

RESEARCH CENTRE

**Inria Centre
at the University of Lille**

IN PARTNERSHIP WITH:
Université de Lille, CNRS

2023

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
4.1 Introduction	7
4.2 Distributed software services for the retail industry	7
4.3 Distributed software services for the digital home	7
5 Social and environmental responsibility	7
5.1 Impact of research results	7
6 Highlights of the year	7
6.1 Awards	8
7 New software, platforms, open data	8
7.1 New software	8
7.1.1 amiunique	8
7.1.2 APISENSE	9
7.1.3 cloudnet	9
7.1.4 PowerAPI	10
8 New results	10
8.1 Design of correct-by-construction self-adaptive cloud applications using formal methods	10
8.2 Improving Web User Privacy Through Content Blocking	11
8.3 Optimization of Software Systems by Configuration	11
9 Bilateral contracts and grants with industry	12
10 Partnerships and cooperations	13
10.1 International initiatives	13
10.1.1 Associate Teams in the framework of an Inria International Lab or in the framework of an Inria International Program	13
10.2 European initiatives	14
10.2.1 Other european programs/initiatives	14
10.3 National initiatives	14
10.3.1 ANR	14
10.3.2 PEPR and PTCC	16
10.4 Regional initiatives	17
10.5 Inria initiatives	18
11 Dissemination	19
11.1 Promoting scientific activities	19
11.1.1 Scientific events: organisation	19
11.1.2 Scientific events: selection	19
11.1.3 Journal	19
11.1.4 Invited talks	20

11.1.5 Leadership within the scientific community	20
11.1.6 Scientific expertise	20
11.1.7 Research administration	21
11.2 Teaching - Supervision - Juries	22
11.2.1 Teaching	22
11.2.2 Supervision	23
11.2.3 Juries	24
11.3 Popularization	25
11.3.1 Articles and contents	25
11.3.2 Education	25
12 Scientific production	25
12.1 Major publications	25
12.2 Publications of the year	26
12.3 Other	29
12.4 Cited publications	29

Project-Team SPIRALS

Creation of the Project-Team: 2015 January 01

Keywords

Computer sciences and digital sciences

- A1.3. – Distributed Systems
- 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, HDR]
- Philippe Merle [INRIA, Senior Researcher, HDR]

Faculty Members

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

Post-Doctoral Fellows

- Mohammed Chakib Belgaid [INRIA]
- Guillaume Fieni [INRIA]
- Imane Fouad [INRIA, Post-Doctoral Fellow]
- Walid Ghandour [UNIV LILLE, Post-Doctoral Fellow, until Feb 2023]

PhD Students

- Wassim Aroui [ORANGE, CIFRE, from Jun 2023]
- Alexandre Bonvoisin [INRIA]
- Tristan Coignon [INRIA]
- Thomas Collignon [Qarnot Computing, CIFRE, from Nov 2023]
- Boubacar Diarra [ORANGE, CIFRE]
- Belkis Djefal [INRIA, from Dec 2023]
- Salman Farhat [INRIA]
- Iliana Fayolle [UNIV LILLE, from Oct 2023]
- Romain Fouquet [INRIA, until Sep 2023]
- Antoine Geimer [UNIV RENNES I]
- Edouard Guégain [UNIV LILLE, until Sep 2023]
- Maxime Huyghe [UNIV LILLE]

- Jean Intumwayase [UNIV LILLE]
- Pierre Jacquet [INRIA]
- Naif Mehanna [UNIV LILLE]
- Hugo Monfleur [INRIA]
- Rémy Raes [INRIA, from Mar 2023]
- Maryam Rahmani [UNIV LILLE]
- Jakub Rozycki [INRIA, until Oct 2023]
- Thibault Simon [ORANGE, CIFRE]
- Yifan Wang [ORANGE, CIFRE]
- Jilai Xu [Inria]

Technical Staff

- Sihem Bouhenniche [INRIA, Engineer]
- Daniel Romero Acero [INRIA, Engineer, SED Permanent Engineer]
- Brell Peclard Sanwouo Chekam [UNIV LILLE, Engineer, from May 2023]
- Jiali Xu [INRIA, Engineer, until Feb 2023]

Interns and Apprentices

- Virginie Amand [INRIA, Apprentice, from Sep 2023]
- Virginie Amand [INRIA, Intern, from Jun 2023 until Aug 2023]
- Baptiste Demoussel [INRIA, Intern, until May 2023]
- Iliana Fayolle [UNIV LILLE, Intern, from Mar 2023 until Sep 2023]
- Arthur Goarant [CENTRALE LILLE, Intern, from Apr 2023 until Jun 2023]
- Nassim Guinet [INRIA, Intern, from Apr 2023 until Aug 2023]
- Kouds Halitim [INRIA, Intern, from Apr 2023 until Aug 2023]
- Laurine Leulliette [INRIA, Intern, from Apr 2023 until Aug 2023]
- Martin Raes [INRIA, Intern, from Jun 2023 until Aug 2023]
- Emie Vincent [Lycée Thérèse d'Avila, Intern, from Jun 2023 until Jun 2023]

Administrative Assistant

- Karine Lewandowski [INRIA]

Visiting Scientists

- Olga Kouchnarenko [UNIV FRANCHE-COMTE, until Feb 2023, Secondment (délégation) Inria]
- Juliette Sénéchal [UNIV LILLE, from Sep 2023, Secondment (délégation SHS) 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 [72], the NESSI Research Priorities for the next Framework Programme for Research and Technological Development FP8 [75], the Roadmap for Advanced Cloud Technologies under H2020 [73], and research roadmaps, such as [58, 61, 71].

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

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

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. [77] 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 [59]. 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 [67] and Facebook [52] 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 [74] 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 [47, 50] 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

4.1 Introduction

Although our research is general enough to be applied to many application domains, we currently focus on applications and distributed services for the retail industry and for the digital home. These two application domains are supported by a strong expertise in mobile computing and in cloud computing that are the two main target environments on which our research prototypes are built, for which we are recognized, and for which we have already established strong collaborations with the industrial ecosystem.

4.2 Distributed software services for the retail industry

This application domain is developed in relation with the **PICOM** (*Pôle de compétitivité Industries du Commerce*) cluster. We have established strong collaborations with local companies in the context of former funded projects, such as Cappucino and **Macchiato**, which focused on the development of a new generation of mobile computing platforms for e-commerce. We are also involved in the Datalyse and OCCIware funded projects that define cloud computing environments with applications for the retail industry. Finally, our activities in terms of crowd-sensing and data gathering on mobile devices with the APISENSE[®] platform share also applications for the retail industry.

