

RESEARCH CENTRE

Lille - Nord Europe

IN PARTNERSHIP WITH:

Université de Lille

2021

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

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	8
6.1 Awards	8
7 New software and platforms	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 Understanding the Impact of Release Policies on Software Development Processes	10
8.2 Software Eco-Design: Investigating and Reducing the Energy Consumption of Software	11
9 Bilateral contracts and grants with industry	11
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
10.1.2 STIC/MATH/CLIMAT AmSud project	13
10.1.3 Participation in other International Programs	14
10.2 European initiatives	14
10.2.1 Other european programs/initiatives	14
10.3 National initiatives	14
10.3.1 ANR	14
10.4 Regional initiatives	17
11 Dissemination	17
11.1 Promoting scientific activities	17
11.1.1 Scientific events: organisation	17
11.1.2 Scientific events: selection	17
11.1.3 Journal	18
11.1.4 Leadership within the scientific community	18
11.1.5 Scientific expertise	19

11.1.6 Research administration	20
11.2 Teaching - Supervision - Juries	20
11.2.1 Teaching	20
11.2.2 Supervision	22
11.2.3 Juries	23
11.3 Popularization	24
12 Scientific production	24
12.1 Major publications	24
12.2 Publications of the year	25
12.3 Cited publications	27

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 Blyudze [Inria, Researcher]
- Pierre Bourhis [CNRS, Researcher]
- Sophie Cerf [Inria, from Oct 2021, Starting Faculty Position]
- Pierre Laperdrix [CNRS, Researcher]
- Clementine Maurice [CNRS, Researcher, from Feb 2021]
- Philippe Merle [Inria, Senior Researcher, HDR]

Faculty Members

- Lionel Seinturier [Team leader, Université de Lille, Professor, HDR]
- Laurence Duchien [Université de Lille, Professor, HDR]
- Clément Quinton [Université de Lille, Associate Professor]
- Romain Rouvoy [Université de Lille, Professor, HDR]
- Walter Rudametkin Ivey [Université de Lille, Associate Professor, HDR]

Post-Doctoral Fellows

- Emile Cadorel [Inria, from Mar 2021]
- Imane Fouad [Université de Lille, from Oct 2021]
- Alexandre Garnier [Inria, until Feb 2021]
- Walid Ghandour [Université de Lille, from Sep 2021]
- Adrien Luxey [Université de Lille]
- Olivier Ruas [Inria]
- Larisa Safina [Inria]

PhD Students

- Zeinab Abou Khalil [Université de Mons - Belgique, until Mar 2021]
- Mohammed Chakib Belgaid [Inria]
- Sacha Brisset [Amaris Consulting-Biot]
- Antonin Durey [Université de Lille]
- Salman Farhat [Inria]
- Guillaume Fieni [Université de Lille]
- Romain Fouquet [Inria]
- Edouard Guegain [Université de Lille]
- Jean Luc Intumwayase [Université de Lille]

- Pierre Jacquet [Inria, from Oct 2021]
- Trinh Le Khanh [Université de Lille]
- Naif Mehanna [Université de Lille]
- Vikas Mishra [Inria]
- Zakaria Ournani [Orange, CIFRE, until Oct 2021]
- Maryam Rahmani [Université de Lille, from Oct 2021]
- Marion Tommasi [Inria, until May 2021]

Technical Staff

- Arthur D'Azémar [Inria, Engineer, until Sep 2021]
- Lauric Desauw [Inria, Engineer, from Sep 2021]
- Jeremie Dusart [Université de Lille, Engineer, from Feb 2021]
- Tetouhe Kilimou [Inria, Engineer, from Jul 2021]
- Timothée Lefebvre [Université de Lille, Engineer]
- Alexis Pernet [Université de Lille, Engineer]
- Remy Raes [Inria, Engineer]

Interns and Apprentices

- Matteo Chancerel [Ecole normale supérieure Paris-Saclay, from Mar 2021 until Jun 2021]
- Theophile Dubuc [Inria, from Mar 2021 until Aug 2021]
- Antoine Geimer [Université de Lille, from Jun 2021 until Aug 2021]
- Bogdan Gorelkin [Inria, from Jun 2021 until Sep 2021]
- Alejandro Kaminetzky Parada [Inria, from Sep 2021 until Nov 2021]
- Lasha Koroshinadze [Inria, from Jul 2021 until Aug 2021]
- Maxime Leger [Inria, from May 2021 until Aug 2021]
- Quentin Roussel [Lycée Sainte Marie, Beaucamps-Ligny, Jun 2021]
- Nicolas Van Sint Jan Campos [Inria, from Sep 2021 until Nov 2021]

Administrative Assistant

- Karine Lewandowski [Inria]

Visiting Scientist

- Olga Kouchnarenko [Université de Franche-Comté, HDR]

External Collaborator

- Martin Monperrus [KTH Royal Institute of Technology, Sweden]

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-healing* and *self-optimization*, 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 [77], the Roadmap for Advanced Cloud Technologies under H2020 [73], and research roadmaps, such as [60, 63, 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. [79] 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 [62]. 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 [68] and Facebook [56] 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 [75] 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 [51, 54] 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. Software-defined power meters are configurable software libraries that can estimate the power consumption of software in real-time. As a middleware toolkit, PowerAPI offers the capability of assembling power meters “à la carte” to accommodate user requirements. The defended PhD thesis of Zakaria Ournani [42] and the ongoing PhD theses of Guillaume Fieni, Chakib Belgaid, and Thibault Simon, contribute to this line of research. With the now defended PhD theses of Maxime Colmant [57] and Adel Nourredine [69], this is the 3rd generation of PhD candidates that work on this topic in the team. We are also involved in industrial collaborations in this context. We work with Orange (defended CIFRE PhD thesis of Zakaria Ournani and ongoing CIFRE PhD thesis of Thibault Simon) and OVHcloud (ongoing bilateral contract, ANR GreenAct) on solutions to measure and optimize the power consumption of virtual machines and servers. We work with the Davidson consulting company (ongoing PhD thesis of Chakib Belgaid) on end-to-end solutions for the design of eco-friendly software systems. We also work with these three companies in the context of the ANR Distiller project. In 2021, we obtained and published several new results on green and power efficient computing [35, 36, 30, 37, 27].

Finally, a new line of research has just been started in the team with the thesis of Romain Fouquet to limit what is currently running in a modern browser. By correctly identifying what is truly needed on a webpage 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

A highlight of the year is the organization by Romain Rouvoy of the "Green Computing meets Green Energy" summer school that was held in Lille from 6 to 10 September. The summer school is one of the outputs of our collaboration with the University of Oslo, the University of Stavanger and TU Munich, in the context of the norwegian-funded PACE (Partnership for joint Curriculum Development and Research in Energy Informatics) project.

