benchmarks confirm that FATS enables fast and thorough unlearning with significantly reduced computation and communication overhead compared to

prior methods. Third, to ensure robustness against unreliable or malicious devices, the dissertation investigates FL under heavy-tailed gradient noise. It first considers scenarios where gradients follow coordinate-wise finite-variance heavy-tailed distributions, and proposes a local soft truncation-based gradient processing strategy that, when combined with appropriate global aggregation, achieves optimal Byzantine resilience in homogeneous data settings. The work is then extended to more realistic conditions involving infinite-variance noise and statistical heterogeneity. Two robust algorithms are developed using gradient and momentum clipping techniques. Additionally, a random projection-based approximate neighbor mixing method is introduced to reduce aggregation overhead in high-dimensional scenarios. On the theoretical side, the dissertation establishes the first high-probability convergence bounds under these relaxed assumptions. Empirically, the proposed algorithms demonstrate strong performance in adversarial and heterogeneous environments where previous approaches fail. Although the dissertation focuses on the FL framework, the underlying algorithmic principles and theoretical insights generalize to other CML paradigms such as split learning, swarm learning, and peer-to-peer learning. This work thus lays a cohesive foundation for trustworthy CML at the edge, supporting the secure and reliable deployment of intelligent edge systems.

https://www.tkn.tu-berlin.de/bib/tao2025trustworthy/

KuVS Newsletter

Project News

FEATHER

At Karlsruhe University of Applied Sciences (HKA), researchers are developing a new AI platform as part of the interdisciplinary project FEATHER (Federated Embodied AI Technology Harmonizing Efficient Robotics), funded by the Carl-Zeiss-Stiftung. The initiative tackles a central challenge in industrial robotics: enabling advanced AI methods without compromising sensitive production data. FEATHER applies a federated learning strategy, allowing robots to train models locally while sharing only the necessary parameters for joint improvement.

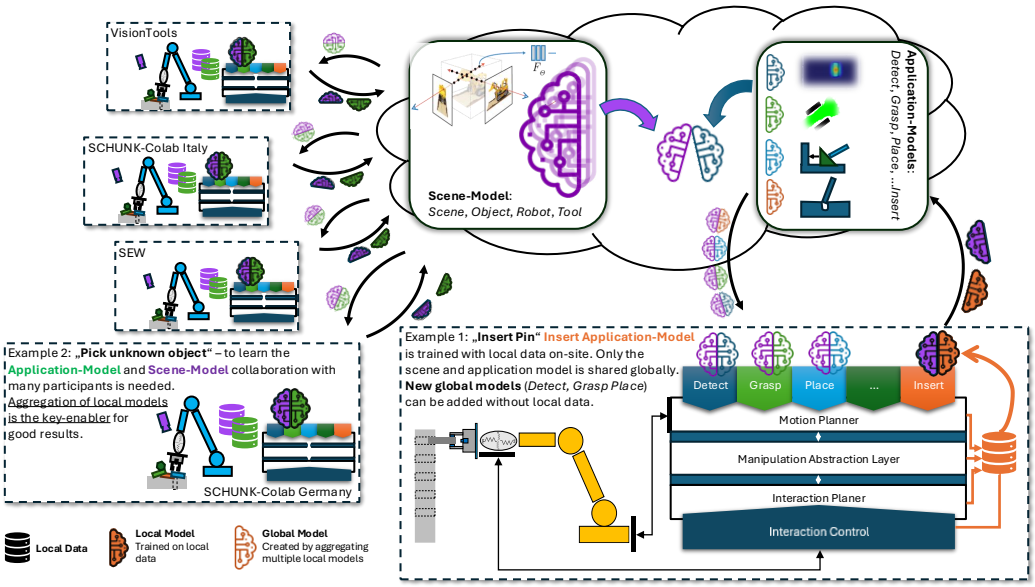

Figure 1: Overview picture of the FEATHER project

FEATHER aims to support a data-protection-compliant approach to AI-driven automation in industrial robotics. By enabling cooperative AI training across company boundaries, the project seeks to advance capabilities such as flexible assembly, autonomous adaptation, and knowledge transfer between different robotic systems.