4.3 Distributed software services for the digital home

We are developing new middleware solutions for the digital home, in particular through our long standing collaboration with Orange Labs. We are especially interested in developing energy management and saving solutions with the POWERAPI software library for distributed environments such the ones that equip digital homes. We are also working to bridge the gap between distributed services hosted on home gateways and distributed services hosted on the cloud to be able to smoothly transition between both environments. This work is especially conducted with the SALOON platform.

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.

In another line of research, the defended PhD thesis of Romain Fouquet [33] contributes to limit what is currently running in a modern browser. By correctly identifying what is truly needed on a web page to guarantee its core functionality, it becomes possible to block unneeded scripts like superficial third-party libraries, trackers or unwanted ads on the web. This leads to direct power savings as less data will be downloaded and less computing resources will be used to render a page.

6 Highlights of the year

In 2023, Spirals developed a collaboration with Juliette Sénéchal, Professor of Private Law at the University of Lille. The collaboration already existed before and was boosted in September 2023 with the secondment ("délégation Inria SHS") that Juliette Sénéchal obtained.

The purpose of the collaboration is to set up a dialog between the disciplines of computer science and law in order to better understand how they cross-fertilize each other, to study the consequences of new

law texts, such as the Digital Service Act, the Data Act, the AI act, on computer science, and, conversely, to study the consequences of the advancement of research in some computer science domains such as cybersecurity, privacy, data and service management, on law texts.

Since September 2023, in the context of this collaboration, the following actions have been carried on or started.

- Juliette Sénéchal, and her colleague Sabine Bernheim-Desvaux, organise a series of seminars on "Towards the emergence of a neuro-ethical law erected as a counterpoint to rights revisited by digital technology? Reflection based on consumer law". Eight workshops are planned in 2023 and 2024 with more than 40 invited speakers, among them persons from Inria (e.g. Natalia Bielova, Claude Castelluccia, Oana Goga, Ludovic Hoyet, Claude Kirchner).
- Juliette Sénéchal coordinated a submission to a call for expression of interest issued by the CNIL on the creation of training databases for machine learning systems. This submission was discussed and made with participants of Inria and of the PEPR Cybersecurity IPoP project.
- Juliette Sénéchal participated to the event "Digital sciences for sport in the society" that was organized by Inria Lille on 15 December and that was conducted in relation with students from Centrale Lille and Science Po Lille.
- A weekly seminar has been set up with the members of Spirals to discuss matters related to law and computer science.

6.1 Awards

- Clémentine Maurice and her co-authors received a most influential paper award from the DIMVA conference for the period 2014-2018 for their paper on a novel attack technique that relies on the execution time of flush instructions in CPU cache [56].
- Clémentine Maurice and her co-authors received the best paper runner-up award from the DIMVA 2023 conference for their paper on using power analysis for fingerprinting machines with identical hardware and software configurations [19].
- Thibault Simon won the jury prize in the contest Ma thèse en trois minutes that was awarded during the Orange Tech Days 2023².

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.

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:

²[Link to the announcement](#)

- 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/>

Authors: Pierre Laperdrix, Antonin Durey, Walter Rudametkin Ivey

Contact: Benoit Baudry

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>

Authors: Antoine Veuiller, Christophe Ribeiro, Julien Duribreux, Romain Rouvoy, Nicolas Haderer, Romain Sommerard, Lakhdar Meftah

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). Three languages are addressed for the moment with the modules:

* Cloudnet TOSCA toolbox for TOSCA including NFV description * cloudnet-hot for HOT (Heat Orchestration Template) from OpenStack * cloudnet-compose for 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: <http://powerapi.org>

Contact: Romain Rouvoy

Participants: Adel Nouredine, Loïc Huertas, Maxime Colmant, Romain Rouvoy, Mohammed Chakib Belgaid, Arthur D'azemar

8 New results

We highlight below three new major results that we obtained in 2023.

8.1 Design of correct-by-construction self-adaptive cloud applications using formal methods

Participants: Simon Bliudze, Trinh Le Khan, Philippe Merle.

In the context of his PhD defended in January, Trinh Le Khanh obtained new results in the domain of correctness-by-construction of self-adaptive applications [35].

Correctly coordinating access to cloud resources across concurrent cloud software components is essential to ensure that they satisfy user and system requirements and avoid operational faults and deadlocks. Cloud systems must be able to self-adapt to changes at runtime without interruption. Traditional

approaches do not separate the code of the component computations from their coordination, making it hard to debug and maintain. Changes in coordination policies not only require reprogramming the components but also affect the other components that interact with them.

In this work, we provide solutions to ensure that concurrent cloud application entities have the correct access to cloud resources with three main contributions. (1) We introduce an ontology-driven specification approach to propose NaturalBIP—a language for specifying functional requirements. This ontology precisely defines concepts and their relationships in a specific domain. Then, the specifications are written in pseudo-natural templates with placeholders restricted by ontology concepts. (2) We provide a compiler to analyze and translate the specifications written in NaturalBIP language into JavaBIP artifacts (i.e., JavaBIP GlueBuilder, data transfers, and safety properties) and BIP connectors. (3) We extend OCCIware with coordination capabilities using JavaBIP. We leverage the Finite State Machine (FSM) specification in OCCIware design to specify the component’s behavior. Then, the coordination between them is established by JavaBIP generated from the NaturalBIP compiler. We also compute BIP models from the BIP connectors and the configuration model to verify the deadlock-free property using iFinder, a tool for the compositional detection of deadlocks at design time. With these contributions, we provide a toolchain to develop correct-by-construction self-adaptive cloud applications and conclude this thesis by presenting future perspectives to improve this work.

8.2 Improving Web User Privacy Through Content Blocking

Participants: Romain Fouquet, Pierre Laperdrix, Romain Rouvoy.

In the context of his PhD defended in May, Romain Fouquet obtained new results in the domain of content blocking for improving privacy on the web [33].

The web has seen steady growth since its inception in 1990. Along with this constant production of content, user tracking has appeared early and seen continuous development. Users are thus followed across websites and their behavior observed on individual web pages. To combat user tracking, browser vendors and extension developers have proposed different strategies, ranging from site isolation to asking the user before using sensitive features. In this work, we focus on one of these client-side privacy protection strategies: content blocking. Content blocking consists in preventing unwanted parts of web pages from being downloaded or executed, thus protecting the user from them. Existing content blocking tools mostly rely on filter lists which specify what parts of web pages to block. They however suffer from several issues, including incomplete coverage and being unable to target certain kinds of resources.

