UMR IRISA

Activity Report 2025

Team SUSHI

SecUrity at the Software-Hardware Interface

Joint team with Centre Inria de l'Université de Rennes

D3 – Architecture

Université de Rennes | ENS rennes | INSA INSTITUT NATIONAL DES SCIENCES APPLIQUÉES RENNES | cnrs | Inria | Inserm | CentraleSupélec | IMT Atlantique Bretagne-Pays de la Loire École Mines-Télécom | Université Bretagne Sud ubs

Project-Team SUSHI

Creation of the Project-Team: 2024 January 01

Keywords

Computer sciences and digital sciences

A1.1.2. – Hardware accelerators (GPGPU, FPGA, etc.)

A1.1.8. – Security of architectures

A1.1.10. – Reconfigurable architectures

A1.1.13. – Virtualization

A2.2.1. – Static analysis

A2.2.5. – Run-time systems

A2.2.6. – GPGPU, FPGA...

A2.2.9. – Security by compilation

A2.6.1. – Operating systems

A2.6.3. – Virtual machines

A4.1.2. – Hardware attacks

A4.4. – Security of equipment and software

A4.5.3. – Program proof

A4.9.1. – Intrusion detection

A4.9.3. – Reaction to attacks

Other research topics and application domains

B6.5. – Information systems

B6.6. – Embedded systems

Contents

Project-Team SUSHI **1**

1 Team members, visitors, external collaborators **4**

2 Overall objectives **5**

3 Research program **5**

4 Application domains **6**

5 Latest software developments, platforms, open data **6**
5.1 Latest software developments . . . 6
5.1.1 HyperSec . . . 6
5.1.2 koika-sushi . . . 6
5.1.3 COQQTL . . . 6
5.1.4 heRVé . . . 7
5.1.5 Knock-Knock . . . 7
5.1.6 Triskel . . . 7
5.1.7 Tansiv . . . 7

6 New results **8**
6.1 Vulnerability identification and security by design . . . 8
6.1.1 Black-Box, Platform-Agnostic DRAM Address-Mapping Reverse Engineering . . . 8
6.1.2 Flush-based Cache Attacks on Modern / Multi-Socket x86 Systems . . . 8
6.1.3 WIP: A Second Look at Port Assignment on Intel CPUs . . . 9
6.1.4 Hardware Security Checkers for the Detection of Microarchitectural Side-Channel Attacks . . . 9
6.1.5 Hardware-based methodologies for the detection of Hardware Trojans . . . 9
6.1.6 Detection of loop-based microarchitectural memory attacks . . . 10
6.1.7 Hardware-based methodologies for the detection of Hardware Trojans . . . 10
6.1.8 Type-based security properties assurance in operating systems . . . 10
6.1.9 AI implementation security . . . 11
6.1.10 Exploiting microarchitectural vulnerabilities from physical side-channel leakage . . . 11
6.1.11 Side-channel vulnerabilities in heterogeneous mixed-signal systems . . . 12
6.2 Reactive security at the host level . . . 12
6.2.1 Time control for stealth analysis sandboxes . . . 12
6.2.2 Hypervisor extension using a Domain Specific Language to detect rootkits . . . 13
6.3 Formal models and proofs for low-level security . . . 13
6.3.1 Formal verification of hardware/software security mechanisms . . . 13
6.4 Binary analysis and vulnerability detection . . . 13
6.4.1 Implementation of an improved algorithm for control-flow graph visualization . . . 13
6.4.2 WIP: Vulnerability detection in macOS kernel extensions through fuzzing . . . 14
6.5 Other topics . . . 14
6.5.1 AI-based Reasoning and Knowledge Formalization using Large Language Models 14
6.5.2 Hardware acceleration of decision-making functions . . . 14
6.5.3 Ethical and legal challenges of AI in judicial decision-making . . . 15
6.5.4 Efficient signal processing for low-delay embedded systems . . . 15
6.5.5 Review on software attacks against Just-in-Time compilers in Virtual Machines . . . 15

7 Bilateral contracts and grants with industry **16**
7.1 Bilateral Grants with Industry . . . 16
7.1.1 ANSSI: . . . 16
7.1.2 ANSSI: . . . 16
7.1.3 DGA: . . . 16
7.1.4 DGA: . . . 16
7.1.5 HP Labs: . . . 16
7.1.6 Thalium: . . . 17

8 Partnerships and cooperations **17**
8.1 International research visitors . . . 17
8.1.1 Visits of international scientists . . . 17
8.2 National initiatives . . . 18
8.2.1 PEPR Cybersécurité: projet SECUREVAL (2022-2028) . . . 18
8.2.2 ANR Project: TrustGW (2021-2026) . . . 18
8.2.3 ANR Project: ATTILA (2022-2025) . . . 18
8.2.4 CMA project : Train-Cyber-Expert (TCE) (2022-2026) . . . 19
8.3 Regional initiatives . . . 19

9 Dissemination **19**
9.1 Promoting scientific activities . . . 19
9.1.1 Scientific events: organisation . . . 19
9.1.2 Scientific events: selection . . . 20
9.1.3 Journal . . . 20
9.1.4 Invited talks . . . 21
9.1.5 Leadership within the scientific community . . . 21
9.1.6 Scientific expertise . . . 21
9.1.7 Research administration . . . 21
9.2 Teaching - Supervision - Juries - Educational and pedagogical outreach . . . 21
9.2.1 Supervision . . . 21
9.2.2 Educational and pedagogical outreach . . . 23
9.3 Popularization . . . 23
9.3.1 Productions (articles, videos, podcasts, serious games, ...) . . . 23
9.3.2 Participation in Live events . . . 23
9.3.3 Others science outreach relevant activities . . . 23

10 Scientific production **23**
10.1 Publications of the year . . . 23
10.2 Cited publications . . . 26

1 Team members, visitors, external collaborators

Faculty Members

- Guillaume Hiet [Team leader, CENTRALESUPELEC, Professor, HDR]
- Alessandro Palumbo [CENTRALESUPELEC, Associate Professor]
- Thomas Rokicki [CENTRALESUPELEC, Associate Professor]
- Ruben Salvador Perea [CENTRALESUPELEC, Associate Professor Delegation, from Sep 2025]
- Ruben Salvador Perea [CENTRALESUPELEC, Associate Professor, until Aug 2025]
- Frédéric Tronel [CENTRALESUPELEC, Associate Professor]
- Yaelle Vincont [ENS RENNES, Associate Professor]
- Pierre Wilke [CENTRALESUPELEC, Associate Professor]

Post-Doctoral Fellows

- Lorenzo Casalino [CENTRALESUPELEC, Post-Doctoral Fellow]
- Quentin Ducasse [CENTRALESUPELEC, Post-Doctoral Fellow]
- Jonas Bror Haglund [CENTRALESUPELEC, Post-Doctoral Fellow, from Dec 2025]

PhD Students

- Zakaria Belkadi [INRIA]
- Erwan Fasquel [CENTRALESUPELEC]
- Leslie Fifanon [CENTRALESUPELEC, from Nov 2025]
- Théo Goureau [UNIV RENNES I, from Oct 2025]
- Lionel Hemmerle [CENTRALESUPELEC]
- Samuel Hiron [CENTRALESUPELEC, from Dec 2025]
- Cyprien Jules [INRIA, from Nov 2025]
- Simon Legros [HP, from Nov 2025]
- Eliott Quere [UNIV RENNES I]
- Jack Royer [THALES, CIFRE, from Jun 2025]

Interns and Apprentices

- Killian Callac [CENTRALESUPELEC, Intern, from Mar 2025 until Aug 2025]
- Hugo Depuydt [ENS RENNES, Intern, until Jul 2025]
- Antoine Doglioli [CENTRALESUPELEC, Intern, from Mar 2025 until Jul 2025]
- Ismael Gaye [CENTRALESUPELEC, Intern, from May 2025 until Jul 2025]
- Samuel Hiron [CENTRALESUPELEC, Intern, until Jul 2025]
- Cyprien Jules [CENTRALESUPELEC, Intern, from Mar 2025 until Oct 2025]
- Antoine Plin [CENTRALESUPELEC, Intern, from Apr 2025 until Jul 2025]

Administrative Assistant

- Lydie Mabil [INRIA]

Visiting Scientist

- Pierre Olivier [UNIV MANCHESTER, from Jun 2025 until Jul 2025]

External Collaborator

- Louis Rilling [DGA-MI]

2 Overall objectives

Computer systems (e.g. personal computers, servers, or embedded systems) rely on computing platforms to execute user applications and host user data. These computing platforms are made of different hardware and system software (i.e., low-level software components such as Operating Systems, hypervisors, and firmware) and tend to grow in complexity. Indeed, they must fulfil various objectives: optimizing performance and offering new services while limiting energy consumption. To achieve these goals, they rely on complex interactions between heterogeneous computation units like CPU, hardware accelerators, FPGA, and system software. This trend increases with growing architectural heterogeneity and emerging computing paradigms, which might bring new yet hidden vulnerabilities.