A major obstacle in industrial AI is the lack of high-quality, labeled training data. Traditional centralized learning approaches are unsuitable for production settings where data privacy is critical. Companies must therefore find ways to jointly improve algorithms without exposing

confidential process data. Furthermore, the integration of flexible, sensor-rich components, such as soft end effectors, introduces additional complexity to motion planning and control.

FEATHER proposes a federated learning platform that enables decentralized AI training across heterogeneous robot systems. The platform integrates cloud, edge, and fog computing paradigms, allowing data processing to occur near its source while maintaining central coordination. Standardized communication protocols and orchestration tools will facilitate seamless collaboration among diverse industrial partners.

The project's research focus lies in:

- Developing distributed AI models for embodied learning on real robotic systems.
- Investigating soft end effectors as an evaluation scenario for adaptive manipulation and physical interaction.
- Designing a secure and scalable software architecture for federated learning in production environments.
- Establishing standards for communication, synchronization, and quality assurance in collaborative AI training.

FEATHER will create a unified and privacy-preserving AI ecosystem that enables robots to learn collaboratively across organizational boundaries. The expected outcomes include more adaptive and intelligent robot systems, enhanced production flexibility, and strengthened data sovereignty for participating companies. By combining technical innovation with robust data protection, the project aims to contribute to the next generation of autonomous, cooperative robotic systems.

Program Committee

- Prof. Dr.-Ing. Christian Friedrich
- Prof. Dr.-Ing. Björn Hein
- Prof. Dr. rer. nat. Oliver P. Waldhorst
- Prof. Dr. rer. nat. Christian Zirpins

DFG SPP 2378 Resilience in Connected Worlds – 2nd Phase

RESILIENT WORLDS

Vision and Goals (recap)

After a very successful first phase, SPP 2378 continues with the second phase for the next three years. The goal of the SPP is to disrupt fundamental limits of connected worlds by adding resilience as a core building block. Resilience is the ability of a system to provide and maintain an acceptable level of secure and safe service delivery, even in case of failure or compromise of some of its components, and also under completely unexpected situations. Machine learning-based solutions help making our complex network infrastructures more resilient but at the cost of reduced controllability – and with reduced abilities of experts to help in critical situations. Thus, we are faced with even more challenges in terms of resilience in critical network infrastructures.

Resilience, as an emerging research field, is strongly required as a core property of the network infrastructure, from the global internet to the internet of things, from connected cars to complex cyber-physical systems; and resilience will be a primary research objective for the coming years. Resilient Worlds will provide resilience throughout the complete protocol stack, from the hardware layer to wireless communications and to applications. We expect that in modern communication networks unknown and unforeseen events could be handled both from the network as well as utilizing external capabilities to prevent a collapse of this critical infrastructure. This requires a holistic approach to resilience, leading to appropriate, understandable, and easily applicable solutions.

In Resilient Worlds, the focus will be on the investigation of a resilience-by-design approach, which is already very challenging; however, adding resilience to (legacy) systems that were not designed for it can be even more demanding. It is therefore the goal of the SPP to address resilience from a new, multi-disciplinary perspective including, but not limited to, communications and networking, semiconductor electronic hardware systems, information security, and machine learning.

Research Program

Following up on the review process that took place in spring 2025, 12 projects have been selected for funding by German Research Foundation, which are about to kick-off early 2026.

- **Resilient Network Virtualization Architecture**
  Pramod Bhatotia (PI), Systems Research Group, TU Munich
  Antoine Kaufmann (PI), The Max Planck Institute for Software Systems

- **Integrated Sensing and Communication: Is This a Key Enabler for Resilience by Design in Future Communication Systems?**
  Holger Boche (PI), Lehrstuhl für Theoretische Informationstechnik, TU Munich
  Rafael Schaefer (PI), Lehrstuhl für Informationstheorie und Maschinelles Lernen, TU Dresden

- **RESERVE: Resilient Hybrid Time-sensitive Networks**
  Falko Dressler (PI), Telekommunication Networks (TKN), TU Berlin
  Doğanalp Ergenç (PI), Telekommunication Networks (TKN), TU Berlin
  Mathias Fischer (PI), Computer Networks (NET), UHH