We present four contributions for improving user privacy by modifying page content: (1) We measure the dependency on common web page elements on JavaScript and how much blocking JavaScript can improve user privacy. We find that 43% of web pages from our sample do not strictly depend on JavaScript and that 67% of pages are likely to be usable when caring only about the main page section. (2) Building on the acquired knowledge of page breakage when blocking JavaScript, we design a set of repairs to repair common page breakage types. We introduce the concept of User Browsing Intent (UBI) and, focusing on the ‘read-only’ UBI, we measure how much these repairs are useful in the case of this UBI. (3) We propose a server-side system to substitute interface page elements usually relying on JavaScript with noscript alternatives, and discuss the benefits of this replacement in particular in terms of device energy savings and security. (4) We devise a signature scheme to generate robust signatures of JavaScript functions, and detect the bundling of tracking functions from known tracking libraries with functional code. We find that 22.7% of domains in our sample bundle such tracking functions with functional code, effectively circumventing existing blocking tools. We propose a technique for blocking these tracking functions while preserving functional code.

This work led to the following publications: [15], [55].

8.3 Optimization of Software Systems by Configuration

Participants: Edouard Guégain, Clément Quinton.

In the context of his PhD defended in September, Edouard Guégain obtained new results in the domain of the optimization of large software systems [34].

The field of software engineering is rapidly evolving, exposing practitioners to a growing collection of tools, languages, frameworks, and paradigms. Thus, designing a new software system consists in selecting components from this collection, which is akin to creating a configuration. Currently, the criterion to configure such systems often revolves around ease of development, which leads to oversized, power-hungry bloatware. This paradigm is not aligned with frugal or environmental concerns.

This work looks into the ability to leverage the configuration of software systems to optimize their performance. In particular, a specific focus is made on the energy consumption and the size of software systems. A prerequisite to optimizing a system is to understand its current performance. To gain a comprehensive understanding, this work begins with a thorough analysis of a configurable software system, considering multiple performance indicators. This analysis reveals that different configurations indeed yield varying performances, and these variations can be attributed to specific options. Leveraging insights gained from performance analysis enables the creation of high-performance configurations for this system. Furthermore, certain performance indicators proved correlated across configurations. Consequently, the optimization objective can be simplified by ignoring redundant performance indicators. The creation of optimized configurations of this system was performed manually, which is only possible for small configuration spaces. To address larger configuration spaces, we propose an algorithm that outlines how to evaluate the performance of each option and subsequently improves a given configuration based on this performance data. However, optimizing a configuration by selecting high-performance options brought out limitations, as options can interact with each other: in some situations, pairing high-performance options may result in subpar performances. Similarly, low-performance options can prove unexpectedly efficient when paired together. We thus designed the algorithm to leverage such interactions. Applying this algorithm to a large set of configurations showed that most of them can reach near-optimal performances, with only a limited set of modifications. However, performance constraints are not limited to a single performance indicator. Depending on the context, the energy consumption of a system may not be the single most impactful indicator to optimize for environmental concerns. Thus, the optimization algorithm must be generalized to accommodate multiple performance indicators. This improved algorithm is validated on a pair of performance indicators: the execution time and the software size. The key finding from this validation is that half of the assessed configurations can reach a local optimum by adjusting just one setting. To automate configuration optimization, this algorithm is implemented as an open-source tool called the ICO tool suite.

This work led to the following publications: [23, 31], [57].

9 Bilateral contracts and grants with industry

ip-label

Participants: Romain Rouvoy (*contact person*).

A software exploitation license (2014–ongoing) of the APISENSE[®] crowd-sensing platform has been sold to the **ip-label** company. They use this platform as a solution to monitor the quality of the GSM signal in the wild. The objective is to provide developers and stakeholders with a feedback on the quality of experience of GSM connection depending on their location.

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 takes place in the context of this collaboration.

First results have been published in [30] and Thibault Simon won the jury prize in the contest Ma thèse en trois minutes that was awarded during the Orange Tech Days 2023³.

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.

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 (see Section 10.5) take place in the context of this collaboration.

10 Partnerships and cooperations

10.1 International initiatives

10.1.1 Associate Teams in the framework of an Inria International Lab or in the framework of an Inria International Program

SusAI

Title: Towards a Sustainable Artificial Intelligence

Duration: 2021 -> 2023

Coordinator: Felipe Tobar (ftobar@dim.uchile.cl)

Partners:

- Centro de Modelamiento Matemático Santiago de Chile (Chile)

Inria contact: Romain Rouvoy

³[Link to the announcement](#)

Summary: The SusAIIn associate-team is composed of a pluridisciplinary group of researchers across France, Chile and Brazil to address the challenge of reducing the power consumption of artificial intelligence algorithms deployed in the context of high performance computing. By combining the unique scientific expertise on this critical topic for the environment, we intend to leverage the existing hardware and software infrastructures to deliver new tools and recommendations for the scientific community on the appropriate design and deployment of sustainable AI algorithms.

10.2 European initiatives

10.2.1 Other european programs/initiatives

ANR France-Germany FACADES

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

FACADES is a 42-month international project (2022–26) with CISP 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 thesis of Iliana Fayolle takes place in the context of this project.

10.3 National initiatives

10.3.1 ANR

ANR ARCHI-SEC

Participants: Pierre Ayoub, Walid Ghandour, Clémentine Maurice (*contact person*).

ANR ARCHI-SEC is a 42-month project (2019–23) funded by ANR. Attacks exploiting micro-architectural vulnerabilities, such as Meltdown, Spectre, Rowhammer, etc., are on the rise. Modern day SoCs "System-on-a-Chip" embed increasingly complex design features, such as branch prediction, Out-of-Order execution, cache coherency protocols, integrated GPUs/ FPGAs, new nonvolatile memories. The security aspect of these new architectures and technologies remains under-studied. The project aims at modeling the architectural problems with a virtual platform based on gem5. It will be used for penetration testing, evaluate the performance cost of countermeasures, anticipate new attacks and propose protections. These latter are validated on platforms based on ARM and RISC-V processors. The major impact of this project will be through the creation of a community around the virtual platform. Other partners include LTCI, LIRMM, LabHC, Secure-IC. The ongoing PhD thesis of Pierre Ayoub contributes to this project. Our first results in the context of this project have been published in [46].

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. Our first results in the context of this project have been published in [49, 51].

ANR Distiller

Participants: Alexandre Bonvoisin, Tristan Coignion, 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.

ANR FP-Locker

Participants: 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. Our first results in the context of this project have been published in [65, 76, 66].

ANR Koala

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

Koala is a 42-month project (2019–23) 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. Our first results in the context of this project have been published in [57, 64].

ANR MIAOUS

Participants: Thomas Rokicki, Clémentine Maurice (*contact person*).

