

RESEARCH CENTRE

**Lille - Nord Europe**

IN PARTNERSHIP WITH:

**Université de Lille**

2020

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

<b>Project-Team SPIRALS</b>	<b>1</b>
<b>1 Team members, visitors, external collaborators</b>	<b>2</b>
<b>2 Overall objectives</b>	<b>3</b>
2.1 Introduction	3
2.2 Scientific Foundations	4
<b>3 Research program</b>	<b>4</b>
3.1 Introduction	4
3.2 Objective #1: Self-healing — Mining software artifacts to automatically evolve systems	4
3.2.1 Learning from software history how to design software and fix bugs	5
3.2.2 Run-time self-healing	6
3.3 Objective #2: Self-optimization — Sharing runtime behaviors to continuously adapt software	6
3.3.1 Monitoring software in the wild	7
3.3.2 Collaborative decision-making approaches	7
3.3.3 Smart reconfigurations in the large	8
<b>4 Application domains</b>	<b>8</b>
4.1 Introduction	8
4.2 Distributed software services for the retail industry	9
4.3 Distributed software services for the digital home	9
<b>5 Social and environmental responsibility</b>	<b>9</b>
5.1 Impact of research results	9
<b>6 Highlights of the year</b>	<b>9</b>
6.1 Awards	9
<b>7 New software and platforms</b>	<b>10</b>
7.1 New software	10
7.1.1 APISENSE	10
7.1.2 PowerAPI	10
7.1.3 amiunique	11
<b>8 New results</b>	<b>11</b>
8.1 Self-Adapting Software Services with Online Reinforcement Learning	11
8.2 Fingerprint-Based Web Security	12
<b>9 Bilateral contracts and grants with industry</b>	<b>12</b>
<b>10 Partnerships and cooperations</b>	<b>14</b>
10.1 International initiatives	14
10.1.1 Inria international partners	14
10.1.2 Participation in other international programs	14
10.2 European initiatives	15
10.2.1 Collaborations in European programs, except FP7 and H2020	15
10.3 National initiatives	15
10.3.1 ANR	15
10.3.2 CNRS Momentum	17
10.4 Regional initiatives	17

<b>11 Dissemination</b>	<b>19</b>
11.1 Promoting Scientific Activities	19
11.1.1 Scientific Events: Selection	19
11.1.2 Journal	20
11.1.3 Invited Talks	20
11.1.4 Leadership within the Scientific Community	20
11.1.5 Scientific Expertise	21
11.1.6 Research Administration	21
11.2 Teaching - Supervision - Juries	21
11.2.1 Teaching	21
11.2.2 Supervision	23
11.2.3 Juries	23
<b>12 Scientific production</b>	<b>24</b>
12.1 Major publications	24
12.2 Publications of the year	25
12.3 Cited publications	28

## Project-Team SPIRALS

*Creation of the Team: 2014 January 01, updated into Project-Team: 2015 January 01*

### Keywords

#### Computer sciences and digital sciences

- A1.3. – Distributed Systems
- A1.3.3. – Blockchain
- A1.3.5. – Cloud
- A1.4. – Ubiquitous Systems
- A1.6. – Green Computing
- A2.3.1. – Embedded systems
- A2.3.2. – Cyber-physical systems
- A2.4.2. – Model-checking
- A2.5. – Software engineering
- A2.5.2. – Component-based Design
- A2.5.3. – Empirical Software Engineering
- A2.5.4. – Software Maintenance & Evolution
- A2.6.2. – Middleware
- A3.1.3. – Distributed data
- A3.1.4. – Uncertain data
- A3.1.5. – Control access, privacy
- A3.1.9. – Database
- A3.2.1. – Knowledge bases
- A3.2.4. – Semantic Web
- A3.2.5. – Ontologies
- A4.4. – Security of equipment and software
- A4.8. – Privacy-enhancing technologies
- A7.2. – Logic in Computer Science
- A9.1. – Knowledge

#### Other research topics and application domains

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

# 1 Team members, visitors, external collaborators

## Research Scientists

- Simon Bliudze [Inria, Researcher]
- Pierre Bourhis [CNRS, Researcher]
- Pierre Laperdrix [CNRS, Researcher]
- Philippe Merle [Inria, 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 [Polytech Lille, Associate Professor]

## Post-Doctoral Fellows

- Alexandre Garnier [Inria]
- Adrien Luxey [Université de Lille, from Jun 2020]
- Hoang Gia Nguyen [Inria, until Apr 2020]
- Olivier Ruas [Inria, from Nov 2020]
- Larisa Safina [Inria, from Mar 2020]

## PhD Students

- Zeinab Abou Khalil [University of Mons, co-supervised with Université de Lille]
- Mohammed Chakib Belgaid [Inria]
- Sacha Brisset [Amaris Consulting, CIFRE]
- Antonin Durey [Université de Lille]
- Salman Farhat [Inria, from Oct 2020]
- Guillaume Fieni [Université de Lille]
- Romain Fouquet [Inria, from Oct 2020]
- Edouard Guegain [Université de Lille, from Jul 2020]
- Sarra Habchi [Inria, until Mar 2020]
- Jean Luc Intumwayase [Université de Lille, from Oct 2020]
- Trinh Le Khanh [Université de Lille]
- Lakhdar Meftah [Université de Lille, until Jul 2020]
- Naif Mehanna [Université de Lille, from Sep 2020]
- Vikas Mishra [Inria]
- Zakaria Ournani [Orange, CIFRE]
- Marion Tommasi [Inria]

## Technical Staff

- Antoine Canda [Inria, Engineer, until Sep 2020]
- Arthur D’Azémar [Inria, Engineer]
- Timothée Lefebvre [Université de Lille, Engineer, from Aug 2020]
- Alexis Pernet [Université de Lille, Engineer, from Aug 2020]
- Remy Raes [Inria, Engineer]
- Larisa Safina [Inria, Engineer, until Feb 2020]

## Interns and Apprentices

- Zakaria Boulkhir [Université de Lille, from Apr 2020 until Aug 2020]
- Tristan Coignon [Inria, from May 2020 until Aug 2020]
- Mamadou Sanoussy Diallo [Inria, from May 2020 until Aug 2020]
- Junior Fungwa Moke [Inria, from May 2020 until Aug 2020]
- Tussank Gupta [Inria, Sep 2020]
- Josue Happe [Inria, from May 2020 until Jul 2020]
- Joseph Marotte [Inria, from Mar 2020 until Aug 2020]
- Isaac Nivet [Inria, from May 2020 until Jun 2020]
- Andres Quintana Vesperinas [Inria, until Apr 2020]

## Administrative Assistant

- Karine Lewandowski [Inria]

## External Collaborator

- Martin Monperrus [KTH, on leave from Université de Lille, HDR]

## 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 work towards two directions: self-healing software systems with data mining solutions, and self-optimizing software systems with context monitoring. These two objectives are applied to two target environments: mobile computing and cloud computing.

## 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.*”<sup>1</sup> 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 [78], the NESSI Research Priorities for the next Framework Programme for Research and Technological Development FP8 [82], the Roadmap for Advanced Cloud Technologies under H2020 [79], and research roadmaps, such as [66, 69, 77].

## 3 Research program

### 3.1 Introduction

Our research program on self-adaptive software targets two key properties that are detailed in the remainder of this section: *self-healing* and *self-optimization*.