6.1 Awards

XScalibur is one of the laureates of the **Orange 2021 VivaTech challenge**. XScalibur is the startup company founded in 2019 by Christophe Gourdin with Philippe Merle. XScalibur is the result of a transfer of our research activities in the domain of cloud computing. Xscalibur offers a visual, clear and ergonomic interface to design and manage all cloud services in real time while taking energy consumption into account to help make more eco-friendly choices.

Pierre Ayoub and Guillaume Didier, PhD students supervised by Clémentine Maurice, won respectively the 2nd and 3rd place at the ICHSA 2021 student competition that is the reference conference in the domain of hardware and side-channel attack.

7 New software and platforms

The following section lists the software systems on which we contributed this year.

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:

- 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://apisense.io>

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: <https://hal.inria.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 Noureddine, Loïc Huertas, Maxime Colmant, Romain Rouvoy, Mohammed Chakib Belgaid, Arthur D'azemar

8 New results

We highlight below two new major results that we obtained in 2021 in the domain of software engineering and in the domain of energy management for software systems.

8.1 Understanding the Impact of Release Policies on Software Development Processes

In the context of her PhD defended in February, Zeinab Abou Khalil obtained new results in the domain of the understanding of the impact of release policies on software development processes [40]. The advent of delivering new features faster has led many software projects to change their development processes towards more rapid release models where releases are shipped using release cycles of weeks or days. The adoption of rapid release practices has significantly reduced the amount of stabilization time, the time it takes for a software product's failure rate to reach close to the steady-state, available for new features. This forces organizations to change their development process and tools to release to the public, in a timely manner and with good quality. Rapid releases are claimed to offer a reduced time-to-market and faster user feedback; end-users benefit of faster access to functionality improvements and security updates and improve turnaround time for fixing bad bugs. Despite these benefits, previous research has shown that rapid releases often come at the expense of reduced software reliability. Despite the increasing adoption of rapid releases in open-source and commercial software, the effects of this practice on the software development process are not well understood. The goal of this work is to provide a deeper understanding of how rapid releases impact different parts of the open-source software development process. We present empirical evidence about the short and long-term impact of rapid releases on the bug handling and testing process in open source organizations; and the plan and tools needed for successful adoption of rapid releases. This work presents an empirical case study of rapid releases in Eclipse and Mozilla Firefox projects. We follow a mixed-methods approach where we analyze software

repositories, containing different types of data such as source code, testing data and software issues; and we conduct a survey with Eclipse developers. This helps understand the evolution and changes of the software development process, the plans and practices needed for successful adoption of rapid releases, and identifying several future research directions calling for further investigation.

8.2 Software Eco-Design: Investigating and Reducing the Energy Consumption of Software

In the context of his PhD defended in November, Zakaria Ournani obtained new results in the domain of the eco-design of software systems [42]. So far, the energy consumption was mainly related to the used hardware, and its capacity to maintain a low power consumption while achieving tasks. However, the running software is in fact as important as hardware, and is as responsible for very substantial gains or drawbacks in energy consumption. We investigate how to help developers and practitioners understand and actively think about green software design in their work, in order to reduce the energy consumption of their software and deliver energy efficient products. We thus contribute to supplement green software design knowledge. To achieve this, we start with conducting a qualitative study with developers, to discuss the multiple hurdles they are facing and their requirements to promote green software design within companies. To reduce software energy consumption, practitioners have to measure it and track its evolution first. In a second step, we investigate the problem of energy consumption variations. We provide guidelines on controllable factors that one could easily tune to reduce this variation and conduct steady and reproducible energy measurements. Once practitioners are able to measure the energy consumption of their software, they can work on reducing it and produce energy efficient software. Thus, we deliver 3 more contributions, focusing on the Java language. The first contribution aims at helping developers choose and configure their execution environment. We identified substantial differences in energy consumption using multiple JVM platforms with different JIT and GC configurations for different use cases. The second and third contributions study the impact on energy consumption of small changes that developers often apply on their source code (code refactoring and API/methods substitutions respectively). We show through these studies that structure oriented code refactorings do not substantially alter software energy consumption. On the other hand, Java I/O methods substitution drastically changed the energy consumption depending on the use case. This work contributes to enrich the knowledge on green software design and provides insights and approaches to enhance the energy efficiency at multiple levels of software development.

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.

Davidson Consulting

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

This collaboration (2017–21) aims at proposing new solutions for optimizing the energy footprint of ICT software infrastructures. We want to be able to measure and assess the energy footprint of ICT systems while preserving various quality of service parameters, such as performance and security. We

aim at proposing a testbed for assessing the energy footprint of various programming languages. This testbed will also incorporate frameworks for web and mobile programming. Finally, we want to be able to issue recommendations to developers in order to assist them in improving the energy footprint of their programs. This collaboration will take advantage of the POWERAPI software library.

The PhD of Mohammed Chakib Belgaid takes place in the context of this collaboration.

Orange # 1

Participants: Philippe Merle (*contact person*), Lionel Seinturier.

This collaboration (2017–22) aims at defining a computational model for software infrastructures layered on top of virtualized and interconnected cloud resources. This computational model provides application programming and management facilities to distributed applications and services [64, 76] and defines a pivot model that enables the interoperability of various existing and future standards for cloud systems such as OCCI and TOSCA. This pivot model is defined with the Alloy specification language [61]. This collaboration takes advantage of the expertise that we are developing since several years on reconfigurable component-based software systems [74], on cloud systems [70], and on the Alloy specification language [65].

This collaboration with Orange Labs is a joint project with Jean-Bernard Stefani from the *Spades* Inria project-team.

Orange # 2

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

This collaboration (2018–21) aims at proposing new solutions for modeling the energy efficiency of software systems and to design and implement new methods for measuring and reducing the energy consumption of software systems at development time. We especially target software systems deployed on cloud environments.

The CIFRE PhD of Zakaria Ournani [42] took place in the context of this collaboration.

Amaris (now Mantu)

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

This collaboration (2018–21) aims at proposing new solutions for automatically spotting and fixing recurrent user experience issues in web applications. We are interested in developing an autonomic framework that learns and classifies the behaviors and figures out causality links between data such as web GUI events, support tickets and user feedback, source version management events (e.g. recent commits). The ultimate objective is to implement an AI-powered recommendation system to guide the maintenance and even to automatically predict and solve user issues.

The CIFRE PhD of Sacha Brisset takes place in the context of this collaboration.

Berger-Levrault

Participants: Tetouhe Kilimou, Philippe Merle (*contact person*).

This collaboration (2020–21) aims at proposing new solutions for steering, orchestrating, and maintaining the software development activities in the stages related to packaging, delivering, and deploying. The objective is to better control these stages and to make them more reliable. Based on several identified use cases, the envisioned solutions take advantage of our expertise and our recent advances in software metamodeling, smart monitoring, and knowledge based building.

