

RESEARCH CENTRE

**Inria Centre at the University of
Lille**

IN PARTNERSHIP WITH:

Université de Lille, CNRS

2024

ACTIVITY REPORT

Project-Team

SPIRALS

**Self-adaptation for distributed services
and large software systems**

IN COLLABORATION WITH: Centre de Recherche en Informatique, Signal
et Automatique de Lille

DOMAIN

**Networks, Systems and Services,
Distributed Computing**

THEME

Distributed Systems and middleware

Inria

Contents

Project-Team SPIRALS	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	4
2.1 Introduction	4
2.2 Scientific Foundations	4
3 Research program	4
3.1 Self-optimization	4
3.2 Self-protection	5
3.3 Foundational elements for self-adaptation	6
4 Application domains	7
5 Social and environmental responsibility	7
5.1 Impact of research results	7
6 Highlights of the year	7
6.1 Awards	7
7 New software, platforms, open data	8
7.1 New software	8
7.1.1 amiunique	8
7.1.2 APISENSE	8
7.1.3 cloudnet	9
7.1.4 PowerAPI	9
8 New results	10
8.1 Investigating Novel Tracking Practices	10
8.2 Safe Dynamic Reconfiguration of Applications with Features	10
8.3 Enhancing IaaS Consolidation with Resource Oversubscription	11
8.4 Exploring the Impact of the User's Browsing Environment on Web Privacy	11
8.5 Sustainable Software Design: Estimation and Reduction of the Environmental Footprint of Software Systems	12
8.6 Privacy Driven Data Collection	13
8.7 Electromagnetic Emanations Side Channels in Embedded Devices	13
9 Bilateral contracts and grants with industry	14
10 Partnerships and cooperations	15
10.1 International research visitors	15
10.1.1 Visits of international scientists	15
10.2 European initiatives	15
10.2.1 Other european programs/initiatives	15
10.3 National initiatives	16
10.3.1 ANR	16
10.3.2 PEPR and PTCC	18
10.4 Public policy support	19
11 Dissemination	20
11.1 Promoting scientific activities	20
11.1.1 Scientific events: organisation	20
11.1.2 Scientific events: selection	20
11.1.3 Journal	20
11.1.4 Invited talks	20

11.1.5 Leadership within the scientific community	21
11.1.6 Scientific expertise	22
11.1.7 Research administration	23
11.2 Teaching - Supervision - Juries	24
11.2.1 Teaching	24
11.2.2 Supervision	25
11.2.3 Juries	26
11.3 Popularization	27
11.3.1 Productions (articles, videos, podcasts, serious games, ...)	27
11.3.2 Participation in Live events	27
12 Scientific production	27
12.1 Major publications	27
12.2 Publications of the year	28
12.3 Cited publications	32

Project-Team SPIRALS

Creation of the Project-Team: 2015 January 01

Keywords

Computer sciences and digital sciences

- A1.3.3. – Blockchain
- A1.3.5. – Cloud
- A1.4. – Ubiquitous Systems
- A1.6. – Green Computing
- A2.3.1. – Embedded systems
- A2.3.2. – Cyber-physical systems
- A2.4.2. – Model-checking
- A2.5. – Software engineering
- A2.5.2. – Component-based Design
- A2.5.3. – Empirical Software Engineering
- A2.5.4. – Software Maintenance & Evolution
- A2.6.2. – Middleware
- A3.1.3. – Distributed data
- A3.1.4. – Uncertain data
- A3.1.5. – Control access, privacy
- A3.1.9. – Database
- A3.2.1. – Knowledge bases
- A3.2.4. – Semantic Web
- A3.2.5. – Ontologies
- A4.4. – Security of equipment and software
- A4.8. – Privacy-enhancing technologies
- A7.2. – Logic in Computer Science
- A9.1. – Knowledge

Other research topics and application domains

- B6.1. – Software industry
- B6.4. – Internet of things
- B6.5. – Information systems
- B6.6. – Embedded systems
- B8.5.2. – Crowd sourcing
- B9.5.1. – Computer science
- B9.5.6. – Data science
- B9.10. – Privacy

1 Team members, visitors, external collaborators

Research Scientists

- Simon Bliudze [INRIA, Researcher]
- Pierre Bourhis [CNRS, Researcher]
- Sophie Cerf [INRIA, ISFP]
- Pierre Laperdrix [CNRS, Researcher]
- Clémentine Maurice [CNRS, Researcher]
- Philippe Merle [INRIA, Senior Researcher]

Faculty Members

- Lionel Seinturier [Team leader, UNIV LILLE, Professor]
- Laurence Duchien [UNIV LILLE, Professor]
- Adrien Luxey-Bitri [UNIV LILLE, Associate Professor]
- Clément Quinton [UNIV LILLE, Associate Professor]
- Romain Rouvoy [UNIV LILLE, Professor]

Post-Doctoral Fellow

- Imane Fouad [INRIA, Post-Doctoral Fellow]

PhD Students

- Virginie Amand [INRIA, from Oct 2024]
- Wassim Aroui [ORANGE, CIFRE, until Mar 2024]
- Alexandre Bonvoisin [INRIA]
- Sihem Bouhenniche [UNIV LILLE]
- Tristan Coignon [INRIA]
- Thomas Collignon [QARNOT COMPUTING, CIFRE]
- Boubacar Diarra [ORANGE, CIFRE]
- Belkis Djefal [INRIA]
- Salman Farhat [UNIV LILLE, ATER, until Aug 2024]
- Nusrat Jahan Farin [INRIA, from Nov 2024]
- Iliana Fayolle [UNIV LILLE]
- Antoine Geimer [UNIV RENNES I]
- Maxime Huyghe [UNIV LILLE]
- Jean Luc Intumwayase [UNIV LILLE, ATER]
- Pierre Jacquet [INRIA, until Sep 2024]

- Naif Mehanna [UNIV LILLE, until Apr 2024]
- Hugo Monfleur [INRIA]
- Remy Raes [INRIA]
- Maryam Rahmani [UNIV LILLE]
- Brell Peclard Sanwouo Chekam [INRIA]
- Thibault Simon [ORANGE, CIFRE, until Oct 2024]
- Yifan Wang [ORANGE, CIFRE]
- Jiali Xu [INRIA]
- Nada Zine [INRIA, from Nov 2024]

Technical Staff

- Mohammed Chakib Belgaid [INRIA, Engineer, until Aug 2024]
- Noe Chachignot [INRIA, Engineer, from Sep 2024]
- Guillaume Fieni [INRIA, Engineer]
- Gaelle Fret [INRIA, Engineer, from Sep 2024]
- Nathan Leblond [INRIA, Engineer, from Oct 2024]
- Daniel Romero Acero [INRIA, Engineer]
- Brice Arleon Zemtsop Ndadji [INRIA, Engineer, from Oct 2024]

Interns and Apprentices

- Virginie Amand [INRIA, Apprentice, until Aug 2024]
- Noe Chachignot [INRIA, Intern, from Apr 2024 until Aug 2024]
- Lucas Deloison [INRIA, Apprentice, from Sep 2024]
- Lucas Deloison [INRIA, Intern, from Apr 2024 until Jul 2024]
- Samuel Dubuisson [INRIA, Apprentice, from Sep 2024]
- Niloofar Etminani [UNIV LILLE, Intern, from Apr 2024 until Jun 2024]
- Gaelle Fret [INRIA, Intern, from Apr 2024 until Aug 2024]
- Corentin Gauquier [INRIA, Intern, from Apr 2024 until Jun 2024]
- Kellian Leveque [INRIA, Apprentice, from Sep 2024]
- Jeremy Woirhaye [INRIA, Apprentice, from Sep 2024]

Administrative Assistants

- Isabelle Aslani [INRIA]
- Nathalie Bonte [INRIA]
- Karine Lewandowski [INRIA]

Visiting Scientist

- Juliette Sénéchal [UNIV LILLE, Délégation Inria]

2 Overall objectives

2.1 Introduction

Our research is based on two complementary fields: distributed systems and software engineering. We aim at introducing more automation in the adaptation processes of software systems, that is, transitioning from the study of adaptive systems to self-adaptive systems. In particular, we target the two key properties of *self-optimization* and *self-protection*, and we study some *foundational elements for self-adaptation*.

2.2 Scientific Foundations

Distributed software services and systems are central to many human activities, such as communication, commerce, education, defense, etc. Distributed software services consist of an ever growing number of devices, often highly heterogeneous, from cloud platforms, sensor networks, to application servers, desktop machines, and mobile devices, such as smartphones. The future of this huge number of interconnected software services has been called the Internet of Services, a vision “*where everything that is needed to use software applications is available as a service on the Internet, such as the software itself, the tools to develop the software, the platform servers, storage and communication to run the software.*”¹ This pervasiveness continuously leads to new usages that in turn foster the emergence of novel requirements and concepts for new software services. Hence, it is necessary to establish new paradigms to design and execute software programs in these highly interconnected and heterogeneous environments, and it is necessary to ensure not only that these software systems can be adapted to new usages, new infrastructures, and new execution environments in the long term, but also that after the adaptation process the services still perform as expected.

This research project focuses on defining *self-adaptive* software services and middleware. From the perspective of the Internet of Services, this project fits in the vision sketched by e.g. the FP8 Expert Group Services in the Future Internet [80], the NESSI Research Priorities for the next Framework Programme for Research and Technological Development FP8 [84], the Roadmap for Advanced Cloud Technologies under H2020 [81], and research roadmaps, such as [67, 73, 79].

3 Research program

Our research program is organized around three axes: self-optimization, self-protection, and foundational elements for self-adaptation. These three axes are detailed below.

3.1 Self-optimization

This research axis aims to tackle the challenges we can observe with the growing adoption of software services in the wild.

Monitoring software in the wild. Software systems are now widely distributed by design, being natively deployed in the very-large scale across several countries and continents. This infrastructure scale and geographical coverage call for the development of novel software monitoring techniques and algorithms that can follow *key performance indicators* (KPI) and report on critical situations where optimizations would be required. Given this context, collecting and processing such data flows to build a holistic view of the distributed system is a key challenge to deliver timely and targeted adaptations. Beyond the middleware challenge of distributed monitoring in the wild, we can also observe the rise of novel KPI aiming to balance performance and environmental metrics to better control the consumption of limited resources for the purpose of a given business. This evolution demonstrates that software developers and

¹Available from the CORDIS FP7 ICT SSAI web page [version archived on 28/11/2015](#)

operators miss more global indicators of the impact of their software on their environment at large to take more optimal decisions.

Collaborative decision-making approaches. To support the above decisions, we believe that very-large-scale distributed systems require to adopt decentralized and collaborative decision-making strategies to implement from local and quick reactions to more global and long-term planning. With this objective in mind, the combination of multiple decision-making techniques, such as control theory, reinforcement learning, constraint solvers, or rule-based approaches, will offer more flexibility to deal with domain-specific adaptations to be performed. Furthermore, our experience on *software product lines* (SPL) will also bring interesting venues to structure and control the control plane of such distributed systems, hence offering some layer of reflexivity for self-adaptive software systems.

3.2 Self-protection

In this axis, our research activities deal with security and privacy especially web privacy. Even if software is a major source of privacy and security threats, hardware and micro-architectural components can also raise some threats that have major consequences in distributed environments, as exemplified by the well-known Spectre and Meltdown vulnerabilities disclosed in 2018. We will then also work on research questions associated with hardware and micro-architectural components. More specifically, we will work at three levels: at the level of web applications, at the level of web browsers, and at the level of hardware and micro-architectural components.

Web and mobile applications. We work on two topics: improving privacy on the web, and protecting users from software vulnerabilities. On the first topic on privacy, we will work on identifying harmful content related to privacy leakage, and on generating automatic shims to replace this harmful content with benign one. The ultimate goal is to have a browser, or an operating system, that does not send identifying information online, can decide which elements in a page are harmful to block, and can repair/augment webpages to assure their proper functioning. The second topic is on debloating which is related to the fact that modern web and mobile applications rely on an impressive list of dependent libraries. While using external libraries eases the development, this is known to be a major source of problems as a package can present a vulnerability or be downright malicious, affecting the integrity of the whole program. The goal here is to design methods and develop tools to have programs that can still benefit from the wealth of existing dependencies that exist in ecosystems, like NPM or Gradle, without the security problems that come with it.

Browsers and hardware. We investigate hardware and software fingerprinting and their associated defenses. While so far we mostly look at fingerprinting the browser, it remains to be seen how much can be fingerprinted on both the software and hardware side. With the arrivals of new APIs like WebGPU, WebUSB or WebXR that rely much more on the hardware, there is a need to understand the privacy problems that can be caused by these APIs to protect users online. The goal is also to infer hardware characteristics through auxiliary time-related channels and micro-benchmarks instead of relying on attributes sent by the browser.

Hardware and micro-architectural components. We work on analyzing attack surface and on improving the reproducibility of micro-architectural attacks. For that, given the lack of documentation of hardware components by the manufacturers, we will work on the reverse-engineering of micro-architectural components. Recent work by Vila et al. [85] shows that it is possible to use automata learning and program synthesis techniques to reverse engineer cache replacement policies from measurements made with performance counters. This approach could be extended to model hardware prefetchers, and perhaps to refine the models proposed on branch prediction units. Concerning reproducibility, it is not uncommon for code that works on one machine to give other results on another machine – and identifying the root cause can be quite complex. We propose the use of the gem5 simulator to overcome some of these problems, and in particular improving visualization techniques, creating a reference benchmark of attacks available to the community, and study countermeasures.