Besides these various objectives, such platforms are critical for user applications and data security. To that end, they form the so-called Trusted Computing Base of computer systems. Consequently, any breach in the TCB will dramatically impact user applications and data. This growing complexity of interactions between software and hardware components raises serious privacy and trust issues in today's computer systems. The main research goal of the SUSHI team is to address these issues by assessing and increasing the security level of existing and future computing platforms at the software/hardware interface.

3 Research program

The goal of the SUSHI team is to thoroughly study and contribute to the security of computing platforms at the software/hardware interface, both from an attack and a defense perspective. We do this by exploring three complementary research axes :

- **Vulnerability identification and security by design.** This axis aims first to identify new vulnerabilities resulting from software/hardware interactions in such complex and heterogeneous platforms and, second, to propose secure-by-design approaches to prevent the exploitation of such vulnerabilities.
- **Reactive security at the host level.** This axis focuses on host-based intrusion detection and reaction by leveraging software/hardware interactions.
- **Formal models and proofs for low-level security.** This axis aims to formally prove the security properties enforced or detected by software/hardware mechanisms.

We propose to decline these research axes on three different levels at the software/hardware interface:

- The hardware architecture and microarchitecture level focuses on the hardware part of the interface, which should provide software with the required services to ensure security;
- The system software level focuses on low-level software, such as OSes or hypervisors, which are heavily tied to hardware interfaces and must use them correctly to achieve security;
- The binary executable analysis and instrumentation level focuses on analysing and modifying binary executables, i.e., sequences of instructions belonging to the ISA.

4 Application domains

We focus on host-based system security but do not consider any specific application or type of system. We are interested in various systems, from tiny IoT devices to high-end workstations or servers. Moreover, we aim to provide trusted software/hardware computing platforms that could be helpful in different application domains such as defence, health, industry, or finance.

5 Latest software developments, platforms, open data

5.1 Latest software developments

5.1.1 HyperSec

Name: A Dedicated Language for Adaptive Rootkit Detection in Virtual Machines

Keywords: Cybersecurity, Anomaly detection, Security

Functional Description: HyperSec is a domain-specific language allowing virtual machines (VMs) to send programs to an hypervisor. These programs leverage the VM's knowledge of its internal OS structures and enable the hypervisor to detect intrusions. To secure this process, we enforce constraints on the set of programs accepted by the hypervisor to prevent the exploitation of vulnerabilities inside the hypervisor.

Contact: Lionel Hemmerle

5.1.2 koika-sushi

Keywords: Formal methods, Coq, CPU, Cybersecurity

Functional Description: This project is the fork of Kôika maintained by the SUSHI team. Kôika is a rule-based Hardware Design Language embedded within Coq. This fork proposes a framework to verify security properties on Kôika circuits using Coq and SMT solvers.

Release Contributions: - Support for more recent Coq versions - Support for SMT powered proofs - Reorganized codebase - Extended RV processor (moved to its own repository)

News of the Year: Support for more recent Coq versions Support for proofs using the progressive rewriting approach described in our CSF paper Support for SMT powered proofs Reorganized codebase Extended RV processor (moved to its own repository)

URL: https://gitlab.inria.fr/SUSHI-public/FMH/koika

Publication: hal-04118645

Contact: Pierre Wilke

Participant: 2 anonymous participants

Partner: EPFL - Ecole Polytechnique Fédérale de Lausanne

5.1.3 COQQTL

Keywords: Formal methods, Coq

Functional Description: Formal semantics in Coq of the FIRRTL intermediate representation of the CHISEL Hardware Description Language.

News of the Year: Creation of the project, definition of the semantics of most of FIRRTL constructs (excluding modules and assertions).

URL: https://gitlab.inria.fr/SUSHI-public/FMH/coqqtl

Contact: Pierre Wilke

Participant: 2 anonymous participants

5.1.4 heRVé

Keywords: Formal methods, Coq, CPU

Functional Description: This project is a fork of the 4-stage RISC-V pipeline processor of the Kôika project. The design has been rewritten to add support for exceptions and interrupts.

News of the Year: Redesign and support for exceptions and interrupts

URL: https://gitlab.inria.fr/SUSHI-public/FMH/herve

Contact: Pierre Wilke

Participant: 2 anonymous participants

Partner: EPFL - Ecole Polytechnique Fédérale de Lausanne

5.1.5 Knock-Knock

Name: Knock-Knock: Black-Box, Platform-Agnostic DRAM Address-Mapping Reverse Engineering

Keywords: Side Channel, Microarchitecture, Reverse engineering

Functional Description: Knock-Knock is a black-box, platform-agnostic DRAM reversing tool. It allows for automatic reverse mapping from physical addresses to DRAM location using the row-buffer side channel and linear algebra.

URL: https://github.com/antpln/Knock-Knock

Contact: Thomas Rokicki

5.1.6 Triskel

Keywords: Binary analysis, Visualization, Control flow graph

Functional Description: Triskel is a control flow graph visualization library. For instance, it can be combined to a binary analysis framework, or directly on LLVM binary files.

Contact: Jack Royer

5.1.7 Tansiv

Name: Time-Accurate Network Simulation Interconnecting Vms

Keywords: Operating system, Virtualization, Cloud, Simulation, Cybersecurity

Functional Description: Tansiv: Time-Accurate Network Simulation Interconnecting Virtual machines (VMs). Tansiv is a novel way to run an unmodified distributed application on top of a simulated network in a time accurate and stealth way. To this aim, the VMs execution is coordinated (interrupted and restarted) in order to garantee accurate arrival and transfer of network packets while ensuring realistic time flow within the VMs. The project can leverage several frameworks for simulating the data (SimGrid or ns-3) and several virtualization solutions to encapsulate the application, intercept the network traffic and enforce the interruption decision (Qemu in emulation mode, KVM and Xen with hardware-accelerated virtualization). Tansiv can be used in various situations: malware analysis (e.g. to defeat malware evasion technique based on network timing measures) or analysis of an application on a geo-distributed context.

Contact: Louis Rilling

Partner: DGA-MI

6 New results

6.1 Vulnerability identification and security by design

6.1.1 Black-Box, Platform-Agnostic DRAM Address-Mapping Reverse Engineering

Participants: Antoine Plin, Lorenzo Casalino, Thomas Rokicki, Ruben Salvador.

Keywords: Microarchitectecture, DRAM security, Side Channels, Reverse engineering.

Modern Systems-on-Chip (SoCs) employ undocumented linear address-scrambling functions to obfuscate DRAM addressing, which complicates DRAM-aware performance optimizations and hinders proactive security analysis of DRAM-based attacks; most notably, Rowhammer. Although previous work tackled the issue of reversing physical-to-DRAM mapping, existing heuristic-based reverse-engineering approaches are partial, costly, and impractical for comprehensive recovery. In consequence, our research [12] establishes a rigorous theoretical foundation and provides efficient practical algorithms for *black-box, complete physical-to-DRAM address-mapping recovery.*

We first formulate the reverse-engineering problem within a linear algebraic model over the finite field GF(2). We characterize the timing fingerprints of row-buffer conflicts, proving a relationship between a bank addressing matrix and an empirically constructed matrix of physical addresses. Based on this characterization, we develop an efficient, noise-robust, and fully platform-agnostic algorithm to recover the full bank-mask basis in polynomial time, a significant improvement over the exponential search from previous works. We further generalize our model to complex row mappings, introducing new hardware-based hypotheses that enable the automatic recovery of a row basis instead of previous human-guided contributions.

Evaluations across embedded and server-class architectures confirm our method's effectiveness, successfully reconstructing known mappings and uncovering previously unknown scrambling functions. Our method provides a 99% recall and accuracy on all tested platforms. Most notably, our pipeline runs in under a few minutes, even on systems with more than 500GB of DRAM, showcasing the scalability of our method. This approach provides an automated, principled pathway to accurate DRAM reverse engineering.

6.1.2 Flush-based Cache Attacks on Modern / Multi-Socket x86 Systems

Participants: Thomas Rokicki.

Keywords: Microarchitectecture, Side Channels, Cache Attacks.

Flush-based cache attacks have been extensively studied and leveraged, yet their behavior on today's complex x86 platforms is not fully understood. Notably, as cache and memory latency depend on the physical layout of cores, caches, and Numa nodes, system topology increasingly influences the latencies underlying such attacks. With Guillaume Didier, from Universität des Saarlandes and Augustin Lucas from ENS Lyon, we thus investigate in [15, 22] the impact of this growing complexity on the effectiveness of flush-based attacks. We present a large-scale, topology-aware study of Flush+Reload and Flush+Flush across 36 Intel and AMD, single- and multi-socket systems, with server and client CPUs. We show topology-induced contributions dominate the latency variations. In particular, Numa's contribution on multi-socket systems makes topology-unaware attacks unreliable. Our topology-aware calibration accounts for topological parameters like attacker/victim cores, memory Numa node, and target address, improving error rates and attack viability. We demonstrate Flush+Flush works on AMD targets, and is often more accurate than Flush+Reload on modern x86 platforms. Finally, we introduce LoadFlush+Reload, a covert-channel

comparing invalid to shared loads, with double the true capacity than Flush+Reload / Flush+Flush. Our results show that topology awareness is required for dependable cache attacks on recent platforms, and provide practical guidance for attackers and defenders.