- **CRUSTmeetResPECT: Collective Resilient Unattended Smart Things (CRUST) meet Resilient Power-Constrained Embedded Communication Terminals (ResPECT)**
  Anna Förster (PI), Sustainable Communication Networks (ComNets), University of Bremen
  Thorsten Herfet (PI), Telecommunications Lab, Saarland University
  Peter Wägemann (PI), Chair of Computer Science 4 (Systemsoftware), FAU Erlangen

- **ReNO-2: Supporting Human Network Operators with ML - A Cognitive Approach**
  Oliver Hohlfeld (PI), Distributed Systems, University of Kassel
  Stefan Schmid (PI), Internet Network Architectures and Management, TU Berlin

- **Resilience of homomorphically encrypted cloud-computations against integrity attacks: Challenges and countermeasures (REACT)**
  Tibor Jager (PI), Chair for IT Security and Cryptography, University of Wuppertal
  Moritz Schulze Darup (PI), Chair of Control and Cyberphysical Systems, TU Dortmund

- **Towards Network Resilience via Collective Counterfactual Explanations and Interventions**
  Setareh Maghsudi (PI), Lernende Technische Systeme, Ruhr-University Bochum

- **Mastering Offloading and Correlated failures for resilient CommunicAtion networks (MOCCA)**
  Michael Menth (PI), Kommunikationsnetze, University of Tübingen
  Björn Scheuermann (PI), Kommunikationsnetze, TU Darmstadt

- **Cross-layer Resilience for Criticality-aware Wireless Networks**
  Amr Rizk (PI), Distributed Real-Time Systems Group, Leibniz Universität Hannover
  Aydin Sezgin (PI), Digitale Kommunikationssysteme, Ruhr-UniversitV§t Bochum

- **Resilient Networking On the Move**
  Anke Schmeink (PI), Lehrstuhl für Informationstheorie und Datenanalytik, RWTH Aachen University

- **CAREFUL - Complexity-aware symbolic fault analysis**
  Klaus Wehrle (PI), Communication and Distributed Systems Group, RWTH Aachen University

- **Coordination**
  Falko Dressler (PI), Telekommunication Networks (TKN), TU Berlin

Coordination Team

- Falko Dressler (Spokesperson), Telecommunication Networks, TU Berlin
- Matthias Hollick, Secure Mobile Networking Lab (SEEMOO), TU Darmstadt
- Anna Förster, Sustainable Communication Networks (ComNets), University of Bremen
- Frank H. P. Fitzek, Deutsche Telekom Chair of Communication Networks, TU Dresden
- Doğanalp Ergenç (Coordination), Telecommunication Networks, TU Berlin

For more information, please visit the program's web site https://www.resilient-worlds.org

KuVS Newsletter

Event Reports

3rd Workshop on Machine Learning in Networking (MaLeNe)

Michael Seufert, Andreas Blenk, Björn Richerzhagen
MaLeNe 2025 Workshop Co-Chairs

The Third International Workshop on Machine Learning in Networking (MaLeNe, https://www.uni-augsburg.de/en/fakultaet/fai/informatik/prof/netcom/events/malene2025/) was a successful half day event held at Technical University of Ilmenau on September 1, 2025, where it was co-located with the Conference on Networked Systems (NetSys 2025). It follows the history of MaLeNe KuVS workshops (Fachgespräch) and MaLeNe NetSys workshops as the 6th edition of the MaLeNe series, and was organized by workshop co-chairs Michael Seufert (University of Augsburg, Germany), Andreas Blenk (Siemens AG, Germany), and Björn Richerzhagen (Siemens AG, Germany). The workshop accepted 7 full papers for presentation.

On the day of the workshop, the co-chairs welcomed 34 registered participants. The workshop started with an industry keynote given by Philippe Buschmann (Siemens AG, Germany) who elaborated on "Recent AI Trends in Industrial Network Environments". He explored industrial AI and its application in smart mobility, smart infrastructure, and manufacturing. He especially highlighted the challenges and opportunities of AI adoption in these settings through the perspective of Siemens.