3.3 Foundational elements for self-adaptation

In this axis, our research activities deal with the definition of formal and rigorous foundations for self-adaptive software systems. As opposed to simple programs that compute a function, software systems are *structured assemblies* of interacting components that *coordinate their behavior* to perform a function based on continuous observation of *data*—both their internal state and the data provided by their environment. We will leverage formal methods, machine learning, database theory, and knowledge representation to consolidate the foundations of self-adaptation. Thus, we plan to work on developing and consolidating the formal foundations of self-adaptation in three complementary directions: structure, behaviour, and data. This will enable a holistic coverage of the different facets of self-adaptive software systems. Of course, these facets are not isolated from each other with connections having to be established among them. This axis is transversal and applies both to self-optimization and self-protection.

Structure. Software systems are commonly assembled from numerous components, each narrowly focused on its provided functionality. Among others, this allows structural variability—often offering a myriad of configuration options—which is crucial for these systems and must be handled from design-to run-time. Two fundamental aspects to be addressed here are the specification, analysis and implementation (1) of the interaction among the components and (2) of the *(re-)configuration* of component assemblies. W.r.t. configuration specifications, we particularly focus on approaches where these features are described in a so-called feature model [70]. We will continue exploring theoretical and implementation aspects of the underlying interaction and reconfiguration mechanisms to provide developers with appropriate modelling abstractions. Since it is not always possible to entirely explore the whole configuration space relating each configuration to the proper non-functional and functional requirements, we will apply machine learning techniques to predict the properties of configurations, identify influential options, ensure non-regression and select the “best” configuration. Furthermore, since the configuration space of the system is likely to evolve over time and due to external factors, the model will have to be learned again anyway to stay consistent with the system and its functional and performance properties. The learning process can therefore be considered itself as a configurable system, which can thus be fine-tuned with respect to what can be learned, how it can be learned and when it can be learned. We will thus investigate solutions for the self-adaptive learning of configuration spaces.

Behaviour. The ability to provide satisfaction guarantees of behavioural properties, such as deadlock-freedom, safety, some aspects of security and privacy, is rapidly becoming crucial for modern software systems due to increasing societal awareness, as attested by the growing use of formal methods by software giants such as Amazon [78] and Facebook [58] and to legislative evolution, e.g., GDPR. This requires precise and formal behavioural models allowing reasoning for proof and analysis of such properties. In particular, the Rigorous System Design (RSD) approach [82] strives for correctness by construction through enforcing multiple levels of separation of concerns. We plan to work on (1) transforming high-level models commonly defined by a “user”, such as feature models mentioned above (in relation with the structural modelling) or requirements into (Java)BIP [54, 56] behavioural models, (2) learning behavioural models of existing software by static (source code) or dynamic (execution traces) analysis, (3) maintaining coherency between models and code in presence of evolutions and adaptations, and (4) formalisation of domain-specific knowledge to generate efficient distributed code while respecting the behavioural semantics of the system and taking into account its structural constraints.

Data. The self-adaptation mechanisms that are studied in the team are always triggered by data: either data gathered online in distributed systems, or data mined offline from repositories of metadata associated with software systems. In this context, being able to reason, query, and manage data and metadata associated with software systems is a central, yet complex and difficult task that raises many challenges. In relation with our expertise in database theory, we want to focus on two main challenges: (1) how data management and knowledge representation techniques can be used for improving the self-adaptation of software systems, and (2) how to take into account the data management concern along the development of self-adaptive software systems.

4 Application domains

Our research activities are general enough to encompass a wide range of application domains in relation with distributed systems and software engineering. We currently focus on sustainability, safety, and cybersecurity, that are three crosscutting properties of software systems.

From an industrial perspective, the application domains of our research activities include cloud and telecom operators, IT services companies and software editors, cybersecurity agencies and companies. Examples of companies we are working with include the OVHcloud, Scalair and Qarnot Computing cloud operators, the Orange telecom company, the Davidson consulting company.

In terms of sustainability, we aim at better understanding how computing and software activities can better take into account the physical limits imposed by the environmental constraints of our world. Instead of measuring a posteriori the impact of digitalization, we aim at taking into account the environmental limits of our world in the very first stage of the design of software systems and services.

In terms of safety, we aim at building better and safer digital services and systems that can be proven safe by construction. We also want to study the tradeoff between safety and self-adaptation that are too often contradictory properties.

In terms of cybersecurity, we aim at improving privacy and preventing vulnerabilities. We are studying tracking techniques with a focus on fingerprinting to better understand how users can be identified on the Internet and how one can protect against such mechanisms. In terms of vulnerabilities we especially apply our research to microarchitectural components and study how one can automate as much as possible countermeasures to these vulnerabilities.

5 Social and environmental responsibility

5.1 Impact of research results

Some parts of our research activities deal with green and power efficient computing. We are especially working on **PowerAPI** that is a middleware toolkit for building software-defined power meters. The results from this research activity have a potential high impact to go towards more sustainable software systems. The impact is important for the IT industry, from software editors to consulting company to telco and cloud operators. We aim at being able to issue recommendations for a greener design and development of software systems and to enable a finer measurement of energy consumption in modern distributed systems that can unlock important power savings.

The defended PhD theses of Pierre Jacquet [45] and Thibault Simon [49] contributes to this objective by proposing, respectively, novel oversubscription techniques to lower the power consumption of data centers, and a methodology to identify levers that can reduce the environmental footprint of digital services.

6 Highlights of the year

6.1 Awards

Naif Mehanna, Pierre Laperdrix, and their co-authors, obtained a best paper award at the MADWeb workshop on measurements, attacks, and defenses for the Web in the context of the NDSS Network and Distributed System Security Symposium for their work on a longitudinal study analyzing the stability, security, and potential manipulation of free web proxies [34].

Juliette Sénéchal and her co-authors obtained the "Prix CNIL-Inria pour la protection de la vie privée" for their study on the targeting mechanisms and legal aspects of marketing to children through online targeted advertising [74].

Rémy Raes, Adrien Luxey-Bitri, Romain Rouvoy, and their co-author, obtained a best paper award at the DAIS International Conference on Distributed Applications and Interoperable Systems for their work on the compact storage of data streams in mobile devices [37].

Martin Monperrus, a past member of the Spirals project-team, and his co-authors, obtained a 10-year most influential paper the ASE IEEE/ACM International Conference on Automated Software Engineering,

one of the two best conferences in software engineering, for their work on fine-grained and accurate source code differencing that was published in 2014 [59].

7 New software, platforms, open data

7.1 New software

7.1.1 amiunique

Name: amiunique

Keywords: Privacy, Browser fingerprinting

Scientific Description: The amiunique web site has been deployed in 2014 in the context of the DiverSE team research activities on browser fingerprinting to understand how software diversity can be leveraged to mitigate the impact of fingerprinting on the privacy of users. In 2018, it was migrated to the Spirals team where the research on browser fingerprinting still continues to this day. In 2024, a novel mobile app was deployed to investigate the fingerprintability of mobile devices.

The web site has yielded multiple datasets of genuine fingerprints to understand the multiple facets of browser fingerprinting and how they can be used on the web to reinforce security. The web site presents regular updates to include the latest development in web technology and understand their impact of users' privacy.

The whole source code of amiunique is open source and is distributed under the terms of the MIT license.

Main innovative features:

- canvas fingerprinting
- WebGL fingerprinting
- advanced JS features (platform, DNT, etc.)

Impact: The website has been visited by more than 3,000,000 unique visitors since its creation and it has been showcased in several professional forums and tutorial sessions over the years. It produced multiple datasets over the years that were used in articles published in top-tier conferences. Amiunique has received in 2018 the prize “**Protection de la vie privée**” granted by Inria and the CNIL. The research around fingerprints in amiunique has also been a **source of influence for the Brave web browser**.

Functional Description: This web site aims at informing visitors about browser fingerprinting and possible tools to mitigate its effect, as well as at collecting data about the fingerprints that can be found on the web. It collects browser fingerprints with the explicit agreement of the users (they have to click on a button on the home page). Fingerprints are composed of 17 attributes, which include regular HTTP headers as well as the most recent state of the art techniques (canvas fingerprinting, WebGL information).

URL: <https://amiunique.org/>

Contact: Pierre Laperdrix

Partners: INSA Rennes, Université de Lille

7.1.2 APISENSE

Keywords: Mobile sensing, Crowd-sensing, Mobile application, Crowd-sourcing, Android

Functional Description: APISENSE platform is a software solution to collect various contextual information from Android devices (client application) and automatically upload collected data to a server (deployed as a SaaS). APISENSE is based on a Cloud computing infrastructure to facilitate datasets collection from significant populations of mobile users for research purposes.

URL: <https://github.com/APISENSE>

Contact: Romain Rouvoy

Partner: Université de Lille

7.1.3 cloudnet

Name: Cloudnet

Keywords: Cloud configuration, Tosca, Docker Compose, Heat Orchestration Template, Alloy

Scientific Description: The multiplication of models, languages, APIs and tools for cloud and network configuration management raises heterogeneity issues that can be tackled by introducing a reference model. A reference model provides a common basis for interpretation for various models and languages, and for bridging different APIs and tools. The Cloudnet Computational Model formally specifies, in the Alloy specification language, a reference model for cloud configuration management. The Cloudnet software formally interprets several configuration languages in it, including the TOSCA configuration language, the OpenStack Heat Orchestration Template and the Docker Compose configuration language.

The use of the software shoes, for examples, how the Alloy formalization allowed us to discover several classes of errors in the OpenStack HOT specification.

Functional Description: Application of the Cloudnet model developed by Inria to software network deployment and reconfiguration description languages.

The Cloudnet model allows syntax and type checking for cloud configuration templates as well as their visualization (network diagram, UML deployment diagram). Four languages are addressed for the moment: The OASIS's TOSCA specification, the ETSI's NFV SOL001 specification, OpenStack's HOT (Heat Orchestration Template), and Docker Compose.

We can use directly the software from an Orange web portal: <https://toscatoolbox.orange.com>

URL: <https://github.com/Orange-OpenSource/Cloudnet-TOSCA-toolbox>

Publication: [hal-02940938v1](https://hal.archives-ouvertes.fr/hal-02940938v1)

Contact: Philippe Merle

Participants: Philippe Merle, Jean-Bernard Stefani, Roger Pissard-Gibollet, Souha Ben Rayana, Karine Guillouard, Meryem Ouzzif, Frédéric Klamm, Jean-Luc Coulin

Partner: Orange Labs

7.1.4 PowerAPI

Keywords: Energy efficiency, Energy management

Functional Description: PowerAPI is a library for monitoring the energy consumption of software systems.

PowerAPI differs from existing energy process-level monitoring tools in its software orientation, with a fully customizable and modular solution that lets the user precisely define what he/she wants to monitor. PowerAPI is based on a modular and asynchronous event-driven architecture using the Akka library. PowerAPI offers an API which can be used to define requests about energy spent by a process, following its hardware resource utilization (in term of CPU, memory, disk, network, etc.).

URL: <https://powerapi.org/>

Publications: [hal-01069142v1](#), [hal-00912996v2](#), [hal-00772454v1](#), [hal-00912613v1](#), [hal-00681560v3](#), [hal-00715331v1](#), [hal-01827132v2](#), [hal-01439889v1](#), [hal-01403486v1](#), [hal-01130030v1](#), [hal-03173410v1](#), [hal-02470128v1](#), [hal-02904300v1](#)

Contact: Romain Rouvoy

Participants: Adel Nouredine, Loic Huertas, Maxime Colmant, Romain Rouvoy, Mohammed Chakib Belgaid, Arthur D'azemar, Guillaume Fieni, Daniel Romero Acero

8 New results

In 2024, we obtained new results in the domains of security and privacy (see Sections 8.1, 8.4, 8.6, and 8.7), the safe reconfiguration of software systems (see Section 8.2), the optimization of the energy consumption of software systems (see Sections 8.3 and 8.5).

8.1 Investigating Novel Tracking Practices

Participants: Pierre Laperdrix, Naif Mehanna, Walter Rudametkin.

In the context of his PhD defended in May, Naif Mehanna obtained new results in the domain of web privacy, especially based on tracking techniques on the web [46].

In this work, we introduce DrawnApart, a technique that leverages unique properties in the GPU stack to identify individual devices. Using DrawnApart, we show that we are capable of differentiating between identical devices, with success rates reaching over 95% in a controlled environment. Through a crowd-sourced data collection on the AmIUnique platform, we evaluated our technique on 2,550 unique devices and 370,392 fingerprints and found that it can significantly extend the tracking time of the state-of-the-art fingerprint-based tracking algorithm, FP-Stalker. In addition, as smartphones have found their way into the pockets of the vast majority of the world's population, we investigated the tracking ecosystem of games on Android. Our analysis of a dataset comprised of 6,355 free games and 396 paid games shows that paid games are effectively less prone to online tracking compared to their free counterpart, but were still subject to a significant amount of tracking in some cases. We also investigated the Teachers Approved program and found that it significantly reduced the number of trackers and advertising in games, highlighting the importance of enforcement on the privacy ecosystem.