6.1.3 WIP: A Second Look at Port Assignment on Intel CPUs

Participants: Thomas Rokicki.

Keywords: Microarchitectecture, Side channel, Port Assignment.

Modern Intel CPUs use a superscalar, out-of-order execution (OoOE) pipeline that maximizes instruction throughput through instruction-level parallelism (ILP). A critical component affecting system performance is the process in which µ-ops are mapped to execution ports, which lays at the core of the ability to perform OoOE. Significant effort has been undertaken in the attempt to reverse-engineer and model this process.

Our work in [9] revisits this question from a *security* perspective, focusing on potential vulnerabilities stemming from µ-op port assignment and execution. Using carefully designed code gadgets, we expose behaviors that contradict state-of-the-art models.

A key observation, which underlies our entire work, is that *all proposed models fail to capture significant aspects of µ-op port assignment and execution*. For most Intel architectures the µ-op port assignment policy exhibits significant irregularities and dynamics, that are inconsistent with state-of-the-art models.

In this Work In Progress, we ask: *Which undocumented behaviors in Intel's port assignment algorithm affect instruction scheduling?* By extending the micro-benchmarks used by Abel and Reineke, we identify corner cases that strongly deviates from the common case, where the CPU adopts at least three different port assignment strategies. These depend not only on the instruction and the microarchitectural generation, but, surprisingly, also on immediate operands.

To make these findings fully accessible, we release an interactive online database showcasing all our experiments and results. We envision this resource as a long-term reference point for researchers exploring Intel port assignment.

6.1.4 Hardware Security Checkers for the Detection of Microarchitectural Side-Channel Attacks

Participants: Alessandro Palumbo.

Keywords: Hardware Security, Microarchitectural Side-Channel Attacks, RISC-V.

Microarchitectural Side-Channel Attacks represent significant challenges to the security and reliability of modern microprocessor-based systems. The poster presented at the RISC-V Summit Europe 2025 discusses a hardware-based approach to enhance microprocessor security detecting Microarchitectural Side-Channel Attacks by employing hash functions and Machine Learning [29], [18].

6.1.5 Hardware-based methodologies for the detection of Hardware Trojans

Participants: Alessandro Palumbo, Ruben Salvador.

Keywords: Hardware Security, Hardware Trojans, RISC-V.

Hardware Trojans represent another significant challenge to the security and reliability of modern microprocessor-based systems. The manuscripts presented at the RISC-V Summit Europe 2025 [30], [19] report two complementary approaches to enhance security at hardware level. First, programmable Hardware Security Modules are proposed to detect Hardware Trojan Horses by monitoring instruction-fetch activities,

identifying malicious interferences, and preventing software-exploitable Hardware Trojan activations. Second, a methodology based on side-channel analysis is proposed to verify the integrity of FPGA bitstreams, allowing the identification of tampered configurations through the extraction and classification of both high- and low-level features.

The paper presented at IEEE IOLTS 20255 [20], goes into more details of the first methodology, leveraging a Hardware Security Checker integrated into a RISC-V microprocessor, featuring Error Correction Codes, and in particular Hamming Single Error Correction (HSEC) architectures, to identify malicious instruction injections that disrupt the normal execution flow and trigger unauthorized programs. Experimental results show that this solution achieves a 100% detection rate with no false positives, while incurring minimal hardware overhead and no performance degradation.

6.1.6 Detection of loop-based microarchitectural memory attacks

Participants: Alessandro Palumbo, Ruben Salvador.

Keywords: Hardware Security, Count-Min Sketche, XOR Filter, Rowhammer, Spectre.

With the miniaturization of DRAM technology, memory devices have become increasingly susceptible to hardware-based attacks that exploit physical proximity between rows to alter or extract data. Notable among these are loopbased attacks such as Rowhammer, which induces bit flips in adjacent memory rows. Other known attacks, like Spectre, leverage shared memory to infer confidential data by cache access time estimations. To address these threats, which share similar loop-based code patterns, an integrated detection system based on XOR Filters (XFs) and Count-Min Sketches (CMSs) has been presented at IEEE DFT 2025 [8]. The proposed security system requires only three memory accesses per activation of a memory row to detect loop attacks, filtering most malicious loops in a first stage (i.e., the XF) and monitoring the remaining ones in a second stage (i.e., the CMS). An alert is triggered when a loop exceeds a defined activation threshold. Experimental results demonstrate that the proposed approach effectively detects not only loop attacks but also other anomalous instructions that generate suspicious loops, achieving low undetected attack rates (on average below 0.4%), while consuming minimal memory and computational resources.

6.1.7 Hardware-based methodologies for the detection of Hardware Trojans

Participants: Alessandro Palumbo, Ruben Salvador.

Keywords: Microarchitectural Side-Channel Attacks, RISC-V, gem5, Machine Learning, Hardware Trojans.

A simulation-based framework relying on the gem5 platform was developed to support the detection of hardware-level attacks on RISC-V microprocessors. The framework enables the automated extraction of a large set of Hardware Performance Counter features from simulated executions, facilitating the construction of datasets for Machine Learning-based detection.

This tool was applied to both the detection of software-exploitable Hardware Trojans and Microarchitectural Side-Channel Attacks, demonstrating the effectiveness of simulation-based approaches for evaluating and comparing detection strategies [11], [10].

6.1.8 Type-based security properties assurance in operating systems

Participants: Zakaria Belkadi, Killian Callac, Hugo Depuydt, Guillaume Hiet, Louis Rilling, Frédéric Tronel.

Keywords: Operating systems, Rust, security assurance.

With the PhD thesis of Zakaria Belkadi we have started the study of a new trade-off between getting assurance on security properties of an operating system and the development and maintenance cost that this assurance involves. Based on using the strong typing system of memory-safe system languages like Rust, this new trade-off is expected to provide better assurance than traditional, manual check-based development and test-based assurance evaluation, as well as sustainable costs contrary to state-of-the-art formal proof methods. To verify these expectations, we revisit the implementation of operating system services to encode low-level security properties with types that represent only the authorized entities for which operations can be performed. A simple case study on the operations authorized for file descriptors (read and write), which are set when creating the descriptor and must be enforced for the lifetime of the descriptor, has been started on the Redox operating system and is being extended to the recently published Asterinas operating system. Compared to traditional manual check-based development, the preliminary results obtained suggest a negligible performance overhead and a comparable development cost. Preliminary results were presented at GT SSLR days.

6.1.9 AI implementation security

Participants: Ruben Salvador.

Keywords: IP Piracy, Reverse engineering, Deep Neural Network, Side-Channel Analysis.

Dataflow neural network accelerators efficiently process AI tasks on FPGAs, with deployment simplified by readyto-use frameworks and pre-trained models. However, this convenience makes them vulnerable to malicious actors seeking to reverse engineer valuable Intellectual Property (IP) through Side-Channel Attacks (SCA). This paper [7] proposes a methodology to recover the hardware configuration of dataflow accelerators generated with the FINN framework. Through unsupervised dimensionality reduction, we reduce the computational overhead compared to the state-of-the-art, enabling lightweight classifiers to recover both folding and quantization parameters. We demonstrate an attack phase requiring only 337 ms to recover the hardware parameters with an accuracy of more than 95% and 421 ms to fully recover these parameters with an averaging of 4 traces for a FINN-based accelerator running a CNN, both using a random forest classifier on side-channel traces, even with the accelerator dataflow fully loaded. This approach offers a more realistic attack scenario than existing methods, and compared to SoA attacks based on tsfresh, our method requires less time for preparation and attack phases, respectively, and gives better results even without averaging traces.

6.1.10 Exploiting microarchitectural vulnerabilities from physical side-channel leakage

Participants: Eliott Quéré, Ruben Salvador.

Keywords: FPGA security, side-channel attacks, DRAM access patterns, cache-miss leakage, power analysis, PCIe contention, RDMA timing attacks, CPU fingerprinting, embedded sensors, cloud computing security..

The widespread adoption of FPGA-accelerated computing in embedded and cloud environments introduces new side-channel threats due to shared hardware resources. This work [21] investigates DRAM access patterns as a leakage source to fingerprint CPU activity, examining both SoC-FPGA and cloud-based co-processor models. In SoC environments, cache-missinduced DRAM activity generates measurable power fluctuations that can be remotely observed. While previous research has detected these fluctuations using external electromagnetic probes for side-channel-based disassembly, we assess whether embedded FPGA sensors can achieve similar results, enabling attackers to infer CPU operations without physical access. However, in cloud-based co-processor models, where FPGA-CPU interactions occur over PCIe and RDMA, large-scale

power management appears to significantly lower the Signal-to-Noise Ratio (SNR), potentially making power side channels more challenging to exploit compared to SoC-FPGAs. Given this uncertainty, we investigate the feasibility of power-based leakage while also exploring timing-based side channels leveraging PCIe contention and RDMA latency variations, which have been shown to reveal workload characteristics. By evaluating both power and timing leakage across these architectures, we comprehensively assess side-channel risks in FPGA-accelerated platforms and emphasize the need for stronger isolation mechanisms.