Afterwards, the first technical session started. Katharina Dietz (University of Würzburg, Germany) evaluated the consensus of XAI -based network intrusion detection and how to improve consensus with feature selection methods. Nasim Nezhadsistani (University of Zurich, Switzerland) presented an asynchronous consensus-driven multi-agent approach to decentralized federated learning for intrusion detection in 5G networks. Zineddine Bettouche (Deggendorf Institute of Technology, Germany) discussed spatiotemporal machine learning techniques for cellular traffic forecasting with a particular focus on Mamba deep learning architectures.

After the coffee break, Christian Maier (Salzburg Research, Austria) elaborated on predicting performance metrics in edge-cloud networks using GNNs. Timothy Harrison (University of Hagen, Germany) presented work on state cloning for high order Markov chains with the GraphLearner. Yanakorn Ruamsuk (University of Hagen, Germany) talked about emotion-controlled communication in agent networks. Finally, Alexander Niedermayer (Karlsruhe University of Applied Sciences, Germany) presented an approach for client-agnostic continuous authentication via keystroke-induced traffic patterns.

The workshop co-chairs closed the day with a short recap and thanked all speakers and participants who engaged in the fruitful discussions. As the workshop has proven to foster active collaborations in the research community, another edition will be considered in the future.

We would like to thank all the authors, reviewers, and attendants for their precious contributions towards the successful organization of the workshop!

KuVS Newsletter

IEEE LCN 2025 – Celebrating 50 Years of Computer Networking

Florian Tschorsch and Burkhard Stiller

The IEEE Conference on Local Computer Networks (LCN) celebrated its 50th anniversary in 2025, marking a historic milestone for the networking community.

Founded in 1976, LCN emerged in the same year that the seminal paper on Ethernet by Bob Metcalfe and David Boggs was published, way before the Internet became mainstream. Over five decades, LCN has evolved from its initial focus on local networks to encompass all aspects of computer networking, making it the longest continuously running academic conference in the field. It remains large enough to sustain a vibrant international community while serving as an attractive entry point for young researchers.

The anniversary edition, held in Sydney, Australia, featured a dedicated celebration session. Overall highlights included keynote talks by Grenville Armitage (Netflix) and Elaine Wong (University of Melbourne), as well as commemorative speeches by Hironori Washizaki (IEEE CS President 2025) and Geoff Huston (APNIC). The proceedings also contain greetings from Bob Metcalfe, who shared his preference for LCN over LAN.[1] And the event received additional coverage by the IEEE Computer Society, reflecting on its legacy.[2]

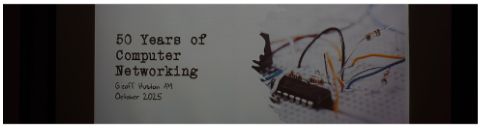

The next edition of LCN in 2026 will be hosted in Coimbra, Portugal. The Call for Papers will be available soon. In the meantime, feel free to explore more of LCN's history at ieeelcn.org.

[1] https://doi.org/10.1109/LCN65610.2025.11146361
[2] https://www.computer.org/press-room/ieee-cs-legacy-lcn

1st Workshop on Resilient Networks and Systems (ReNeSys)

in conjunction with NetSys 2025

The first Workshop on Resilient Networks and Systems (ReNeSys) brought together researchers working on network and system resilience in September in Ilmeanu. Held in conjunction with NetSys 2025, the workshop was organized by early-career researchers of the DFG priority program SPP 2378 *Resilient Worlds*: Dr. Doğanalp Ergenç, Dr. Shadi Attarha, and Dr. Peter Wägemann.

Photo (right): Peter Amthor, CC BY 4.0

The workshop featured eight accepted papers, each presented in 20-minute talks covering a wide range of topics in network and system resilience. In addition, it hosted two keynote speakers, Prof. Dr. Guevara Noubir and Dr. Joachim Sachs, who delivered inspiring talks on secure and robust wireless communication. The Best Paper Award was presented to Sascha Rösler and his co-authors for their work **"***Robust LoRa via Repetition in Frequency through Signal Emulation using WiFi."*

The accepted papers are available in open access through TU Berlin's repository at: https://doi.org/10.14279/depositonce-24399.