OVHcloud

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

This collaboration (2020–21) aims at improving the monitoring of the power consumption of virtual machines hosted in datacenters. The objective is to propose new solutions to be able to continuously, in real time, monitor the power consumption of IaaS solutions such as OpenStack. This will contribute to a better understanding of the power models that can be used online in datacenters. This collaboration will take advantage of the POWERAPI software library.

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

SusAIIn

Title: Towards a Sustainable Artificial Intelligence

Duration: 2021 -> 2023

Coordinator: Romain Rouvoy

Partners:

- Inria Spirals
- Inria Chile
- Univ. Chile
- LNCC Brazil

Inria contact: Romain Rouvoy

Summary: The project addresses the challenge of reducing the power consumption of artificial intelligence algorithms deployed in the context of high performance computing. The research program is organized around a set of interdisciplinary themes: better/custom hardware; energy-aware high-performance and cloud computing; smarter experimentation and hyperparameter search; Model self-adaptation, transfer learning and domain adaptation; active learning and sample efficiency.

10.1.2 STIC/MATH/CLIMAT AmSud project

CLIMAT AmSud Green AI

Participants: Mohammed Chakib Belgaid, Guillaume Fieni, Zakaria Ournani, Romain Rouvoy (*contact person*), Lionel Seinturier.

CLIMAT AmSud **GREEN AI** is a 2-year (2020–21) project that aims to design a systemic and multi-component approach to the problem of the Artificial Intelligence's ecological impact. The project focuses on cloud and mobile computing, transfer learning, model reuse, active learning and evolutionary computing, among others. Partners are: Inria Chile, Universidad de Chile, Universidad de la Republica (Uruguay), Universidad Nacional de Asuncion (Paraguay), Pontificia Universidade Catolica de Rio de Janeiro (Brazil).

10.1.3 Participation in other International Programs

Partnership for joint Curriculum Development and Research in Energy Informatics (PACE)

Participants: Mohammed Chakib Belgaid, Arthur D’Azémar, Guillaume Fieni, Zakaria Ournani, Clément Quinton, Romain Rouvoy (*contact person*), Lionel Seinturier.

PACE is a 3-year (2019–21) project funded by the Research Council of Norway. The goal of the project is to establish a sustained education and research-oriented collaboration between four partner universities in energy informatics and green computing that will strengthen quality academic relations and mutually improve each other’s quality of research and researcher training both at PhD and master level. Partner universities are: University of Oslo (Norway), University of Stavanger (Norway), TU Munich (Germany), Université de Lille.

10.2 European initiatives

10.2.1 Other european programs/initiatives

DG ECHO AIDERS

Participants: Timothée Lefebvre, Alexis Pernet, Romain Rouvoy (*contact person*), Lionel Seinturier.

AIDERS is an European project (2020–22) funded by the directorate general for European Civil Protection and Humanitarian Aid Operations of the EU. The partners are the University of Cyprus, the Cyprus Civil Defense, the Entente pour la Forêt Méditerranéenne french public organisation, the Corpo nazionale vigili del fuoco italian state organisation, the Center for Security Studies greek think thank, and the Spirals project-team. The AIDERS project aims at developing application-specific algorithms and novel mapping platform that will harness the large volume of data that first responders are now able to collect through heterogeneous sensors (including visual, thermal and multispectral cameras, LIDAR, CBRN sensors, etc.) on-board RPAS units, and converting that data into actionable decisions for improved emergency response.

DG ECHO ARTION

Participants: Timothée Lefebvre, Alexis Pernet, Romain Rouvoy (*contact person*), Lionel Seinturier.

ARTION is a 18-month project (2021—22) that envisions to become a world-class network for knowledge sharing in the area of artificial intelligence for disaster management that will guide the development and use of AI tools by first responders across Europe. By achieving its strategic objectives ARTION aims to bridge the gap between AI scientists and disaster management experts, build capacity and competency of first responders in the use of AI technology, share knowledge and data, and stimulate further AI research towards application-specific challenges faced throughout the disaster management cycle.

10.3 National initiatives

10.3.1 ANR

ANR ARCHI-SEC

Participants: 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-Chips" 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 [23].

ANR CQFD

Participants: Pierre Bourhis (*contact person*).

CQFD is a 48-month project (2018–22) 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 [53, 55].

ANR Delta

Participants: Pierre Bourhis (*contact person*).

Delta is a 48-month project (2016–21) funded by ANR. The project focuses on the study of logic, transducers and automata. In particular, it aims at extending classical framework to handle input/output, quantities and data. This project is coordinated by M. Zeitoun from LaBRI. Other partners include LIF (Marseille), IRIF (Paris-Diderot), and D. Gallois from the Inria Lille Links team. Several results and publications have been obtained in the context of this project [52, 48, 49, 50].

ANR FP-Locker

Participants: Naif Mehanna, Vikas Mishra, Walter Rudametkin Ivey (*contact person*).

FP-Locker is a 42-month project (2019–23) 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 [34, 34] and [78, 67].

ANR Headwork

Participants: Pierre Bourhis (*contact person*), Marion Tommasi.

Headwork is a 48-month project (2016–21) funded by ANR. The main objective of Headwork is to develop data-centric workflows for programming crowd sourcing systems in a flexible declarative manner. The problem of crowd sourcing systems is to fill a database with knowledge gathered by thousands or more human participants. A particular focus is to be put on the aspects of data uncertainty and for the representation of user expertise. This project is coordinated by D. Gross-Amblard from the Druid Team (Rennes 1). Other partners include the Dahu team (Inria Saclay), Sumo (Inria Bretagne), and Links (Inria Lille) with J. Nierhen and M. Sakho. Our results from this project have been published in the following paper [47].

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 [30] and [66].

ANR MIAOUS

Participants: Clémentine Maurice (*contact person*).

ANR 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 theses of Guillaume Didier and Thomas Rokicki contribute to this project. Our first results in the context of this project have been published in [25, 38].

ANR SATAS

Participants: Alexandre Garnier, Philippe Merle (*contact person*), Romain Rouvoy, Lionel Seinturier.

SATAS is a 66-month project (2015–21) funded by ANR. SATAS aims to advance the state of the art in massively parallel SAT solving with a particular eye to the applications driving progress in the field. The final goal of the project is to be able to provide a “pay as you go” interface to SAT solving services, with a particular focus on their power consumption. This project will extend the reach of SAT solving technologies, daily used in many critical and industrial applications, to new application areas, which were previously considered too hard, and lower the cost of deploying massively parallel SAT solvers on the cloud. Our results from this project have been published in the following papers [58, 59].

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 [38].

BRiCoS

Participants: Simon Bliudze (*contact person*), Larisa Safina.