6.1.11 Side-channel vulnerabilities in heterogeneous mixed-signal systems

Participants: Ruben Salvador.

Keywords: Side-Channel Attacks, EM Side Channels, Screaming-Channel Attacks, Multi-Channel Attacks..

Side-channel attacks consist of retrieving internal data from a victim system by analyzing its leakage, which usually requires proximity to the victim in the range of a few millimetres. Screaming channels are EM side channels transmitted at a distance of a few meters. They appear on mixed-signal devices integrating an RF module on the same silicon die as the digital part. Consequently, the side channels are modulated by legitimate RF signal carriers and appear at the harmonics of the digital clock frequency. While initial works have only considered collecting leakage at these harmonics, late work has demonstrated that the leakage is also present at frequencies other than these harmonics. This result significantly increases the number of available frequencies to perform a screaming-channel attack, which can be convenient in an environment where multiple harmonics are polluted. This work [24] studies how this diversity of frequencies carrying leakage can be used to improve attack performance. We first study how to combine multiple frequencies. Second, we demonstrate that frequency combination can improve attack performance and evaluate this improvement according to the performance of the combined frequencies. Finally, we demonstrate the interest of frequency combination in attacks at 15 and, for the first time to the best of our knowledge, at 30 meters. One last important observation is that this frequency combination divides by 2 the number of traces needed to reach a given attack performance.

This work is currently under evaluation (major revision) at IEEE TIFS (Transactions on Information Forensics and Security).

6.2 Reactive security at the host level

6.2.1 Time control for stealth analysis sandboxes

Participants: Louis Rilling.

Keywords: Evasive malware, Virtualization, Network simulation, Introspection.

Virtual Machine Introspection (VMI) is used by sandbox-based dynamic malware detection and analysis frameworks to observe malware samples while staying isolated and stealthy. Sandbox detection and evasion techniques based on hypervisor introspection are becoming less of an issue since running server and workstation environments on hypervisors is becoming standard and high-end sandboxes manipulate virtual clocks to mask VM execution pauses caused by VMI. However masking VM execution pauses on multicore VMs assumes that pauses involve all virtual cores, which breaks introspection practices observed with popular tools like libVMI. Indeed some introspection procedures require to pause a virtual core while waiting for an event on another virtual core. Moreover, despite virtualization not being a reliable hint for a malware to be under analysis, single-core VMs are less and less common in production workloads and having only one core is a reliable hint of being analyzed. To lower the VMI detection ability of malware on multi-core VMs, we first proposed several metrics to characterize the side effects caused by a VMI monitor in multi-core VMs. Then we introduced a new strategy to conceal VMI pauses in multi-core guests, and asserted its performance

compared to the regular approaches used by sandboxes. Our results show that this new strategy improves the stealthiness of VMI pauses on a multi-core system for some metrics. This work is published in [3].

This work on hiding VMI pauses was also presented in details in the more general framework of Léo Cosseron's PhD thesis defended in December 2025, where controlling the flow of time in VMs allows to interconnect VMs with a network simulator. This combination of hardware virtualization and network simulation provides both a tool to analyze real distributed applications in reproducible setups and the infrastructure for a performance-wise stealth network environment in malware analysis sandboxes.

6.2.2 Hypervisor extension using a Domain Specific Language to detect rootkits

Participants: Lionel Hemmerlé, Frédéric Tronel, Pierre Wilke, Guillaume Hiet.

Keywords: Virtualization, Introspection, Intrusion Detection System, Rootkits.

Endpoint Detection Reaction (EDR) needs to be protected against attackers who manage to gain access to a high privilege level. We propose using virtual extensions for this purpose: the protected system is placed in a VM, and the EDR is implemented in the hypervisor [6]. This ensures that the EDR remains functional even if the VM is fully compromised. However, in such setup, the EDR loses the operating system abstraction provided by the VM. To reduce this semantic gap, we developed a new language that allows a VM to write and send programs to the hypervisor. Since these programs are produced by the VM, we can legitimately assume they have the necessary knowledge about the internal structure of the VM's operating system. The hypervisor runs these programs to detect intrusions. To secure this process, the hypervisor must enforce strict security constraints to ensure that these programs do not introduce new vulnerabilities. Most tested rootkits could be detected with only an acceptable overhead added to the VM's execution. Furthermore, we demonstrate that this overhead can be further reduced by executing native binaries instead of relying on an interpreter. A paper about these results has been submitted in early 2025 and is under review.

6.3 Formal models and proofs for low-level security

6.3.1 Formal verification of hardware/software security mechanisms

Participants: Guillaume Hiet, Pierre Wilke, Cyprien Jules.

Keywords: Formal Methods, Hardware Description Language.

The security of computer systems eventually relies on the security of the underlying hardware and on the security mechanisms that it offers. We extend a simple RISC-V pro- cessor written in the Kôika Hardware Description Language (HDL) with security mechanisms, and prove that these mech- anisms are correct [28, 25]. We do so by compiling Kôika circuits into a representation well-suited for automated theorem provers, therefore delegating the proof construction to an SMT solver. We also strive to make our proofs modular, i.e., function by function or rule by rule.

6.4 Binary analysis and vulnerability detection

6.4.1 Implementation of an improved algorithm for control-flow graph visualization

Participants: Jack Royer, Frédéric Tronel, Yaelle Vincont.

Keywords: Reverse engineering, control-flow graph, visualization.

Human reverse engineers rely on automatic disassembly tools to make sense of the binary code they are analyzing. One view usually offered by such tools is the control-flow graph : by showing relations between

code blocks, it helps the engineer visualize how the code is organized. This is particularly useful to detect common patterns, such as if-then-else constructions, loops, etc.

However, most of these tools use general graph layout algorithm, which do not account for the specificities of control-flow graphs. After studying the state-of-the-art, we proposed a modified version of popular algorithms, tailored for control-flow graphs and which improves readability by showcasing common patterns. This work was presented at the BAR workshop in 2025 [14] and an open-source tool was published (5.1.6).

6.4.2 WIP: Vulnerability detection in macOS kernel extensions through fuzzing

Participants: Erwan Fasquel, Frédéric Tronel, Yaelle Vincent.

Keywords: Vulnerability discovery, fuzzing, macOS.

Fuzzing is a widely used technique to detect vulnerabilities in non-malicious code, usually binary programs as run on computers. Its application to closed-source code such as macOS kernel extensions is more recent, as it presents new challenges. In particular, kernel extensions are drivers for the macOS kernel, which has recently transitioned to a new processor architecture. As such, fuzzing them requires specialized binary analysis, along with other custom techniques. In this research, we have studied the state-of-the art in order to compile a comprehensive guide to macOS kernel fuzzing. Preliminary results were presented through a poster [27].

6.5 Other topics

6.5.1 AI-based Reasoning and Knowledge Formalization using Large Language Models

Participants: Alessandro Palumbo.

Keywords: Large Language Models, Reasoning, Knowledge Formalization, Human-in-the-loop.

This research line investigates the capabilities and limits of Artificial Intelligence (AI) systems in reasoning and knowledge formalization. A first contribution, presented at JURIX 2025 [16], proposes a structured and replicable framework for empirically evaluating the argumentative and interpretative capabilities of Large Language Models in judicial contexts. The methodology combines an abstract assessment of legal concept understanding with an applied evaluation on real judicial decisions, using Italian case law as a demonstrative use case.

A complementary contribution addresses the formalization of legal knowledge through computational languages. A human-in-the-loop methodology has been presented at ISDFS 2025 [4] to translate normative texts and fundamental case-law principles into Attempto Controlled English, highlighting the challenges posed by the ambiguity of natural legal language and the syntactic constraints of controlled languages.

A further contribution, published at Journal of Inclusive Methodology and Technology in Learning and Teaching [2], introduces an AI-based decision support system aimed at assisting preliminary assessments in criminal proceedings. The proposed system analyzes a structured set of descriptive variables extracted from case scenarios and processes them through a Boolean logic function to evaluate the configurability of the offence.

6.5.2 Hardware acceleration of decision-making functions

Building the software-level implementation of decision-making logic function proposed in [32], the contribution presented at IADIS 2025 [17] investigates its hardware acceleration.

6.5.3 Ethical and legal challenges of AI in judicial decision-making

The integration of predictive technologies into justice system represents a transformative opportunity to enhance judicial efficiency, reduce case processing times, and ensure greater consistency in decisions. These tools, powered by AI, have the potential to optimize case management and resource allocation, particularly in complex areas such as road homicides, where objective and verifiable factors dominate legal reasoning. However, this shift also raises critical ethical and legal challenges, including risks of algorithmic bias, opacity in decision-making processes, and the potential erosion of judicial autonomy. By focusing on issues such as decision-making consistency, the preservation of judicial discretion, and the legal responsibility associated with AI outputs, the work presented at ISDFS 2025 [5] reports a comprehensive analysis of the ethical and legal implications of adopting AI in judicial proceedings. In particular, it explores the benefits and risks of AI-based techniques in the Italian criminal context, emphasizing the need for their implementation to uphold fundamental legal principles such as transparency, accountability, and human-centered decision-making.