More details about the workshop are available on the program website: https://www.resilient-worlds.org/renesys/program/.

WueWoWAS'25: With momentum and vigor into future 6G networks

The Würzburg Workshop on Next-Generation Communication Networks (WueWoWAS'25) took place from October 6-8, 2025. About 45 participants from academia and industry discussed about 6G networks, covering advanced architectures, technologies and concepts that will enable the next generation of connectivity, intelligence and automation across multiple domains.

The tradition and intention of the WueWoWAS workshop series is to foster communication among researchers from industry, universities, and other research institutes. To that end, the program featured **technical talks** on current research, peer-reviewed from an open call, **keynote talks** by experts of outstanding repute, and **tutorial presentations**. The technical sessions covered a wide range of subjects addressing various next-generation network paradigms. Some of these works used artificial intelligence (AI) and machine learning (ML) for 6G communication networks and, in particular, RAN Intelligent Controller (RIC) telemetry, AI agents and agentic workflows to tackle misconfiguration in campus networks, or AI assisted intent based network management systems. *Henning Sanneck (Apple)* provided a keynote from the UE perspective on native AI/ML in 6G and especially current efforts in UE mobility, utilizing ML for the air interface and RAN (3GPP release 20), RSRP predictions, as well as predictive mobility in a flexible multi-connectivity and cell-free environment. This included an outlook on the emergence of advanced design patterns that integrate AI-native architectures and semantic communications towards self-evolving, context-aware networks beyond traditional connectivity.

From the perspective of verticals, industrial networking drives 6G toward deterministic, AI-enabled communication frameworks that integrate sensing, control, and connectivity to meet the stringent demands of industry applications. *Johannes Riedl (Siemens)* highlighted in his keynote the requirements of industrial networks and verticals. The vision is to enable vertical industries to focus on their application and innovation rather than operating the infrastructure, through zero-effort networking and providing connectivity as a service to industries. Key components are intent-based network orchestration using AI and ML, as well as technologies supporting deterministic networking like time-sensitive networking (TSN). The technical sessions included also methods of how to improve time synchronization accuracy for 5G links or time-aware shaping in TSN based on Centralized Network Controller (CNC) for dynamic reconfiguration to meet industrial requirements.

Several contributions addressed O-RAN, which is an open, interoperable architecture for mobile networks that disaggregates traditional RAN components and introduces open interfaces. Approaches were presented on how to optimize the RAN lifecycle management for industrial 5G using NetDevOps-driven deployment framework for industrial open RAN that replaces complex orchestration with infrastructure as code (IaC). Also, mobility was discussed in that context and, in particular, a UE-centric handover mechanism tailored for industrial Open RAN deployments that proactively initiates transitions before connectivity degrades below acceptable levels. *Stanislav Lange (NTNU)* provided a tutorial to the attendees on Open RAN and 5G campus networks setup. Key building blocks for setting up a private 5G testbed using Software-Defined Radios (SDRs), open source software components such as OpenAirInterface and srsRAN, and off-the-shelf modems / phones.

Several technical contributions emphasized the evaluation of network performance and quality of service (QoS), but also reliability and resilience, as well as power consumption and sustainability. One study presented a detailed assessment of power consumption and energy efficiency of fiber optic SFP modules. *Frank Loh (University of Würzburg)* gave a tutorial on energy measurements, metrics, and models, with a focus on methods for measuring power consumption in communication networks as well as software-based approaches for approximating the power consumption. Energy efficiency metrics from literature were revisited and pitfalls in their usage were discussed. *Carmen Mas Machuca (Universität der Bundeswehr München)* gave a tutorial on resilience metrics and models for communication networks. A variety of resilience metrics and their computation were discussed, providing meaningful guide for research. In the technical sessions, flow availabilities were also analyzed by investigating the correlation between minimal cut sets and flow availabilities, which was presented by *Shakthivelu Janardhanan (TU Munich)* who received the best student paper award.