MIAOUS is a 42-month project (2019–23) funded by ANR in the context of the JCJC program. The project aims to propose a generic framework to provide a better understanding of the attack surface for microarchitectural attacks, both on the hardware and on the software side, and the tools to close the attack surface. Hardware is often considered as an abstract layer that behaves correctly, executing instructions and giving an output. However, side effects due to software implementation and its execution on actual hardware can cause information leakage from side channels, resulting in critical vulnerabilities impacting both the security and privacy of these systems. The project targets in particular information leakage that does not require any physical proximity to devices and that is due to processor microarchitecture, as well as the constructions of novel countermeasures. The ongoing PhD thesis of Guillaume Didier and the defended PhD thesis of Thomas Rokicki contribute to this project. Several results and publications have been obtained in the context of this project [68, 70, 69, 54, 53, 48, 62, 60, 63].

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 Monfleur contributes to this project. First results have been published in [29].

10.3.2 PEPR and PTCC

PEPR Cloud computing Taranis

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

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.

PEPR Cloud computing CARECloud

Participants: Pierre Bourhis, Sophie Cerf, Clément Quinton, Adrien Luxey-Bitri, Romain Rouvoy (*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.

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: 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*), 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 Regional initiatives

ASCOT

Participants: Imane Fouad, Pierre Laperdrix (*contact person*).

ASCOT is a 36-month (2020–23) project funded in the context of the STaRS program of Hauts-de-France region. The goal of the project is to improve web security and privacy and put back control into users hands by blocking unwanted trackers. The project will combine information flow analysis, machine learning and deobfuscation to detect a wide range of trackers. The project will also identify page breakage when blocking online trackers. All in all, this project aims to advance the actual techniques that protect users online while maintaining a high level of usability that is key to offer a comfortable browsing experience. First results have been published in [69].

10.5 Inria initiatives

Inria Défi Federated Machine Learning over the Internet

Participants: Adrien Luxey, Rémy Raes, Romain Rouvoy (*contact person*), Lionel Seinturier.

Federated Machine Learning over the Internet (Fed-Malin) is a 48-month Défi (2021–25) funded by Inria. The goal of the project is to push federated learning research and to address a number of challenges that arise when it is deployed over the Internet, including privacy & fairness, energy consumption, personalization, and location/time dependencies. 10 Inria project-teams participate to this Défi with the support of the Groupe La Poste. The PhD thesis of Rémy Raes takes place in the context of this project.

Inria Défi Towards a More Frugal Cloud

Participants: Romain Rouvoy (*contact person*), Lionel Seinturier.

Towards a More Frugal Cloud is a 48-month (2021–25) Défi funded by Inria. The goal is to investigate new solutions for designing cloud-based digital services that can be more frugal in terms of energy consumption and that can reduce the environment impact of these environments. 5 Inria project-teams participate to this Défi along with the OVHcloud company.

Inria Défi Pushing carbon-neutral services towards the edge

Participants: Mohamed Chakib Belgaid, Thomas Collignon, Sophie Cerf, Daniel Romero, Romain Rouvoy (*contact person*), Lionel Seinturier.

Pushing carbon-neutral services towards the edge (PULSE) is a 48-month (2021–25) Défi funded by Inria. The goal is to develop and promote best practices in geo-repaired hardware and software infrastructures for more environmentally friendly intensive computing. 6 Inria project-teams participate to this Défi along with the Qarnot Computing company. This Défi is led by Romain Rouvoy for Inria and Rémi Bouzel for Qarnot Computing. The CIFRE PhD thesis of Thomas Collignon takes place in the context of this context.

11 Dissemination

11.1 Promoting scientific activities

11.1.1 Scientific events: organisation

Member of the organizing committees

Sophie Cerf: 7th GDR RSD / ASF Winter School on Distributed Systems & Networks.

Romain Rouvoy: 7th GDR RSD / ASF Winter School on Distributed Systems & Networks.

11.1.2 Scientific events: selection

Chair of conference program committees

Pierre Laperdrix: Vice-chair of the 7th CNIL-Inria Privacy Award.

Member of the conference program committees

Simon Bliudze: International Conference on Formal Methods in Software Engineering (FormaliSE), ACM International Systems and Software Product Line Conference (SPLC, Journal First track), IPM International Conference on Fundamentals of Software Engineering (FSEN).

Pierre Bourhis: ProvenanceWeek, ACM Special Interest Group on Management of Data (SIGMOD), Knowledge Representation (KR), International Joint Conferences on Artificial Intelligence.

Sophie Cerf: International Conference on Software Engineering for Adaptive and Self-Managing Systems (SEAMS), International Symposium On Benchmarking, Measuring And Optimizing (Bench).

Pierre Laperdrix: USENIX Security, Conference on Computer and Communications Security (CCS), "Security, Privacy, and Trust" track of The Web Conference.

Clément Quinton: International Systems and Software Product Line Conference (SPLC), International Working Conference on Variability Modelling of Software-Intensive Systems (VaMoS).

Romain Rouvoy: International Middleware Conference (Middleware), International Conference on Mobile Software Engineering and Systems (MobileSoft), Industry Track of the Automated Software Engineering Conference (ASE), International Conference on Distributed Applications and Interoperable Systems (DAIS).

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).

11.1.3 Journal

Member of the editorial boards

Sophie Cerf: Guest Editor for ACM Transactions on Autonomous and Adaptive Systems Special Issue on Control for Computing Systems.

Laurence Duchien: Special Issue co-editor for Journal of System and Software (JSS).

Philippe Merle: Member of the editorial board of MDPI IoT.

Adrien Luxey-Bitri: Guest Editor for AMUE Collection Numérique edition 29 entitled "Urgence sur les sobriétés numériques!".

Reviewer - reviewing activities

Simon Bliudze: International Journal of Foundations of Computer Science (IJFCS).

Pierre Bourhis: Theoretical Computer Science.

Sophie Cerf: IEEE Transactions on Dependable and Secure Computing, American Control Conference (ACC), International Conference on Control, Decision and Information Technologies (CoDIT).

Pierre Laperdrix: IEEE Transactions on Reliability, Annals of Telecommunications, Elsevier Computers & Security.

Clément Quinton: Journal of Systems and Software (JSS), Annals of Mathematics and Artificial Intelligence (AMAI), ACM Transactions on Software Engineering and Methodology (TOSEM), International Systems and Software Product Line Conference (SPLC), International Working Conference on Variability Modelling of Software-Intensive Systems (VaMoS).

Romain Rouvoy: Sustainable Computing: Informatics and Systems (SusCom).

Lionel Seinturier: Springer The Journal of Supercomputing.

11.1.4 Invited talks

Simon Bliudze: Invited lecture (6 hours) at the EuroProofNet Summer School on Verification Technology, Systems & Applications (VTSA 2023).

Pierre Bourhis: Invited talk at the Simon Institute Workshop Logic and Algebra for Query Evaluation.