6.5.4 Efficient signal processing for low-delay embedded systems

Participants: Ruben Salvador.

Negative Group Delay (NGD) remains a little-known concept in embedded digital signal processing, and most electronic designers are unfamiliar with its design and analysis. This paper [13] presents a clear, step-by-step theoretical framework and a practical methodology for implementing baseband NGD using a low-order finite impulse response (FIR) filter. All FIRbased NGD parameters and characteristics can be defined under specific conditions based solely on the desired NGD value and the sampling frequency. Our synthesis results confirm that the realized time-advance increases with normalized frequency. We validate the approach through simulation and then demonstrate a proof-of-concept FPGA implementation. In transient experiments using a 3 kHz Gaussian pulse, the NGD circuit produces an advance of approximately 13 µs relative to the input signal and the NGD execution time is about 2µs.

6.5.5 Review on software attacks against Just-in-Time compilers in Virtual Machines

Participants: Quentin Ducasse.

Programming Language Virtual Machines (VMs) are composed of several components that together execute and manage languages efficiently. They are deployed in virtually all computing systems through modern web browsers. However, vulnerabilities in any VM component pose a significant threat to security and privacy. In this article [1], we present a survey of software attacks on Just-In-Time (JIT) compilers, which dynamically produce optimized code at run time. We first present an overview and categorization of software attacks and their vectors as presented in the literature, identifying three main attack classes: code injection, code-reuse, and data-only attacks. We show how each can lead to arbitrary code execution. Next, we present a comprehensive taxonomy of defenses, including diversification, strict memory permissions and capability containment. While some were integrated in modern VMs, we draw recommendations for future protections. Securing JIT compilers remains challenging due to inherent conflicts with security principles, such as W `xor` X (Writable exclusive-or eXecutable), and the complexity of JIT optimizations. Finally, we examine how newer architectures, like ARMv8 and RISC-V, face similar threats. With RISC-V's open architecture offering a promising platform for prototyping VM-specific protections and custom security instructions, we discuss hardware-assisted runtime protections and RISC-V extensions that could enhance VM security.

7 Bilateral contracts and grants with industry

7.1 Bilateral Grants with Industry

7.1.1 ANSSI:

Participants: Théo Goureau, Thomas Rokicki, Guillaume Hiet.

Théo Goureau started his PhD in October 2025 in the context of a collaboration between Inria and the ANSSI. In this project, we want to tackle the discovery and analysis of microarchitectural side channels on modern ARM devices, mainly new smartphone and computer chips. In particular, we would like to study the threat surface of microarchitectural and transient execution attacks on increasingly complex ARM systems.

7.1.2 ANSSI:

Participants: Zakaria Belkadi, Louis Rilling, Frédéric Tronel.

Zakaria Belkadi started his PhD in October 2024 in the context of a collaboration between Inria and the ANSSI. In this project, we explore using types in operating system source code as a means to get assurance on security properties. With the rise of memory-safe languages for system programming like Rust, type-based techniques in operating system sources have recently started being investigated to assure functional correctness. With security properties, considering the whole program at once instead of individual functions or modules is an additional challenge.

7.1.3 DGA:

Participants: Jack Royer, Frédéric Tronel.

Jack Royer was a research engineer from November 2024 to May 2025. This position was funded by a DGA grant. In this project, we are interested in deobfuscating binaries protected by anti-debugging technical measures. This project led to the publication of a paper [14] about the vizualisation of control flow graphs.

7.1.4 DGA:

Participants: Quentin Ducasse, Guillaume Hiet.

Quentin Ducasse started his PostDoc position in October 2024 thanks to a grant from DGA. This project aims to propose an intrusion detection approach for hybrid applications by distributing the detection process across different monitors dedicated to each execution unit while minimizing the performance impact on the monitored system and ensuring the protection of the monitors. We propose to use heterogeneous applications developed with High-Level Synthesis (HLS) for platforms utilizing FPGA SoCs (e.g., Xilinx UltraScale+). The implementation involves integrating the monitor generation process into Xilinx's HLS tool based on the LLVM compiler. This approach allows for the insertion of compiler passes capable of analyzing and modifying the intermediate code generated by the compiler [23].

7.1.5 HP Labs:

Participants: Simon Legros, Guillaume Hiet.

Simon Legros started a PhD position in october 2025 in the context of a CIFRE contract between HP and CentraleSupélec. This PhD project investigates the design of an innovative cyber-resilience architecture aimed at ensuring the continuity of critical services in the presence of cyberattacks. Leveraging virtualization technologies, the work focuses on hypervisors as privileged components for security, capable of detecting, analyzing, and responding to attacks even when guest operating systems are compromised. The research explores the use of virtual machine introspection (VMI) and fine-grained control mechanisms to enable automated, contextual detection, response, and recovery.

The objective goes beyond intrusion detection by developing targeted response strategies that can temporarily restrict or repair affected services while preserving the integrity of critical data. The project follows a co-design approach involving application services, standard operating systems (Linux/Windows), and the hypervisor, without requiring a complete redesign of existing systems.

7.1.6 Thalium:

Participants: Jack Royer, Frédéric Tronel, Yaelle Vincont.

Jack Royer started a PhD position in June 2025 in the context of a CIFRE contract between Thales and CentraleSupélec. The goal of this PhD is to measure up against today"s leading commercial obfuscators by delivering deobfuscation methods that are not only effective but also industrial-ready: they must be easy for analysts to adopt, robust across diverse target environments, and scalable to large code bases.

To achieve this, we will architect a modular, multi-architecture toolchain. Existing frameworks such as [33] and other generic dynamic binary analysis suites suffer from performance bottlenecks and limited extensibility, steering us toward a more specialized platform. While [31] follows a similar design philosophy, its closed-source nature impedes community-driven improvements.

8 Partnerships and cooperations

8.1 International research visitors

8.1.1 Visits of international scientists

Other international visits to the team

Pierre Olivier

Status Senior lecturer

Institution of origin: University of Manchester

Country: United Kingdom

Dates: June 2-July 4

Context of the visit: Visiting Professor (CentraleSupélec grant).

Mobility program/type of mobility: research stay.

During this visit, collaboration opportunities about Pierre Olivier’s works on software compartmentalization and operating system design for security in SUSHI were explored. An investigation on input validation in Linux ioctl system calls was started and, based on the internship results of Killian Callac, led to the proposal of a PhD topic for which funding has been requested.

Research stays abroad

- Ruben Salvador got a "**Severo Ochoa**" grant (EUR 1 500) from Barcelona Supercomputing Center in 2025 for a 1-month visit to be completed in 2026.

8.2 National initiatives

8.2.1 PEPR Cybersécurité: projet SECUREVAL (2022-2028)

Participants: Frédéric Tronel, Guillaume Hiet, Pierre Wilke, Erwan Fasquel.

The security assessment of digital systems relies on compliance and vulnerability analyses to provide recognized cybersecurity assurances. The SECUREVAL project of PEPR Cybersecurity aims to design new tools around new digital technologies to verify the absence of hardware and software vulnerabilities and achieve the required compliance proofs. These developments are based on a double approach, first theoretical and founded on the French school of symbolic reasoning, then applied and anchored in the practice of tool development and security assessment techniques. In addition, by exploring new methods for security assessments, this project will also allow France to remain at the top of the world in assessment capabilities by anticipating the evolution of international certification schemes. Within this project's framework, our contribution concerns tasks 4.4 Formal analysis and models at the software-hardware boundary (led by Guillaume Hiet) and 3.2 Vulnerability analysis tools in binary codes (led by Frédéric Tronel). So far, we have started two PhD positions for the task 3.2 about vulnerability analysis (Erwan Fasquel and Samuel Hiron).

8.2.2 ANR Project: TrustGW (2021-2026)

Participants: Guillaume Hiet, Frédéric Tronel, Pierre Wilke, Lionnel Hemmerlé.

In the ANR TrustGW project, we consider a system composed of IoT objects connected to a gateway. This gateway is, in turn, connected to one or more cloud servers. The architecture of the gateway, which is at the heart of the project, is heterogeneous (software/hardware), composed of a baseband processor, an application processor, and hardware accelerators implemented on an FPGA. A hypervisor allows sharing of these resources and allocating them to different virtual machines. TrustGW is a collaborative project between the ARCAD team from Lab-STICC, the ASIC team from IETR, and the SUSHI team from IRISA. The project addresses three main challenges: (1) to define a heterogeneous, dynamically configurable and trusted gateway architecture, (2) to propose a trusted hypervisor allowing the deployment of virtual machines on a heterogeneous software-hardware architecture with virtualization of the whole resources and (3) to secure the applications running on the gateway. Within this project's framework, the SUSHI team's contribution focuses mainly on the last challenge, particularly through the PhD of Lionel Hemmerlé (2022-2025). Guillaume Hiet is the director of this PhD, co-supervised by Frédéric Tronel, Pierre Wilke and Jean-Christophe Prévotet. We will also explore hardware-assisted Dynamic Information Flow Tracking approaches for hybrid applications, which offload part of their computation to an FPGA.