BRiCoS is 24-month (2021–22) project funded by the I-Site ULNE Foundation in the framework of its “Support for young researchers 2020” programme. The project addresses two fundamental challenges for the implementation of the Rigorous System Design approach to general-purpose software: 1) obtaining behavioural models of the coordinated software entities and 2) detecting the deviations between these models and the corresponding executable code in the face of software evolution.

11 Dissemination

11.1 Promoting scientific activities

11.1.1 Scientific events: organisation

Romain Rouvoy organized the summer school “Green Computing meets Green Energy” that was held in Lille from 6 to 10 September.

11.1.2 Scientific events: selection

Chair of conference program committees

Clémentine Maurice: USENIX Security Symposium 2022 Artifact Evaluation co-chair.

Simon Bliudze: International Conference on Formal Methods in Software Engineering (FormaliSE).

Member of the conference program committees

Simon Bliudze: International Conference on Formal Methods for Industrial Critical Systems (FMICS), International Colloquium on Theoretical Aspects of Computing (ICTAC), International Conference on Coordination Models and Languages (Coordination) PC & Steering committee, International Conference on Formal Aspects of Component Software (FACS).

Laurence Duchien: European Conference on Software Architecture (ECSA) PC & Steering committee, International Conference on Software Architecture, Systems and Software Product Line Conference (SPL) PC & Steering Committee.

Pierre Laperdrix: The Web Conference 2022, Designing Security for the Web (SecWeb), International Symposium on Research in Attacks, Intrusions and Defenses (RAID), European Workshop on Systems Security (EuroSec), USENIX Security, Conference on Detection of Intrusions and Malware & Vulnerability Assessment (DIMVA), IEEE Workshop on Offensive Technologies (WOOT).

Clémentine Maurice: USENIX Security Symposium, IEEE Symposium on Security and Privacy (S&P) 2021 & 2022, Workshop on Hardware and Architectural Support for Security and Privacy (HASP), Workshop on DRAM Security (DRAMSec), Steering committee of the IEEE Workshop on Offensive Technologies (WOOT), Steering committee of the Conference on Detection of Intrusions and Malware & Vulnerability Assessment (DIMVA).

Philippe Merle: International Conference on Adaptive and Self-Adaptive Systems and Applications (ADAPTIVE), International Conference on Big Data, Machine Learning and IoT (BMLI), Conférence en Ingénierie du Logiciel (CIEL), International Conference on Cloud Computing, GRIDs, and Virtualization (CLOUD COMPUTING), Future Internet Services and Applications (FISA@WETICE), International Conference on Evolving Internet (INTERNET), International Conference on Advanced Service Computing (SERVICE COMPUTATION).

Clément Quinton: ACM Symposium on Applied Computing (SAC) track SA-TTA, ACM International Systems and Software Product Line Conference (SPLC), tracks Challenges, Solutions, Artifact, International Working Conference on Variability Modelling of Software-Intensive Systems, ACM Student Research Competition (SRC) @ MODELS.

Romain Rouvoy: IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGrid), ACM Symposium on Applied Computing (SAC) track DADS, ACM/IFIP International Middleware Conference (Middleware), IEEE/ACM International Conference on Mobile Software Engineering and Systems (MOBILESoft), ACM/IEEE International Conference on Mining Software Repositories (MSR) registered reports, ACM International Conference on Distributed and Event-Based Systems (DEBS) industrial track, IFIP International Conference Distributed Applications and Interoperable Systems (DAIS).

Lionel Seinturier: ACM Symposium on Applied Computing (SAC) tracks DADS, OS, SA-TTA, International Conference on Service Oriented Computing (ICSOC), ACM International Conference on Management of Emergent Digital EcoSystems (MEDES), Workshop of Formal Approaches for Advanced Computing Systems (FAACS) @ ECSA, International Workshop on Models@run.time (MRT) @ MODELS.

11.1.3 Journal

Member of the editorial boards

Laurence Duchien: Guest editor for a special issue on Configurable Systems of Empirical Software Engineering (EMSE), Special Issue co-editor for Journal of System and Software (JSS).

Reviewer - reviewing activities

Philippe Merle: Future Generation Computer Systems (FGCS), Software Practice and Experiments (SPE).

Lionel Seinturier: Elsevier Science of Computer Programming.

11.1.4 Leadership within the scientific community

Simon Bliudze

- Co-head of the YODA (trustworthyY and Optimal Dynamic Adaptation) working group of the GDR GPL

Laurence Duchien

- President of the scientific council of IRT SystemX
- Scientific advisor at INS2I
- Elected member of CoNRS section 6 (until June)

Philippe Merle

- Elected member of the Inria scientific board (CS)

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

11.1.5 Scientific expertise**Simon Bliudze**

- South Africa's National Research Foundation (NRF), evaluating the quality, impact and standing of the research portfolio of a professor from a leading South African university (expert)

Laurence Duchien

- recruitment committee for an associate professor position, Grenoble INP
- recruitment committee Linnaeus University, Sweden
- recruitment committee for the program Helmholtz Distinguished Professorship – Funding for the recruitment of leading international women scientists (W3), Germany
- scientific advisory board of Labex CIMI-Toulouse
- scientific advisory board of IMT Atlantique

Clémentine Maurice

- recruitment committee for an associate professor position, ENS Rennes
- recruitment committee for an associate professor position, Université de Lorraine
- secrétaire du prix de thèse Gilles Kahn de la Société Informatique de France

Philippe Merle

- recruitment committee for an associate professor position, ENSIIE

Clément Quinton

- recruitment committee for an associate professor position, Université de Lorraine - LORIA

Romain Rouvoy

- recruitment committee for an associate professor position, IMT Atlantique - LS2N
- recruitment committee for an associate professor position, Université de Bordeaux - LaBRI

Lionel Seinturier

- Institut Universitaire de France (IUF) Jury senior
- expert for the Icelandic Research Fund
- recruitment committee for a professor position, Sorbonne Université

11.1.6 Research administration

Simon Bliudze

- Elected member of the comité de centre Inria Lille Nord Europe

Laurence Duchien

- in charge of the Career development & Intersectoral secondments in the PEARL Project ("Programme for EARly-stage Researchers in Lille") at I-Site Université Lille Nord Europe
- member of the council of the Faculty of science and technology of the University of Lille

Antonin Durey

- Elected member of the CRISAL laboratory council

Philippe Merle

- Elected member of the Inria technical committee (CTI)
- Elected member of the Inria national committee on "hygiène, de sécurité et des conditions de travail" (CNHSCT)
- President of the CUMI (Comité des Utilisateurs des Moyens Informatiques)
- Permanent secretary of the CLHSCT (Comité Local d'Hygiène, de Sécurité et de Conditions de Travail)
- Elected member of the centre committee for the Inria Lille - Nord Europe research center
- Member of the steering committee of the Inria continuous integration service

Lionel Seinturier

- President of section 27 (Informatique) of Conseil national des universités (CNU)
- Member of the scientific council of University of Lille

Walter Rudametkin

- Member of the CDT (Comité de Développement Technologique) of the Inria Lille - Nord Europe research center
- Elected member of the CRISAL laboratory council

11.2 Teaching - Supervision - Juries

11.2.1 Teaching

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

- INF411: Les bases de la programmation et de l'algorithmique, 40h, 2nd year of the Engineering cycle
- INF442: Traitement des données massives, 40h, 2nd year of the Engineering cycle

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.