### 3.2 Objective #1: Self-healing — Mining software artifacts to automatically evolve systems

Software systems are under the pressure of changes all along their lifecycle. Agile development blurs the frontier between design and execution and requires constant adaptation. The size of systems (millions of lines of code) multiplies the number of bugs by the same order of magnitude. More and more systems, such as sensor network devices, live in “surviving” mode, in the sense that they are neither rebootable nor upgradable.

Software bugs are hidden in source code and show up at development-time, testing-time or worse, once deployed in production. Except for very specific application domains where formal proofs are achievable, bugs can not be eradicated. As an order of magnitude, on the 16<sup>th</sup> of December 2011, the Eclipse bug repository contained 366 922 bug reports. Software engineers and developers work on bug fixing on a daily basis. Not all developers spend the same time on bug fixing. In large companies, this is sometimes a full-time role to manage bugs, often referred to as *Quality Assurance* (QA) software engineers. Also, not all bugs are equal, some bugs are analyzed and fixed within minutes, others may take months to be solved [85].

In terms of research, this means that: (i) one needs means to automatically adapt the design of the software system through automated refactoring and API extraction, (ii) one needs approaches to automate the process of adapting source code in order to fix certain bugs, (iii) one needs to revisit the notion of error-handling so that instead of crashing in presence of errors, software adapts itself to continue with its execution, e.g. in degraded mode.

There is no one-size-fits-all solution for each of these points. However, we think that novel solutions can be found by using **data mining and machine learning techniques tailored for software engineering** [86]. This body of research consists of mining some knowledge about a software system by analyzing the source code, the version control systems, the execution traces, documentation and all kinds of software

---

<sup>1</sup><https://web.archive.org/web/20170619225302/http://cordis.europa.eu/fp7/ict/ssai/>

development and execution artifacts in general. This knowledge is then used within recommendation systems for software development, auditing tools, runtime monitors, frameworks for resilient computing, etc.

The novelty of our approach consists of using and tailoring data mining techniques for analyzing software artifacts (source code, execution traces) in order to achieve the **next level of automated adaptation** (e.g. automated bug fixing). Technically, we plan to mix unsupervised statistical learning techniques (e.g. frequent item set mining) and supervised ones (e.g. training classifiers such as decision trees). This research is currently not being performed by data mining research teams since it requires a high level of domain expertise in software engineering, while software engineering researchers can use off-the-shelf data mining libraries, such as Weka [63].

We now detail the two directions that we propose to follow to achieve this objective.

### 3.2.1 Learning from software history how to design software and fix bugs

The first direction is about mining techniques in software repositories (e.g. CVS, SVN, Git). Best practices can be extracted by data mining source code and the version control history of existing software systems. The design and code of expert developers significantly vary from the artifacts of novice developers. We will learn to differentiate those design characteristics by comparing different code bases, and by observing the semantic refactoring actions from version control history. Those design rules can then feed the test-develop-refactor constant adaptation cycle of agile development.

**Fault localization of bugs reported in bug repositories.** We will build a solid foundation on empirical knowledge about bugs reported in bug repository. We will perform an empirical study on a set of representative bug repositories to identify classes of bugs and patterns of bug data. For this, we will build a tool to browse and annotate bug reports. Browsing will be helped with two kinds of indexing: first, the tool will index all textual artifacts for each bug report; second it will index the semantic information that is not present by default in bug management software, i.e. “contains a stacktrace”). Both indexes will be used to find particular subsets of bug reports, for instance “all bugs mentioning invariants and containing a stacktrace”. Note that queries with this kind of complexity and higher are mostly not possible with the state-of-the-art of bug management software. Then, analysts will use annotation features to annotate bug reports. The main outcome of the empirical study will be the identification of classes of bugs that are appropriate for automated localization. Then, we will run machine learning algorithms to identify the latent links between the bug report content and source code features. Those algorithms would use as training data the existing traceability links between bug reports and source code modifications from version control systems. We will start by using decision trees since they produce a model that is explicit and understandable by expert developers. Depending on the results, other machine learning algorithms will be used. The resulting system will be able to locate elements in source code related to a certain bug report with a certain confidence.

**Automated bug fix generation with search-based techniques.** Once a location in code is identified as being the cause of the bug, we can try to automatically find a potential fix. We envision different techniques: (1) infer fixes from existing contracts and specifications that are violated; (2) infer fixes from the software behavior specified as a test suite; (3) try different fix types one-by-one from a list of identified bug fix patterns; (4) search fixes in a fix space that consists of combinations of atomic bug fixes. Techniques 1 and 2 are explored in [59] and [84]. We will focus on the latter techniques. To identify bug fix patterns and atomic bug fixes, we will perform a large-scale empirical study on software changes (also known as changesets when referring to changes across multiple files). We will develop tools to navigate, query and annotate changesets in a version control system. Then, a grounded theory will be built to master the nature of fixes. Eventually, we will decompose change sets in atomic actions using clustering on changeset actions. We will then use this body of empirical knowledge to feed search-based algorithms (e.g. genetic algorithms) that will look for meaningful fixes in a large fix space. To sum up, our research on automated bug fixing will try not only to point to source code locations responsible of a bug, but to search for code patterns and snippets that may constitute the skeleton of a valid patch. Ultimately, a blend of expert heuristics and learned rules will be able to produce valid source code that can be validated by developers and committed to the code base.



### 3.2.2 Run-time self-healing

The second proposed research direction is about inventing a self-healing capability at run-time. This is complementary to the previous objective that mainly deals with development time issues. We will achieve this in two steps. First, we want to define frameworks for resilient software systems. Those frameworks will help to maintain the execution even in the presence of bugs, i.e. to let the system survive. As exposed below, this may mean for example to switch to some degraded modes. Next, we want to go a step further and to define solutions for automated runtime repair, that is, not simply compensating the erroneous behavior, but also determining the correct repair actions and applying them at run-time.

**Mining best effort values.** A well-known principle of software engineering is the “fail-fast” principle. In a nutshell, it states that as soon as something goes wrong, software should stop the execution before entering incorrect states. This is fine when a human user is in the loop, capable of understanding the error or at least rebooting the system. However, the notion of “failure-oblivious computing” [76] shows that in certain domains, software should run in a resilient mode (i.e. capable of recovering from errors) and/or best-effort mode, i.e. a slightly imprecise computation is better than stopping. Hence, we plan to investigate data mining techniques in order to learn best-effort values from past executions (i.e. somehow learning what is a correct state, or the opposite what is not a completely incorrect state). This knowledge will then be used to adapt the software state and flow in order to mitigate the error consequences, the exact opposite of fail-fast for systems with long-running cycles.

**Embedding search based algorithms at runtime.** Harman recently described the field of search-based software engineering [64]. We believe that certain search based approaches can be embedded at runtime with the goal of automatically finding solutions that avoid crashing. We will create software infrastructures that allow automatically detecting and repairing faults at run-time. The methodology for achieving this task is based on three points: (1) empirical study of runtime faults; (2) learning approaches to characterize runtime faults; (3) learning algorithms to produce valid changes to the software runtime state. An empirical study will be performed to analyze those bug reports that are associated with runtime information (e.g. core dumps or stacktraces). After this empirical study, we will create a system that learns on previous repairs how to produce small changes that solve standard runtime bugs (e.g. adding an array bound check to throw a handled domain exception rather than a spurious language exception). To achieve this task, component models will be used to (1) encapsulate the monitoring and reparation meta-programs in appropriate components and (2) support runtime code modification using scripting, reflective or bytecode generation techniques.