However, balancing sustainability, performance and resilience may introduce a fundamental research challenge for communication networks. One work considered to design photovoltaic-powered server clusters that operate energy-efficiently with minimal environmental impact while maintaining high performance and resilience under varying workloads and fault conditions. Novel approaches addressed island-ready 6G communication, which envisions resilient, self-sustaining networks capable of maintaining connectivity in isolated or infrastructure-limited environments. To this end, the conceptual model of island connectivity was introduced and results on the users' perspectives presented, which triggered discussions on the design implications for operators, developers, and authorities.

The next highlight was the keynote by *Rastin Pries (Nokia)* on spatial computing in the 6G era. The paradigm of spatial computing refers to technologies and concepts that merge the physical world, the digital world, and the human world. Spatial computing represents a paradigm shift in the manner in which human-computer interactions function, moving towards a more bidirectional model. The device serves not only to augment the user's perception of the world with digital information, but also to provide contextual information like the user's current physical context and gestures to the application. This requires advanced sensing, AI-driven data fusion, and low-latency processing to ensure accurate spatial awareness and seamless synchronization between physical and digital environments for real-time digital twins.

The paradigm of quantum communication was introduced by *Karim Elsayed (Leibniz University Hannover)* in a tutorial. Quantum communication forms the foundation for connecting multiple quantum computers to collaboratively solve complex problems. This approach is driven by the difficulty of sustaining large-scale quantum systems. Unlike classical communication, quantum communication is governed by unique quantum properties, requiring fundamentally new concepts and techniques. To this end, the tutorial introduced the core principles of quantum communication, its unique features, and evaluated the performance of the communication system for different approaches in light of these distinctive features.

A key part of the WueWoWAS workshops is the integration of PhD students into the research community and the exchange between academia and industry. This included the technical mentorship through focused tutorials, but also a dedication **speedmentoring session**. In that speedmentoring event, PhD students were connected with leading experts from academia and industry for short, targeted discussions on research topics, emerging trends, and career development. The discussions continued at the social event in a traditional Würzburg restaurant, which served Bavarian cuisine and specialties like Franconian wine soup in a relaxed atmosphere.

To integrate PhD students also in the academic review process, experienced PhD students were invited to participate as reviewers. A **reviewer reflection forum session** was part of the WueWoWAS'25 event in order to foster transparency and understanding of community standards from the perspective of reviewers and authors: what reviewers and authors value in the review process; what should be addressed in a submission and in a review; what practices to avoid; what is the hardest part as reviewer. Such community-oriented efforts are essential and highly relevant, as they cultivate an inclusive and supportive research ecosystem where early-career scientists can learn, connect, and contribute meaningfully. This was followed by open discussions on the focus theme for the next WueWoWAS workshop in 2026. Integrating PhD students into mentorship, peer reviews, and open discussions is key for their academic maturation and for strengthening the sense of belonging within the community and the vitality of the research community.

More information in the workshop report, available online: https://doi.org/10.25972/OPUS-42709

WueWoWAS'25 would like to thank our industrial partners Infosim and EMnify for providing financial support and the supporting projects ORIGAMI and SUSTAINET-Advance, without whom the workshop would not be possible. The event was technically co-sponsored by the VDE ITG expert group KT 2 on “Communication Networks and Systems” as well as the GI/ITG special interest group “Communication and Distributed Systems'” (KuVS).

Finally, we thank all authors, contributors, reviewers, the keynote speakers, the tutorial speakers and students from the University of Würzburg. David Raunecker's tremendous support in the organization and efforts in the social activities are deserving of particular recognition. Their efforts made WueWoWAS'25 a success. We will see you at the next edition in 2026.

Tobias Hoßfeld (University of Würzburg,
Co-Chair of KT 2 “Communication Networks and Systems)

Amr Rizk (Leibniz Universität Hannover,
GI/ITG Special Interest Group KuVS)

KuVS Newsletter

KuVS Calls and Announcements

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

Calls for Papers and Presentations

7th KuVS Fachgespräch on Machine Learning in Networking (MaLeNe 2026)

Call For Papers

The 7th edition of the KuVS Fachgespräch "Machine Learning in Networking (MaLeNe)" will be held on March 19–20, 2026 at the University of Augsburg in Augsburg, Germany. The workshop is also technically sponsored by VDE ITG KT2.