Pierre Laperdrix: Invited talk on online tracking for the Summer School "Cyber in Sophia 2023".

11.1.5 Leadership within the scientific community

Simon Bliudze

- Co-head of the YODA (trustworthY 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

Romain Rouvoy

- Elected member of CoNRS section 6 (since September)
- Elected member of the "bureau" 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 "Génie Logiciel pour les Systèmes Cyberphysiques" working group of the GDR GPL

11.1.6 Scientific expertise

Simon Bliudze

- Member of the recruitment committee for Junior researcher positions (CRCN/ISFP) at the Centre Inria de l'Université de Lille

Pierre Bourhis

- Member of the jury of agrégation of informatique

Sophie Cerf

- Member of the Inria secondment campaign committee

Laurence Duchien

- Member of the scientific advisory board of Labex CIMI-Toulouse

- Member of the scientific advisory board of IMT Atlantique

Romain Rouvoy

- Member of the jury of agrégation of informatique

Lionel Seinturier

- Member of the recruitment committee for an associate professor position, Université Paris 10 Nanterre

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 technical committee (CTI)
- Deputy secretary of the Inria national committee on "hygiène, de sécurité et des conditions de travail" (CNHSCT)
- Permanent secretary of the "Comité Local d'Hygiène, de Sécurité et de Conditions de Travail" (CLHSCT)
- Elected member of the centre committee for the Inria Center of University of Lille
- Elected member of the Commission d'évaluation Inria (since July)

Lionel Seinturier

- President of section 27 (Informatique) of Conseil national des universités (CNU) (until November)
- 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, is part-time Assistant Professor at École Polytechnique, Palaiseau, France, in the Department of Computer Sciences (DIX).

- CSE202: 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), 14h, Level M1-2, École Centrale Lille
- Commande des Systèmes (basics of linear feedback control), 24h, Level L2 and L3, ITEEM
- Projet d'intégration en Robotique (robot building and control), 12h, Level M1, ITEEM
- Enjeux Environnementaux du Numérique (Introduction à l'Analyse du Cycle de Vie), 4h, Level M2, Master of Computer Science, with A. Luxey-Bitri

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

- Javascript, 18h, 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
- Distributed Systems, 24h, Level M2, Master MIAGE
- 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: Clémentine Maurice, Micro-architectural side channels: Studying the attack surface from hardware to browsers, defended on 24 May 2023, [36].
- PhD defended: Trình Lê Khánh, Design of Correct-by-Construction Self-Adaptive Cloud Applications using Formal Methods, defended on 29 January 2023, supervised by Philippe Merle & Simon Bliudze, [35].
- PhD defended: Romain Fouquet, Improving Online Privacy through Content Blocking and Information Restriction, defended on 31 May 2023, supervised by Romain Rouvoy & Pierre Laperdrix, [33].
- PhD defended: Edouard Guégain, Configuration of Large Scale Fog Environments, defended on 29 September 2023, supervised by Clément Quinton, [34].
- PhD in progress: Wassim Aroui, Modeling and scaling of 5G+/6G network functions on a large-scale edge infrastructure, since 1 June 2023, supervised by Philippe Merle.
- PhD in progress: Pierre Ayoub, IoT devices security inside 5G networks through side-channel analysis, since October 2020, supervised by Clémentine Maurice.
- PhD in progress: Dejfal Belkis, Optimization of Databases in Cloud Environments in a Context of Green Computing, since December 2023, supervised by Romain Rouvoy & Pierre Bourhis.
- PhD in progress: Alexandre Bonvoisin, Reducing the Energy Consumption of Software Stacks, since October 2022, supervised by Romain Rouvoy & Clément Quinton.
- PhD in progress: Tristan Coignon, Environmental Impact of Development Assistants, since October 2022, supervised by Romain Rouvoy & Clément Quinton.
- PhD in progress: Thomas Collignon, Control of high performance computing tasks emissions, since November 2023, supervised by Sophie Cerf & Lionel Seinturier.
- PhD in progress: Boubacar Diarra, Formal modeling and reliability of cloud network configurations, since 1 January 2023, supervised by Philippe Merle.

- PhD in progress: Salman Farhat, Safe Dynamic Reconfiguration of Cloud Applications, since October 2020, supervised by Laurence Duchien & Simon Bliudze.
- PhD in progress: Iliana Fayolle, Side channel vulnerabilities in web environments, since 1 October 2023, 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, supervised by Sandrine Blazy & Clémentine Maurice.
- PhD in progress: Maxime Huyghe, Automated Software Testing to Improve the Privacy of Browsers, since October 2022, supervised by Lionel Seinturier & Clément Quinton & Walter Rudametkin.
- PhD in progress: Jean Luc Intumwayase, Improving Online Privacy through Content Blocking and Information Restriction, since November 2020, supervised by Romain Rouvoy & Pierre Laperdrix.
- PhD in progress: Pierre Jacquet, Fostering the frugal design of cloud services, since October 2021, supervised by Romain Rouvoy in collaboration with Thomas Ledoux (Inria Stack).
- PhD in progress: Naif Mehanna, Hardening web authentication with browser fingerprinting, since October 2020, supervised by Walter Rudametkin.
- PhD in progress: Hugo Monfleur, Design and implementation of a micro-services oriented language, since 1 October 2022, supervised by Philippe Merle.
- PhD in progress: Rémy Raes, Distributed machine learning in ubiquitous environments using location dependant models, since 1 March 2023, supervised by Romain Rouvoy & Adrien Luxey-Bitri.
- PhD in progress: Maryam Rahmani, Multi-scale processing of spatio-temporal data applied to air quality in urban areas, since October 2021, supervised by Romain Rouvoy
- PhD in progress: Thibault Simon, Sustainable software engineering, since November 2021, supervised by Romain Rouvoy
- PhD in progress: Yifan Wang, Autonomous Management of Database Systems, since October 2022, supervised by Romain Rouvoy & Pierre Bourhis.
- PhD in progress: Jiali Xu, Behavior characterisation in smart and deep systems with heterogeneous and reprogrammable nodes, supervised by Romain Rouvoy & Valéria Loscri.

11.2.3 Juries

Simon Bliudze

- PhD Maria Pittou (Aristotle University of Thessaloniki, Greece), examiner

Pierre Bourhis

- PhD Gaston Zanitti (Ecole Polytechnique), member
- PhD Alejandro Grez (Catholic University of Chili), reviewer

Sophie Cerf

- PhD Samuele Zoboli (Université de Lyon), member

Lionel Seinturier

- HDR Sylvain Malacria (U Lille), president

11.3 Popularization

11.3.1 Articles and contents

Pierre Laperdrix was interviewed for the "CNRS Le Journal" on the dangers of online tracking in mobile ([Link](#)).