### 3.3 Objective #2: Self-optimization — Sharing runtime behaviors to continuously adapt software

Complex distributed systems have to seamlessly adapt to a wide variety of deployment targets. This is due to the fact that developers cannot anticipate all the runtime conditions under which these systems are immersed. A major challenge for these software systems is to develop their capability to continuously reason about themselves and to take appropriate decisions and actions on the optimizations they can apply to improve themselves. This challenge encompasses research contributions in different areas, from environmental monitoring to real-time symptoms diagnosis, to automated decision making. The variety of distributed systems, the number of optimization parameters, and the complexity of decisions often resign the practitioners to design monolithic and static middleware solutions. However, it is now globally acknowledged that the development of dedicated building blocks does not contribute to the adoption of sustainable solutions. This is confirmed by the scale of actual distributed systems, which can, for example, connect several thousands of devices to a set of services hosted in the Cloud. In such a context, the lack of support for smart behaviors at different levels of the systems can inevitably lead to its instability or its unavailability. In June 2012, an outage of Amazon’s Elastic Compute Cloud in North Virginia has taken down Netflix, Pinterest, and Instagram services. During hours, all these services failed to satisfy their millions of customers due to the lack of integration of a self-optimization mechanism going beyond the boundaries of Amazon.

The research contributions we envision within this area will therefore be organized as a reference model for engineering **self-optimized distributed systems** autonomously driven by *adaptive feedback control loops*, which will automatically enlarge their scope to cope with the complexity of the decisions to

be taken. This solution introduces a multi-scale approach, which first privileges local and fast decisions to ensure the homeostasis<sup>2</sup> property of a single node, and then progressively propagates symptoms in the network in order to reason on a longer term and a larger number of nodes. Ultimately, domain experts and software developers can be automatically involved in the decision process if the system fails to find a satisfying solution. The research program for this objective will therefore focus on the study of mechanisms for **monitoring, taking decisions, and automatically reconfiguring software at runtime and at various scales**. As stated in the self-healing objective, we believe that there is no one-size-fits-all mechanism that can span all the scales of the system. We will therefore study and identify an optimal composition of various adaptation mechanisms in order to produce long-living software systems.

The novelty of this objective is to exploit the wisdom of crowds to define new middleware solutions that are able to continuously adapt software deployed in the wild. We intend to demonstrate the applicability of this approach to distributed systems that are deployed from mobile phones to cloud infrastructures. The key scientific challenges to address can be summarized as follows: *How does software behave once deployed in the wild? Is it possible to automatically infer the quality of experience, as it is perceived by users? Can the runtime optimizations be shared across a wide variety of software? How optimizations can be safely operated on large populations of software instances?*

The remainder of this section further elaborates on the opportunities that can be considered within the frame of this objective.

### 3.3.1 Monitoring software in the wild

Once deployed, developers are generally no longer aware of how their software behave. Even if they heavily use testbeds and benchmarks during the development phase, they mostly rely on the bugs explicitly reported by users to monitor the efficiency of their applications. However, it has been shown that contextual artifacts collected at runtime can help to understand performance leaks and optimize the resilience of software systems [87]. Monitoring and understanding the context of software at runtime therefore represent the first building block of this research challenge. Practically, we intend to investigate crowd-sensing approaches, to smartly collect and process runtime metrics (e.g. request throughput, energy consumption, user context). Crowd-sensing can be seen as a specific kind of **crowdsourcing** activity, which refers to the capability of lifting a (large) diffuse group of participants to delegate the task of retrieving trustable data from the field. In particular, crowd-sensing covers not only *participatory sensing* to involve the user in the sensing task (e.g. surveys), but also *opportunistic sensing* to exploit mobile sensors carried by the user (e.g. smartphones).

While reported metrics generally enclose raw data, the monitoring layer intends to produce meaningful indicators like the *Quality of Experience* (QoE) perceived by users. This QoE reflects representative symptoms of software requiring to trigger appropriate decisions in order to improve its efficiency. To diagnose these symptoms, the system has to process a huge variety of data including runtime metrics, but also history of logs to explore the sources of the reported problems and identify opportunities for optimizations. The techniques we envision at this level encompass **machine learning, principal component analysis**, and fuzzy logic [75] to provide enriched information to the decision level.

### 3.3.2 Collaborative decision-making approaches

Beyond the symptoms analysis, decisions should be taken in order to improve the *Quality of Service* (QoS). In our opinion, collaborative approaches represent a promising solution to effectively converge towards the most appropriate optimization to apply for a given symptom. In particular, we believe that exploiting the **wisdom of the crowd** can help the software to optimize itself by sharing its experience with other software instances exhibiting similar symptoms. The intuition here is that the body of knowledge that supports the optimization process cannot be specific to a single software instance as this would restrain the opportunities for improving the quality and the performance of applications. Rather, we think that any software instance can learn from the experience of others.

With regard to the state-of-the-art, we believe that a multi-levels decision infrastructure, inspired from distributed systems like Spotify [61], can be used to build a decentralized decision-making algorithm

---

<sup>2</sup>Homeostasis is the property of a system that regulates its internal environment and tends to maintain a stable, relatively constant condition of properties [Wikipedia].

involving the surrounding peers before requesting a decision to be taken by more central control entity. In the context of collaborative decision-making, peer-based approaches therefore consist in quickly reaching a consensus on the decision to be adopted by a majority of software instances. Software instances can share their knowledge through a micro-economic model [56], that would weight the recommendations of experienced instances, assuming their age reflects an optimal configuration.

Beyond the peer level, the adoption of algorithms inspired from evolutionary computations, such as **genetic programming**, at an upper level of decision can offer an opportunity to test and compare several alternative decisions for a given symptom and to observe how does the crowd of applications evolves. By introducing some diversity within this population of applications, some instances will not only provide a satisfying QoS, but will also become naturally resilient to unforeseen situations.

### 3.3.3 Smart reconfigurations in the large

Any decision taken by the crowd requires to propagate back to and then operated by the software instances. While simplest decisions tend to impact software instances located on a single host (e.g. laptop, smartphone), this process can also exhibit more complex reconfiguration scenarios that require the orchestration of various actions that have to be safely coordinated across a large number of hosts. While it is generally acknowledged that centralized approaches raise scalability issues, we think that self-optimization should investigate different reconfiguration strategies to propagate and apply the appropriate actions. The investigation of such strategies can be addressed in two steps: the consideration of *scalable data propagation protocols* and the identification of *smart reconfiguration mechanisms*.

With regard to the challenge of scalable data propagation protocols, we think that research opportunities encompass not only the exploitation of gossip-based protocols [60], but also the adoption of publish/subscribe abstractions [70] in order to decouple the decision process from the reconfiguration. The fundamental issue here is the definition of a communication substrate that can accommodate the propagation of decisions with relaxed properties, inspired by *Delay Tolerant Networks* (DTN), in order to reach weakly connected software instances. We believe that the adoption of asynchronous communication protocols can provide the sustainable foundations for addressing various execution environments including harsh environments, such as developing countries, which suffer from a partial connectivity to the network. Additionally, we are interested in developing the principle of *social networks of applications* in order to seamlessly group and organize software instances according to their similarities and acquaintances. The underlying idea is that grouping application instances can contribute to the identification of optimization profiles not only contributing to the monitoring layer, but also interested in similar reconfigurations. Social networks of applications can contribute to the anticipation of reconfigurations by exploiting the symptoms of similar applications to improve the performance of others before that problems actually happen.