This new result has been published in [34]. Previous results had been presented in [75, 71, 72].

8.2 Safe Dynamic Reconfiguration of Applications with Features

Participants: Salman Farhat, Simon Bliudze, Laurence Duchien.

In the context of his PhD defended in July, Salman Farhat obtained new results in the domain of the safe reconfiguration of software systems [43].

In this work, we propose an approach named FeCo4Reco that leverages feature models to automatically generate, in a component-based formalism called JavaBIP, component-based run-time variability models that respect the feature model constraints. These component-based run-time variability models are executable and can be used at runtime to enforce the variability constraints, that is, to ensure the (partial) validity of all reachable configurations. As complex systems' architectures may evolve at run-time by acquiring new functionalities while respecting new constraints, we define composition operators for component-based run-time variability models that not only encode these feature model composition operators, but also ensure safe run-time reconfiguration. To prove the correctness and compositionality properties, we propose a novel multi-step UP-bisimulation equivalence and use it to show that the component-based run-time variability models preserve the semantics of the composed

feature models. For the experimental evaluation, we demonstrated the applicability of our approach in real-world scenarios by generating a run-time model based on the feature model of the Heroku cloud platform using our approach. This model is then used to deploy a real-world web application on the Heroku platform. Furthermore, we measured the time and memory overheads induced by the generated run-time models on systems involving up to 300 features. The results show that the overheads are negligible, demonstrating the practical interest of our approach.

This new result has been published in [25]. Previous results had been presented in [60, 61].

8.3 Enhancing IaaS Consolidation with Resource Oversubscription

Participants: Pierre Jacquet, Romain Rouvoy.

In the context of his PhD defended in July, Pierre Jacquet obtained new results in the domain of the optimization of energy consumption for cloud computing environments [45]. This work has been conducted with Thomas Ledoux from the Inria Nantes Stack project-team.

In this work we propose to examine the usage rates of servers in cloud computing environments along four complementary contributions: 1) The creation of realistic controlled experiments in an Infrastructure-as-a-Service (IaaS) context. While platforms supporting Cloud infrastructures are extensively studied, generating realistic workloads is crucial. As each Cloud Provider has its characteristics (distribution of VM sizes, individual usage rates), we propose a tool to generate these workloads. 2) The calculation of individual server oversubscription ratio. By considering the individual stability of servers, it is possible to fine-tune the calculation of this ratio without causing additional violations. 3) The introduction of a new oversubscription paradigm. By first demonstrating that Virtual Machine (VM) vCPUs are not uniformly used in a real-world context, we expose to VMs cores of different powers (by oversubscribing them to different amounts) and demonstrate that this paradigm can improve overall performance. 4) The complementarity of oversubscription techniques to reduce unallocated resources. Comparing so-called premium VMs and oversubscribed VMs identifies that they tend to saturate their hosts' resources differently. By hosting them on the same servers, it is thus possible to benefit from synergies and reduce the number of servers by up to 9.6%.

This new result has been published in [14, 32, 33]. Previous results had been presented in [69, 68].

8.4 Exploring the Impact of the User's Browsing Environment on Web Privacy

Participants: Jean Luc Intumwayase, Pierre Laperdrix, Romain Rouvoy.

In the context of his PhD defended in December, Jean Luc Intumwayase obtained new results in the domain of web privacy [44].

In this work our aim is to enhance web privacy by minimizing unnecessary user's browsing environment (UBE) information exposure while maintaining website functionality. We provide three contributions for understanding the interplay between the user's browsing environment and web privacy:

1. We present a framework to determine the relevance of information within the UBE in relation to website functionality when accessed by websites. We categorize UBE information into geolocation, device, and browser attributes, and simulate website visits using various UBE constructs through browser instrumentation. This approach allows us to systematically restrict specific UBE attributes and observe the resulting website behavior. Our methodology involves designing a web crawler that uses these different UBE constructs and collects data from websites. To quantify the impact of modifying UBE attributes, we introduce SimilarityRadar, a multidimensional tool that compares web pages and computes similarity scores. We use SimilarityRadar to compare website behavior in normal versus restricted environments to identify discrepancies that indicate the relevance of specific UBE attributes. This framework enables the safe restriction of nonessential UBE information, thereby enhancing user privacy.

2. We use SimilarityRadar to investigate the impact of the User Agent (UA) string and associated device information on serving its original purpose of tailoring website content to various browsers. Similar to the comparison between standard and restricted UBE described previously, we crawl 270,048 web pages across 11,252 domains using three standard browsers and a set of ‘None’ browsers—clones of the standard browsers with their UA and associated information set to ‘None’. We then compare the similarity between the standard and ‘None’ browsers. Before JavaScript (JS) execution, we observe 100% similarity, indicating that UA headers are irrelevant to website functionality today. However, after JS execution, 8.4% of crawled web pages change. These changes are primarily driven by third-party scripts related to advertising, bot detection, and content delivery networks. Further investigating, we classify these changes based on levels of usability severity, with most being minor CSS adjustments and a smaller percentage causing significant usability issues. This indicates the obsolescence of the UA string and associated device information amidst its significance in browser fingerprinting.

3. We investigate the enforcement of cookie consent notices across different continents, focusing on how visibility and impact vary by geographic location. We use a novel automated visual detection technique based on SimilarityRadar to efficiently detect consent notices. For each of the five countries—Brazil, France, Japan, South Africa, and the United States—we analyzed 14,078 websites, repeating this process three times under different interaction scenarios: first without interacting with cookie consent notices, then by accepting them, and finally by rejecting them. The results reveal disparities in consent notice prevalence, with France showing the highest visibility at 69% and Japan the lowest at 27%. We also observe that third-party cookies increase after users interact with consent notices, particularly in the US, while France maintains the lowest number of third-party cookies due to its stringent regulations. This indicates a correlation between consent notice visibility and user tracking practices.

This work led to the following publication: [66].

8.5 Sustainable Software Design: Estimation and Reduction of the Environmental Footprint of Software Systems

Participants: Thibault Simon, Romain Rouvoy.

In the context of his PhD defended in December, Thibault Simon obtained new results in the domain of sustainable software design [49].

This work combines different approaches from various research fields to identify significant levers to reduce the environmental footprint of software. First, we assess the advantages and disadvantages of top-down analyses to assess the carbon footprint of the Information and Communication Technologies (ICT) sector, and we demonstrate their interest in assessing the sector’s impact on other impact categories, notably on metals and minerals. Building on the observed upward trend, we then adopt a bottom-up methodology to develop tools and methodologies to assess and identify reduction levers in various aspects of the environmental footprint of digital services. More specifically, we assess the environmental footprint of cloud services and user devices with a lifecycle approach, while proposing a new methodology to systematically track uncertainties arising from reference sources and modeling choices within these estimates.

To go beyond the usage phase and the energy footprint of software, we also propose a methodology and an associated tool to holistically assess the impacts occurring throughout the software lifecycle. However, impact estimation is only the first step in software ecodesign. We therefore explore the different responsibilities of software components, and introduce a conceptual model to help different software stakeholders define metrics to reduce the environmental footprint of software, in their areas of responsibility. Within this conceptual framework, we introduce a new architectural quality metric that focuses on minimizing the resource waste induced by the software architecture, as a simple and implementable solution. Furthermore, we propose a practical approach for software stakeholders to strive to achieve proportionality between their environmental footprint and the evolution of the functional unit over time.

This work was conducted in the context of a collaboration with Orange and funded by the Cifre scheme. This new result has been published in [28, 38]. Previous results had been presented in [83].

8.6 Privacy Driven Data Collection

Participants: Maryam Rahmani, Romain Rouvoy.

In the context of her PhD defended in December, Maryam Rahmani obtained new results in the domain of privacy driven data collection techniques, applied to the domain of air quality prediction [48]. This is a multidisciplinary work with Suzanne Crumeyrolle, a researcher in physics at the University of Lille (LOA laboratory) who studies the chemical and physical properties of aerosols.

This work presents a novel approach to air pollution prediction in urban areas by integrating Artificial Intelligence (AI), spatiotemporal modeling, and privacy-preserving data collection techniques. The first major contribution of this research is the development of PMFORECAST, a temporal prediction model specifically designed to forecast PM_{2.5} levels. By utilizing advanced machine learning techniques and temporal attention mechanisms, PMFORECAST effectively captures temporal dependencies in pollutant concentrations, leading to highly accurate predictions for both short-term and long-term forecasting. Additionally, the model demonstrates significant multi-tasking capabilities. It achieves a notable prediction accuracy of 99.7% for 1-hour forecasts and 73.5% for 12-hour forecasts, representing substantial improvements over existing models in terms of precision and computational efficiency.

Spatial and temporal data from underground sensor networks to predict PM_{2.5} concentrations across different geographic regions. The Graph Temporal LSTM (GT-LSTM) model employs Graph Convolutional Networks (GCNs) to capture the complex interactions between pollution sources and atmospheric conditions at ground level, while utilizing Long Short-Term Memorys (LSTMs), as described in the previous contribution, to model temporal dependencies. This approach provides a more refined understanding of pollutant dispersion over time and space. By operating with fixed zone resolutions corresponding to available data resources, the model ensures accurate and localized predictions.

The third contribution is the design of a federated learning architecture called FEDAIRNET, aimed at enhancing air quality prediction using mobile sensor data while safeguarding user privacy. Traditional air quality monitoring stations are often constrained by limited spatial coverage and high costs, whereas mobile sensors offer a more flexible and granular data source. However, the collection of mobile sensor data introduces privacy concerns. FEDAIRNET addresses these challenges by distributing the learning process across multiple devices, ensuring that sensitive data remains on local devices while still contributing to global model updates. This decentralized approach not only improves prediction accuracy but also mitigates risks associated with centralized data collection, such as point-of-interest (PoI) attacks.

The models presented in this thesis have been rigorously tested in real-world environments, demonstrating their potential to transform air pollution monitoring systems. The PMFORECAST model provides robust predictions of PM_{2.5} concentrations, making it valuable for public health interventions and environmental policies. The Spatiotemporal Model adds a critical layer of understanding by analyzing how pollutants behave across spatial and temporal dimensions, while the FEDAIRNET architecture ensures that privacy is protected as the use of mobile sensors becomes more prevalent.

This research represents significant advancements in air pollution prediction by integrating AI-driven insights with privacy-preserving data collection techniques. Future work should focus on incorporating additional data sources, and refining hybrid models that combine temporal, spatial, and mobile sensing data. These innovations will contribute to more accurate, timely, and secure air pollution forecasting systems, ultimately helping to mitigate the harmful effects of air pollution on human health and the environment.

This new result has been published in [16]. The code and data developed in the context of this work are available online on [this site](#).

8.7 Electromagnetic Emanations Side Channels in Embedded Devices

Participants: Pierre Ayoub, Clémentine Maurice.

In the context of his PhD defended in December, Pierre Ayoub obtained new results in the domain of compromising electromagnetic emanations, and in particular in embedded devices [53].

In this work, we focus on new security risks impacting microcontrollers linked to unexpected interactions between heterogeneous digital and analog embedded modules. In particular, we analyze threats related to cross-layer interactions that can be exploited through electromagnetic side-channel analysis and propose two main contributions. First, despite a flourishing literature on electromagnetic side channels, the modulation of leaked signals remains a complex phenomenon which is not fully understood. In this context, the first major contribution of our work aims at enhance understanding of how digital activity modulate leaked electromagnetic signals for both offensive and defensive applications. More precisely, electromagnetic side channels typically focus on the amplitude of leaked signals, neglecting the potential interest of other modulation types from a security perspective. Our first contribution, *PhaseSCA*, uncovers unintended phase modulation in leaked signals as a novel source of side-channel. Second, the applicability of newly discovered side-channel attacks regarding modern communication protocols is not systematically evaluated. An illustration of this is the *Screaming Channels* attack, which exploits a phenomenon of intermodulation between the leakage from a digital activity and the carrier of a radio transceiver in mixed-signal chips. Modern protocols only enable the radio transceiver for a short duration, introducing a serious limitation since Screaming Channels exploits leakage broadcasted through the radio transceiver. Our second contribution explores the impact of this threat on Bluetooth Low Energy protocol. We highlight how an attacker can manipulate the protocol parameters through traffic injection, forcing a victim to transmit during sensitive operations, demonstrating the threat introduced by Screaming Channels for the Internet of Things ecosystem.

Overall, this work underscores the need to systematically analyze the unexpected interactions between hardware and software to design secure chips and protocols, as well as highlight emerging security threats and their implications for wireless communications.

This work was conducted in the context of the ANR MobiS5. This new result has been published in [20] and [11].

9 Bilateral contracts and grants with industry

Orange # 1

Participants: Thibault Simon, Romain Rouvoy (*contact person*).

This collaboration (2021–24) aims at working on the sustainable design of software systems. The purpose is especially to characterize the quality of software components from an environmental point of view to go towards the production of low environment footprint software systems.

The CIFRE PhD thesis of Thibault Simon was defended in December in the context of this collaboration. Results have been published in [28, 38][83]. Thibault Simon won the jury prize in the contest Ma thèse en trois minutes that was awarded during the Orange Tech Days 2023².

Orange # 2

Participants: Yifan Wang, Pierre Bourhis (*contact person*), Romain Rouvoy.

This collaboration (2022–25) aims at working on the self-optimization of database management systems.