11.3.2 Education

Sophie Cerf co-organized the "Rendez-vous des jeunes mathématiciennes et informaticiennes" (RJMI) ([web page of the event](#)), a scientific event for high school girls. She participated to the [Chiche!](#) program, and made presentations for pupils in high schools.

Pierre Laperdrix talked about online tracking and fingerprinting during the "Journée de l'Enseignement de l'Informatique et de l'Algorithmique 2023" ([JEIA](#)) at the University of Lille aimed at teachers from a wide range of background.

Rémy Raes participated to three events in relation with dissemination activities towards high school pupils in general and girls in particular. On 21 June, he participated to an academic forum to tell high school teachers about what can be done with their students and computer science labs of the university. From 26 to 27 October, he participated to the so-called "Rendez-vous des jeunes mathématiciennes et informaticiennes" (RJMI), an event to welcome high school girls in the lab during two days to make them work on scientific projects ([web page of the event](#)). On 17 November, he participated to an event to welcome secondary school pupils to Inria Lille Interface showcase room to present his research career and computer science in general.

12 Scientific production

12.1 Major publications

- [1] M. Benedikt, P. Bourhis, L. Jachiet and E. Tsamoura. 'Balancing expressiveness and inexpressiveness in view design'. In: *KR 2020 - 17th International Conference on Principles of Knowledge Representation and Reasoning*. Ed. by D. Calvanese, E. Erdem and M. Thielscher. Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, Rhodes, Greece: IJCAI Organization, Sept. 2020, pp. 109–118. DOI: [10.24963/kr.2020/12](https://doi.org/10.24963/kr.2020/12). URL: <https://hal.telecom-paris.fr/hal-02964692>.
- [2] 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>.
- [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] 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>.
- [5] L. Meftah, R. Rouvoy and I. Chrismet. '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>.
- [6] 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>.

- [7] 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>.
- [8] 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>.
- [9] A. Vastel, P. Laperdrix, W. Rudametkin and R. Rouvoy. ‘FP-STALKER: Tracking Browser Fingerprint Evolutions’. In: *IEEE S&P 2018 - 39th IEEE Symposium on Security and Privacy*. Ed. by B. Parno and C. Kruegel. Proceedings of the 39th IEEE Symposium on Security and Privacy (S&P). San Francisco, United States: IEEE, May 2018, pp. 728–741. DOI: [10.1109/SP.2018.00008](https://doi.org/10.1109/SP.2018.00008). URL: <https://hal.inria.fr/hal-01652021>.
- [10] Z. Yu, C. Bai, L. Seinturier and M. Monperrus. ‘Characterizing the Usage, Evolution and Impact of Java Annotations in Practice’. In: *IEEE Transactions on Software Engineering* (2019). DOI: [10.1109/TSE.2019.2910516](https://doi.org/10.1109/TSE.2019.2910516). URL: <https://hal.inria.fr/hal-02091516>.

12.2 Publications of the year

International journals

- [11] A. Auvoilat, T. Goldoin and A. Luxey-Bitri. ‘Du concept d’entr’hébergement à Garage, solution technique pour le stockage réparti’. In: *1024 : Bulletin de la Société Informatique de France* 22 (Nov. 2023), pp. 171–183. DOI: [10.48556/SIF.1024.22.171](https://doi.org/10.48556/SIF.1024.22.171). URL: <https://cnrs.hal.science/hal-04387879>.
- [12] S. Brisset, R. Rouvoy, L. Seinturier and R. Pawlak. ‘SFTM: Fast Matching of Web Pages using Similarity-based Flexible Tree Matching’. In: *Information Systems* (Feb. 2023). DOI: [10.1016/j.is.2022.102126](https://doi.org/10.1016/j.is.2022.102126). URL: <https://hal.science/hal-03774245>.
- [13] S. Cerf, B. Robu, N. Marchand and S. Bouchenak. ‘Privacy protection control for mobile apps users’. In: *Control Engineering Practice* 134.May (May 2023), p. 105456. DOI: [10.1016/j.conengprac.2023.105456](https://doi.org/10.1016/j.conengprac.2023.105456). URL: <https://hal.science/hal-03977386>.
- [14] B. Danglot, J.-R. Falleri and R. Rouvoy. ‘Can We Spot Energy Regressions using Developers Tests?’ In: *Empirical Software Engineering* (2023). URL: <https://hal.science/hal-04286574>.
- [15] R. Fouquet, P. Laperdrix and R. Rouvoy. ‘Breaking Bad: Quantifying the Addiction of Web Elements to JavaScript’. In: *ACM Transactions on Internet Technology* 23.1 (23rd Feb. 2023), pp. 1–28. DOI: [10.1145/3579846](https://doi.org/10.1145/3579846). URL: <https://hal.science/hal-03936503>.
- [16] R. Guerraoui, A.-M. Kermarrec, G. Niot, O. Ruas and F. Taïani. ‘GoldFinger: Fast & Approximate Jaccard for Efficient KNN Graph Constructions’. In: *IEEE Transactions on Knowledge and Data Engineering* 35.11 (1st Nov. 2023), pp. 11461–11475. DOI: [10.1109/TKDE.2022.3232689](https://doi.org/10.1109/TKDE.2022.3232689). URL: <https://inria.hal.science/hal-04394851>.
- [17] E. Molina, M. Fiacchini, S. Cerf and B. Robu. ‘React to the Worst: Lightweight and proactive protection of location privacy’. In: *IEEE Control Systems Letters* 7 (2023), pp. 2371–2376. DOI: [10.1109/LCSYS.2023.3286989](https://doi.org/10.1109/LCSYS.2023.3286989). URL: <https://hal.science/hal-04128118>.

International peer-reviewed conferences

- [18] S. Bliudze, P. van den Bos, M. Huisman, R. Rubbens and L. Safina. ‘JavaBIP meets VerCors: Towards the Safety of Concurrent Software Systems in Java’. In: *FASE 2023 - 26th International Conference on Fundamental Approaches to Software Engineering*. Vol. 13991. Lecture Notes in Computer Science. Paris, France: Springer Nature Switzerland, 20th Apr. 2023, pp. 143–150. DOI: [10.1007/978-3-031-30826-0_8](https://doi.org/10.1007/978-3-031-30826-0_8). URL: <https://inria.hal.science/hal-03911393>.