With regard to the challenge of smart reconfiguration mechanisms, we are interested in building on our established experience of adaptive middleware [80] in order to investigate novel approaches to efficient application reconfigurations. In particular, we are interested in adopting seamless micro-updates and micro-reboot techniques to provide in-situ reconfiguration of pieces of software. Additionally, the provision of safe and secured reconfiguration mechanisms is clearly a key issue that requires to be carefully addressed in order to avoid malicious exploitation of dynamic reconfiguration mechanisms against the software itself. In this area, although some reconfiguration mechanisms integrate transaction models [68], most of them are restricted to local reconfigurations, without providing any support for executing distributed reconfiguration transactions. Additionally, none of the approached published in the literature include security mechanisms to preserve from unauthorized or malicious reconfigurations.

## 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<sup>®</sup> 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 ongoing PhD theses of Guillaume Fieni, Zakaria Ournani, and Chakib Belgaid, contribute to this line of research. With the now defended PhD theses of Maxime Colmant [57] and Adel Nourredine [73], this is the 3<sup>rd</sup> 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 (ongoing CIFRE PhD thesis of Zakaria Ournani) and OVHcloud (ongoing bilateral contract) 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. In 2020, we obtained and published several new results on green and power efficient computing [29, 36, 37].

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

## 6.1 Awards

Our research activities in 2020 led to four awards that are highlighted in this section.

Sarra Habchi won one of the two accessits to the Prix de thèse Gilles Kahn 2020 for her work on understanding mobile-specific code smells. Sarra Hachi defended her PhD thesis [62] in 2019 under the supervision of Romain Rouvroy. She is now a postdoctoral researcher at the University of Luxembourg.

The work of Sarra contributed to a better understanding and some solutions to prevent bad design and development practices commonly observed in software systems, especially those for mobile applications.

Walter Rudametkin, Romain Rouvoy, and their co-authors, won a best paper award for their paper entitled “FP-Crawlers: Studying the Resilience of Browser Fingerprinting to Block Crawlers” at the workshop on Measurements, Attacks, and Defenses for the Web (MADWeb) in the context of the The Network and Distributed System Security Symposium 2020 [38].

Pierre Bourhis and his co-authors were awarded a SIGMOD research highlight award for their paper entitled “Constant-Delay Enumeration for Nondeterministic Document Spanner” [13].

Clément Quinton and his co-authors won a best paper award for their paper entitled “Feature-Model-Guided Online Reinforcement Learning for Self-Adaptive Services” at the 18th International Conference on Service Oriented Computing (ICSOC 2020) [6].

## 7 New software and platforms

### 7.1 New software

#### 7.1.1 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

#### 7.1.2 PowerAPI

**Keywords:** Energy efficiency, Energy management

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

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

**URL:** <http://powerapi.org>

**Contact:** Romain Rouvoy

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

### 7.1.3 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

## 8 New results

In 2020, we obtained new results in the domains of self-adaption with online reinforcement learning, and fingerprint-based web privacy.

### 8.1 Self-Adapting Software Services with Online Reinforcement Learning

A self-adaptive service can maintain its QoS requirements in the presence of dynamic environment changes. To develop a self-adaptive service, service engineers have to create self-adaptation logic encoding when the service should execute which adaptation actions. However, developing self-adaptation logic may be difficult due to design time uncertainty; e.g. anticipating all potential environment changes at design time is in most cases infeasible. Online reinforcement learning addresses design time uncertainty

by learning suitable adaptation actions through interactions with the environment at runtime. To learn more about its environment, reinforcement learning has to select actions that were not selected before, which is known as exploration. How exploration happens has an impact on the performance of the learning process. We focus on two problems related to how a service's adaptation actions are explored: (1) Existing solutions randomly explore adaptation actions and thus may exhibit slow learning if there are many possible adaptation actions to choose from. (2) Existing solutions are unaware of service evolution, and thus may explore new adaptation actions introduced during such evolution rather late. We propose novel exploration strategies that use feature models (from software product line engineering) to guide exploration in the presence of many adaptation actions and in the presence of service evolution. Experimental results for a self-adaptive cloud management service indicate an average speed-up of the learning process of 58.8% in the presence of many adaptation actions, and of 61.3% in the presence of service evolution. The improved learning performance in turn led to an average QoS improvement of 7.8% and 23.7% respectively.

This result has been published in [6]. This paper has obtained the best paper award at the 18th International Conference on Service Oriented Computing (ICSOC 2020). This work is a collaboration with the University of Duisburg-Essen and Politecnico di Milano.

## 8.2 Fingerprint-Based Web Security

Data available on the Web, such as financial data or public reviews, provides a competitive advantage to companies able to exploit them. Web crawlers, a category of bot, aim at automating the collection of publicly available Web data. While some crawlers collect data with the agreement of the websites being crawled, most crawlers do not respect the terms of service. CAPTCHAs and approaches based on analyzing series of HTTP requests classify users as humans or bots. However, these approaches require either user interaction or a significant volume of data before they can classify the traffic. In this paper, we study browser fingerprinting as a crawler detection mechanism. We crawled the Alexa top 10K and identified 291 websites that block crawlers. We show that fingerprinting is used by 93 (31.96%) of them and we report on the crawler detection techniques implemented by the major fingerprinters. Finally, we evaluate the resilience of fingerprinting against crawlers trying to conceal themselves. We show that although fingerprinting is good at detecting crawlers, it can be bypassed with little effort by an adversary with knowledge on the fingerprints collected.

This result has been published in [38]. This paper has obtained the best paper award at the workshop on Measurements, Attacks, and Defenses for the Web (MADWeb) in the context of the The Network and Distributed System Security Symposium 2020. This work is a collaboration with the Datadome web security company and the University of Bordeaux.

## 9 Bilateral contracts and grants with industry

### ip-label

**Participants** Romain Rouvoy (*contact person*).

A software exploitation license (2014–ongoing) of the APISENSE<sup>®</sup> 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 [71, 81]. This computational model 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 [67]. This collaboration takes advantage of the expertise that we are developing since several years on reconfigurable component-based software systems [80], on cloud systems [74], and on the Alloy specification language [72].

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 takes 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** 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 Inria international partners

#### Informal international partners

- University of Oslo: green computing.
- University of Mons: software engineering.
- TU Berlin & University of Bayreuth.
- Pontifical Catholic University of Chile.
- University of Göttingen: cloud computing.
- CETIC: cloud computing.

### 10.1.2 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, Alexandre Garnier, 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 Collaborations in European programs, except FP7 and H2020

#### DG ECHO Aiders

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

**Aiders** is an European project (2020–21) 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.

## 10.3 National initiatives

### 10.3.1 ANR

#### 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, 65].

#### 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 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 [16] [48, 49, 50].