What is a Fachgespräch?

Fachgespräche are a low-key version of workshops, organized (among others) by the Communication and Distributed Systems group of the Gesellschaft für Informatik. They foster community building by organizing meetings on topics of current interest, with a low entry barrier. The intended audience primarily comprises advanced Master's students and early-stage PhD researchers, but of course everybody is welcome. Participants are invited to submit full papers, early-stage research ideas, extended abstracts, or work in progress.

What is MaLeNe?

MaLeNe is one such Fachgespräch, dedicated to the interaction of machine learning and networking. Earlier editions can be found here: 2020a, 2020b, 2021 (1st NetSys workshop), 2022, 2023 (2nd NetSys workshop), 2025 (3rd NetSys workshop).

In recent years, communication networks have become highly flexible through the employment of virtualization and softwarization paradigms. Still, networks are highly complex, dynamic, and time-varying systems, such that the statistical properties of networks and network traffic cannot be easily understood and modeled. Furthermore, the interplay between networking and the dynamic and heterogeneous requirements, expectations, and experiences of applications and users increases the complexity of the systems, which makes fault, configuration, performance, and security management in networks a hard problem. As observed in other disciplines, the successful application of machine learning can help to overcome these issues by following a more data-driven approach. In the networking domain as well, technological advancements in the area of machine learning, the increasing availability of network analytics data, and the flexibility of programmable networks and virtualized network resources have made this approach applicable, which creates exciting new opportunities.

MaLeNe 2026 aims to provide a forum for researchers addressing emerging concepts and challenges related to machine learning in networking. The Fachgespräch will address opportunities where machine learning can bring benefits to networking in different facets, such as network monitoring, management, and security. Together with flexible and programmable networks, this paves the way toward a more proactive and autonomous network design and "self-driving" networks. The long-term vision is that configuration decisions can be made in real time in an automated fashion before service and experience degradation occurs. The Fachgespräch will feature original paper presentations that foster discussions and joint work among participants.

Topics of Interest

Authors are invited to submit extended abstracts, early-stage research ideas, or work in progress related to the topic areas listed below:

- **Methodology**
  - Data sets for benchmarking, verification, proof of concept
  - Data augmentation; dimensionality reduction (e.g., autoencoder); prediction and generation (e.g., GANs)
  - Performance evaluation methodology (best practices)
- **Artificial Intelligence and Machine Learning Methods**
  - Classical and deep learning methods for supervised, unsupervised, and reinforcement learning
  - Advanced methods (e.g., self-supervised learning) and models (e.g., adversarial, agentic, generative)
  - Large language models and foundation models; tiny models
- **Generalizability**
  - Transfer of trained models (e.g., small to large networks, enterprise to data center)
  - Federated learning (combine models trained for different data sets) and privacy
  - Context drift, catastrophic forgetting, and machine unlearning
- **Explainability**
  - Explainable Artificial Intelligence (XAI) and visualization
  - Understanding decisions of ML-based systems (management, traffic engineering, etc.)
  - Game-theory-based approaches to approximate guarantees
- **Networking for Machine Learning and AI**
  - Network architectures, applications, and use cases (data center, enterprise, etc.)
  - Network resource management (algorithms, schedulers, etc.)
  - In-network computation and processing
- **Applications in Networking**
  - Network monitoring, especially from encrypted traffic (e.g., traffic classification, QoE)
  - Network configuration (e.g., suggest optimal configurations, “spell-check” text-based configuration data)
  - Network planning (e.g., reconfigurable data centers, job placement)
  - Network management (e.g., autonomous management, self-driving networks)

- Network security (e.g., intrusion detection, covert channels, firewall)
- Advanced networks (e.g., 5G to 6G, industry, slicing)

