

ABSTRACTS

CyberSweden 2025



Accelerating Transformer-Based Network Traffic Classification: Hybrid Architectures for Real-Time Performance

Mohamed Hashim Changrampadi
Chalmers University of Technology

Supervisor: Magnus Almgren

Abstract:

Accurate traffic analysis is critical for network security, enabling detection of malicious behavior and enforcement of defense policies. Traditional rule-based methods struggle with the growing complexity and diversity of network traffic. Transformer-based models improve classification accuracy but are often too resource-intensive for real-time use. We evaluate trade-offs between accuracy and system cost using ET-BERT, a leading BERT-based traffic classifier. While ET-BERT delivers strong performance, it requires substantial computing. Simpler models with 200x fewer FLOPs achieve comparable accuracy—within 0.5%—and 20x faster inference, offering a more practical path for real-time, security-oriented deployment.

Dynamic Graph Neural Networks for Advanced Persistent Threat Detection

Zhenlu Sun
Uppsala University

Supervisors: André Teixeira, Salman Toor

Abstract:

This project explores using Dynamic Graph Neural Networks (DGNNs) to detect evolving Advanced Persistent Threats (APTs). APTs are stealthy and complex, making early detection essential. Provenance graphs capture system interactions, but most existing methods lack early and explainable detection. DGNNs learn spatial and temporal patterns in these graphs, enabling early alerts and evolution analysis of APTs. The proposed method supports both graph- and node-level detection. Extensive experiments on public APT datasets demonstrate the model's effectiveness in early detection, explainability, and robustness against adversarial attacks.



Automating Automotive IDS

Torbjörn Livén

Chalmers University of Technology

Supervisor: Magnus Almgren

Abstract:

Deploying IDS systems and handling intrusions in in-vehicle CAN networks requires significant human effort. Within the scope of the Vinnova FFI MAGIC project, we identify bottlenecks and propose a framework to reduce the need for manual human intervention. The framework consists of two main components, a federated learning scheme that continually trains personalized models for each individual vehicle while maintaining privacy, and an intrusion handling system that safely handles incidents and generates a human-readable report for root-cause analysis. The system deployment could significantly reduce the need for human intervention in IDS tuning and root-cause analysis.

Federated Learning with Efficient Privacy-Preserving Training

Yenan Wang

Chalmers University of Technology, University of Gothenburg

Supervisor: Elad Michael Schiller

Abstract:

Federated learning enables collaborative machine learning model training without sharing sensitive local data, but robust privacy measures like homomorphic encryption or differential privacy incur heavy resource overheads. We introduce a novel federated learning framework that alternates rounds with privacy protections on authentic data and rounds on synthetic data without the need to use privacy protection. The new framework balances model accuracy against computational and communication costs by alternating authentic and synthetic rounds. Experimental results and analysis show that the proposed framework defends against data leakage attacks while significantly reducing privacy overhead and maintaining strong privacy guarantees.



Human Factors of CTI Adoption in Organizations

Patrick Shirazi

University of Skövde

Supervisor: Ali Padyab

Abstract:

Cyber Threat Intelligence (CTI) is increasingly critical for organizations to remain resilient. While the technical aspects of CTI have been widely studied, the human factors influencing its adoption remain underexplored. This qualitative study investigates how individual expertise, team dynamics, and communication practices impact CTI adoption across organizations. Data collection is complete, and analysis is in its final stages. Preliminary findings reveal key socio-organizational drivers that shape CTI success beyond technical capacity. Findings will contribute to both academic discourse and practical strategies for advancing human-centered CTI approaches.