- Système de Transport Autonome, 12h, 2nd year of the Engineering cycle, École Centrale Lille
- Commande des Systèmes, 24h, Level L2, ITEEM
- Projet d'intégration en Robotique, 10h, Level M1, ITEEM

Laurence Duchien teaches at the Université de Lille in the FST faculty. She is project leader for doctoral studies at Université de Lille.

- Software engineering project, 60h, Level M2, Master MIAGE FI
- Software engineering project, 50h, Level M2, Master MIAGE FC/FA
- Research initiation, 20h, Level M2, Master of Computer Science

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

- Introduction to Computer Science, 46.5h, Level L1, Licence of Computer Science
- Object-oriented programming, 36h, Level L2, Licence of Computer Science
- Object-oriented design, 42h, Level L3, Licence of Computer Science
- Design of distributed applications, 42h, Level M1, Master of Computer Science
- Advanced design of distributed applications, 37.5h, Level M2, Master MIAGE
- Infrastructure and frameworks for the Internet, 33.75h, Level M2, Master of Computer Science
- Software product lines, 7.5h, Level M2, Master of Computer Science
- Suivi de stages et de projets, 30h, Licence and Master of Computer Science

Romain Rouvroy teaches at the Université de Lille in the FST faculty. He heads the **Master of Computer Science** program at the Université de Lille.

- 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

Walter Rudametkin Ivey teaches at the Polytech Lille engineering school.

- GIS4 Programmation par Objets, 32h
- GIS4 Architectures Logicielles, 26h
- GIS2A3 (apprentissage) Projet programmation par Objet, 24h
- IMA2A4 (apprentissage) Conception Modélisation Objet, 24h
- IMA3 Programmation Avancée, 62h
- GBIAAL4 Bases de données, 22h
- GIS5 Suivi de projets, 42h

- GIS2A (apprentissage) Suivi d'apprentis, 28h

Lionel Seinturier teaches at the Université de Lille in the FST faculty. Until July 2019, he headed the Computer Science Department at the Faculty of Science and Technology of the Université de Lille.

- Conception d'applications réparties, 48h, Level M1, Master MIAGE
- Systèmes répartis avancés 1, 52h, Level M2, Master of Computer Science

11.2.2 Supervision

- PhD defended: Zeinab Abou Khalil, Understanding the impact of release policies on software development processes, defended on 26 February 2021, supervised by Laurence Duchien & Clément Quinton, co-supervision with Tom Mens (University of Mons), [40].
- PhD defended: Zakaria Ournani, Software Eco-Design: Investigating and Reducing the Energy Consumption of Software, defended on 8 November 2021, supervised by Romain Rouvoy, [42].
- 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: Mohammed Chakib Belgaid, Développement durable des logiciels vers une optimisation énergétique de bout en bout des systèmes logiciels, since January 2018, supervised by Romain Rouvoy & Lionel Seinturier.
- PhD in progress: Sacha Brisset, Automatic Spotting and fixing of Recurrent user Experience issues. Detecting and Fixing Anomalies by applying Machine Learning on user Experience Data, since November 2018, supervised by Lionel Seinturier & Romain Rouvoy & Renaud Pawlak.
- PhD in progress: Guillaume Didier, Security of microarchitectures, since October 2019, supervised by Clémentine Maurice.
- PhD in progress: Antonin Durey, Leveraging Browser Fingerprinting to Fight Fraud on the Web, since October 2018, supervised by Romain Rouvoy & Walter Rudametkin.
- PhD in progress: Salman Farhat, Safe Dynamic Reconfiguration of Cloud Applications, since October 2020, supervised by Laurence Duchien & Simon Bliudze.
- PhD in progress: Guillaume Fieni, GreenData : Vers un traitement efficient et éco-responsable des grandes masses de données numériques, since October 2017, supervised by Romain Rouvoy & Lionel Seinturier.
- PhD in progress: Romain Fouquet, Improving Online Privacy through Content Blocking and Information Restriction, since October 2020, supervised by Romain Rouvoy & Pierre Laperdrix.
- PhD in progress: Omid Gheibi, Dependable Adaptive Software Systems for the Digital World, since October 2020, supervised by Laurence Duchien & Clément Quinton, co-supervision with Danny Weyns (KU Leuven).
- PhD in progress: Edouard Guegain, Configuration of Large Scale Fog Environments, since October 2020, supervised by Clément Quinton.
- 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: Trình Lê Khánh, Design of Correct-by-Construction Self-Adaptive Cloud Applications using Formal Methods, since October 2019, supervised by Philippe Merle & Simon Bliudze.

- PhD in progress: Naif Mehanna, Hardening Web Authentication with Browser Fingerprinting, since October 2020, supervised by Walter Rudametkin.
- PhD in progress: Vikas Mishra, Collaborative Strategies to Protect Against Browser Fingerprinting, since October 2018, supervised by Romain Rouvoy & Walter Rudametkin.
- 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: Thomas Rokicki, Detection and exploitation of side-channel vulnerabilities, since October 2019, supervised by Clémentine Maurice.
- PhD in progress: Thibault Simon, Sustainable software engineering, since November 2021, supervised by Romain Rouvoy

11.2.3 Juries

Laurence Duchien

- Zeinab Abou Khalil (U Lille & U Mons, Belgium), co-director
- Zeina Houmani (ENS Lyon), reviewer
- Kenza Kraibi (U Lille), president
- Julien Leveau (U Bordeaux), president
- Hugo Martin (U Rennes 1), reviewer
- Geoffrey Pruvost (U Lille), president

Pierre Laperdrix

- Imane Fouad (U Côte d'Azur), examiner

Clémentine Maurice

- Cesar Pereida García (Tampere U, Finland), dissertation pre-examiner
- Mathieu Escouteloup (U Rennes 1), examiner
- Georges-Axel Jaloyan (ENS Paris), examiner
- Thomas Troughkine (U Grenoble Alpes), examiner

Philippe Merle

- Antoine Cheron (U Rennes 1), reviewer

Clément Quinton

- Zeinab Abou Khalil (U Lille & U Mons, Belgium), invitee

Romain Rouvoy

- Hamidreza Arkian (U Rennes 1), reviewer
- Thibault Béziers La Fosse (U Nantes), reviewer
- Patrick Lavoisier Wapet (INP Toulouse), reviewer
- Dimitri Saingre (IMT Atlantique), reviewer
- Zakaria Ournani (U Lille), supervisor

- HDR Mathieu Acher (U Rennes 1), reviewer

Lionel Seinturier

- Mirko D'Angelo (Linnaeus U, Sweden), opponent
- Maxime Savary-Leblanc (U Lille), president
- HDR Jérémy Buisson (U Bretagne Sud), reviewer
- HDR François Trahay (IP Paris Telecom SudParis), reviewer

11.3 Popularization

Lionel Seinturier participates to the Chiche program by giving talks to high school pupils in the context of the SNT Sciences numériques et technologie courses. Talks given at Lycée Queneau, Villeneuve d'Ascq, 15 March, 16 March, 7 December.