### 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.

### 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.

### 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.

### 10.3.2 CNRS Momentum

#### Manage Your Data Without Information Leakage

**Participants** Pierre Bourhis (*contact person*).

“Gérer vos données sans fuite d’information” is a 3-year (2018–20) project granted in the context of the CNRS-Momentum call for projects. Data manipulated by modern applications are stored in large databases. To protect these pieces of data, security policies limit a user’s access to what she is allowed to see. However, by using the semantics of the data, a user can deduce information that she was not supposed to have access to. The goal of this project is to establish methods and tools for understanding and detecting such data leaks. Several results and publications have been obtained in the context of this project [55, 51, 52].

## 10.4 Regional initiatives

### CIRRUS

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

CIRRUS is a 3-year (2017–20) joint team with the Scalair cloud operator and architect company funded by the Hauts-de-France region. The CIRRUS joint team is developing novel solutions in the domains of the on demand configuration of heterogeneous cloud resources, the management of cloud elasticity for all deployed services (SaaS, PaaS, IaaS) in order to guarantee quality of service and user quality of experience, and the taming of financial costs of cloud infrastructures.

### Rigorous Component-Based Design of Correct-by-Construction Software and Systems: Application to Cloud Computing

**Participants** Simon Bliudze (*contact person*), Trinh Le Khanh, Larisa Safina.

This 24-month (2019–20) project is funded in the context of the STaRS program of Hauts-de-France region. It aims at the development of methods and tools for the rigorous design of cloud computing platforms and applications, which can be proven to be correct by construction. First results have been published in [14] [53, 54].

### **BRiCoS: Building Blocks for Rigorous Design of Concurrent Software**

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

This 24-month (2021–22) project — funded by the I-SITE ULNE Foundation in the framework of its “Support for young researchers 2020” programme — 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.

### **ASCOT: Advanced Script analysis to prevent Online Tracking**

**Participants** Romain Fouquet, Jean Luc Intumwayase, Pierre Laperdrix (*contact person*), Romain Rouvoy, Walter Rudametkin Ivey.

This 36-month (2020–23) project is 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.

### **Indoor Analytics**

**Participants** Pierre Bourhis, Remy Raes, Romain Rouvoy (*contact person*), Lionel Seinturier.

Indoor Analytics is a 32-month (2019–21) project funded in the context of CPER Data program. Indoor Analytics aims at collaborating with the Mapwize company on the development of novel analytics for indoor location systems. In particular, Mapwize and Spirals target the joint delivery of an open-source software solution devoted to the acquisition, storage and processing of location events at scale.

### **COMMODE**

**Participants** Pierre Bourhis, Laurence Duchien, Clément Quinton (*contact person*).

COMMODE (Knowledge COMpilation for feature MODEls) is a 24-month (2019–21) project funded in the context of CPER Data program. COMMODE aims at using techniques from knowledge compilation, a subarea of artificial intelligence, for feature models, a representation of software products used in software engineering.

## ADT FingerKit

**Participants** Antoine Canda, Antonin Durey, Pierre Laperdrix, Walter Rudametkin Ivey (*contact person*).

ADT FingerKit (2018–20) is a technology development initiative supported by the Inria Lille - Nord Europe Center that focuses on the design and development of a new and enhanced version of the **AmIUnique** platform. AmIUnique is a data collection and analysis platform to better understand, analyze and vulgarize the uses and threats of browser fingerprinting. This initiative led by Inria is a key asset to better understand novel techniques that threatens the user privacy on Internet. This ADT builds on our first results with the PhD thesis of Antoine Vastel [83].

## ADT e-Lens

**Participants** Arthur d’Azémar, Guillaume Fieni, Romain Rouvoy (*contact person*).

ADT e-Lens (2018–21) is a technology development initiative supported by the Inria Lille - Nord Europe Center that aims at extending the **PowerAPI** energy monitoring library that we develop in the team since 2011. The extension deals with the integration of new power models (for GPU, disk, network interface), the implementation of a self-optimization algorithm, the port of the platform to embedded systems running with Raspberry Pi, ROS and Android, and the implementation of an active learning algorithm for power models. This ADT builds on our results with the defended PhD theses of Adel Nouredine [73] and Maxime Colmant [57], and with the ongoing PhD thesis of Guillaume Fieni.

# 11 Dissemination

## 11.1 Promoting Scientific Activities

### 11.1.1 Scientific Events: Selection

#### Chair of Conference Program Committees

**Simon Bliudze:** International Conference on Coordination Models and Languages (COORDINATION).

#### Member of the Conference Program Committees

**Simon Bliudze:** International Conference on Formal Methods in Software Engineering (FormaliSE), NASA Formal Methods Symposium (NFM), International Colloquium on Theoretical Aspects of Computing (ICTAC), IFIP International Conference on Topics in Theoretical Computer Science (TTCS, Track B), International Conference on Microservices, International Conference on Verification and Evaluation of Computer and Communication Systems (VECoS).

**Pierre Bourhis:** International Conference on Database Theory (ICDT), ACM SIGMOD/PODS International Conference on Management of Data.

**Laurence Duchien:** ACM International Systems and Software Product Line Conference (SPLC), European Conference on Software Architecture (ECSA), European Conference on Software Architecture Doctoral Symposium, IEEE International Conference on Software Architecture (ICSA), International Workshop on Software Engineering Research and Industrial Practice (SER&IP).

**Pierre Laperdrix:** USENIX Security, Conference on Detection of Intrusions and Malware & Vulnerability Assessment (DIMVA), USENIX Workshop on Offensive Technologies (WOOT), Workshop on Privacy in the Electronic Society (WPES).

**Philippe Merle:** International Conference on Cooperative Information Systems (CoopIS), IEEE International Conference on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE) - Future Internet Services and Applications track (FISA), International Conference on Mobile Ubiquitous Computing, Systems, Services and Technologies (UBICOMM), International Conference on Adaptive

and Self-Adaptive Systems and Applications (ADAPTIVE), International Conference on Evolving Internet (INTERNET), International Conference on Cloud Computing, GRIDs, and Virtualization (CLOUD COMPUTING), International Conference on Advanced Service Computing (SERVICE COMPUTATION), International Symposium on Security in Computing and Communications (SSCC), International Conference on Applied Soft Computing and Communication Networks (ACN), Conférence en Ingénierie du Logiciel (CIEL), International Workshop on Formal Approaches for Advanced Computing Systems (FAACS).

**Clément Quinton:** ACM Symposium on Applied Computing (SAC) track SA-TTA, ACM International Systems and Software Product Line Conference (SPLC) tracks RA DT.

**Romain Rouvoy:** ACM Symposium on Applied Computing (SAC) track SA-TTA, ACM/IFIP International Middleware Conference (Middleware), IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGrid), International Conference on Software and Systems Reuse (ICSR), International Conference on Distributed Applications and Interoperable Systems (DAIS), ACM International Conference on Distributed and Event-based Systems (DEBS) industry track.

**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).

**Walter Rudametkin:** Workshop on Measurements, Attacks, and Defenses for the Web (MADWeb).

### 11.1.2 Journal

#### Member of the Editorial Boards

**Philippe Merle:** IoT — Open Access Journal (MDPI).

**Walter Rudametkin:** Springer Annals of Telecommunications.

#### Reviewer – Reviewing Activities

**Simon Bliudze:** Journal of Software: Evolution and Process (formerly The Journal of Software Maintenance and Evolution; JSME)

**Pierre Bourhis:** ACM Transactions on Database Systems (TODS).

**Philippe Merle:** ACM Computing Surveys (CSUR), IEEE Access, Springer Software and Systems Modeling (SoSyM), IEEE Software, Wiley Software: Practice and Experience (SPE), ACM Transactions on Internet Technology (TOIT), Wiley Concurrency and Computation: Practice and Experience (CCPE), Journal of Robotics (Hindawi).