A Taxonomy of Cybersecurity Economic Models

Aous Al Salek

University of Skövde

Supervisors: Ali Padyab, Martin Lundgren, Rose-Mharie Åhlfeldt

Abstract:

Organizations face challenges in selecting suitable cybersecurity economic models due to limited awareness and contextual guidance. This study addresses the issue through a systematic literature review, analyzing 60 models to develop a structured taxonomy. Using a validated methodology, seven dimensions and 30 characteristics were identified, enabling effective classification and comparison of 38 unique objects. The resulting taxonomy assists stakeholders in choosing appropriate models and uncovers gaps in current research. This contribution supports more informed cybersecurity investment decisions and lays the groundwork for future model development.



An exploratory study on cybersecurity specialists' workflow and views on AI

Alexander Lindström
Karlstad University

Abstract:

There is much ongoing research for using generative AI in cybersecurity, with the aim of simplifying and making processes, such as penetration testing more efficient. There is, however, a lack of research on what cybersecurity professionals think about generative AI, what risks and benefits they see with it, and what requirements they might have for using these tools in their work. To fill this gap, we conducted a semi-structured interview study with cybersecurity professionals. The goal of the study was to answer these questions, and the findings from the study will later be used to develop a tool using generative AI for cybersecurity defense.

Innovation and improvement of management systems for IT and information security

Safia El Moutaouakil
Luleå University of Technology

Supervisor: John Lindström, Karl Andersson, LTU

Abstract:

Digitization has shifted how organizations operate and the data being shared necessitates effective data management. The implementation of IT and Information Security Management Systems (ISMS) is becoming crucial. Management systems are driving organizational processes and operations to via seamless business processes achieve organizations' strategic objectives efficiently. However, these management systems are not immune to misfunction with evolving technologies and cyber threat landscape. Thus, continuous performance monitoring and alignment with technology innovations are needed. Organizations use, e.g., governance, risk management and incident response planning, for assessing their ISMS's performance and overall effectiveness in preventing, detecting and responding to cyber threats.



Trusted Execution Environment (TEE) for Low-Power RISC-V Devices

Asrin Abdollahi
RISE, Mälardalen University

Supervisor: Shahid Raza

Abstract:

Securing IoT devices can be challenging due to their small size, limited battery power, and resource constraints. Protecting RISC-V devices at the hardware level is necessary to reduce the attack surface. RISC-V offers an open-source architecture, license-free, and customizability, making it ideal for small IoT devices and custom designs. A lightweight TEE can protect sensitive data and code by creating a private operating environment within the processor, isolated from the main OS, and preventing unauthorized access and exploitation. The work focuses on low-level hardware programming using Renode. The secure design shall demonstrate the integrity and confidentiality within the TEE.

TinyTor: Enabling Onion Routing for the IoT

Rikard Höglund
RISE, Uppsala University

Supervisor: Shahid Raza

Abstract:

This work presents an approach for implementing Tor-like onion routing for the Internet of Things (IoT), utilizing a stack of communications and security protocols adapted to constrained IoT devices. We analyze the challenges of adapting onion routing to the constraints and requirements of IoT systems. We propose a modified onion routing solution that is tailored for IoT scenarios, detailing its design and impact on IoT security. We outline the architectural modifications needed and the appropriate usage of protocols suitable for constrained IoT devices. Furthermore, we highlight considerations and design choices that enable assuring privacy and security of the proposed solution.



How Feasible are Passive Network Attacks on 5G Networks and Beyond? A Survey.

Atmane Ayoub Mansour Bahar
Chalmers University of Technology

Abstract:

Passive network attacks (PNAs) exploit encrypted traffic patterns to infer sensitive user data without detection. While widely demonstrated in Wi-Fi and LTE settings, their feasibility in 5G remains unclear due to link-layer encryption, beamforming, and massive MIMO. This survey evaluates the reproducibility of PNAs in 5G and beyond, focusing on device and app fingerprinting, system identification, and user geolocation. Findings indicate that although 5G raises significant barriers to passive surveillance, advanced attackers may still pose a risk in specific contexts. This is the first systematic evaluation of PNA feasibility in next-generation cellular networks.