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. Chrisment. 'FOUGERE: User-Centric Location Privacy in Mobile Crowdsourcing Apps'. In: *19th IFIP International Conference on Distributed Applications and Interoperable Systems (DAIS)*. Ed. by J. Pereira and L. Ricci. Vol. LNCS-11534. Distributed Applications and Interoperable Systems. Kongens Lyngby, Denmark: Springer International Publishing, 2019, pp. 116–132. DOI: [10.1007/978-3-030-22496-7_8](https://doi.org/10.1007/978-3-030-22496-7_8). URL: <https://hal.inria.fr/hal-02121311>.
- [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] Z. Abou khalil, E. Constantinou, T. Mens and L. Duchien. ‘On the impact of release policies on bug handling activity: A case study of Eclipse’. In: *Journal of Systems and Software* 173 (Mar. 2021), p. 110882. DOI: [10.1016/j.jss.2020.110882](https://doi.org/10.1016/j.jss.2020.110882). URL: <https://hal.inria.fr/hal-03035765>.
- [12] M. Ahmed-Nacer, S. Kallel, F. Zalila, P. Merle and W. Gaaloul. ‘Model driven simulation of elastic OCCI cloud resources’. In: *The Computer Journal* (2022), pp. 1–22. DOI: [10.1093/comjnl/bxaa159](https://doi.org/10.1093/comjnl/bxaa159). URL: <https://hal.inria.fr/hal-02976775>.
- [13] M. Benedikt, P. Bourhis, B. Ten Cate, G. Puppis and M. Vanden Boom. ‘INFERENCE FROM VISIBLE INFORMATION AND BACKGROUND KNOWLEDGE’. In: *ACM Transactions on Computational Logic* (21st June 2021). URL: <https://hal.archives-ouvertes.fr/hal-03454215>.
- [14] S. Brisset, R. Rouvoy, L. Seinturier and R. Pawlak. ‘ERRATUM: Leveraging Flexible Tree Matching to Repair Broken Locators in Web Automation Scripts’. In: *Information and Software Technology* 144 (Apr. 2022). DOI: [10.1016/j.infsof.2021.106754](https://doi.org/10.1016/j.infsof.2021.106754). URL: <https://hal.inria.fr/hal-03377001>.
- [15] S. Challita, F. Korte, J. Erbel, F. Zalila, J. Grabowski and P. Merle. ‘Model-Based Cloud Resource Management with TOSCA and OCCI’. In: *Software and Systems Modeling* (26th Feb. 2021), pp. 1–23. DOI: [10.1007/s10270-021-00869-y](https://doi.org/10.1007/s10270-021-00869-y). URL: <https://hal.archives-ouvertes.fr/hal-03122452>.
- [16] Y. Al-Dhuraibi, F. Zalila, N. Djarallah and P. Merle. ‘Model-Driven Elasticity Management with OCCI’. In: *IEEE transactions on cloud computing* 9.4 (Oct. 2021), pp. 1549–1562. DOI: [10.1109/TCC.2019.2923686](https://doi.org/10.1109/TCC.2019.2923686). URL: <https://hal.inria.fr/hal-02375362>.
- [17] L. Duchien, P. Grünbacher and T. Thüm. ‘Foreword to the Special Issue on Configurable Systems’. In: *Empirical Software Engineering* 26 (17th June 2021), pp. 1–4. DOI: [10.1007/s10664-021-09964-6](https://doi.org/10.1007/s10664-021-09964-6). URL: <https://hal.archives-ouvertes.fr/hal-03263567>.
- [18] 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>.
- [19] J. Leveau, X. Blanc, L. Réveillère, J.-R. Falleri and R. Rouvoy. ‘Fostering the Diversity of Exploratory Testing in Web Applications’. In: *Journal of: Software Testing, Verification and Reliability* (2022). URL: <https://hal.inria.fr/hal-03536908>.
- [20] S. Mengel, A. Amarilli, P. Bourhis and M. Niewerth. ‘Constant-Delay Enumeration for Nondeterministic Document Spanners’. In: *ACM Transactions on Database Systems* 46.1 (Apr. 2021), pp. 1–30. DOI: [10.1145/3436487](https://doi.org/10.1145/3436487). URL: <https://hal.archives-ouvertes.fr/hal-03402446>.
- [21] Y. Wu, J. He, S. Yan, J. Wu, T. Yang, O. Ruas, G. Zhang and B. Cui. ‘Elastic Bloom Filter: Deletable and ExpandableFilter Using Elastic Fingerprints’. In: *IEEE Transactions on Computers* (22nd Mar. 2021), pp. 1–1. DOI: [10.1109/TC.2021.3067713](https://doi.org/10.1109/TC.2021.3067713). URL: <https://hal.inria.fr/hal-03347690>.
- [22] 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* 47.5 (2021), pp. 969–986. DOI: [10.1109/TSE.2019.2910516](https://doi.org/10.1109/TSE.2019.2910516). URL: <https://hal.inria.fr/hal-02091516>.