**Clément Quinton:** Journal of Systems and Software (JSS).

**Lionel Seinturier:** Elsevier Future Generation Computing Systems (FGCS), Springer Empirical Software Engineering (EMSE), Elsevier Internet of Things, Springer Software and Systems Modeling (SoSyM).

### 11.1.3 Invited Talks

#### 11.1.4 Leadership within the Scientific Community

**Simon Bliudze** is co-head of the YODA (trustworthy and Optimal Dynamic Adaptation) working group of the GDR GPL. He was guest editor for the Springer Innovations in Systems and Software Engineering journal special issue with selected papers presented at VECoS '18, the Springer Software Tools for Technology Transfer journal special issue with selected papers presented at MeTRiD '18–'19, and the CCSD Logical Methods in Computer Science journal special issue with selected papers presented at COORDINATION '20. Since June 2020, he is a member of the Steering Committee of COORDINATION.

**Laurence Duchien** is a member of the scientific councils of IMT Atlantique (Nantes), Labex CIMI (Toulouse), IRT SystemX (Saclay). She is member of the steering committees of the European Conference on Software Architecture (ECSA) and of the Systems and Software Product Lines Conference (SPLC). She was guest editor for the Empirical Software Engineering Journal for a special issue on software configuration.

**Philippe Merle** is an elected member of the Inria scientific board (CS).

**Romain Rouvoy** is an elected member of the “bureau” of the French chapter of the ACM Special Interest Group in Operating Systems (SIGOPS / ASF) and elected member of the administrative council of

SpecifCampus. He is co-head of the “Groupe de Travail Génie Logiciel pour les Systèmes Cyber-physiques” of the GDR GPL.

### 11.1.5 Scientific Expertise

**Simon Bliudze** was member of the recruitment committee of junior researchers (CRCN and ISFP) at Inria Lille. He was a member of the chair professor recruitment committee at UCLouvain.

**Laurence Duchien** was member of the recruitment committee for two researchers at the University of Luxembourg. She was president of the recruitment committee for Professor in Computer Science at ENS Paris. She was president of the CNRS Research Scientists committee for section CoNRS 6. She was member of the jury for the GDR GPL Prix de thèse. She was expert for MESRI DGRI and evaluated projects for the PHC and STIC-Asie programs.

**Romain Rouvoy** was member of scientific committee 25 — Infrastructures (networks, high performance computing and storage), software sciences and technologies of the French Research Agency (ANR). He was scientific expert for DRRT IDE.

**Lionel Seinturier** was president of the recruitment committee for Professor in Computer Science at the University of Lille, and member of the recruitment committee for Associate Professor in Computer Science at the University of Paris. He was member of the Institut Universitaire de France (IUF) Senior Jury. He was scientific expert for the Swiss National Science Foundation, DRRT IDE.

### 11.1.6 Research Administration

**Simon Bliudze** is an elected member of the comité de centre Inria Lille Nord Europe.

**Laurence Duchien** is member of the CNRS CoNRS section 6 committee, and of the “bureau” of this committee. She is 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.

**Antonin Durey** is elected member of the CRISAL laboratory council.

**Philippe Merle** is an 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), and elected member of the centre committee for the Inria Lille – Nord Europe research center. He is member of the steering committee of the Inria’s continuous integration service. He is member of SPI doctoral school council of University of Lille.

**Walter Rudametkin** is an member of the CDT (Comité de Développement Technologique) of the Inria Lille – Nord Europe research center. He is an elected member of the CRISAL laboratory council.

**Lionel Seinturier** is a member of the BSC (Bureau Scientifique du Centre) for the Inria Lille – Nord Europe research center. He is chair of CNU27, the Computer Science section of the “Conseil National des Universités”.

## 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, 2<sup>nd</sup> year of the Engineering cycle
- INF442: Traitement des données massives, 40h, 2<sup>nd</sup> 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

**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 Rouvoy** 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, 50h, Level M1, Master MIAGE
- Infrastructures et Frameworks Internet, 70h, Level M2 E-Services IAGL TIIR, Master of Computer Science

### 11.2.2 Supervision

- PhD in progress: Zeinab Abou Khalil, Understanding the impact of release policies in software development processes, November 2017, Laurence Duchien & Clément Quinton, co-supervision with Tom Mens (University of Mons).
- PhD in progress: Mohammed Chakib Belgaid, Développement durable des logiciels vers une optimisation énergétique de bout en bout des systèmes logiciels, January 2018, 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, November 2018, Lionel Seinturier & Romain Rouvoy & Renaud Pawlak & Yoann Couillec.
- PhD in progress: Antonin Durey, Leveraging Browser Fingerprinting to Fight Fraud on the Web, October 2018, Romain Rouvoy & Walter Rudametkin & Lionel Seinturier.
- PhD in progress: Salman Farhat, Safe Dynamic Reconfiguration of Cloud Applications, October 2020, Laurence Duchien & Simon Bliudze.
- PhD in progress: Guillaume Fieni, GreenData : Vers un traitement efficient et éco-responsable des grandes masses de données numériques, October 2017, Romain Rouvoy & Lionel Seinturier.
- PhD in progress: Romain Fouquet, Improving Online Privacy through Content Blocking and Information Restriction, October 2020, Romain Rouvoy & Pierre Laperdrix.
- PhD in progress: Omid Gheibi, Dependable Adaptive Software Systems for the Digital World, October 2020, Laurence Duchien & Clément Quinton, co-supervision with Danny Weyns (KU Leuven).
- PhD in progress: Edouard Guegain, Configuration of Large Scale Fog Environments, October 2020, Clément Quinton.
- PhD in progress: Jean-Luc Intumwayase, Improving Online Privacy through Content Blocking and Information Restriction, November 2020, Romain Rouvoy & Pierre Laperdrix.
- PhD in progress: Trinh Lê Khánh, Design of Correct-by-Construction Self-Adaptive Cloud Applications using Formal Methods, October 2019, Philippe Merle & Simon Bliudze.
- PhD in progress: Naif Mehanna, Hardening Web Authentication with Browser Fingerprinting, October 2020, Walter Rudametkin.
- PhD in progress: Vikas Mishra, Collaborative Strategies to Protect Against Browser Fingerprinting, October 2018, Romain Rouvoy, Walter Rudametkin & Lionel Seinturier.
- PhD in progress: Zakaria Ournani, Eco-conception des logiciels : modélisation de l'efficacité énergétique des logiciels et conception d'outils pour mesurer et réduire leur consommation d'énergie, November 2018, Romain Rouvoy.
- PhD in progress: Marion Tommasi, Collaborative Data-centric Workflows: Towards Knowledge Centric Workflows and Integrating Uncertain Data, October 2018, Pierre Bourhis & Lionel Seinturier.

### 11.2.3 Juries

PhD juries, or otherwise specified such as HDR.