8.2.3 ANR Project: ATTILA (2022-2025)

Participants: Ruben Salvador.

ATTILA tackles the interplay between security and Approximate Computing (AxC) in the context of DNN accelerator security. In particular, it studies the threats posed to such accelerators when built using AxC techniques. We build on the hypothesis of hidden side-channel vulnerabilities that might be due to AxC and on the possibility of leveraging AxC itself to create countermeasures. Specifically, the objectives are:

1. to study power/EM side-channel vulnerabilities of approximate DNN accelerators and the impact of AxC on leakage behaviour and SCA resistance;
2. to build more secure implementations leveraging on DSE and Pareto fronts to facilitate trading-off SCA resistance with inference quality for different approximations;
3. to evaluate AxC and intelligent run-time managers as countermeasures that enable self-adaptation through the Pareto front and beyond to render SCA attacks more difficult;
4. to extend current SCA practices for DNN implementations towards more powerful ML-based techniques.

Ruben Salvador is the PI of ATTILA, which runs in collaboration with the ASIC team from IETR. The project employs 1 PhD student directed by Jean-Christophe Prévotet (INSA Rennes/IETR) and co-supervised by Ruben Salvador and Maria Mendez Real (UBS/Lab-STICC).

8.2.4 CMA project : Train-Cyber-Expert (TCE) (2022-2026)

Participants: Yohann Rio, Frédéric Tronel.

As part of the France 2030 recovery plan, the SUSHI team participates in the Train-Cyber-Expert (TCE) project, funded by the CMA (Competences and Jobs of the Future) call for projects. TCE is a collaborative project involving several academic partners to develop educational resources in the form of digital content and technological platforms, organized into skill blocks, focusing on modularity, reusability, and competency-based pedagogy leading to certifications. We are involved in this project together with the Inria PIRAT team. Our goal is to propose pedagogical resources in the field of system security. We are currently creating a set of lectures about memory attacks and defences based on existing lectures at CentraleSupélec.

8.3 Regional initiatives

Alessandro Palumbo got an **Allocation d'Installation Scientifique (AIS)** from Rennes Métropole, a regional competitive funding (EUR 10 000) awarded to support newly recruited researchers in the establishment and structuring of their research activities.

9 Dissemination

Participants: Guillaume Hiet, Ruben Salvador, Frédéric Tronel, Louis Rilling, Yaëlle Vinçont, Alessandro Palumbo, Lionel Hemmerle, Thomas Rokicki, Lorenzo Casalino, Zakaria Belkadi.

9.1 Promoting scientific activities

9.1.1 Scientific events: organisation

General chair, scientific chair

- Guillaume Hiet was the co-chair of the HS3 workshop at ESORICS 2025
- Louis Rilling is co-organizer of the CREACH LABS SoSySec seminar (Software and Systems Security).

Member of the organizing committees

- Ruben Salvador was workshop chair at ACM Computing Frontiers 2025 and publicity chair at SAMOS 2025

9.1.2 Scientific events: selection

Member of the conference program committees

- Guillaume Hiet was part of the program committees of the following conferences and workshops: EAI SecureComm 2025, NSS 2025, VERDI@DSN 2025, IEEE AIoT 2025, SIF annual congress and HS3.
- Rubén Salvador was part of the program committees of the following conferences: IEEE ISVLSI, IEEE ISCAS, SAMOS, IEEE LASCAS, DASIP. He also served in the program committee of the workshops PARMA-DITAM@HiPEAC, MAL-IoT@ACM Computing Frontiers.
- Yaëlle Vinçont was part of the program committee for SSTIC 2025.
- Rubén Salvador is part of the scientific committee of the CREACH LABS SemSecuElec seminar (security of embedded electronic systems).
- Guillaume Hiet and Louis Rilling are part of the scientific committee of the SoSySec seminar.
- Lorenzo Casalino was part of the program committee of the conference IEEE LATS and the workshop PARMA-DITAM.
- Pierre Wilke was part of the PC for the HS3 workshop.

Reviewer

- Rubén Salvador served as reviewer for the following conferences: IEEE ISCAS, IEEE ISVLSI, IEEE ETS, IEEE LASCAS, SAMOS, DASIP. He served as reviewer in the following workshops: MAL-IoT@Computing Frontiers, PARMA-DITAM@HiPEAC.
- Alessandro Palumbo served as reviewer for the following conferences: IEEE ISCAS and IEEE LASCAS
- Lorenzo Casalino served as reviewer for the following conferences: DASIP, IEEE ISVLSI, SAMOS.
- Pierre Wilke served as a reviewer for the HS3 workshop
- Guillaume Hiet served as a reviewer for the following conferences and workshop: DATE, HS3, VERDI, NSS, SecureComm, and AIoT.
- Quentin Ducasse served as a reviewer for the following conferences: IEEE ISCAS and IEEE ISVLSI.

9.1.3 Journal

Member of the editorial boards

- Rubén Salvador is Associate Editor for the journal IEEE Embedded Systems Letters (from 2022).

Reviewer - reviewing activities

- Guillaume Hiet served as reviewer for IEEE Transactions on Software Engineering
- Alessandro Palumbo served as reviewer for the following journals: Journal of Systems Architecture (JSA), IEEE Transactions on VLSI, IEEE Sensors, IEEE Embedded Systems Letters and the Journal of Supercomputing
- Rubén Salvador served as reviewer for the following journals: IEEE Embedded Systems Letters, IEEE TVLSI, IEEE TCAD, IEEE TDMR, Elsevier Computers & Security, Elsevier JSA, Springer Genetic Programming and Evolvable Machines, and IACR TCHES (as subreviewer)
- Lorenzo Casalino served as review for the journals IEEE ACCESS and IEEE TODAES. Pierre Wilke served as a reviewer for STVR (Software Testing, Verification and Reliability).

Member	Licence-level	Master-level	CS	Univ. of Rennes	ENS	Amout (h eqTD)
Guillaume Hiet	✓	✓	✓			237
Alessandro Palumbo	✓	✓	✓	✓		203.125
Thomas Rokicki	✓	✓	✓			135
Ruben Salvador	✓	✓	✓	✓		233
Frédéric Tronel	✓	✓	✓			215
Yaëlle Vinçont		✓			✓	144
Pierre Wilke	✓	✓	✓			40
Lorenzo Casalino	✓	✓	✓	✓		57.5

Table 1: Summary of teaching effort (eqTD)

9.1.4 Invited talks

- Alessandro Palumbo gave a seminar talk titled “Hardware Trojan Horses and Microarchitectural Side-Channel Attacks: Detection and Mitigation via Hardware-based Methodologies” as part of the SemSecuElec seminar series.

- Frédéric Tronel, Lionel Hemmerlé, and Louis Rilling delivered a lecture titled Defending Against the Undetectable: Advanced Threat Detection with eBPF and Hypervisors at the third EUR CyberSchool winter school in February 2025.

9.1.5 Leadership within the scientific community

- Guillaume Hiet is the co-chair of the Systems, Software and Network Security working group of the GDR Sécurité Informatique.

9.1.6 Scientific expertise

Guillaume Hiet served in the evaluation comitee of the *Conseil d’évaluation externe du métier CYBER* of the DGA

9.1.7 Research administration

- Guillaume Hiet was a member of the recruitment committees for Assistant Professor positions at CentraleSupélec and at Sorbonne Université.

9.2 Teaching - Supervision - Juries - Educational and pedagogical outreach

Several team members are involved in initial and continuing education at CentraleSupélec, a French institute of research and higher education in engineering and science, ENS of Rennes and University of Rennes.

In these institutions, Guillaume Hiet is responsible for the new engineering curriculum in Cybersecurity at CentraleSupélec. Yaelle Vincont is responsible for the "préparation agrégation" (intermediate year leading to the agrégation competitive exam) at ENS as of September 2025. Pierre Wilke was in parental leave from September 2024 to March 2025 included.

The teaching duties are summed up in table 1.