PKI for 6G Satellite Communication

Arlette Houndji
RISE, Karlstad University

Supervisors: Simone Fischer-Hübner, Shahid Raza

Abstract:

The rapid growth of the "New Space" sector and the deployment of Low Earth Orbit (LEO) satellite constellations are revolutionizing satellite communications (SATCOM), making it an integral part of the anticipated 6G infrastructure. However, integrating SATCOM with terrestrial networks introduces complex security and privacy challenges that conventional Public Key Infrastructure (PKI) solutions cannot fully address. This is due to the unique characteristics of space-based networks, including highly dynamic topologies, mobility, and resource constraints. PKI serves as a cornerstone of internet security and has, for decades, ensured robust authentication and message confidentiality and integrity across terrestrial networks. While proven effective for terrestrial networks, traditional PKI faces limitations in SATCOM deployments due to the unique characteristics of space-based communication systems. Specifically, SATCOM networks introduce a highly distributed and mobile infrastructure, exhibit significant latency variability, and operate under strict resource constraints, which collectively necessitate the development of optimized, lightweight security solutions. We study the



challenges and requirements for deploying PKI in a space environment, for a safer 6G.

Evaluating the Viability of Computational Offloading for Vehicles Under Adverse Network Conditions

Samuel Bach
AstaZero, Karlstad University

Supervisor: Anna Brunström

Abstract:

The safe and efficient operation of automated vehicles requires processing massive amounts of sensor data. However, the computational capabilities of vehicles are often limited. This work evaluates the viability of utilizing computational offloading under adverse network conditions. To emulate such conditions, we use synthetic network interference such as packet loss, throughput rate limiting, packet corruption, and RF attenuation. The results indicate that network conditions can significantly impact performance. The median for false detections running 5% packet corruption reached as high as 20%. The results emphasize the need for resilience and robustness to poor network conditions when designing computational offloading strategies.

EagerTAP: Pre-evaluation for data minimization on TAPs

Daniel Freiermuth
Chalmers University of Technology

Supervisor: Andrei Sabelfeld

Abstract:

Trigger-Action-Platforms (TAPs) enable automations between different services. Typically, TAPs have access to more data than necessary for the task at hand. As a result, an honest-but-curious TAP has rich access to data and they are lucrative targets for attackers. Automations can be customized by small JS snippets (filter code) which typically only use a part of the data available. Our approach introduces a pre-computation step at the information provider that limits the data sent to the TAP. This improves on previous research in a setting with multiple information providers and when only parts of the attributes are used.



Multi-party Hybrid Homomorphic Encryption

Anton Israelsson
RISE, Cybercampus Sweden

Supervisor: Shahid Raza

Abstract:

Collaborative data analysis and machine learning are often challenged by technical complexity and privacy, regulation, or proprietary constraints. Multi-Party Computation (MPC) offers a solution but is commonly limited by client-side overhead and complexity. We propose an approach alleviating these issues using Fully Homomorphic Encryption (FHE)-based MPC. This method shifts the primary computational load to a powerful, untrusted server, reducing client-side complexity and overhead. Our scheme combines hybrid homomorphic encryption (transcipherring) with multi-party FHE. This allows multiple resource-limited parties to perform privacy-preserving computations while encrypting their data using symmetric encryption, with a central server managing the complex operations, fostering simplified collaboration.

Secure access control and blockchain

Israel Ehile Ehile
Luleå University of Technology

Supervisor: John Lindström, Karl Andersson (LTU)

Abstract:

Access control systems regulate resource usage and assign permissions for specific actions. The current approaches have limited fine-grained control, are vulnerable to single points of failure, and have concerns related to confidentiality, integrity, and availability. Blockchain-based access control mechanisms attract significant interest, but missing is a comprehensive analysis of the current empirical research. With features like decentralization, immutability, transparency, and peer-to-peer architecture, blockchain offers a trusted and secure foundation for applications in security, data management, and access control. This poster presents a review of blockchain-based access control systems, platforms used, blockchain properties applied, model types, and associated tools and testbeds.