The CIFRE PhD thesis of Yifan Wang takes place in the context of this collaboration. First results have been published in [40].

²[Link to the announcement](#)

Orange # 3

Participants: Boubacar Diarra, Philippe Merle (*contact person*).

This collaboration (2023–26) aims at working on the formal verification of cloud native applications. The CIFRE PhD thesis of Boubacar Diarra takes place in the context of this collaboration. First results have been published in [24].

Qarnot Computing

Participants: Sophie Cerf (*contact person*), Thomas Collignon, Lionel Seinturier.

This collaboration (2023–26) aims at working on the control of high performance computing tasks emissions. In tomorrow’s computing services, users will have objectives in terms of quality of service and cost, to which will be added new objectives related to the energy impact and environmental footprint of their calculations. As these objectives are contradictory, there will necessarily be a compromise to ensure. This collaboration focuses on the realization of this trade-off at the level of a computational task, or a cluster of computational tasks of a same user, the user deciding on the relative importance to be brought to the parameters mentioned above.

The CIFRE PhD thesis of Thomas Collignon and the Inria Défi PULSE take place in the context of this collaboration. First results have been published in [50].

10 Partnerships and cooperations

10.1 International research visitors

10.1.1 Visits of international scientists

In the framework of the SmartCloud ANR project, Giuseppe De Palma, a PhD student from the University of Bologna working on cloud and serverless computing, visited us for 6 months from January to June 2024.

10.2 European initiatives

10.2.1 Other european programs/initiatives

ANR France-Germany FACADES

Participants: Iliana Fayolle, Sihem Bouhenniche, Pierre Laperdrix (*contact person*), Clémentine Maurice, Romain Rouvroy.

FACADES is a 42-month international project (2022–26) with CISPA and Saarland University funded in the context of the German-French ANR joint call. The project investigates fingerprinting and CPU attack and defense exploration from browser scripts. The aim is to analyze the security implications of new features in web browsers (WebAssembly, WebGPU, WebUSB, etc.) that provide direct or indirect access to low-level hardware features.

The PhD theses of Sihem Bouhenniche and Iliana Fayolle take place in the context of this project. First results have been published in [26].

10.3 National initiatives

10.3.1 ANR

ANR ADAPT

Participants: Simon Bliudze (*contact person*), Sophie Cerf, Adrien Luxey-Bitri, Daniel Romero, Lionel Seinturier.

ANR ADAPT is a 48-month project (2024–28) funded by ANR. The project’s objective is to develop a formal framework and tools that support hierarchical modeling of self-adaptive systems and integrate adaptive control to allow dynamic reconfigurations. To this end, we intend to 1) extend component-based models with hierarchical motifs to allow specifying modular robots (MR) functions and constraints on resources; 2) define appropriate control mechanisms, based on hierarchical control motifs, for adapting and reconfiguring hierarchies of motifs, utilising aggregate measures over components; 3) implement control mechanisms over extended component-based models; and use them to 4) simulate and validate adaptations, before deploying them on the real MRs, such as Blinky Blocks.

ANR CQFD

Participants: Pierre Bourhis (*contact person*).

CQFD is a 72-month project (2018–24) funded by ANR. The project focuses on the complex ontological queries over federated heterogeneous data. The project targets to set the foundations, to provide efficient algorithms, and to provide query rewriting oriented evaluation mechanisms, for ontology-mediated query answering over heterogeneous data models. This project is coordinated by Federico Ulliana from Inria Sophia Antipolis. Other partners include LaBRI, Inria Saclay, IRISA, LTCI, and LIG.

First results have been published in [55, 57].

ANR Distiller

Participants: Alexandre Bonvoisin, Tristan Coignon, Clément Quinton, Romain Rouvoy (*contact person*), Lionel Seinturier.

Distiller is a 36-month project (2021–24) funded by ANR. The project intends to better assist practitioners by delivering software artifacts recommendations to promote sustainable cloud native software. Distiller will rely on a software sustainability index. The core idea is to evaluate the sustainability of different cloud-native services. Based on those evaluations and your project requirements, Distiller will recommend a more sustainable next technical stack.

The ongoing PhD theses of Alexandre Bonvoisin and Tristan Coignon take place in the context of this project. First results have been published in [21, 23].

ANR FP-Locker

Participants: Pierre Laperdrix, Naif Mehanna, Walter Rudametkin (*contact person*).

FP-Locker is a 46-month project (2019–24) funded by ANR in the context of the JCJC program. This project proposes to investigate advanced browser fingerprinting as a configurable authentication mechanism. We argue that it has the potential to be the only authentication mechanism when used in very low-security, public websites; it can be used to block bots and other fraudulent users from otherwise open websites. It also has the potential to be used as a second factor authentication mechanism, or as

an additional factor in Multi-Factor Authentication (MFA) schemes. Besides strengthening a session's initial authentication, it can also be used for continuous session authentication to protect against session hijacking. In many contexts, fingerprinting is fully transparent to users, meaning that contrary to authentication processes that rely on external verification cards, code generating keys, special apps, SMS verification codes, users do not have to do anything to improve their security. In more restricted contexts, administrators can enforce different policies, for example, enrolling fingerprints from devices that connect from trusted IP addresses (e.g., an internal network), and then verifying these fingerprints when the same users connect from untrusted IP addresses. Consequently, we plan to design an architecture and implement it to be able to plug the browser fingerprinting authentication process to an existing authentication system.

The PhD thesis of Naif Mehanna has been defended in May in the context of this project [46]. This work led to the following publications: [34][75, 71, 72]. A best paper award at the MADWeb 2024 workshop on measurements, attacks, and defenses for the Web in the context of the NDSS Network and Distributed System Security Symposium for their work on a longitudinal study analyzing the stability, security, and potential manipulation of free web proxies [34].

ANR Koala

Participants: Pierre Bourhis, Edouard Guegain, Clément Quinton (*contact person*).

Koala is a 46-month project (2019–24) funded by ANR in the context of the JCJC program. The project aims to deliver a series of innovative tools, methods and software to deal with the complexity of fog computing environments configurations and adaptations. In particular, we take a step back on the current limitations of existing approaches (e.g., lack of expressiveness and scalability) and address them placing knowledge as a first-class citizen. We plan to tackle configuration issues from a novel perspective in the field of variability management, using recent techniques from the area of knowledge compilation. Specifically, we will investigate the best-suited d-DNNF representation for each reasoning operation, and we plan to provide new variability modeling mechanisms (e.g., dimensions, priorities and scopes) required in a fog context. Regarding adaptation concerns, we want to leverage machine learning techniques to improve adaptation management and evolution under uncertainty, relying on a continuously enriched and reusable knowledge base. In particular, we plan to propose an approach for suggesting evolution scenarios in a predictive manner, relying on an evolution-aware knowledge base acquired at run-time through machine learning feedback.

The PhD thesis of Edouard Guegain has been defended in December 2023 in the context of this project [62]. This work led to the following publications: [65, 76, 63, 64].

ANR RAISIN

Participants: Sophie Cerf (*contact person*).

RAISIN is a 48-month project (2024–28) funded by ANR. The project intends to develop techniques enabling resource-aware conservative static analyses. Complexity evaluation of static analysis will allow to estimate the analysis time of a program, then the objective is to develop dynamic, control-based techniques that change the precision of a given analysis during the analysis itself. The project is led by the Inria SyCoMoRES project-team.

ANR SCALER

Participants: Philippe Merle (*contact person*), Hugo Monfleur, Romain Rouvoy.

SCALER is a 42-month project (2022–26) funded by ANR. The project aims to optimize the scaling of microservice-based networked services while satisfying their stringent IT and telco requirements. Especially the objectives are to identify relevant metrics to characterize microservices, to define microservices integration patterns, and to design smart management strategies. Partners are the University of Grenoble, the Orange and Eolas companies.

The ongoing PhD thesis of Hugo Monfleuret takes place in this project. First results have been published in [77].

ANR SmartCloud

Participants: Simon Bliudze (*contact person*), Philippe Merle, Arleon Zemtsop.

SmartCloud is a 42-month project (2024–27) funded by ANR. The project goal is to develop a flexible infrastructure for smart and coordinated dynamic reconfiguration of Cloud Computing systems. Formal Methods techniques will be used to allow explicit specification of the structural and behavioural constraints of the system to provide formal correctness guarantees and allow proactive adaptivity in a coordinated manner. Distributed monitoring and online optimisation techniques will be used for dynamic adaptation, aiming to optimise the efficiency of resource usage in a scalable manner. Partners are the Scalair company and the Inria OLAS project-team.

10.3.2 PEPR and PTCC

PEPR Cloud Taranis

Participants: Philippe Merle (*contact person*), Clément Quinton, Brell Sanwouo.

Taranis is a 7-year project (2023–29) funded in the context of PEPR Cloud Computing framework. New infrastructures, such as Edge Computing or the Cloud-Edge-IoT computing continuum, make cloud issues more complex as they add new challenges related to resource diversity and heterogeneity (from small sensor to data center/HPC, from low power network to core networks), geographical distribution, as well as increased dynamicity and security needs, all under energy consumption and regulatory constraints. In order to efficiently exploit new infrastructures, the project proposes a strategy based on a significant abstraction of the application structure description to further automate application and infrastructure management. Thus, it will be possible to globally optimize the resources used with respect to multi-criteria objectives (price, deadline, performance, energy, etc.) on both the user side (applications) and the provider side (infrastructures). This abstraction also includes the challenges related to the abstraction of application reconfiguration and to automatically adapt the use of resources.

The ongoing PhD thesis of Brell Sanwouo takes place in the context of this project. First results have been published in: [17].

PEPR Cloud CARECloud

Participants: Djefal Belkis, Pierre Bourhis, Sophie Cerf, Clément Quinton, Adrien Luxey-Bitri, Romain Rouvroy (*contact person*).

CARECloud is a 7-year project (2023–29) funded in the context of PEPR Cloud Computing framework. The first objective of the project is to understand how cloud infrastructures consume energy in order to identify sources of waste and to design new models and metrics to qualify energy efficiency. The second objective focuses on the energy efficiency of cloud infrastructures, i.e., optimizing their consumption during the usage phase. In particular, this involves designing resource allocation and energy lever orchestration strategies: mechanisms that optimize energy consumption (sleep modes, dynamic adjustment of the size of virtual resources, optimization of processor frequency, etc.). Finally, the third

objective targets digital sobriety in order to sustainably reduce the environmental impact of clouds and aims to design infrastructures that are more energy and IT resource efficient, resilient to electrical intermittency, adaptable to the production of electricity from renewable energy sources and tolerant of the disconnection of a highly decentralized part of the infrastructure.

The PhD thesis of Djeflal Belkis takes place in the context of this project.

PEPR Cybersecurity IPoP

Participants: Pierre Bourhis, Imane Fouad, Clémentine Maurice, Pierre Laperdrix (*contact person*), Romain Rouvoy, Juliette Sénéchal.

IPoP is a 72-month project (2022–28) funded in the PEPR Cybersecurity framework. The objectives of the IPoP (Interdisciplinary Project on Privacy) project are to study the threats on privacy that have been introduced by these new digital technologies, and to conceive theoretical and technical privacy-preserving solutions that are compatible with French and European regulations, that preserve the quality of experience of the users. Spirals is leader of WP2 on new forms of personal data gathering and their associated threats for privacy.

PEPR Cybersecurity REV

Participants: Nusrat Jahan Farin, Clémentine Maurice (*contact person*), Pierre Laperdrix.

REV is a 72-month project (2022–28) funded in the PEPR Cybersecurity framework. The objectives of the REV (Recherche et Exploitation de Vulnérabilités) project are to study the presence of vulnerabilities in modern devices by attacking multiple layers of a system at the same time from the software, hardware and its interfaces. This broad-spectrum analysis is key to understand how far an attack can go since the deployment of numerous protections prevents a device from being compromised through a single vector. The goals of the project are multiple from finding novel attacks to determining the best course for correction while considering how potential findings fit into the legal framework. Spirals is leader of WP2 focused on low level attacks with one task in WP4 being dedicated to web vulnerabilities.

PTCC SWHSec

Participants: Pierre Laperdrix (*contact person*), Clémentine Maurice, Lionel Seinturier.

SWHSec is a 36-month project (2023–26) funded in the context of the CampusCyber. The objective of the SWHSec project is to assess and control the software vulnerabilities that can be maliciously introduced in large code bases. The project leverages the existing massive Software Heritage archive. The purpose is to enrich the archive with security-relevant information.

10.4 Public policy support

Participants: Pierre Bourhis, Pierre Laperdrix, Lionel Seinturier, Juliette Sénéchal (*contact person*).

We contributed to two public consultations launched by the CNIL on:

- the building of datasets for machine learning based artificial intelligence systems,
- the opening and reuse of data published on the Internet.

These contributions have been coordinated with the ones made by the IPoP project (PEPR Cybersecurity program).

11 Dissemination

11.1 Promoting scientific activities

11.1.1 Scientific events: organisation

Member of the organizing committees

Sophie Cerf: 8th GDR RSD / ASF Winter School on Distributed Systems & Networks.

Rémy Raes: 8th GDR RSD / ASF Winter School on Distributed Systems & Networks.

Romain Rouvoy: 8th GDR RSD / ASF Winter School on Distributed Systems & Networks.