9.2.1 Supervision

PhD students of the team:

- Lionel Hemmerlé (in progress), *Conception et implémentation d’un langage dédié à l’introspection de machine virtuelle*, funded by ANR TrustGW, started November 2022, supervised by Guillaume Hiet (25%, director), Pierre Wilke (25%), Frédéric Tronel (25%), and Jean-Christophe Prévotet (25%)

- Zakaria Belkadi (in progress), *Type-based security properties assurance in operating systems*, funded by ANSSI, started December 2024, supervised by Louis Rilling (25%), Frédéric Tronel (25%), Guillaume Hiet (25%, director), and Florence Schadle (25%).
- Erwan Fasquel (in progress), *Fuzzing for automatic vulnerabilities discovery in closed-source operating systems*, funded by PEPR SECUREVAL, started in May 2024, supervised by Frédéric Tronel (50%, director) and Yaelle Vincont (50%).
- Eliott Quere (in progress), *Towards Secure FPGA-Accelerated Clouds: Identification, Exploitation and Detection of Remote Side-Channel Leakage Sources*, started October 2024, funded by Univ. of Rennes, supervised by Ruben Salvador (75% director, 25% supervisor), Lilian Bossuet (25% director, 25% supervisor), Alessandro Palumbo (25%), and Maria Méndez-Real (25%).
- Théo Goureau (in progress), *Arm Microarchitectural Attacks iscovery and Analysis*, started October 2025, funded by ANSSI and Inria, supervised by Daniel De Almeida Braga (25%), Pierre-Alain Fouque (25%, director), Guillaume Hiet (25%, director), and Thomas Rokicki (25%).
- Leslie Fifanon (in progress), *Design and Formal Verification of Hardware/Software Security Mechanisms*, funded by CEA "Programme Audace !", TwinSec project, started November 2025, supervised by Guillaume Hiet (25%, director), Mathieu Jan (25%, director), Damien Couroussé (25%) and Pierre Wilke (25%)
- Jack Royer (in progress), *Development of a modular, multi-architecture binary deobfuscation tool for industrial-scale use.* funded by a CIFRE contract between Thales and CentraleSupélec, started in June 2025, supervised by Guillaume Hiet (34%, director), Frédéric Tronel (33%) and Yaelle Vincont (33%).
- Samuel Hiron (in progress), *Automatic discovery of vulnerabilities in binary executables through the use of hybrid approaches mixing static and dynamic analysis*, funded by PEPR SECUREVAL, started in December 2025, supervised by Ludovic Mé (34%, director), Frédéric Tronel (33%) and Yaelle Vincont (33%).
- Simon Legros (in progress), *Dynamically ensuring the cyber resilience of OS applications and services through a hypervisor-based architecture*, funded by a CIFRE contract between HP and CentraleSupélec, started in June 2025, supervised by Guillaume Hiet (50%, director) and Boris Balacheff from HP (50%).
- Cyprien Jules (in progress), *Formal verification of security mechanisms in a RISC-V processor*, funded by Défi Inria CocoRISCo, supervised by Pierre Wilke (50%, director), Guillaume Hiet (25%), Simon Rokicki (25%)

Supervision of PhD students in other teams:

- Seungah Lee (defended in January 2025), *Efficient designs of On-Board heterogeneous Embedded Systems for Space Applications*, funded by CNES, started October 2021, supervised by Rubén Salvador (35%), Angeliki Kritikakou (35%), and Emmanuel Casseau (30%, director).
- Léo Cosseron (defended in December 2025), *Time-Accurate Network Simulation Interconnecting VMs with Hardware Virtualization Towards Stealth Analysis*, funded by DGA via CREACH LABS, started October 2022, supervised by Louis Rilling (50%), Martin Quinson (25%, director), and Matthieu Simonin (25%).
- Guillaume Lomet (in progress), *Guess What I'm Learning: Side-Channel Analysis of Edge AI Training Accelerators*, started October 2022, supervised by Rubén Salvador (35%), Olivier Sentieys (30%, codirector), and Cédric Killian (35%, co-director)
- Valentin Abgrall (in progress), Vulnerability analysis, fault modeling, and countermeasures towards dependable real-time computing in safety-critical drone systems, started December 2025, funded by Univ. of Rennes, supervised by Angeliki Kritikakou (25% director), Marcello Traiola (25%), Rubén Salvador (25%), Alessandro Palumbo (25%)

9.2.2 Educational and pedagogical outreach

- Thomas Rokicki gave a talk on Sushi's research topics to M1 level students at the Research Track of ESILV.
- Yaelle Vincont gave a talk on "Teaching (and convincing) students to test their code" at the Algorithms and Programming conference for CPGE teachers, held at the CIRM in May 2025.

9.3 Popularization

9.3.1 Productions (articles, videos, podcasts, serious games, ...)

Alessandro Palumbo got a AAP Oser - Université Paris-Saclay Grant (EUR 5 000) awarded for the development of innovative teaching material, focusing on the use of comic-style illustrations to explain microprocessor security attacks. The project was carried out in collaboration with the artist Lorenzo Colangeli.

9.3.2 Participation in Live events

- Lorenzo Casalino and Ruben Salvador Perea participated at the workshop *WISG 2025* (Paris, France), presenting a poster [26] on the ANR JCJC project *ATTILA*.
- Pierre Wilke participated at the workshop PriSC 2025 (Rennes, France) colocated with the POPL conference, giving a talk about the team's work on Kôika.
- Guillaume Hiet delivered invited presentations about the research activities of the SUSHI team at the annual congres of the SIF, the *café science* at CentraleSupélec, and the Sci-Rennes seminar at IRISA.
- Guillaume Hiet took part in a roundtable discussion during the Cross-Border Cybersecurity Tour in Metz and Saarbrücken on cybersecurity training. He also chaired a roundtable on hardware roots of trust at the RISC-V Exploration of Industrial Synergies event.

9.3.3 Others science outreach relevant activities

- Lorenzo Casalino presented part of his research works during the *GT SSLR 2025*.
- Zakaria Belkadi presented his research work during the *GT SSLR 2025*.
- Jack Royer presented his research work during the *GT SSLR 2025*.
- Pierre Wilke presented Lionel Hemmerle's work at *GDR RSD* and the work around Kôika at the *GT LVP* of the *GDR GPL*

10 Scientific production

10.1 Publications of the year

International journals

[1] Q. Ducasse, P. Cotret and L. Lagadec. 'War on JITs: Software-Based Attacks and Hybrid Defenses for JIT Compilers - A Comprehensive Survey'. In: *ACM Computing Surveys* 57.9 (2nd Apr. 2025), pp. 1–36. DOI: 10.1145/3731598. URL: https://hal.science/hal-05041484 (cit. on p. 15).

[2] G. Garzo and A. Palumbo. 'Come l'Intelligenza Artificiale può Supportare le Decisioni Giuridiche: Un Caso di Studio sugli Omicidi Stradali'. In: *Journal of Inclusive Methodology and Technology in Learning and Teaching* (5th May 2025), pp. 1–8. URL: https://hal.science/hal-05056369. In press (cit. on p. 14).

International peer-reviewed conferences

[3] L. Cosseron, L. Rilling and M. Quinson. 'Mitigation of the impact of Virtual Machine Introspection Pauses on Multi-core Virtual Machines'. In: *Lecture Notes in Computer Science (LNCS)*. HS3 2025 - 1st Workshop on Hardware-Supported Software Security. Toulouse, France, 2025, pp. 1–20. URL: https://inria.hal.science/hal-05253492 (cit. on p. 13).

[4] G. Garzo and A. Palumbo. 'Human-in-the-Loop: Legal Knowledge Formalization in Attempto Controlled English'. In: ISDFS - 13th International Symposium on Digital Forensics and Security. Boston, United States, 2025, pp. 1–6. DOI: 10.1109/isdfs65363.2025.11011971. URL: https://hal.science/hal-05021540 (cit. on p. 14).

[5] G. Garzo and A. Palumbo. 'Legal & Ethical Implications of Predictive Digital Techniques in the Judicial Criminal Proceedings'. In: ISDFS 2025 - 13th International Symposium on Digital Forensics and Security. Boston, United States, 2025, pp. 1–6. DOI: 10.1109/isdfs65363.2025.11012003. URL: https://hal.science/hal-05021547 (cit. on p. 15).

[6] L. Hemmerlé, G. Hiet, F. Tronel, P. Wilke and J.-C. Prévotet. 'HYPERSEC: An Extensible Hypervisor-Assisted Framework for Kernel Rootkit Detection'. In: *Information Security - 28th International Conference, ISC 2025, Seoul, South Korea, October 20-22, 2025, Proceedings.Lecture Notes in Computer Science 16186, Springer 2026, ISBN 978-3-032-08123-0*. ISC 2025 - 28th Information Security Conference. Vol. 16186. Lecture Notes in Computer Science. Seoul, South Korea: Springer, 22nd Oct. 2025, pp. 431–451. DOI: 10.1007/978-3-032-08124-7_25. URL: https://inria.hal.science/hal-05469459 (cit. on p. 13).

[7] G. Lomet, R. Salvador, B. Colombier, V. Grosso, O. Sentieys and C. Killian. 'Side-Channel Extraction of Dataflow AI Accelerator Hardware Parameters'. In: *2025 IEEE 30th International Symposium on On-Line Testing and Robust System Design (IOLTS)*. 2025 IEEE 31st International Symposium on On-Line Testing and Robust System Design (IOLTS). Ischia, Italy, 2025, pp. 1–7. DOI: 10.1109/IOLTS65288.2025.11117043. URL: https://inria.hal.science/hal-05120223 (cit. on p. 11).