Digital Sovereignty for Collaborative AI Engineering

Venkata Satya Sai Ajay Daliparth
Blekinge Institute of Technology

Supervisor: Kurt Tutschku

Abstract:

Collaborative AI engineering enables the exchange of artifacts among multiple stakeholders. However, establishing digital sovereignty is essential for gaining stakeholder trust to collaborate. This work investigates the methods and tools for implementing digital sovereignty for collaborative AI engineering use cases in AI marketplaces, under Horizon 2020 projects - BonsApps and dAIEdge. The results include (i) a decentralized data marketplace that enables owners with sovereignty over maintaining data artifacts, (ii) A license management system for defining and enforcing artifact usage rights. This work contributes by translating digital sovereignty definitions and requirements into the context of collaborative AI, as well as designing and implementing control mechanisms that enable stakeholders with digital sovereignty.

Quantum random number generation and generation of shared keys using quantum key distribution

Joakim Argillander
Linköping University

Supervisor: Mikael Asplund

Abstract:

My research interests lie in the field of experimental quantum communication, with a particular focus on the development of photonic quantum random number generators (QRNGs) and integration of them. Throughout my PhD studies, I have worked extensively on designing and implementing novel QRNG architectures capable of certifying the generated randomness as both unpredictable and private. These devices leverage fundamental quantum processes to produce entropy that is provably resistant to adversarial influence, making them essential components for future-proof cryptographic protocols.

In parallel with this work, I have a strong interest in the supporting electronics and photonic instrumentation required to realize high-performance, low-



latency quantum communication systems. This includes custom hardware design for fast signal acquisition, real-time processing, and synchronization using field-programmable gate arrays (FPGAs). My expertise in this area has enabled me to develop efficient, scalable solutions that bridge a gap in the market for lab-grade instrumentation.

More recently, I have contributed significantly to the National Quantum Communication Infrastructure in Sweden (NQCIS) project, where we are building the first large-scale quantum communication network in the country. As part of this initiative, I have been directly involved in the deployment, characterization, and operation of quantum key distribution (QKD) systems over metropolitan and inter-city fiber links, including a 300 km long connection between Linköping and Stockholm. This work integrates optical communication engineering, system-level integration, and experimental quantum cryptography, and contributes toward the realization of national large-scale quantum networks.

As part of my work on QKD technology, I am also exploring the concept of vendor-agnostic QKD device certification, which aims at providing a method for validating the security of QKD systems independently of the manufacturers - effectively mitigating one of the biggest arguments against the widespread adoption of QKD technology. Lastly, drawing from my background in both quantum communication and computer engineering, I have recently started exploring applications for bandwidth-efficient classical communication protocols that can utilize ultra-low-bandwidth one-time-pad-encrypted channels established by QKD systems.



Making Moving Target Defense Practical

Syed Umer Bukhari

Chalmers University of Technology

Abstract:

Since the seminal work by Ristenpart et al.[1] in 2009, an ever-increasing number of side-channel and covert-channel attacks have been shown to affect cloud systems. We introduce a practical approach to Moving Target Defense (PMTD), designed to create a dynamic and complex attack surface against side-channel and covert-channel attacks, as well as to counter application-level threats, while being practical. Our solution relies on using multiple state-of-the-art system techniques to eliminate performance overheads, including techniques for image reduction using debloating and better container provisioning, minimizing the overhead and latency while keeping the overall system well secure. We analyze a large number of state-of-the-art side-channel attacks with respect to their time of collocation requirements, and show how PMTD can be used against many of these attacks. While attackers aim to maximize co-location time and opportunities, PMTD introduces stochasticity in the system, reducing the probability of an attacker being co-located with a victim, even when the attacker has a coordinated distributed attack with multiple attack containers. Through simulation and real time experiments, we show the efficacy of PMTD in reducing the probability and time of collocation by up to 66% compared to naive MTD, while reducing the total number of swaps by up to 90%.