11.1.2 Scientific events: selection

Chair of conference program committees

Pierre Laperdrix: Vice Program Chair for Privacy Enhancing Technologies Symposium (PETS).

Member of the conference program committees

Simon Bliudze: International Conference on Formal Aspects of Component Software (FACS), International Conference on Fundamentals of Software Engineering (FSEN), International Conference on Formal Methods in Software Engineering (FormalISE).

Sophie Cerf: International Conference for High Performance Computing, Networking, Storage, and Analysis (SC) Reproducibility Initiative, International Conference on Software Engineering for Adaptive and Self-Managing Systems (SEAMS), BenchCouncil International Symposium On Benchmarking, Measuring And Optimizing (Bench), EuroSys Doctoral Workshop (EuroDW).

Clémentine Maurice: International Symposium on Research in Attacks, Intrusions and Defenses (RAID).

Clément Quinton: International Working Conference on Variability Modelling of Software-Intensive Systems (VAMOS), International Systems and Software Product Lines Conference (SPLC) - Artefact Evaluation track.

Romain Rouvoy: International Conference on Software Engineering (ICSE) - Industry challenge track, ACM Symposium on Applied Technology (SAC) DADS track, International Middleware Conference, International Conference on Dependable Systems and Networks (DSN).

Lionel Seinturier: European Conference on Software Architecture (ECSA), ACM Symposium on Applied Technology (SAC) SA-TTA track, International Conference on Service Computing (ICSOC), IEEE International Conference on Web Services (ICWS).

Pierre Laperdrix: ACM Conference on Computer and Communications Security (CCS) - Track "Web Security", USENIX Security Symposium, USENIX WOOT Conference on Offensive Technologies (WOOT).

11.1.3 Journal

Member of the editorial boards

Laurence Duchien: special issues co-editor for Journal of Systems and Software (JSS).

Reviewer - reviewing activities

Simon Bliudze: Science of Computer Programming (SciCo).

Sophie Cerf: ACM/IFIP International Middleware Conference, American Control Conference (ACC), Springer Journal of Supercomputing.

Clément Quinton: ACM Transactions on Autonomous and Adaptive Systems (TAAS), ACM Transactions on Software Engineering and Methodology (TOSEM), Journal of Systems and Software (JSS).

11.1.4 Invited talks

Sophie Cerf gave the following invited talk.

- A talk on "Automatique et IA : des opportunités pour maîtriser l'utilisation des ressources" in the Workshop frugalité en IA et en statistique on Oct. 4.

Pierre Laperdrix gave the following invited talks.

- A keynote on the topic of web evolution at the MADWeb workshop on March 1st.
- A talk with Vincent Rocca on "Web, smartphone, AdTech: the privacy viewpoint" in the 2024 Winter School of the PEPR Cybersecurity on 29 January [18].

Clémentine Maurice gave the following invited talks.

- A masterclass on microarchitectural attacks at the Forum InCyber (FIC), on March 26,
- "A Systematic Evaluation of Automated Tools for Side-Channel Vulnerabilities Detection in Cryptographic Libraries" at the SWHSec Conference, on June 6,
- "Side-channel-free software, are we there yet?" at the Max Planck Institute for Security and Privacy Symposium on System Security on June 17,
- "Side-channel-free software, are we there yet?" at the Journées Scientifiques Inria on August 29.

Clément Quinton gave the following invited talks.

- "Generation of Software Variants" in the Défi Inria LLM4Code conference on July 5,
- "Optimizing Configurable Software Systems" in the Scientific day on optimization at CNRS on October 4.

Juliette Sénéchal gave the following invited talks.

- "AI Alignement through IA Act" in Paris Conference on Artificial Intelligence and Digital Ethics, Sciences Po Paris, 7 June,
- "DSA and AI Act" in MDD summer school, Ceillac, 24 June,
- "DSA and AI Act" in CAp-RFIAP Conference, 1 July,
- "L'AI Act dans son rapport aux valeurs humaines et aux droits fondamentaux, un enchevêtrement pragmatique d'approches normatives modernes, post-moderne et non moderne" in Séminaire "Sociologie du numérique : d'internet à l'intelligence artificielle" at EHESS,
- "Le contrat conclu avec un service d'intermédiation en ligne (plateforme)" in "Comparaison du droit français et du droit roumain des contrats spéciaux", Université de Bucarest & Université Paris I Panthéon Sorbonne (co-org.), 19 September,
- "Vers l'émergence d'un droit neuro-éthique érigé en contrepoint des droits revisités par le numérique ? Réflexion à partir du droit de la consommation" in Troisième journée scientifique de l'Arcom en partenariat avec l'ENS Paris Saclay, 14 November,
- "Pratiques d'IA influençant et manipulant les consommateurs : quelles réglementations ?" in "Encadrement éthique et juridique de l'IA : bilan 2 ans après ChatGPT", Digital Dauphine Days, 19 November.

11.1.5 Leadership within the scientific community

Simon Bliudze

- Co-head of the YODA (trustworthyY and Optimal Dynamic Adaptation) working group of the GDR GPL
- Member of the Formal Methods Europe (FMEurope) **Book Review Committee** (the aim of the committee is to provide the formal methods community, and the scientific community in general, with high-quality reviews of books on topics of interest to the community)

Sophie Cerf

- Elected member of the "bureau" of the French chapter of the ACM Special Interest Group in Operating Systems (SIGOPS / ASF)

Laurence Duchien

- President of the scientific council of IRT SystemX

Clémentine Maurice

- Member of the "Software and source codes college" of the Committee for Open Science of the Higher Education Ministry
- Member of the SSLR ("Sécurité des systèmes, des logiciels et des réseaux") working group of the GDR SI

Romain Rouvoy

- Elected member of CoNRS section 6 (since September 2021)
- President of the French chapter of the ACM Special Interest Group in Operating Systems (SIGOPS / ASF)
- Elected member of the administrative council of Specif Campus
- Co-head of the "Logiciel Eco-Responsable" working group of the GDR GPL

Pierre Laperdrix

- Assistant director of the GDR Sécurité Informatique (since September 2024)

11.1.6 Scientific expertise**Pierre Bourhis**

- Member of the jury of agrégation of informatique

Sophie Cerf

- Member of the recruitment committee for two associate professor position, at Université de Lille and Université Grenoble Alpes
- Member of the Inria national committee for "détachement"

Laurence Duchien

- Member of the scientific advisory board of Labex CIMI-Toulouse
- Member of the scientific advisory board of IMT Atlantique
- President of the recruitment committee for a professor position at ENS Lyon

Clément Quinton

- Vice-President of the recruitment committee for an associate professor position at Université de Lille

Romain Rouvoy

- Member of the jury of agrégation of informatique

Lionel Seinturier

- President of the recruitment committee for a professor position at Université Grenoble Alpes

11.1.7 Research administration

Simon Bliudze

- Elected member of the Centre Committee of the Centre Inria de l'Université de Lille
- Member of the Gender Parity committee of the CRISAL laboratory
- Co-organiser of the joint Inria/CRISAL Mentoring programme

Pierre Bourhis

- Referent of Communication of the CRISAL laboratory
- Animator of the communication committee of the CRISAL laboratory
- Member of the Gender Parity committee of the CRISAL laboratory

Sophie Cerf

- Member of the Gender Parity committee of the CRISAL laboratory
- Mediation scientific co-referent for Inria Center of the University of Lille

Laurence Duchien

- Dean of the Faculty of science and technology at the University of Lille

Maxime Huyghe

- Elected member of the Council of the doctoral school MADIS

Clémentine Maurice

- Member of the Gender Parity committee of the CRISAL laboratory

Philippe Merle

- Elected member of the Inria scientific board (CS)
- Elected member of the Inria evaluation committee (CE) (since July)
- Elected member of the Inria social administration committee (CSA)
- Member of the Inria national committee on "hygiène, de sécurité et des conditions de travail" (FS CSA)
- Member of the "Comité Local d'Hygiène, de Sécurité et de Conditions de Travail" (FSS)
- Elected member of the centre committee for the Inria Center of University of Lille

Lionel Seinturier

- ICT Panel Scientific Conselor for the Research Evaluation Department of Hcéres (since September)
- Member of the Scientific council of the University of Lille

11.2 Teaching - Supervision - Juries

11.2.1 Teaching

Simon Bliudze is, in addition to his tenure Junior Researcher position at Inria, part-time Assistant Professor (part-time Associate Professor since September 2024) at École Polytechnique, Palaiseau, France, in the Department of Computer Sciences (DIX).

- CSC_2F002_EP: Design and Analysis of Algorithms, 28h, Bachelor cycle
- INF361: Introduction to Computer Science, 40h, Cycle Polytechnique

Pierre Bourhis is, in addition to his tenure Junior Researcher position at CNRS, part-time Associate Professor of Data Science at École Polytechnique, Palaiseau, France, in the Department of Computer Sciences (DIX).

- Info553: Bases de données, 18h, Cycle Polytechnique
- Modal Graphe Géant, 36h
- INF517: Projet de Recherche Data Science, 20h
- INF583: System for Big Data, 20h

Sophie Cerf teaches at Centrale Lille Institute.

- Modélisation et commande de systèmes (state feedback, observation, discrete control), 12h, Level M1-2, École Centrale Lille
- Commande des Systèmes (basics of linear feedback control), 16h, Level L2 and L3, ITEEM
- Projet d'intégration en Robotique (robot building and control), 24h, Level M1, ITEEM

Laurence Duchien teaches at the Université de Lille in the FST faculty.

- Software engineering project, 64h, Level M2, Master MIAGE FI/FC/FA

Adrien Luxey-Bitri teaches at the Université de Lille in the FST faculty.

- Informatics, 36h, Level L1, Licence of Computer Science
- Computer networks, 20h, Level L3, Licence of Computer Science
- Software Engineering, 18h, Level L3, Licence of Computer Science
- Distributed systems, 48h, Level M1, Master of Computer Science
- Advanced distributed systems, 24h, Level M2, Master of Computer Science
- Suivi de stages, projets et mémoires, 30h, Licence and Master of Computer Science
- Enjeux Environnementaux du Numérique (Introduction à l'Analyse du Cycle de Vie), 4h, Level M2, Master of Computer Science, with S. Cerf

Philippe Merle teaches at the Université de Lille in the FST faculty.

- Software Configuration, 12h, Level M2, Master of Computer Science

Clément Quinton teaches at the Université de Lille in the FST faculty.

- Object-oriented programming, 36h, Level L2, Licence of Computer Science
- Software project, 36h, Level L2, Licence of Computer Science
- Object-oriented design, 31h, Level L3, Licence of Computer Science

- Distributed systems, 24h, Level M1, Master of Computer Science
- Software product lines, 24h, Level M2, Master of Computer Science
- Suivi de stages et de projets, 30h, Licence and Master of Computer Science

Romain Rouvoy teaches at the Université de Lille in the FST faculty.

- Design of distributed applications, 12h, Level M1, Master of Computer Science
- Object-oriented design, 4h, Level L3, Licence of Computer Science
- Suivi de projets, 20h, Level M2, Master of Computer Science

Lionel Seinturier teaches at the Université de Lille in the FST faculty.

- Conception d'applications réparties, 48h, Level M1, Master MIAGE
- Systèmes répartis avancés 1, 24h, Level M2, Master of Computer Science

11.2.2 Supervision

- HDR defended: Clément Quinton, Evolving, Adapting and Optimizing Configurable Software Systems, defended on 26 January 2024, University of Lille, [47].
- HDR defended: Simon Bliudze, Contributions to the Rigorous Design of Concurrent Component-Based Software and Systems Using BIP, defended on 8 March 2024, University of Lille, [41].
- HDR defended: Pierre Bourhis, Dialog between Logic, Trees, and Circuits, defended on 18 September 2024, University of Lille, [42].
- PhD defended: Naif Mehanna, Intrinsically Inseparable: Investigating Novel Tracking Practices and Assessing the Carbon Footprint of Ads, defended on 30 May 2024, University of Lille, [46].
- PhD defended: Salman Farhat, Safe Dynamic Reconfiguration of Applications with Features, defended on 11 July 2024, University of Lille, [43].
- PhD defended: Pierre Jacquet, Enhancing IaaS Consolidation with Resource Oversubscription, defended on 19 July 2024, University of Lille, [45].
- PhD defended: Pierre Ayoub, Compromising Electromagnetic Emanations: Side-Channel Leakages in Embedded Devices, defended on 3 December 2024, Eurecom.
- PhD defended: Jean Luc Intumwayase, Exploring the Impact of the User's Browsing Environment on Web Privacy, defended on 5 December 2024, University of Lille, [44].
- PhD defended: Thibault Simon, Sustainable Software Design: Estimation and Reduction of the Environmental Footprint of Software Systems, defended on 6 December 2024, University of Lille, [49].
- PhD defended: Maryam Rahmani, Next-Generation Air Pollution Forecasting: Integrating AI, Spatiotemporal Dynamics, and Privacy-Ensuring Approaches for Urban Areas, defended on 11 December 2024, University of Lille, [48].
- PhD in progress: Virginie Amand, Designing Sustainable Software Services with Large Language Models, since October 2024, University of Lille, supervised by Romain Rouvoy & Clément Quinton.
- PhD in progress: Alexandre Bonvoisin, Reducing the Energy Consumption of Software Stacks, since October 2022, University of Lille, supervised by Romain Rouvoy & Clément Quinton.
- PhD in progress: Sihem Bouenniche, Security Analysis of Current and Future Web Standards, since January 2024, University of Lille, supervised by Clémentine Maurice & Walter Rudametkin.