- [19] M. Botvinnik, T. Laor, T. Rokicki, C. Maurice and Y. Oren. ‘The Finger in the Power: How to Fingerprint PCs by Monitoring their Power Consumption’. In: DIMVA 2023 - 20th Conference on Detection of Intrusions and Malware & Vulnerability Assessment. Hamburg, Germany, 12th July 2023. URL: <https://inria.hal.science/hal-04153854>.
- [20] P. Bourhis, A. Boussidan and P. Gambette. ‘On Distances between Words with Parameters’. In: *LIPICs*. CPM 2023. Vol. 259. Proceedings of the 34th Annual Symposium on Combinatorial Pattern Matching. Champs-sur-Marne, Marne-la-Vallée, France: Schloss Dagstuhl, 2023, 6:1–6:23. DOI: 10.4230/LIPICs.CPM.2023.16. URL: <https://hal.science/hal-04080842>.
- [21] S. Cerf and E. Rutten. ‘Combining neural networks and control: potentialities, patterns and perspectives’. In: *Proceedings of The 22nd World Congress of the International Federation of Automatic Control*. IFAC 2023 - 22nd World Congress of the International Federation of Automatic Control. Yokohama, Japan, July 2023. URL: <https://inria.hal.science/hal-04060379>.
- [22] 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. Lisbon, Portugal, 22nd Apr. 2023. URL: <https://inria.hal.science/hal-04098074>.
- [23] E. Guégain, A. Taherkordi and C. Quinton. ‘Configuration Optimization with Limited Functional Impact’. In: CAISE 2023 - 35th International Conference on Advanced Information Systems Engineering. Advanced Information Systems Engineering. Zaragoza, Spain, 12th June 2023. URL: <https://hal.science/hal-04034888>.
- [24] 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, 26th Nov. 2023. DOI: 10.1145/3603216.3624958. URL: <https://hal.science/hal-04271971>.
- [25] P. Jacquet, T. Ledoux and R. Rouvoy. ‘CLOUDFACTORY: An Open Toolkit to Generate Production-like Workloads for Cloud Infrastructures’. In: IC2E 2023 - 11th IEEE International Conference on Cloud Engineering. Boston, Massachusetts, USA, United States: IEEE, 25th Sept. 2023, pp. 1–11. URL: <https://inria.hal.science/hal-04168667>.
- [26] G. KP, G. Pierre and R. Rouvoy. ‘Studying the Energy Consumption of Stream Processing Engines in the Cloud’. In: IC2E 2023 - 11th IEEE International Conference on Cloud Engineering. Boston (MA), United States: IEEE, 25th Sept. 2023, pp. 1–9. URL: <https://inria.hal.science/hal-04164074>.
- [27] 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)*, The Web Conference 2023. Austin (TX), United States, 2023, pp. 1–9. URL: <https://hal.science/hal-04084097>.
- [28] E. Molina, M. Fiacchini, S. Cerf and B. Robu. ‘Optimal privacy protection of mobility data: a predictive approach’. In: IFAC WC 2023 - 22nd IFAC World Congress. Yokohama, Japan, 9th July 2023. URL: <https://hal.science/hal-04040962>.
- [29] H. Monfleur and P. Merle. ‘Towards Concern-Oriented Microservice Architecture’. In: *Proceedings of 2023 International Conference on Microservices*. 2023 International Conference on Microservices. Pise, Italy, 10th Oct. 2023. URL: <https://inria.hal.science/hal-04201189>.
- [30] 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, 5th June 2023. URL: <https://inria.hal.science/hal-04082263>.

Conferences without proceedings

- [31] 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, 11th Sept. 2023. URL: <https://hal.science/hal-04213194>.

Edition (books, proceedings, special issue of a journal)

- [32] *Urgence sur les sobriétés numériques !* La collection numérique 29 (31st Oct. 2023). URL: <https://hal.science/hal-04269072>.

Doctoral dissertations and habilitation theses

- [33] R. Fouquet. 'Improving Web User Privacy Through Content Blocking'. Université de Lille, 31st May 2023. URL: <https://theses.hal.science/tel-04123409>.
- [34] E. Guégain. 'Configuration-driven Software Optimization'. Université de lille, 29th Sept. 2023. URL: <https://theses.hal.science/tel-04350545>.
- [35] T. Le Khanh. 'Design of correct-by-construction self-adaptive cloud applications using formal methods'. Université de Lille, 27th Jan. 2023. URL: <https://theses.hal.science/tel-04136838>.
- [36] C. Maurice. 'Micro-architectural side channels: Studying the attack surface from hardware to browsers'. Université de Lille, 24th May 2023. URL: <https://inria.hal.science/tel-04180074>.

Reports & preprints

- [37] A. Amarilli, P. Bourhis, F. Capelli and M. Monet. *Ranked Enumeration for MSO on Trees via Knowledge Compilation*. 1st Oct. 2023. DOI: [10.4230/LIPIcs](https://doi.org/10.4230/LIPIcs). URL: <https://inria.hal.science/hal-04377344>.
- [38] L. Desauw, A. Luxey-Bitri, R. Raes, R. Rouvoy, O. Ruas and W. Rudametkin. *A critical review of mobile device-to-device communication*. 19th Sept. 2023. URL: <https://inria.hal.science/hal-04198528>.
- [39] S. Farhat, S. Bliudze, L. Duchien and O. Kouchnarenko. *Run-time Coordination of Reconfiguration Requests in Cloud Computing Systems*. RR-9504. Inria, 28th Apr. 2023. URL: <https://inria.hal.science/hal-04085278>.
- [40] D. Gross-Amblard, M. Tommasi, I. Rakotoniaina, C. Thierry, R. Singh and L. Jacoboni. *HEADWORK: a Data-centric Crowdsourcing Platform for Complex Tasks and Participants: (Submitted to EDBT 2024)*. 5th July 2023. URL: <https://inria.hal.science/hal-04162652>.
- [41] E. Guégain, A. Taherkordi and C. Quinton. *The ICO Tool Suite: Optimizing Highly Configurable Systems*. 16th Mar. 2023. URL: <https://hal.science/hal-03874051>.

Other scientific publications

- [42] K. Halitim. 'Enhancing Efficiency through Control theory in Compute-Intensive Applications'. Lille, France.: Spirals, Inria Lille, 11th Oct. 2023, p. 73. URL: <https://inria.hal.science/hal-04357812>.
- [43] D. Romero and R. Rouvoy. *La fabrique du numérique frugal*. 1st Apr. 2023. URL: <https://inria.hal.science/hal-04056838>.

12.3 Other

Scientific popularization

- [44] P. Jacquet, R. Rouvoy and T. Ledoux. ‘La chasse au gaspillage dans le cloud et les data centers’. In: *The Conversation France* (29th Jan. 2023). URL: <https://inria.hal.science/hal-04168662>.
- [45] A. Luxey-Bitri, D. Mallarino and B. Ninassi. ‘La sobriété numérique : un impératif environnemental ? : Point de situation, contexte et enjeux, l’état des lieux est sans appel : les limites sont dépassées !’ In: *Collection numérique de l’AMUE, Agence de mutualisation des universités et établissements d’enseignement supérieur* 29 (9th Oct. 2023), p. 2. URL: <https://inria.hal.science/hal-04390215>.