Detecting Prototype Pollution in Client-side JavaScript

Samuel Kajava
Chalmers University of Technology

Supervisor: Andrei Sabelfeld

Abstract:

Prototype pollution is a vulnerability that exploits the internals of the inheritance mechanism in JavaScript (JS). Given the widespread use of this programming language, detecting and mitigating the threats introduced by the vulnerability is critical. We address this issue by using state-of-the-art tools for static analysis, high coverage source code retrieval, and readily available means to verify identified vulnerabilities.

Our approach introduces a multistage framework for detecting prototype pollution vulnerabilities utilizing web-crawling, static analysis with CodeQL, static payload generation, and dynamic gadget detection.

Our preliminary results demonstrate the viability of our approach, detecting 28 new cases of prototype pollution, two of which are vulnerable to XSS.

Time Predictability vs. Security: Randomizing Schedules for Real-Time DAG Tasks

Xiuqi Zhang
Chalmers University of Technology

Abstract:

The deterministic nature of real-time DAG schedules creates security vulnerabilities exploitable by timing-based side-channel attacks. We introduce a novel framework that enhances security by strategically inserting and parameterizing 'fake' subtasks into a DAG, provably maintaining schedulability under any work-conserving global scheduler. We propose two new metrics, System Threat and Task Distribution Entropy, to quantify security provided by randomness. Our experiments demonstrate significant threat reduction and reveal a key insight: more computational resources do not necessarily improve security, highlighting the need for co-designed, resource-aware randomization.



A Practical Guideline and Taxonomy to LLVM's Control Flow Integrity

Sabine Houy
Umeå University

Supervisor: Alexandre Bartel

Abstract:

Memory corruption vulnerabilities enable arbitrary code execution by hijacking control flow, posing a severe threat to modern systems. While Control Flow Integrity (CFI) offers promising protection, practical guidance for real-world deployment is lacking. We established a taxonomy that maps LLVM's forward-edge CFI variants to four high-impact memory corruption vulnerabilities, providing developers with actionable guidance. We evaluate LLVM's CFI against one CVE of each category and explain why CFI blocks exploitation in two cases while failing in the remaining cases. Our findings support informed deployment decisions and provide a foundation for improving the practical use of CFI in production systems.

FAST: Rapid Prototyping Framework for Secure IoT Development and Deployment

Hendarmawan
RISE

Supervisor: Shahid Raza

Abstract:

The rapid development of the Internet of Things (IoT) presents significant challenges in securing devices, particularly in resource-constrained environments such as edge computing. Designing secure hardware accelerators for IoT systems, primarily using Field Programmable Gate Arrays (FPGAs), is a complex and time-consuming task. To address these challenges, we introduce the FAST Framework, a platform designed to simplify the rapid prototyping of secure hardware accelerators for FPGA-based edge computing. It provides an integrated hardware/software (HW/SW) co-design solution, enabling IoT developers to efficiently create, optimize, and deploy secure hardware designs. The framework streamlines the development of secure and high-performance IoT systems by



automating design generation and incorporating integrated security features, such as cryptographic functions and memory protection, and reducing the complexity typically associated with FPGA design, allowing software developers to create custom hardware intellectual property (IP) and accelerators without requiring extensive hardware expertise. In addition, the framework emphasizes resource efficiency and scalability, enabling developers to deploy optimized IoT hardware that can scale up and adapt to evolving security requirements. FAST enables faster prototyping and improved performance through flexible design modules and streamlined deployment processes, empowering developers to meet IoT systems' security and throughput demands. Our proposal shows significant improvements in throughput (up to 8x) for hashing algorithm speedup compared to traditional processors and reductions in development time and complexity. This work paves the way for a more accessible and scalable approach to secure IoT deployment, enhancing the overall security of the IoT ecosystem.