- PhD in progress: Tristan Coignon, Environmental Impact of Development Assistants, since October 2022, University of Lille, supervised by Romain Rouvoy & Clément Quinton.
- PhD in progress: Thomas Collignon, Control of high performance computing tasks emissions, since November 2023, University of Lille, supervised by Sophie Cerf & Lionel Seinturier.
- PhD in progress: Belkis Dejffal, Optimization of Databases in Cloud Environments in a Context of Green Computing, since December 2023, University of Lille, supervised by Romain Rouvoy & Pierre Bourhis.
- PhD in progress: Boubacar Diarra, Formal modeling and reliability of cloud network configurations, since January 2023, University of Lille, supervised by Philippe Merle.
- PhD in progress: Nusrat Jahan Farin, since November 2024, University of Lille, supervised by Clémentine Maurice.
- PhD in progress: Iliana Fayolle, Side channel vulnerabilities in web environments, since 1 October 2023, University of Lille, supervised by Clémentine Maurice.
- PhD in progress: Antoine Geimer, Détection et correction automatique de vulnérabilités par canaux auxiliaires dans les bibliothèques cryptographiques, since October 2022, University of Rennes, supervised by Clémentine Maurice & Sandrine Blazy (Inria Epicure project-team).
- PhD in progress: Maxime Huyghe, Automated Software Testing to Improve the Privacy of Browsers, since October 2022, University of Lille, supervised by Lionel Seinturier & Clément Quinton & Walter Rudametkin (Inria Diverse project-team).
- PhD in progress: Hugo Monfleur, Design and implementation of a micro-services oriented language, since 1 October 2022, University of Lille, supervised by Philippe Merle.
- PhD in progress: Rémy Raes, Distributed machine learning in ubiquitous environments using location dependant models, since 1 March 2023, University of Lille, supervised by Romain Rouvoy & Adrien Luxey-Bitri.
- PhD in progress: Brell Sanwouo, Artificial Intelligence applied to Self-Adaptive Systems, since March 2024, University of Lille, supervised by Clément Quinton & Paul Temple (Inria Diverse project-team).
- PhD in progress: Yifan Wang, Autonomous Management of Database Systems, since October 2022, University of Lille, supervised by Romain Rouvoy & Pierre Bourhis.
- PhD in progress: Jiali Xu, Behavior characterisation in smart and deep systems with heterogeneous and reprogrammable nodes, University of Lille, supervised by Romain Rouvoy & Valéria Loscri (Inria Fun project-team).
- PhD in progress: Nada Zine, since November 2024, University of Lille, supervised by Romain Rouvoy & Clément Quinton.

11.2.3 Juries

Simon Bliudze

- PhD Paulina Paraponiari (Aristotle University of Thessaloniki, Greece), examiner
- PhD Perla Tannoury (Univ. Franche-Comté), reviewer

Clémentine Maurice

- PhD Vincent Giraud (ENS Paris), reviewer
- PhD Marco Casagrande (EURECOM), reviewer

Romain Rouvoy

- Yassine El Amraoui (Univ. Côte d'Azur), examiner
- David Delande (Univ. Toulouse), examiner
- Milad Jamalzadeh (Univ. Lille), president
- Adrien Berthelot (ENS Lyon), examiner
- Daniel Lezcano (Univ. Toulouse), reviewer
- Chih-Kai Huang (Univ. Rennes), reviewer
- HDR Antoine Boutet (Insa Lyon), reviewer
- Théophile Dubuc (ENS Lyon), reviewer

Lionel Seinturier

- HDR Samy Yanguï (INP Toulouse), president
- HDR Nawal Guermouche (INP Toulouse), examiner

11.3 Popularization

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

Clémentine Maurice wrote an article with David Monniaux on "Les arcanes des processeurs mis à l'épreuve" that has been published in July in the magazine La Recherche [52].

11.3.2 Participation in Live events

Adrien Luxey-Bitri worked with the communication service of Inria Lille to set up the exhibit "Le numérique en eaux troubles". The goal is to raise awareness among citizens about the impacts of digital development on water resources. The exhibit was displayed from 11 October to 25 November in Forum départemental des sciences in Villeneuve d'Ascq in the context of Fête de la science.

Clémentine Maurice and Iliana Fayolle contributed to Les rendez-vous des jeunes mathématiciennes et informaticiennes (RJMI) on 21 and 22 October in Inria Lille by giving a talk on cybersecurity and animating workshops.

Iliana Fayolle and Tristan Coignon animated workshops for middle-school pupils during their Stage de découverte de 3ème at Inria Lille on December 20.

12 Scientific production

12.1 Major publications

- [1] Y. Al-Dhuraibi, F. Paraiso, N. Djarallah and P. Merle. 'Elasticity in Cloud Computing: State of the Art and Research Challenges'. In: *IEEE Transactions on Services Computing (TSC)* 11.2 (Mar. 2018), pp. 430–447. DOI: [10.1109/TSC.2017.2711009](https://doi.org/10.1109/TSC.2017.2711009). URL: <https://hal.inria.fr/hal-01529654>.
- [2] A. Geimer, M. Vergnolle, F. Recoules, L.-A. Daniel, S. Bardin and C. Maurice. 'A Systematic Evaluation of Automated Tools for Side-Channel Vulnerabilities Detection in Cryptographic Libraries'. In: *CCS 2023 - ACM SIGSAC Conference on Computer and Communications Security*. Copenhagen, Denmark: ACM, 2023, pp. 1690–1704. DOI: [10.1145/3576915.3623112](https://doi.org/10.1145/3576915.3623112). URL: <https://inria.hal.science/hal-04474774>.
- [3] S. Habchi, N. Moha and R. Rouvoy. 'Android Code Smells: From Introduction to Refactoring'. In: *Journal of Systems and Software* 177 (July 2021). DOI: [10.1016/j.jss.2021.110964](https://doi.org/10.1016/j.jss.2021.110964). URL: <https://hal.inria.fr/hal-03182277>.

- [4] T. Laor, N. Mehanna, A. Durey, V. Dyadyuk, P. Laperdrix, C. Maurice, Y. Oren, R. Rouvoy, W. Rudametkin and Y. Yarom. ‘DRAWNAPART: A Device Identification Technique based on Remote GPU Fingerprinting’. In: *Network and Distributed System Security Symposium*. San Diego, United States, 27th Feb. 2022. DOI: [10.14722/ndss.2022.24093](https://doi.org/10.14722/ndss.2022.24093). URL: <https://inria.hal.science/hal-03526240>.
- [5] P. Laperdrix, N. Bielova, B. Baudry and G. Avoine. ‘Browser Fingerprinting: A Survey’. In: *ACM Transaction on the Web* 14.2 (Apr. 2020), pp. 1–33. DOI: [10.1145/3386040](https://doi.org/10.1145/3386040). URL: <https://hal.archives-ouvertes.fr/hal-02864872>.
- [6] L. Meftah, R. Rouvoy and I. Chrisment. ‘FOUGERE: User-Centric Location Privacy in Mobile Crowdsourcing Apps’. In: *19th IFIP International Conference on Distributed Applications and Interoperable Systems (DAIS)*. Ed. by J. Pereira and L. Ricci. Vol. LNCS-11534. Distributed Applications and Interoperable Systems. Kongens Lyngby, Denmark: Springer International Publishing, 2019, pp. 116–132. DOI: [10.1007/978-3-030-22496-7_8](https://doi.org/10.1007/978-3-030-22496-7_8). URL: <https://hal.inria.fr/hal-02121311>.
- [7] A. Metzger, C. Quinton, Z. Adam-Mann, L. Baresi and K. Pohl. ‘Feature Model-Guided Online Reinforcement Learning for Self-Adaptive Services’. In: *International Conference on Service Oriented Computing*. Proceedings of the 18th International Conference on Service-Oriented Computing. Dubai, United Arab Emirates, Dec. 2020. URL: <https://hal.archives-ouvertes.fr/hal-02982029>.
- [8] Z. Ournani, M. C. Belgaid, R. Rouvoy, P. Rust and J. Penhoat. ‘Evaluating the Impact of Java Virtual Machines on Energy Consumption’. In: *15th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)*. Bari, Italy, 11th Oct. 2021. URL: <https://hal.inria.fr/hal-03275286>.
- [9] E. Stachtari, A. Mavridou, P. Katsaros, S. Bliudze and J. Sifakis. ‘Early validation of system requirements and design through correctness-by-construction’. In: *Journal of Systems and Software* 145 (2018), pp. 52–78. DOI: [10.1016/j.jss.2018.07.053](https://doi.org/10.1016/j.jss.2018.07.053). URL: <https://hal.archives-ouvertes.fr/hal-01873999>.
- [10] Y. Wang, P. Bourhis, R. Rouvoy and P. Royer. ‘Challenges and Opportunities in Automating DBMS: A Qualitative Study’. In: *ASE ’24: 39th IEEE/ACM International Conference on Automated Software Engineering*. Sacramento - California, United States: ACM; ACM, 24th Oct. 2024, pp. 2013–2023. DOI: [10.1145/3691620.3695264](https://doi.org/10.1145/3691620.3695264). URL: <https://hal.science/hal-04771192>.

12.2 Publications of the year

International journals

- [11] P. Ayoub, A. Hernandez, R. Cayre, A. Francillon and C. Maurice. ‘PhaseSCA: Exploiting Phase-Modulated Emanations in Side Channels’. In: *IACR Transactions on Cryptographic Hardware and Embedded Systems* (2025). URL: <https://hal.science/hal-04726109>. In press (cit. on p. 14).
- [12] B. Danglot, J.-R. Falleri and R. Rouvoy. ‘Can We Spot Energy Regressions using Developers Tests?’ In: *Empirical Software Engineering* 29.121 (25th July 2024). DOI: [10.1007/s10664-023-10429-1](https://doi.org/10.1007/s10664-023-10429-1). URL: <https://hal.science/hal-04286574>. In press.
- [13] G. Fieni, D. R. Acero, P. Rust and R. Rouvoy. ‘PowerAPI: A Python framework for building software-defined power meters’. In: *Journal of Open Source Software* 9.98 (4th June 2024), p. 6670. DOI: [10.21105/joss.06670](https://doi.org/10.21105/joss.06670). URL: <https://hal.science/hal-04601379>.
- [14] P. Jacquet, T. Ledoux and R. Rouvoy. ‘ScroogeVM: Boosting Cloud Resource Utilization with Dynamic Oversubscription’. In: *IEEE Transactions on Sustainable Computing* (2024), pp. 1–13. DOI: [10.1109/TSUSC.2024.3369333](https://doi.org/10.1109/TSUSC.2024.3369333). URL: <https://hal.science/hal-04466538> (cit. on p. 11).
- [15] E. Molina, M. Fiacchini, A. Goarant, R. Raes, S. Cerf and B. Robu. ‘Application of a predictive method to protect privacy of mobility data’. In: *Control Engineering Practice* 156 (Mar. 2025), p. 106223. DOI: [10.1016/j.conengprac.2024.106223](https://doi.org/10.1016/j.conengprac.2024.106223). URL: <https://hal.science/hal-04885266>.

- [16] M. Rahmani, S. Crumeyrolle, N. Allegri-Martiny, A. Taherkordi and R. Rouvoy. ‘PmForecast: Leveraging Temporal LSTM to Deliver In situ Air Quality Predictions’. In: *Environmental Science and Pollution Research* (10th Aug. 2024). DOI: [10.1007/s11356-024-34623-w](https://doi.org/10.1007/s11356-024-34623-w). URL: <https://hal.science/hal-04847894> (cit. on p. 13).
- [17] B. Sanwouo, C. Quinton and R. Rouvoy. ‘TS-Pothole: Automated Imputation of Missing Values in Univariate Time Series’. In: *Neural Computing and Applications* (30th Sept. 2024). URL: <https://hal.science/hal-04703773> (cit. on p. 18).

Invited conferences

- [18] V. Roca and P. Laperdrix. ‘Web, smartphone, AdTech: the privacy viewpoint’. In: Winter school 2024 of the PEPR Cybersecurity. Autrans, France, 29th Jan. 2024. URL: <https://inria.hal.science/hal-04550725> (cit. on p. 21).