International peer-reviewed conferences

- [23] 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, 26th Apr. 2021. DOI: [10.1145/3447852.3458715](https://doi.org/10.1145/3447852.3458715). URL: <https://hal.inria.fr/hal-03215326>.
- [24] P. Bourhis, A. Grez, L. Jachiet and C. Riveros. ‘Ranked enumeration of MSO logic on words’. In: 24th International Conference on Database Theory, {ICDT} 2021, March 23-26, 2021, Nicosia, Cyprus. Vol. 186. Nicosia / Virtual, Cyprus: Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 23rd Mar. 2021, 20:1–20:19. DOI: [10.4230/LIPIcs.ICDT.2021.20](https://doi.org/10.4230/LIPIcs.ICDT.2021.20). URL: <https://hal.telecom-paris.fr/hal-03223242>.
- [25] 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, 14th July 2021. URL: <https://hal.inria.fr/hal-03267431>.
- [26] A. Durey, P. Laperdrix, W. Rudametkin and R. Rouvoy. ‘FP-Redemption: Studying Browser Fingerprinting Adoption for the Sake of Web Security’. In: International Conference on the Detection of Intrusions and Malware & Vulnerability Assessment (DIMVA). Lisboa, Portugal, 14th July 2021. URL: <https://hal.archives-ouvertes.fr/hal-03212726>.
- [27] G. Fieni, R. Rouvoy and L. Seinturier. ‘SEFWATTS: On-the-fly Selection of Performance Events to Optimize Software-defined Power Meters’. In: CCGRID 2021 - 21th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing. Melbourne, Australia, 10th May 2021. URL: <https://hal.inria.fr/hal-03173410>.
- [28] Q. Forcioli, J.-L. Danger, C. Maurice, L. Bossuet, F. Bruguier, M. Mushtaq, D. Novo, L. France, P. Benoit, S. Guilley and T. Perianin. ‘Virtual Platform to Analyze the Security of a System on Chip at Microarchitectural Level’. In: IEEE European Symposium on Security and Privacy Workshops (EuroS&PW 2021). Vienne, Austria, 2021, pp. 96–102. DOI: [10.1109/EuroSPW54576.2021.00017](https://doi.org/10.1109/EuroSPW54576.2021.00017). URL: <https://hal.archives-ouvertes.fr/hal-03353878>.
- [29] G. Giakkoupis, A.-M. Kermarrec, O. Ruas and F. Taïani. ‘Cluster-and-Conquer: When Randomness Meets Graph Locality’. In: ICDE 2021 - IEEE 37th International Conference on Data Engineering. Chania, Greece: IEEE, 19th Apr. 2021, pp. 2027–2032. DOI: [10.1109/ICDE51399.2021.00195](https://doi.org/10.1109/ICDE51399.2021.00195). URL: <https://hal.inria.fr/hal-03346860>.
- [30] 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, 6th Sept. 2021, pp. 89–99. DOI: [10.1145/3461001.3471142](https://doi.org/10.1145/3461001.3471142). URL: <https://hal.inria.fr/hal-03269168>.
- [31] T. Laor, N. Mehanna, A. Durey, V. Dyadyuk, P. Laperdrix, C. Maurice, Y. Oren, R. Rouvoy, W. Rudametkin and Y. Yarom. ‘DRAWNAPART: A Device Identification Technique based on Remote GPU Fingerprinting’. In: Network and Distributed System Security Symposium. San Diego, United States, 27th Feb. 2022. DOI: [10.14722/ndss.2022.24093](https://doi.org/10.14722/ndss.2022.24093). URL: <https://hal.inria.fr/hal-03526240>.
- [32] P. Laperdrix, O. Starov, Q. Chen, A. Kapravelos and N. Nikiforakis. ‘Fingerprinting in Style: Detecting Browser Extensions via Injected Style Sheets’. In: 30th USENIX Security Symposium. Virtual, France, 11th Aug. 2021. URL: <https://hal.archives-ouvertes.fr/hal-03152176>.
- [33] G. Meiser, P. Laperdrix and B. Stock. ‘Careful Who You Trust: Studying the Pitfalls of Cross-Origin Communication’. In: ASIACCS 2021 - 16th ACM Asia Conference on Computer and Communications Security. 16th ACM Asia Conference on Computer and Communications Security. Hong Kong / Virtual, China, 7th June 2021. URL: <https://hal.archives-ouvertes.fr/hal-03021256>.
- [34] 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, 12th July 2021. URL: <https://hal.inria.fr/hal-03017222>.

- [35] 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>.
- [36] Z. Ournani, R. Rouvoy, P. Rust and J. Penhoat. ‘Comparing the Energy Consumption of Java I/O Libraries and Methods’. In: ICSME 2021 - 37th International Conference on Software Maintenance and Evolution. Proceedings of the 37th International Conference on Software Maintenance and Evolution (ICSME). Luxembourg / Virtual, Luxembourg, 27th Sept. 2021. URL: <https://hal.inria.fr/hal-03269129>.
- [37] Z. Ournani, R. Rouvoy, P. Rust and J. Penhoat. ‘Tales from the Code #1: The Effective Impact of Code Refactorings on Software Energy Consumption’. In: ICSOFT 2021 - 16th International Conference on Software Technologies. Virtual, France, 6th July 2021. URL: <https://hal.archives-ouvertes.fr/hal-03202437>.
- [38] 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, 6th Sept. 2021. URL: <https://hal.inria.fr/hal-03215569>.

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

- [39] S. Bliudze and L. Semini, eds. *2021 IEEE/ACM 9th International Conference on Formal Methods in Software Engineering (FormalISE 2021)*. 2021 IEEE/ACM 9th International Conference on Formal Methods in Software Engineering (FormalISE). Madrid: IEEE, 2021. DOI: [10.1109/FormalISE52586.2021](https://doi.org/10.1109/FormalISE52586.2021). URL: <https://hal.archives-ouvertes.fr/hal-03321958>.

Doctoral dissertations and habilitation theses

- [40] Z. Abou khalil. ‘Understanding the Impact of Release Policies on Software Development Processes.’ Université de Lille; Université de Mons, 26th Feb. 2021. URL: <https://tel.archives-ouvertes.fr/tel-03156973>.
- [41] A. Durey. ‘Leveraging browser fingerprinting to strengthen web authentication’. Université de Lille, 14th Jan. 2022. URL: <https://tel.archives-ouvertes.fr/tel-03544475>.
- [42] Z. Ournani. ‘Software Eco-Design: Investigating and Reducing the Energy Consumption of Software’. University of Lille, 8th Nov. 2021. URL: <https://tel.archives-ouvertes.fr/tel-03429300>.
- [43] W. Rudametkin. ‘Improving the Security and Privacy of the Web through Browser Fingerprinting’. Université de Lille, 21st June 2021. URL: <https://hal.inria.fr/tel-03370277>.

Reports & preprints

- [44] S. Bird, V. Mishra, S. Englehardt, R. Willoughby, D. Zeber, W. Rudametkin and M. Lopatka. *Actions speak louder than words: Semi-supervised learning for browser fingerprinting detection*. 23rd July 2021. URL: <https://hal.inria.fr/hal-03297181>.
- [45] A. Durey, P. Laperdrix, W. Rudametkin and R. Rouvoy. *An iterative technique to identify browser fingerprinting scripts*. 29th Apr. 2021. URL: <https://hal.archives-ouvertes.fr/hal-03212729>.
- [46] O. Le Goaer, A. Noureddine, F. Barbier, R. Rouvoy and F. Maraninchi. *Vers des Logiciels Éco-responsables: Le génie logiciel au défi de la sobriété écologique*. 13th Sept. 2021. URL: <https://hal.archives-ouvertes.fr/hal-03264970>.