#### **Laurence Duchien**

- Rabin Sahu (Université de Lille), president
- Cagan Arslan (Université de Lille), president
- Maverick Chardet (IMT Atlantique), president

- Lorenzo Pagliani (University of L'Aquila, Italy), reviewer
- Neil Ayebe (Université Grenoble Alpes), reviewer
- Sami Lazreg (Université de Nice), reviewer
- Matthieu Amy (Université de Toulouse), examiner
- Hussein Mhanna (Université Paris-Saclay), examiner

#### **Philippe Merle**

- Hayet Brabra (Télécom Sud Paris), president

#### **Romain Rouvoy**

- Yewan Wang (IMT Atlantique), reviewer & president
- Jean-Émile Dartois (Université de Rennes 1), reviewer
- Dalimir Orfanus (University of Oslo, Norway), opponent
- Dorian Burihabwa (University of Neuchâtel, Switzerland), reviewer
- Emile Cadorel (IMT Atlantique), reviewer
- Florentin Delaine (INSA Lyon), reviewer
- Louison Gitzinger (Université de Rennes 1), reviewer

#### **Lionel Seinturier**

- HDR, Stéphanie Chollet (Université Grenoble Alpes), reviewer
- HDR, Anne-Cécile Orgerie (ENS Rennes), examiner
- Soulaïmane Guedria (Université Grenoble Alpes), reviewer

#### **Pierre Laperdrix**

- Nampoina Andriamilanto (Université de Rennes 1), examiner

## **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, X. Blanc and R. Rouvoy. 'On Adopting Linters to Deal with Performance Concerns in Android Apps'. In: *ASE18 - Proceedings of the 33rd IEEE/ACM International Conference on Automated Software Engineering*. Vol. 11. Montpellier, France: ACM Press, Sept. 2018. DOI: [10.1145/3238147.3238197](https://doi.org/10.1145/3238147.3238197). URL: <https://hal.inria.fr/hal-01829135>.

- [4] P. Laperdrix, N. Bielova, B. Baudry and G. Avoine. ‘Browser Fingerprinting: A Survey’. In: *ACM Transaction on the Web* 14.2 (Apr. 2020), pp. 1–33. DOI: [10.1145/3386040](https://doi.org/10.1145/3386040). URL: <https://hal.archives-ouvertes.fr/hal-02864872>.
- [5] L. Meftah, R. Rouvoy and I. Chrismet. ‘FOUGERE: User-Centric Location Privacy in Mobile Crowdsourcing Apps’. In: *19th IFIP International Conference on Distributed Applications and Interoperable Systems (DAIS)*. Ed. by J. Pereira and L. Ricci. Vol. LNCS-11534. Distributed Applications and Interoperable Systems. Kongens Lyngby, Denmark: Springer International Publishing, 2019, pp. 116–132. DOI: [10.1007/978-3-030-22496-7\\_8](https://doi.org/10.1007/978-3-030-22496-7_8). URL: <https://hal.inria.fr/hal-02121311>.
- [6] **Best Paper**  
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] M. Naseri, N. P. Borges Jr., A. Zeller and R. Rouvoy. ‘AccessiLeaks: Investigating Privacy Leaks Exposed by the Android Accessibility Service’. In: *PETS 2019 - The 19th Privacy Enhancing Technologies Symposium*. Stockholm, Sweden, July 2019. URL: <https://hal.inria.fr/hal-01929049>.
- [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* (2021), pp. 1–22. DOI: [10.1093/comjnl/bxaa159](https://doi.org/10.1093/comjnl/bxaa159). URL: <https://hal.inria.fr/hal-02976775>.
- [13] A. Amarilli, P. Bourhis, S. Mengel and M. Niewerth. ‘Constant-Delay Enumeration for Nondeterministic Document Spanners’. In: *SIGMOD record* 49.1 (4th Sept. 2020), pp. 25–32. DOI: [10.1145/3422648.3422655](https://doi.org/10.1145/3422648.3422655). URL: <https://hal.archives-ouvertes.fr/hal-03046848>.
- [14] E. Baranov and S. Bliudze. ‘Expressiveness of component-based frameworks: A study of the expressiveness of BIP’. In: *Acta Informatica* 57 (Dec. 2020), pp. 761–800. DOI: [10.1007/s00236-019-00337-7](https://doi.org/10.1007/s00236-019-00337-7). URL: <https://hal.archives-ouvertes.fr/hal-02149934>.
- [15] S. Basu Roy, L. Chen, A. Morishima, J. A. Monedero, P. Bourhis, F. Charoy, M. Danilevsky, G. Das, G. Demartini, A. Dubey, S. Elbassuoni, D. Gross-Amblard, E. Hoareau, M. Inoguchi, J. Kenworthy, I. Kitahara, D. Lee, Y. Li, R. M. Borromeo, P. Papotti, R. Rao, S. Roy, P. Senellart, K. Tajima, S. Thirumuruganathan, M. Tommasi, K. Umemoto, A. Wiggins, K. Yoshida and S. Amer-Yahia. ‘Making AI Machines Work for Humans in FoW’. In: *SIGMOD record* 49.2 (9th Dec. 2020), pp. 30–35. DOI: [10.1145/3442322.3442327](https://doi.org/10.1145/3442322.3442327). URL: <https://hal.inria.fr/hal-03103700>.

- [16] 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>.
- [17] P. Bourhis, J. L. Reutter and D. Vrgoč. ‘JSON: Data model and query languages’. In: *Information Systems* 89 (Mar. 2020). DOI: [10.1016/j.is.2019.101478](https://doi.org/10.1016/j.is.2019.101478). URL: <https://hal.archives-ouvertes.fr/hal-03094643>.
- [18] 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>.
- [19] B. Danglot, M. Monperrus, W. Rudametkin and B. Baudry. ‘An approach and benchmark to detect behavioral changes of commits in continuous integration’. In: *Empirical Software Engineering* 25.4 (July 2020), pp. 2379–2415. DOI: [10.1007/s10664-019-09794-7](https://doi.org/10.1007/s10664-019-09794-7). URL: <https://hal.inria.fr/hal-03121735>.
- [20] S. Habchi, N. Moha and R. Rouvoy. ‘Android Code Smells: From Introduction to Refactoring’. In: *Journal of Systems and Software* (2021). URL: <https://hal.inria.fr/hal-03182277>.
- [21] L. Meftah, R. Rouvoy and I. Chrismet. ‘Empowering Mobile Crowdsourcing Apps with User Privacy Control’. In: *Journal of Parallel and Distributed Computing* (1st Aug. 2020), p. 15. DOI: [10.1016/j.jpdc.2020.07.011](https://doi.org/10.1016/j.jpdc.2020.07.011). URL: <https://hal.inria.fr/hal-02910246>.
- [22] C. Quinton, M. Vierhauser, R. Rabiser, L. Baresi, P. Grünbacher and C. Schuhmayer. ‘Evolution in Dynamic Software Product Lines’. In: *Journal of Software: Evolution and Process* (21st May 2020). DOI: [10.1002/smr.2293](https://doi.org/10.1002/smr.2293). URL: <https://hal.archives-ouvertes.fr/hal-02952741>.