International peer-reviewed conferences

- [19] A. Amarilli, P. Bourhis, F. Capelli and M. Monet. ‘Ranked Enumeration for MSO on Trees via Knowledge Compilation’. In: *27th International Conference on Database Theory (ICDT 2024)*. International Conference on Database Theory (ICDT 2024). Vol. 290. 25. Paestum, Italy, 14th Mar. 2024, 5:1–25:18. DOI: [10.4230/LIPIcs.ICDT.2024.25](https://doi.org/10.4230/LIPIcs.ICDT.2024.25). URL: <https://inria.hal.science/hal-04377344>.
- [20] P. Ayoub, R. Cayre, A. Francillon and C. Maurice. ‘BlueScream: Screaming Channels on Bluetooth Low Energy’. In: *Proceedings of the 40th Annual Computer Security Applications Conference (ACSAC '24)*. 40th Annual Computer Security Applications Conference (ACSAC '24). Waikiki, Honolulu, Hawaii, United States, 9th Dec. 2024. URL: <https://hal.science/hal-04725668> (cit. on p. 14).
- [21] A. Bonvoisin, C. Quinton and R. Rouvoy. ‘Understanding the Performance-Energy Tradeoffs of Object-Relational Mapping Frameworks’. In: SANER'24 - 31th IEEE International Conference on Software Analysis, Evolution and Reengineering. Rovaniemi, Finland: IEEE, 12th Mar. 2024, p. 11. URL: <https://inria.hal.science/hal-04401643> (cit. on p. 16).
- [22] P. Bourhis, A. Boussidan, C. Fournial and P. Gambette. ‘Detecting semantic or structural similarities for theater play comparison’. In: *17e Journées internationales d'Analyse statistique des Données Textuelles*. JADT 2024. Proceedings of the 17th International Conference on the Statistical Analysis of Textual Data. Bruxelles (BE), Belgium, 2024, pp. 139–148. URL: <https://univ-eiffel.hal.science/hal-04722464>.
- [23] T. Coignion, C. Quinton and R. Rouvoy. ‘A Performance Study of LLM-Generated Code on Leetcode’. In: EASE'24 - 28th International Conference on Evaluation and Assessment in Software Engineering. Proceedings of the 28th International Conference on Evaluation and Assessment in Software Engineering (EASE'24). Salerno, Italy, 18th June 2024. DOI: [10.1145/3661167.3661221](https://doi.org/10.1145/3661167.3661221). URL: <https://hal.science/hal-04525620> (cit. on p. 16).
- [24] B. Diarra, K. Guillouard, M. Ouzzif, P. Merle and J.-B. Stefani. ‘In-depth analysis of Kubernetes manifest verification tools for robust CNF deployment’. In: ICIN 2024 - Conference on Innovation in Clouds, Internet and Networks. Paris, France, 2024, pp. 1–8. URL: <https://inria.hal.science/hal-04421758> (cit. on p. 15).
- [25] S. Farhat, S. Bliudze, L. Duchien and O. Kouchnarenko. ‘Composing Run-Time Variability Models’. In: *Software Engineering and Formal Methods*. SEFM 2024. Vol. 15280. Lecture Notes in Computer Science. Aveiro, Portugal: Springer Nature Switzerland, 26th Nov. 2024, pp. 234–252. DOI: [10.1007/978-3-031-77382-2_14](https://doi.org/10.1007/978-3-031-77382-2_14). URL: <https://hal.science/hal-04855648> (cit. on p. 11).
- [26] I. Fayolle, J. Wichelmann, A. Köhl, W. Rudametkin, T. Eisenbarth and C. Maurice. ‘Semi-Automated and Easily Interpretable Side-Channel Analysis for Modern JavaScript’. In: CANS 2024 - 23rd International Conference on Cryptology And Network Security. Cambridge, United Kingdom, 2024, pp. 1–22. URL: <https://hal.science/hal-04652991> (cit. on p. 15).

- [27] I. Fouad, C. Santos and P. Laperdrix. ‘The Devil is in the Details: Detection, Measurement and Lawfulness of Server-Side Tracking on the Web’. In: *Privacy Enhancing Technologies*. 24th Privacy Enhancing Technologies Symposium (PETS 2024). Vol. 2024. 4. Bristol, United Kingdom, 15th July 2024. URL: <https://hal.science/hal-04617727>.
- [28] É. Guégain, T. Simon, A. Rahier and R. Rouvoy. ‘Managing Uncertainties in ICT Services Life Cycle Assessment using Fuzzy Logic’. In: *ICT4S’24 - 10th International Conference on ICT for Sustainability*. Vol. Proceedings of the 10th International Conference on ICT for Sustainability (ICT4S’24). Stockholom, Sweden: IEEE, 24th June 2024. URL: <https://hal.science/hal-04532041> (cit. on pp. 12, 14).
- [29] Q. Guilloteau, S. Cerf, R. Bleuse, B. Robu and E. Rutten. ‘Under Control: A Control Theory Introduction for Computer Scientists’. In: *ACSOS 2024 - 5th IEEE International Conference on Autonomic Computing and Self-Organizing Systems (ACSOS 2024)*. Aarhus, Denmark: IEEE, 2024, pp. 1–10. URL: <https://hal.science/hal-04666859>.
- [30] Q. Guilloteau, S. Cerf, E. Rutten, R. Bleuse and B. Robu. ‘Tutorial: "Under-Control: A Control Theory Introduction for Computer Scientists"’. In: *FlexScience 2024 - 14th workshop on AI and Scientific Computing at Scale using Flexible Computing Infrastructures*. Pise, Italy, 3rd June 2024, pp. 1–2. URL: <https://hal.science/hal-04602298>.
- [31] M. Huyghe, C. Quinton and W. Rudametkin. ‘Taming the Variability of Browser Fingerprints’. In: *SPLC’24 - 28th ACM International Systems and Software Product Lines Conference*. Luxembourg, Luxembourg, 2nd Sept. 2024, pp. 1–6. URL: <https://hal.science/hal-04622269>.
- [32] P. Jacquet, T. Ledoux and R. Rouvoy. ‘SLACKVM: Packing Virtual Machines in Oversubscribed Cloud Infrastructures’. In: *2024 CLUSTER - IEEE International Conference on Cluster Computing*. Vol. Proceedings of the IEEE International Conference on Cluster Computing (CLUSTER). Kobe, Japan: IEEE, 2024, pp. 1–12. DOI: [10.1109/CLUSTER59578.2024.00024](https://doi.org/10.1109/CLUSTER59578.2024.00024). URL: <https://hal.science/hal-04636648> (cit. on p. 11).
- [33] P. Jacquet, T. Ledoux and R. Rouvoy. ‘SweetspotVM: Oversubscribing CPU without Sacrificing VM Performance’. In: *Proceedings of the 24th IEEE/ACM international Symposium on Cluster, Cloud and Internet Computing (CCGrid’24)*. CCGrid’24 - 24th IEEE/ACM international Symposium on Cluster, Cloud and Internet Computing. Philadelphia, United States: IEEE, 2024, pp. 1–10. DOI: [10.1109/CCGrid59990.2024.00026](https://doi.org/10.1109/CCGrid59990.2024.00026). URL: <https://hal.science/hal-04454043> (cit. on p. 11).
- [34] N. Mehanna, W. Rudametkin, P. Laperdrix and A. Vastel. ‘Free Proxies Unmasked: A Vulnerability and Longitudinal Analysis of Free Proxy Services’. In: *Workshop on Measurements, Attacks, and Defenses for the Web (MADWeb’24)*. MADWeb 2024 - Workshop on Measurements, Attacks, and Defenses for the Web. San Diego (CA), United States, 2024, pp. 1–12. DOI: [10.14722/madweb.2024.23035](https://doi.org/10.14722/madweb.2024.23035). URL: <https://hal.science/hal-04489166> (cit. on pp. 7, 10, 17).
- [35] R. Pagano, S. Cerf, B. Robu, Q. Guilloteau, R. Bleuse and E. Rutten. ‘Making Control in High Performance Computing for Overload Avoidance Adaptive in Time and Job Size’. In: *CCTA 2024 - 8th IEEE Conference on Control Technology and Applications*. Newcastle Upon Tyne, United Kingdom: IEEE, 2024, pp. 1–8. URL: <https://hal.science/hal-04669743>.
- [36] R. Raes, A. Luxey-Bitri, R. Rouvoy, D. Frey and F. Taïani. ‘Venice: eschewing the cloud by leveraging local communication channels’. In: *ICT4S 2024 - International Conference on Information and Communications Technology for Sustainability*. Stockholm, Sweden, 2024, pp. 1–4. URL: <https://hal.science/hal-04576743>.
- [37] R. Raes, O. Ruas, A. Luxey-Bitri and R. Rouvoy. ‘Compact Storage of Data Streams in Mobile Devices’. In: *Distributed Applications and Interoperable Systems*. DAIS’24 - 24th International Conference on Distributed Applications and Interoperable Systems. Vol. 14677. Proceedings of the 24th International Conference on Distributed Applications and Interoperable Systems (DAIS’24). Groningen, Netherlands: LNCS, 17th June 2024, p. 45. URL: <https://hal.science/hal-04535716> (cit. on p. 7).

- [38] T. Simon, D. Ekchajzer, A. Berthelot, E. Fourboul, S. Rince and R. Rouvoy. ‘BoaviztAPI: a bottom-up model to assess the environmental impacts of cloud services’. In: HotCarbon’24 - 3rd Workshop on Sustainable Computer Systems. Santa Cruz, United States, 9th July 2024. URL: <https://hal.science/hal-04621947> (cit. on pp. 12, 14).
- [39] Q. Wang, H. Dai, Y. Zhao, M. Zhang and S. Bliudze. ‘Enabling Behaviour Tree Verification via a Translation to BIP’. In: *Lecture Notes in Computer Science*. Formal Aspects of Component Software. FACS 2024. Vol. LNCS - 15189. Lecture Notes in Computer Science. Milan, Italy: Springer Nature Switzerland, 8th Sept. 2024, pp. 3–20. DOI: 10.1007/978-3-031-71261-6_1. URL: <https://hal.science/hal-04860800>.
- [40] Y. Wang, P. Bourhis, R. Rouvoy and P. Royer. ‘Challenges and Opportunities in Automating DBMS: A Qualitative Study’. In: ASE ’24: 39th IEEE/ACM International Conference on Automated Software Engineering. ASE2024 Industry Track. Sacramento - Californie, United States: ACM; ACM, 24th Oct. 2024, pp. 2013–2023. DOI: 10.1145/3691620.3695264. URL: <https://hal.science/hal-04771192> (cit. on p. 14).

Doctoral dissertations and habilitation theses

- [41] S. Bliudze. ‘Contributions to the Rigorous Design of Concurrent Component-Based Software and Systems Using BIP’. Université de Lille, 8th Mar. 2024. URL: <https://hal.univ-lille.fr/tel-04567479> (cit. on p. 25).
- [42] P. Bourhis. ‘Dialog Between Logic, Trees and Circuits’. Université de Lille, 18th Sept. 2024. URL: <https://hal.science/tel-04743648> (cit. on p. 25).
- [43] S. Farhat. ‘Safe Dynamic Reconfiguration of Applications with Features: Run-time Coordination of Reconfiguration Requests’. Université de Lille, 11th July 2024. URL: <https://theses.hal.science/tel-04657867> (cit. on pp. 10, 25).
- [44] J. L. Intumwayase. ‘Exploring the Impact of the User’s Browsing Environment on Web Privacy’. Université de Lille, 5th Dec. 2024. URL: <https://hal.science/tel-04837269> (cit. on pp. 11, 25).
- [45] P. Jacquet. ‘Enhancing IaaS Consolidation with Resource Oversubscription’. Université de Lille, 19th July 2024. URL: <https://theses.hal.science/tel-04685771> (cit. on pp. 7, 11, 25).
- [46] N. Mehanna. ‘Intrinsically Inseparable: Investigating Novel Tracking Practices and Assessing the Carbon Footprint of Ads’. Université de Lille, 30th May 2024. URL: <https://theses.hal.science/tel-04647367> (cit. on pp. 10, 17, 25).
- [47] C. Quinton. ‘Evolving, Adapting and Optimizing Configurable Software Systems’. University of Lille, 26th Jan. 2024. URL: <https://theses.hal.science/tel-04424609> (cit. on p. 25).
- [48] M. Rahmani. ‘Next-Generation Air Pollution Forecasting: Integrating AI, Spatiotemporal Dynamics, and Privacy-Ensuring Approaches for Urban Areas’. Université de Lille, 11th Dec. 2024. URL: <https://theses.hal.science/tel-04851216> (cit. on pp. 13, 25).
- [49] T. Simon. ‘Software Ecodesign: Estimating and Reducing Software Environmental Footprint’. Université de Lille, 6th Dec. 2024. URL: <https://theses.hal.science/tel-04875910> (cit. on pp. 7, 12, 25).

Other scientific publications

- [50] T. Collignon and A. van Kempen. ‘Using Control Theory to Reduce Disk Congestion Caused by Unpredictable I/O in Cloud Computing’. In: Performance and Scalability of Storage Systems (Per3S). Ed. by S. Cerf. Paris, France, 28th May 2024. URL: <https://hal.science/hal-04679808> (cit. on p. 15).
- [51] R. Raes, R. Rouvoy, A. Luxey-Bitri, D. Romero, D. Frey and F. Taïani. ‘Eschewing the Cloud by leveraging Local Communication Channels: Angus & Bob want to share cool things’. In: ICT4S 2024 - International Conference on ICT for Sustainability. Stockholm, Sweden, 2024, pp. 1–1. URL: <https://hal.science/hal-04633670>.

Scientific popularization

- [52] D. Monniaux and C. Maurice. ‘Les arcanes des processeurs mis à l’épreuve’. In: *La Recherche* 578 (July 2024). URL: <https://hal.science/hal-04643335> (cit. on p. 27).