12.3 Cited publications

- [47] S. Abiteboul, P. Bourhis and V. Vianu. ‘Explanations and Transparency in Collaborative Workflows’. In: *PODS 2018 - 37th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles Of Database Systems*. Houston, Texas, United States, June 2018. URL: <https://hal.inria.fr/hal-01744978>.

- [48] A. Amarilli, P. Bourhis, S. Mengel and M. Niewerth. ‘Constant-Delay Enumeration for Nondeterministic Document Spanners’. In: *ICDT*. 22nd International Conference on Database Theory (ICDT 2019). 25 pages including 17 pages of main material. Integrates all reviewer feedback. Outside of possible minor formatting differences, this paper is exactly the same as the ICDT’19 paper except that it contains 6 pages of technical appendix. Lisbon, Portugal, Mar. 2019. DOI: [10.4230/LIPIcs.ICDT.2019.19](https://doi.org/10.4230/LIPIcs.ICDT.2019.19). URL: <https://hal.telecom-paris.fr/hal-02338344>.
- [49] A. Amarilli, P. Bourhis, S. Mengel and M. Niewerth. ‘Enumeration on Trees with Tractable Combined Complexity and Efficient Updates’. In: *PODS*. Proceedings of the 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems. Amsterdam, France: ACM Press, June 2019, pp. 89–103. DOI: [10.1145/3294052.3319702](https://doi.org/10.1145/3294052.3319702). URL: <https://hal.archives-ouvertes.fr/hal-02326243>.
- [50] A. Amarilli, P. Bourhis, M. Monet and P. Senellart. ‘Evaluating Datalog via Tree Automata and Cycluits’. In: *Theory of Computing Systems* 63.7 (2019), pp. 1620–1678. DOI: [10.1007/s00224-018-9901-2](https://doi.org/10.1007/s00224-018-9901-2). URL: <https://hal.inria.fr/hal-01891811>.
- [51] 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).
- [52] M. Benedikt, P. Bourhis, G. Gottlob and P. Senellart. ‘Monadic Datalog, Tree Validity, and Limited Access Containment’. In: *ACM Transactions on Computational Logic* 21.1 (2020), 6:1–6:45. DOI: [10.1145/3344514](https://doi.org/10.1145/3344514). URL: <https://hal.inria.fr/hal-02307999>.
- [53] 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, Macau 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>.
- [54] 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).
- [55] 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>.
- [56] C. Calcagno, D. Distefano, J. Dubreil, D. Gabi, P. Hooimeijer, M. Luca, P. O’Hearn, I. Papakonstantinou, J. Purbrick and D. Rodriguez. ‘Moving Fast with Software Verification’. In: *NASA Formal Methods*. Ed. by K. Havelund, G. Holzmann and R. Joshi. Cham: Springer International Publishing, 2015, pp. 3–11.
- [57] M. Colmant. ‘Multi-Dimensional Analysis of Software Power Consumptions in Multi-Core Architectures’. Theses. Université Lille 1 - Sciences et Technologies, Nov. 2016. URL: <https://tel.archives-ouvertes.fr/tel-01403559>.
- [58] M. Colmant, P. Felber, R. Rouvoy and L. Seinturier. ‘WattsKit: Software-Defined Power Monitoring of Distributed Systems’. In: *17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*. Ed. by F. Capello, G. Fox and J. Garcia-Blas. Proceedings of the 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid). Madrid, Spain: IEEE, Jan. 2017, p. 10. URL: <https://hal.inria.fr/hal-01439889>.
- [59] A. Havet, V. Schiavoni, P. Felber, M. Colmant, R. Rouvoy and C. Fetzer. ‘GenPack: A Generational Scheduler for Cloud Data Centers’. In: *5th IEEE International Conference on Cloud Engineering (IC2E)*. Ed. by I. Gupta and J. Liu. Proceedings of the 5th IEEE International Conference on Cloud Engineering (IC2E). Vancouver, Canada: IEEE, Apr. 2017, p. 10. URL: <https://hal.inria.fr/hal-01403486>.
- [60] 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>.

- [61] D. Jackson. 'Alloy: A Lightweight Object Modelling Notation'. In: *ACM Transactions on Software Engineering and Methodology* 11.2 (Apr. 2002), pp. 256–290.
- [62] 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.
- [63] 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>.
- [64] P. Merle, A. Ndeye Sylla, M. Ouzzif, F. Klamm and K. Guillooard. 'A Lightweight Toolchain to Validate, Visualize, Analyze, and Deploy ETSI NFV Topologies'. In: *NetSoft 2019 - The 5th IEEE International Conference on Network Softwarization*. Paris, France, June 2019. DOI: [10.1109/NETSOFT.2019.8806632](https://doi.org/10.1109/NETSOFT.2019.8806632). URL: <https://hal.archives-ouvertes.fr/hal-02124164>.
- [65] P. Merle and J.-B. Stefani. *A formal specification of the Fractal component model in Alloy*. Research Report 6721. INRIA, Nov. 2008, p. 44. URL: <http://hal.inria.fr/inria-00338987>.
- [66] 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>.
- [67] 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>.
- [68] 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>.
- [69] A. Nouredine. 'Towards a Better Understanding of the Energy Consumption of Software Systems'. Theses. Université des Sciences et Technologie de Lille - Lille I, Mar. 2014. URL: <https://tel.archives-ouvertes.fr/tel-00961346>.
- [70] F. Paraiso, P. Merle and L. Seinturier. 'soCloud: A service-oriented component-based PaaS for managing portability, provisioning, elasticity, and high availability across multiple clouds'. In: *Springer Computing* 98.5 (May 2016), pp. 539–565. DOI: [10.1007/s00607-014-0421-x](https://doi.org/10.1007/s00607-014-0421-x). URL: <https://hal.inria.fr/hal-01019385>.
- [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] L. Seinturier, P. Merle, R. Rouvoy, D. Romero, V. Schiavoni and J.-B. Stefani. ‘A Component-Based Middleware Platform for Reconfigurable Service-Oriented Architectures’. English. In: *Software: Practice and Experience* 42.5 (May 2012), pp. 559–583. DOI: [10.1002/spe.1077](https://doi.org/10.1002/spe.1077). URL: <http://hal.inria.fr/inria-00567442>.
- [75] 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>.
- [76] A. N. Sylla, K. Guillouard, F. Klamm, M. Ouzzif, P. Merle, S. Ben Rayana and J.-B. Stefani. ‘Formal Verification of Orchestration Templates for Reliable Deployment with OpenStack Heat’. In: *CNSM 2019 - 15th International Conference on Network and Service Management*. Halifax, Canada, Oct. 2019, pp. 1–5. URL: <https://hal.inria.fr/hal-02375386>.
- [77] 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.
- [78] 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>.
- [79] P. Vila, P. Ganty, M. Guarnieri and B. Köpf. ‘CacheQuery: learning replacement policies from hardware caches’. In: *PLDI*. 2020.