[8] R. Martínez, A. Palumbo, P. Reviriego, R. Salvador and D. Larrabeiti. 'LAD-IXoC: Loop-based Attack Detection with Integrated Xor Filter and CMS'. In: DFT 2025 38th IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems. Barcelone, Spain: IEEE, 2025, pp. 1–5. URL: https://hal.science/hal-05195618 (cit. on p. 10).

[9] Y. Oziel, T. Laor, S. Levy, C. Maurice, Y. Oren, T. Rokicki and G. Scalosub. 'WIP: A Second Look at Port Assignment on Intel CPUs'. In: uASC 2026 - 2nd Microarchitecture Security Conference. Leuven, Belgium, 2026. URL: https://hal.science/hal-05337434 (cit. on p. 9).

[10] A. Palumbo. 'Machine Learning-Based Detection of Microarchitectural Attacks on RISC-V via Gem5'. In: TechDefense 2025 - IEEE International Workshop on Technologies for Defense and Security. Rome, Italy: IEEE, 2025, pp. 1–6. URL: https://hal.science/hal-05288651 (cit. on p. 10).

[11] A. Palumbo and R. Salvador. 'Leveraging gem5 for Hardware Trojan Research: Simulation for Machine-Learning-Based Detection'. In: 10th International Workshop on Malicious Software and Hardware in the Internet of Things (MAL-IoT) 2025 - Proceedings of the 22nd ACM International Conference on Computing Frontiers: Workshops and Special Sessions. Cagliari, Italy, 2025, pp. 9–16. DOI: 10.1145/3706594.3728869. URL: https://hal.science/hal-05044677 (cit. on p. 10).

[12] A. Plin, L. Casalino, T. Rokicki and R. Salvador. 'Knock-Knock: Black-Box, Platform-Agnostic DRAM Address-Mapping Reverse Engineering'. In: uASC 2026 - 2nd Microarchitecture Security Conference. Leuven, Belgium, 2026. URL: https://hal.science/hal-05273255 (cit. on p. 8).

[13] R. Randriatsiferana, R. Salvador and O. Tamarin. 'Digital NGD Design for FPGA-based Application'. In: INSCIT 2025 - 9th International Symposium on Instrumentation Systems, Circuits and Transducers. Manaus, Brazil: IEEE, 2025, pp. 1–5. DOI: 10.1109/INSCIT66472.2025.11181019. URL: https://ird.hal.science/ird-05448291 (cit. on p. 15).

[14] J. Royer, F. Tronel and Y. Vinçont. 'Towards Better CFG Layouts'. In: BAR 2025 - Workshop on Binary Analysis Research. San Diego (CA), United States, 2025, pp. 1–17. DOI: 10.14722/bar.2025.23011. URL: https://hal.science/hal-04996939 (cit. on pp. 14, 16).

Conferences without proceedings

[15] G. Didier, A. Lucas and T. Rokicki. 'Cache Attacks in Modern/Multi-Socket x86 Systems (Work in Progress)'. In: HS3 2025 - 1st Workshop on Hardware-Supported Software Security. Toulouse, France, 26th Sept. 2025, pp. 1–10. URL: https://hal.science/hal-05249476 (cit. on p. 8).

[16] G. Garzo and A. Palumbo. 'Does ChatGPT Understand the Law? A Case Study on Road Homicide in Italy'. In: JURIX 2025 - 38th International Conference on Legal Knowledge and Information Systems. Turin, Italy, Sept. 2025, pp. 1–12. URL: https://hal.science/hal-05325502 (cit. on p. 14).

[17] G. Garzo and A. Palumbo. 'The First Hardware Circuit Emulating Italian Road Homicides Legal Logic, DAJE!' In: IS 2025 - 18th IADIS International Conference Information Systems. Madeira island, Portugal, 13th Mar. 2025, pp. 1–8. URL: https://hal.science/hal-04990194 (cit. on p. 14).

[18] A. Palumbo. 'Detecting Microarchitectural Side-Channel Attacks via Hardware Security Checkers'. In: RISC-V Summit Europe 2025. Paris, France, 2025, pp. 1–2. URL: https://hal.science/hal-05287163 (cit. on p. 9).

[19] A. Palumbo. 'Tackling Hardware Trojans via Hardware-based Methodologies'. In: RISC-V Summit Europe 2025. Paris, France, 2025, pp. 1–2. URL: https://hal.science/hal-05287160 (cit. on p. 9).

[20] A. Palumbo and R. Salvador. 'Detecting Hardware Trojans in Microprocessors via Hardware Error Correction Code-based Modules'. In: 2025 IEEE 31st International Symposium on On-Line Testing and Robust System Design (IOLTS). Naples, Italy, 2025, pp. 1–7. DOI: 10.1109/IOLTS65288.2025.11116960. URL: https://hal.science/hal-05055204 (cit. on p. 10).

[21] E. Quéré, M. Mendez Real, A. Palumbo, L. Bossuet and R. Salvador. 'Side-Channel Exploitation of DRAM Access Patterns for Fingerprinting FPGA-CPU Environments'. In: RESSI 2025 - Rendez-Vous de la Recherche et de l'Enseignement de la Sécurité des Systèmes d'Information. Lannion, France, 2025. URL: https://hal.science/hal-05266486 (cit. on p. 11).

Reports & preprints

[22] G. Didier, T. Rokicki and A. Lucas. *Flush-based Cache Attacks on Modern / Multi-Socket x86 Systems*. 19th Dec. 2025. URL: https://hal.science/hal-05424273 (cit. on p. 8).

[23] Q. Ducasse, G. Hiet, V. Stolz and P. Wilke. *WIP: InSight -A CoreSight Trace Interpreter for Dynamic Information Flow Tracking*. CentraleSupélec, 25th Sept. 2025. URL: https://hal.science/hal-05472331 (cit. on p. 16).

[24] J. Guillaume, M. Pelcat, A. Nafkha and R. Salvador. *Multi-Screaming-Channel Attacks: Frequency Diversity for Enhanced Attacks*. 2025. DOI: 10.48550/arXiv.2504.02979. URL: https://inria.hal.science/hal-05024950 (cit. on p. 12).

[25] C. Jules, P. Wilke and G. Hiet. *Modular and automatic formal verification of a RISC-V processor with security mechanisms*. 11th Jan. 2026. URL: https://hal.science/hal-05472602 (cit. on p. 13).

Other scientific publications

[26] L. Casalino, M. Mendez Real, J.-C. Prévotet and R. Salvador. 'ATTILA – Addressing security threats to artificial intelligence in approximate computing systems'. In: WISG 2025 - Workshop interdisciplinaire sur la sécurité globale. Paris, France, 2025, pp. 1–1. URL: https://hal.science/hal-05262917 (cit. on p. 23).

[27] E. Fasquel, F. Tronel and Y. Vinçont. 'Fuzzing macOS Kernel Extensions'. In: RESSI 2025 - Rendez-vous de la Recherche et de l'Enseignement de la Sécurité des Systèmes d'Information. Lanniron, France, 2025. URL: https://inria.hal.science/hal-05173342 (cit. on p. 14).

[28] C. Jules, P. Wilke, G. Hiet and G. Desfrene. *heRVé: towards a formally verified RISC-V processor with security mechanisms*. Toulouse, France, 2025. URL: https://inria.hal.science/hal-05274107 (cit. on p. 13).

[29] A. Palumbo. 'Detecting Microarchitectural Side-Channel Attacks via Hardware Security Checkers'. In: RISC-V Summit Europe 2025. Paris, France, 2025, pp. 1–1. URL: https://hal.science/hal-05287161 (cit. on p. 9).

[30] A. Palumbo. 'Tackling Hardware Trojan Horses via Hardware-based Methodologies'. In: RISC-V Summit Europe 2025. Paris, France, 2025, pp. 1–1. URL: https://hal.science/hal-05287158 (cit. on p. 9).

10.2 Cited publications

[31] P. Garba and M. Favaro. 'SATURN - Software Deobfuscation Framework Based On LLVM'. In: *Proceedings of the 3rd ACM Workshop on Software Protection*. SPRO'19. New York, NY, USA: Association for Computing Machinery, Nov. 2019, pp. 27–38. DOI: 10.1145/3338503.3357721. URL: https://doi.org/10.1145/3338503.3357721 (visited on 03/07/2024) (cit. on p. 17).

[32] G. Garzo, S. Ribes and A. Palumbo. 'Opening the Black Box: How Boolean AI can Support Legal Analysis'. In: *CCAI 2024 - 4th International Conference on Computer Communication and Artificial Intelligence*. Xi'an, China, May 2024, pp. 269–272. DOI: 10.1109/ccai61966.2024.10603017. URL: https://hal.science/hal-04685601 (cit. on p. 14).

[33] J. Salwan, S. Bardin and M.-L. Potet. 'Symbolic Deobfuscation: From Virtualized Code Back to the Original'. en. In: *Detection of Intrusions and Malware, and Vulnerability Assessment*. Ed. by C. Giuffrida, S. Bardin and G. Blanc. Cham: Springer International Publishing, 2018, pp. 372–392. DOI: 10.1007/978-3-319-93411-2_17 (cit. on p. 17).