Paper Submission

Authors may submit extended abstracts, early-stage research ideas, or works in progress presenting original research, ongoing studies, or practical experiences. Previously published work or work currently under review elsewhere cannot be submitted. Each submission must be written in English, accompanied by a 75 to 200 word abstract and a list of up to 5 key words. There is a length limitation of 2 A4 (210 mm × 297 mm) pages for full papers (including title, abstract, figures, tables) plus 1 page for references. Submissions must be in 2-column IEEE conference style with a minimum font size of 10 pt. Papers exceeding these limits, multiple submissions, and self-plagiarized papers will be rejected without further review. Authors should submit their papers electronically via the EasyChair online submission system under the following link: https://easychair.org/conferences?conf=malene2026

Important Dates

- Paper submission deadline: January 15, 2026
- Paper notification: January 31, 2026
- Camera-ready deadline: February 16, 2026

Proceedings

All papers accepted for MaLeNe 2026 will be included in proceedings, which are published open access via OPUS. We reserve the right to remove any paper from the proceedings if the paper is not presented at the workshop.

Organizers

- Michael Seufert (University of Augsburg)
- Andreas Blenk (Siemens AG)
- Björn Richerzhagen (Siemens AG)

For any questions please contact Michael Seufert (michael.seufert@uni-a.de).

KuVS Newsletter

Open Positions

Ph.D. Position in collaboration between Deutsche Telekom Chair for Communication Networks and BMW Group: Non-Terrestrial 6G Networks (NTN-6G)

https://cn.ifn.et.tu-dresden.de/chair/career/

- Focus on the development and improvement of communication protocols for satellite networks, especially for LEO satellites, in order to close existing research gaps in satellite-based applications.
- Design and implementation of communication systems for UAVs, HAPs, and satellites with seamless integration into terrestrial networks.
- Research into advanced network technologies, such as the application of machine learning (ML) and artificial intelligence (AI), for network and resource optimization.
- Development of mobility management and multi-path solutions to improve communication in the automotive sector and its connection to NTNs.
- Collaboration in an interdisciplinary team, close cooperation with specialist departments and universities, and participation in international standardization and conference activities (e.g., 3GPP).

Ph.D. Position collaboration between Deutsche Telekom Chair for Communication Networks and BMW Group: Ultra-reliable 6G communications fot the automotive

https://cn.ifn.et.tu-dresden.de/chair/career/

- Focus on research into the development and improvement of communication protocols for vehicle-to-everything communication, whether direct or network-based.
- Research into advanced network technologies, such as the application of machine learning (ML) and artificial intelligence (AI), for network and resource optimization.
- Research into multipath communication and network coding in densely meshed V2V scenarios to enable extremely robust communication links for V2X applications.
- Direct connection to the host university in Dresden, collaboration in an interdisciplinary team, close cooperation with specialist departments and universities, and participation in international standardization and conference activities (e.g., 3GPP).

KuVS Newsletter

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:* $\sqrt{a\ trE}$
- *Playing with Capital Letters:* **XLNt**
- *Advanced Persons:*
  the H of the B was < 1 wEk
- *Experts:*
  V stRTd 2 0 this 1derful vU from the moun10 2 the C
  [hint: 0 pronounced "love" (as in tennis)]
- *Geniuses:*
  {(T • T), (T • T • T)}

Fig. 1: Illustration to assist the reader in solving the third 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):

- root of a tree [because: *root of*-a-tr-*E*]
- excellent [because: *X*-*L*-*N*-t]
- the age of the bee was less than one week [because: the-*H*-of-the-*B*-was-*less than*-*one*-w-*E*-k]
- we started to love this wonderful view from the mountain to the sea [because: *V*-st-*R*-*T*-d-*two*-*love*-this-*one*-derful-v-*U*-from-the-moun-*ten*-*two*-the-*C*]
- set of tea products [because: *set of*-*T*-*products*].

KuVS Newsletter

Next Newsletter - Deadline May 15th

Next newsletter : 06/2026

Deadline for submissions and contributions : 15th May 2026

We ask you for submissions in English. Topics can be from the following time frame: December 2025 - May 2026

- 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:

Oliver Hohlfeld
mailto:oliver.hohlfeld@uni-kassel.de

Bastian Bloessl
mailto:mail@bastibl.net