12.3 Cited publications

- [53] P. Ayoub. ‘Emissions électromagnétique compromettantes : fuites par canaux auxiliaires dans les périphériques embarqués’. PhD thesis. Sorbonne Université, Dec. 2024 (cit. on p. 14).
- [54] A. Basu, B. Bensalem, M. Bozga, J. Combaz, M. Jaber, T.-H. Nguyen and J. Sifakis. ‘Rigorous Component-Based System Design Using the BIP Framework’. In: *IEEE Software* 28.3 (2011), pp. 41–48. DOI: [10.1109/MS.2011.27](https://doi.org/10.1109/MS.2011.27) (cit. on p. 6).
- [55] M. Bienvenu and P. Bourhis. ‘Mixed-World Reasoning with Existential Rules under Active-Domain Semantics’. In: *Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI)*. Macao, Macao SAR China, Aug. 2019. DOI: [10.24963/ijcai.2019/216](https://doi.org/10.24963/ijcai.2019/216). URL: <https://hal.archives-ouvertes.fr/hal-02342129> (cit. on p. 16).
- [56] S. Bliudze, A. Mavridou, R. Szymanek and A. Zolotukhina. ‘Exogenous coordination of concurrent software components with JavaBIP’. In: *Software: Practice and Experience* 47.11 (Nov. 2017), pp. 1801–1836. DOI: [10.1002/spe.2495](https://doi.org/10.1002/spe.2495) (cit. on p. 6).
- [57] P. Bourhis, M. Leclère, M.-L. Mugnier, S. Tison, F. Ulliana and L. Gallois. ‘Oblivious and Semi-Oblivious Boundedness for Existential Rules’. In: *IJCAI, 2019 - International Joint Conference on Artificial Intelligence*. Macao, China, Aug. 2019. URL: <https://hal-lirmm.ccsd.cnrs.fr/lirmm-02148142> (cit. on p. 16).
- [58] C. Calcagno, D. Distefano, J. Dubreil, D. Gabi, P. Hooimeijer, M. Luca, P. O’Hearn, e. Papakonstantinou Iren, J. Purbrick and D. Rodriguez. ‘Moving Fast with Software Verification’. In: *NASA Formal Methods*. Ed. by K. Havelund, G. Holzmann and R. Joshi. Springer International Publishing, 2015, pp. 3–11 (cit. on p. 6).
- [59] J.-R. Falleri, F. Morandat, X. Blanc, M. Martinez and M. Monperrus. ‘Fine-grained and Accurate Source Code Differencing’. In: *Proceedings of the International Conference on Automated Software Engineering*. update for oadoi on Nov 02 2018. Västerås, Sweden, 2014, pp. 313–324. DOI: [10.1145/2642937.2642982](https://doi.org/10.1145/2642937.2642982). URL: <https://hal.science/hal-01054552> (cit. on p. 8).
- [60] S. Farhat, S. Bliudze and L. Duchien. *Safe Dynamic Reconfiguration of Concurrent Component-based Applications*. ICSA 2022 - 19th IEEE International Conference on Software Architecture. Poster. Mar. 2022. URL: <https://hal.science/hal-03585767> (cit. on p. 11).
- [61] S. Farhat, S. Bliudze, L. Duchien and O. Kouchnarenko. ‘Toward Run-time Coordination of Reconfiguration Requests in Cloud Computing Systems’. In: *COORDINATION 2023 - 25th International Conference on Coordination Models and Languages*. Ed. by S.-S. Jongmans. Part 6: Run-Time Changes. Lisbon, Portugal: Springer Nature Switzerland, June 2023, pp. 271–291. URL: <https://inria.hal.science/hal-04098074> (cit. on p. 11).
- [62] E. Guégain. ‘Configuration-driven Software Optimization’. Theses. Université de Lille, Sept. 2023. URL: <https://theses.hal.science/tel-04350545> (cit. on p. 17).
- [63] E. Guégain, A. Taherkordi and C. Quinton. ‘Configuration Optimization with Limited Functional Impact’. In: *Advanced Information Systems Engineering*. Advanced Information Systems Engineering. Zaragoza, Spain, June 2023. URL: <https://hal.science/hal-04034888> (cit. on p. 17).
- [64] E. Guégain, A. Taherkordi and C. Quinton. ‘ICO: A Platform for Optimizing Highly Configurable Systems’. In: *5th International Workshop on Automated and verifiable Software System Development - ASYDE 2023*. Kirchberg, Luxembourg, Sept. 2023. URL: <https://hal.science/hal-04213194> (cit. on p. 17).
- [65] É. Guégain, C. Quinton and R. Rouvoy. ‘On Reducing the Energy Consumption of Software Product Lines’. In: *SPLC ’21: 25th ACM International Systems and Software Product Line Conference*. Vol. A. Leicester, United Kingdom: ACM, Sept. 2021, pp. 89–99. DOI: [10.1145/3461001.3471142](https://doi.org/10.1145/3461001.3471142). URL: <https://hal.inria.fr/hal-03269168> (cit. on p. 17).

- [66] J. L. Intumwayase, I. Fouad, P. Laperdrix and R. Rouvoy. ‘UA-Radar: Exploring the Impact of User Agents on the Web’. In: *WPES '23: 22nd Workshop on Privacy in the Electronic Society*. Vol. Proceedings of the 22nd Workshop on Privacy in the Electronic Society. Copenhagen, Denmark, Nov. 2023. DOI: [10.1145/3603216.3624958](https://doi.org/10.1145/3603216.3624958). URL: <https://hal.science/hal-04271971> (cit. on p. 12).
- [67] V. Issarny, N. Georgantas, S. Hachem, A. Zarras, P. Vassiliadis, M. Autili, M. A. Gerosa and A. Ben Hamida. ‘Service-Oriented Middleware for the Future Internet: State of the Art and Research Directions’. In: *Journal of Internet Services and Applications* 2.1 (May 2011), pp. 23–45. DOI: [10.1007/s13174-011-0021-3](https://doi.org/10.1007/s13174-011-0021-3). URL: <http://dx.doi.org/10.1007/s13174-011-0021-3> (cit. on p. 4).
- [68] P. Jacquet, T. Ledoux and R. Rouvoy. ‘CloudFactory: An Open Toolkit to Generate Production-like Workloads for Cloud Infrastructures’. In: *Proceedings of the 11th IEEE International Conference on Cloud Engineering (IC2E)*. Boston, Massachusetts, United States: IEEE, Sept. 2023, p. 11. DOI: [10.1109/IC2E59103.2023.00017](https://doi.org/10.1109/IC2E59103.2023.00017). URL: <https://inria.hal.science/hal-04168667> (cit. on p. 11).
- [69] P. Jacquet, R. Rouvoy and T. Ledoux. ‘La chasse au gaspillage dans le cloud et les data centers’. In: *The Conversation France* (Jan. 2023). URL: <https://inria.hal.science/hal-04168662> (cit. on p. 11).
- [70] K. C. Kang, S. G. Cohen, J. A. Hess, W. E. Novak and A. S. Peterson. *Feature-Oriented Domain Analysis (FODA) Feasibility Study*. Tech. rep. Carnegie-Mellon University Software Engineering Institute, Nov. 1990 (cit. on p. 6).
- [71] T. Laor, N. Mehanna, A. Durey, V. Dyadyuk, P. Laperdrix, C. Maurice, Y. Oren, R. Rouvoy, W. Rudametkin and Y. Yarom. ‘DRAWNAPART: A Device Identification Technique based on Remote GPU Fingerprinting’. In: *Network and Distributed System Security Symposium*. San Diego, United States, Feb. 2022. DOI: [10.14722/ndss.2022.24093](https://doi.org/10.14722/ndss.2022.24093). URL: <https://inria.hal.science/hal-03526240> (cit. on pp. 10, 17).
- [72] P. Laperdrix, N. Mehanna, A. Durey and W. Rudametkin. ‘The Price to Play: a Privacy Analysis of Free and Paid Games in the Android Ecosystem’. In: *ACM Web Conference 2022*. Lyon, France, Apr. 2022. DOI: [10.1145/3485447.3512279](https://doi.org/10.1145/3485447.3512279). URL: <https://hal.science/hal-03559973> (cit. on pp. 10, 17).
- [73] R. de Lemos, H. Giese, H. A. Muller, M. Shaw, J. Andersson, L. Baresi, B. Becker, N. Bencomo, Y. Brun, B. Cukic, R. Desmarais, S. Dustdar, G. Engels, K. Geihs, K. M. Göschka, A. Gorla, V. Grassi, P. Inverardi, G. Karsai, J. Kramer, M. Litoiu, A. Lopes, J. Magee, S. Malek, S. Mankovskii, R. Mirandola, J. Mylopoulos, O. Nierstrasz, M. Pezzè, C. Prehofs, W. Schäfer, R. Schlichting, B. Schmerl, D. B. Smith, J. P. Sousa, G. Tamura, L. Tahvildari, N. M. Villegas, T. Vogel, D. Weyns, K. Wong and J. Wuttke. ‘Software Engineering for Self-Adaptive Systems: A Second Research Roadmap’. In: *Software Engineering for Self-Adaptive Systems*. Ed. by R. de Lemos, H. Giese, H. Müller and M. Shaw. Vol. 7475. Dagstuhl Seminar Proceedings. Springer, 2013, pp. 1–26. URL: <http://hal.inria.fr/inria-00638157> (cit. on p. 4).
- [74] T. Medjkoune, O. Goga and J. Sénéchal. ‘Marketing to Children Through Online Targeted Advertising: Targeting Mechanisms and Legal Aspects’. In: *The 2023 ACM SIGSAC Conference on Computer and Communications Security*. Copenhagen, Denmark: ACM, Nov. 2023. DOI: [10.1145/3576915.3623172](https://doi.org/10.1145/3576915.3623172). URL: <https://hal.univ-lille.fr/hal-04311598> (cit. on p. 7).
- [75] N. Mehanna and W. Rudametkin. ‘Caught in the Game: On the History and Evolution of Web Browser Gaming’. In: *Companion Proceedings of the ACM Web Conference 2023 (WWW '23 Companion)*. Austin (TX), United States, Apr. 2023, pp. 1–9. URL: <https://hal.science/hal-04084097> (cit. on pp. 10, 17).
- [76] A. Metzger, C. Quinton, Z. Adam-Mann, L. Baresi and K. Pohl. ‘Feature Model-Guided Online Reinforcement Learning for Self-Adaptive Services’. In: *International Conference on Service Oriented Computing*. Proceedings of the 18th International Conference on Service-Oriented Computing. Dubai, United Arab Emirates, Dec. 2020. URL: <https://hal.archives-ouvertes.fr/hal-02982029> (cit. on p. 17).

- [77] H. Monfleur and P. Merle. ‘Towards Concern-Oriented Microservice Architecture’. In: *Proceedings of 2023 International Conference on Microservices*. Universita di Pisa and Microservices Community. Pisa, Italy, Oct. 2023. URL: <https://inria.hal.science/hal-04201189> (cit. on p. 18).
- [78] C. Newcombe, T. Rath, F. Zhang, B. Munteanu, M. Brooker and M. Deardeuff. ‘How Amazon Web Services Uses Formal Methods’. In: *Commun. ACM* 58.4 (Mar. 2015), pp. 66–73. DOI: [10.1145/2699417](https://doi.org/10.1145/2699417). URL: <https://doi.org/10.1145/2699417> (cit. on p. 6).
- [79] M. Salehie and L. Tahvildari. ‘Self-adaptive software: Landscape and research challenges’. In: *ACM Transactions on Autonomous and Adaptive Systems* 4.2 (May 2009), 14:1–14:42. DOI: [10.1145/1516533.1516538](https://doi.org/10.1145/1516533.1516538). URL: <http://doi.acm.org/10.1145/1516533.1516538> (cit. on p. 4).
- [80] H. Schaffers and M. Sharpe. *Services in the Future Internet*. FP8 Expert Group, EU Directorate-General Information Society & Media. Apr. 2011. URL: http://www.eurosfair.prdd.fr/7pc/doc/1303227305_fp8_future_internet_28_02_2011.pdf (cit. on p. 4).
- [81] L. Schubert, K. Jeffery and B. Neidecker-Lutz. *A Roadmap for Advanced Cloud Technologies under H2020*. Dec. 2012. URL: http://ec.europa.eu/information_society/newsroom/cf/dae/document.cfm?doc_id=2165 (cit. on p. 4).
- [82] J. Sifakis. ‘Rigorous System Design’. In: *Foundations and Trends® in Electronic Design Automation* 6.4 (2013), pp. 293–362. DOI: [10.1561/1000000034](https://doi.org/10.1561/1000000034). URL: <http://dx.doi.org/10.1561/1000000034> (cit. on p. 6).
- [83] T. Simon, P. Rust, R. Rouvoy and J. Penhoat. ‘Uncovering the Environmental Impact of Software Life Cycle’. In: *International Conference on Information and Communications Technology for Sustainability*. Rennes, France, June 2023. URL: <https://inria.hal.science/hal-04082263> (cit. on pp. 12, 14).
- [84] J. Urban. *NESSI Research Priorities for the next Framework Programme for Research and Technological Development FP8*. May 2011. URL: http://www.nessi-europe.eu/files/Docs/NESSI%5C%20SRA_update_May_2011_V1-0.pdf (cit. on p. 4).
- [85] P. Vila, P. Ganty, M. Guarnieri and B. Köpf. ‘CacheQuery: learning replacement policies from hardware caches’. In: *PLDI*. 2020 (cit. on p. 5).