12.4 Cited publications

- [46] P. Ayoub and C. Maurice. ‘Reproducing Spectre Attack with gem5: How To Do It Right?’ In: *14th European Workshop on Systems Security (EuroSec’21)*. Edinburgh, United Kingdom, Apr. 2021. DOI: [10.1145/3447852.3458715](https://doi.org/10.1145/3447852.3458715). URL: <https://hal.inria.fr/hal-03215326>.
- [47] 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).
- [48] S. Bhattacharya, C. Maurice, S. Bhasin and D. Mukhopadhyay. ‘Branch Prediction Attack on Blinded Scalar Multiplication’. In: *IEEE Transactions on Computers* 69.5 (May 2020), pp. 633–648. DOI: [10.1109/TC.2019.2958611](https://doi.org/10.1109/TC.2019.2958611). URL: <https://hal.inria.fr/hal-02866753>.
- [49] 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>.
- [50] 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).
- [51] 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>.
- [52] 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.
- [53] G. Didier and C. Maurice. ‘Calibration Done Right: Noiseless Flush+Flush Attacks’. In: *DIMVA 2021 - The 18th Conference on Detection of Intrusions and Malware & Vulnerability Assessment*. Lisboa / Virtual, Portugal, July 2021. URL: <https://hal.inria.fr/hal-03267431>.
- [54] G. Didier, C. Maurice, A. Geimer and W. J. Ghandour. ‘Characterizing Prefetchers using CacheObserver’. In: *IEEE 34th International Symposium on Computer Architecture and High Performance Computing*. Bordeaux, France, Nov. 2022. URL: <https://hal.science/hal-03798500>.
- [55] R. Fouquet, P. Laperdrix and R. Rouvoy. ‘JSRehab: Weaning Common Web Interface Components from JavaScript Addiction’. In: *WWW’22: Companion Proceedings of the Web Conference 2022*. Lyon, France, May 2022. DOI: [10.1145/3487553.3524227](https://doi.org/10.1145/3487553.3524227). URL: <https://hal.science/hal-03604674>.
- [56] D. Gruss, C. Maurice and K. Wagner. ‘Flush+Flush: A Stealthier Last-Level Cache Attack’. In: *CoRR* abs/1511.04594 (2015). arXiv: [1511.04594](https://arxiv.org/abs/1511.04594). URL: <http://arxiv.org/abs/1511.04594>.

- [57] 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>.
- [58] 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>.
- [59] 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.
- [60] 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>.
- [61] 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>.
- [62] M. Lipp, V. Hadžić, M. Schwarz, A. Perais, C. Maurice and D. Gruss. ‘Take A Way: Exploring the Security Implications of AMD’s Cache Way Predictors’. In: *15th ACM Asia Conference on Computer and Communications Security*. Taipei, Taiwan, Oct. 2020. DOI: [10.1145/3320269.3384746](https://doi.org/10.1145/3320269.3384746). URL: <https://hal.inria.fr/hal-02866777>.
- [63] M. Lipp, M. Schwarz, L. Raab, L. Lamster, M. T. Aga, C. Maurice and D. Gruss. ‘Nethammer: Inducing Rowhammer Faults through Network Requests’. In: *EuroS&P Workshops 2020*. Genova, Italy, Sept. 2020. URL: <https://hal.inria.fr/hal-01872588>.
- [64] 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>.
- [65] V. Mishra, P. Laperdrix, W. Rudametkin and R. Rouvoy. ‘Déjà vu: Abusing Browser Cache Headers to Identify and Track Online Users’. In: *PETS 2021 - The 21th International Symposium on Privacy Enhancing Technologies*. Virtual, France, July 2021. URL: <https://hal.inria.fr/hal-03017222>.
- [66] V. Mishra, P. Laperdrix, A. Vastel, W. Rudametkin, R. Rouvoy and M. Lopatka. ‘Don’t count me out: On the relevance of IP addresses in the tracking ecosystem’. In: *The Web Conference 2020*. Ed. by B. Crispo and N. Nikiforakis. Vol. Security, Privacy, and Trust. Proceedings of The Web Conference (WWW’20). Taipei, Taiwan, Apr. 2020. DOI: [10.1145/3366423.3380161](https://doi.org/10.1145/3366423.3380161). URL: <https://hal.inria.fr/hal-02435622>.
- [67] 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>.
- [68] T. Rokicki, C. Maurice, M. Botvinnik and Y. Oren. ‘Port Contention Goes Portable: Port Contention Side Channels in Web Browsers’. In: *ASIA CCS ’22 - ACM Asia Conference on Computer and Communications Security*. Nagasaki / Virtual, Japan: ACM, May 2022, pp. 1182–1194. DOI: [10.1145/3488932.3517411](https://doi.org/10.1145/3488932.3517411). URL: <https://inria.hal.science/hal-03708833>.

- [69] T. Rokicki, C. Maurice and P. Laperdrix. ‘SoK: In Search of Lost Time: A Review of JavaScript Timers in Browsers’. In: *6th IEEE European Symposium on Security and Privacy (EuroS&P’21)*. Vienna, Austria, Sept. 2021. URL: <https://hal.inria.fr/hal-03215569>.
- [70] T. Rokicki, C. Maurice and M. Schwarz. ‘CPU Port Contention Without SMT’. In: *27th European Symposium on Research in Computer Security (ESORICS 2022)*. Vol. 13556. Lecture Notes in Computer Science. Copenhagen, Denmark: Springer Nature Switzerland, Sept. 2022, pp. 209–228. DOI: [10.1007/978-3-031-17143-7_11](https://doi.org/10.1007/978-3-031-17143-7_11). URL: <https://hal.science/hal-03798342>.
- [71] 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>.
- [72] 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.
- [73] 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.
- [74] 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>.
- [75] 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.
- [76] A. Vastel, W. Rudametkin, R. Rouvoy and X. Blanc. ‘FP-Crawlers: Studying the Resilience of Browser Fingerprinting to Block Crawlers’. In: *MADWeb’20 - NDSS Workshop on Measurements, Attacks, and Defenses for the Web*. Ed. by O. Starov, A. Kapravelos and N. Nikiforakis. San Diego, United States, Feb. 2020. DOI: [10.14722/ndss.2020.23xxx](https://doi.org/10.14722/ndss.2020.23xxx). URL: <https://hal.inria.fr/hal-02441653>.
- [77] P. Vila, P. Ganty, M. Guarnieri and B. Köpf. ‘CacheQuery: learning replacement policies from hardware caches’. In: *PLDI*. 2020.