#### International peer-reviewed conferences

- [23] B. Amin Azad, O. Starov, P. Laperdrix and N. Nikiforakis. ‘Taming The Shape Shifter: Detecting Anti-fingerprinting Browsers’. In: DIMVA 2020 - 17th Conference on Detection of Intrusions and Malware & Vulnerability Assessment. Lisboa / Virtual, Portugal, 24th June 2020. URL: <https://hal.archives-ouvertes.fr/hal-02612461>.
- [24] B. Amin Azad, O. Starov, P. Laperdrix and N. Nikiforakis. ‘Web Runner 2049: Evaluating Third-Party Anti-bot Services’. In: DIMVA 2020 - 17th Conference on Detection of Intrusions and Malware & Vulnerability Assessment. Lisboa / Virtual, Portugal: <https://dimva2020.campus.ciencias.ulisboa.pt/>, 24th June 2020. URL: <https://hal.archives-ouvertes.fr/hal-02612454>.
- [25] P. Bourhis, D. Deutch and Y. Moskovitch. ‘Equivalence-Invariant Algebraic Provenance for Hyperplane Update Queries’. In: Sigmod. Portland, United States, 14th June 2020. DOI: [10.1145/3318464.3380578](https://doi.org/10.1145/3318464.3380578). URL: <https://hal.archives-ouvertes.fr/hal-03094605>.
- [26] P. Bourhis, L. Hérouët, Z. Miklos and R. Singh. ‘Data Centric Workflows for Crowdsourcing’. In: Petri Nets 2020 - 41st International Conference on Application and Theory of Petri Nets and Concurrency. Paris, France, 24th June 2020, pp. 1–39. URL: <https://hal.inria.fr/hal-02508838>.
- [27] J. Enes, G. Fieni, R. R. Expósito, R. Rouvoy and J. Tourino. ‘Power Budgeting of Big Data Applications in Container-based Clusters’. In: IEEE Cluster 2020. Kobe, Japan: <https://clustercomp.org/2020/>, 14th Sept. 2020, pp. 281–287. DOI: [10.1109/CLUSTER49012.2020.00038](https://doi.org/10.1109/CLUSTER49012.2020.00038). URL: <https://hal.inria.fr/hal-02904300>.
- [28] 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: <http://cloudbus.org/ccgrid2021/mainpage.html>, 10th May 2021. URL: <https://hal.inria.fr/hal-03173410>.
- [29] G. Fieni, R. Rouvoy and L. Seinturier. ‘SmartWatts: Self-Calibrating Software-Defined Power Meter for Containers’. In: CCGRID 2020 - 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing. Melbourne, Australia: <http://cloudbus.org/ccgrid2020/>, 11th May 2020. DOI: [10.1109/CCGrid49817.2020.00-45](https://doi.org/10.1109/CCGrid49817.2020.00-45). URL: <https://hal.inria.fr/hal-02470128>.

- [30] 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: <https://www.usenix.org/conference/usenixsecurity21>, 11th Aug. 2021. URL: <https://hal.archives-ouvertes.fr/hal-03152176>.
- [31] J. Leveau, X. Blanc, L. Réveillère, J.-R. Falleri and R. Rouvoy. ‘Fostering the Diversity of Exploratory Testing in Web Applications’. In: ICST 2020 - IEEE International Conference on Software Testing, Verification and Validation. Porto, Portugal: <https://icst2020.info/>, 23rd Mar. 2020. DOI: 10.1109/ICST46399.2020.00026. URL: <https://hal.inria.fr/hal-02398969>.
- [32] L. Meftah, R. Rouvoy and I. Chrisment. ‘Capturing Privacy-preserving User Contexts with IndoorHash’. In: DAIS 2020 - 20th IFIP International Conference on Distributed Applications and Interoperable Systems. Vol. 12135. Valletta, Malta: <http://www.discotec.org/2020/dais.html>, 15th June 2020. DOI: 10.1007/978-3-030-50323-9\_2. URL: <https://hal.inria.fr/hal-02541391>.
- [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: <https://asiaccs2021.comp.polyu.edu.hk/>, 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] 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. Vol. Security, Privacy, and Trust. Proceedings of The Web Conference (WWW’20). Taipei, Taiwan: <https://www2020.thewebconf.org/>, 20th Apr. 2020. DOI: 10.1145/3366423.3380161. URL: <https://hal.inria.fr/hal-02435622>.
- [36] Z. Ournani, M. C. Belgaid, R. Rouvoy, P. Rust, J. Penhoat and L. Seinturier. ‘Taming Energy Consumption Variations in Systems Benchmarking’. In: ICPE’2020 - 11th ACM/SPEC International Conference on Performance Engineering. Edmonton, Canada: <https://icpe2020.spec.org/>, 20th Apr. 2020. DOI: 10.1145/3358960.3379142. URL: <https://hal.inria.fr/hal-02403379>.
- [37] Z. Ournani, R. Rouvoy, P. Rust and J. Penhoat. ‘On Reducing the Energy Consumption of Software: From Hurdles to Requirements’. In: ESEM 2020 - ACM/IEEE International Symposium on Empirical Software Engineering and Measurement. Bari, Italy: <https://eseiw2020.di.uniba.it/>, 5th Oct. 2020. DOI: 10.1145/3382494.3410678. URL: <https://hal.inria.fr/hal-02892900>.
- [38] *Best Paper*  
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. San Diego, United States: <https://madweb.work/>, 23rd Feb. 2020. DOI: 10.14722/ndss.2020.23xxx. URL: <https://hal.inria.fr/hal-02441653>.
- [39] D. Zeber, S. Bird, C. Oliveira, W. Rudametkin, I. Segall, F. Wollsen and M. Lopatka. ‘The Representativeness of Automated Web Crawls as a Surrogate for Human Browsing’. In: The Web Conference. Proceedings of The Web Conference (WWW’20). Taipei, Taiwan, 20th Apr. 2020. DOI: 10.1145/3366423.3380104. URL: <https://hal.inria.fr/hal-02456195>.

### Scientific book chapters

- [40] B. Baudry, Y.-D. Bromberg, D. Frey, A. Gómez-Boix, P. Laperdrix and F. Taïani. ‘Profilage de navigateurs : état de l’art et contre-mesures’. In: *Le profilage en ligne : entre libéralisme et régulation*. 15th Oct. 2020. URL: <https://hal.inria.fr/hal-03043187>.

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

- [41] S. Bliudze and M. F. Atig, eds. *Special Issue: 12th International Conference on Verification and Evaluation of Computer and Communication Systems (VECoS 2018)*. June 2020. URL: <https://hal.archives-ouvertes.fr/hal-03197446>.
- [42] S. Bliudze and L. Bocchi, eds. *22nd IFIP WG 6.1 International Conference, COORDINATION 2020, Held as Part of the 15th International Federated Conference on Distributed Computing Techniques, DisCoTec 2020, Valletta, Malta, June 15–19, 2020, Proceedings*. 10th June 2020. DOI: [10.1007/978-3-030-50029-0](https://doi.org/10.1007/978-3-030-50029-0). URL: <https://hal.archives-ouvertes.fr/hal-03168525>.

### Doctoral dissertations and habilitation theses

- [43] 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>.

### Reports & preprints

- [44] P. Bourhis, M. Leclère, M.-L. Mugnier, S. Tison, F. Ulliana and L. Galois. *Oblivious and Semi-Oblivious Boundedness for Existential Rules*. LIRMM (UM, CNRS), 15th June 2020. URL: <https://hal-lirmm.ccsd.cnrs.fr/lirmm-02920624>.
- [45] G. Giakkoupis, A.-M. Kermarrec, O. Ruas and F. Taïani. *Cluster-and-Conquer: When Randomness Meets Graph Locality*. 21st Oct. 2020. URL: <https://hal.inria.fr/hal-02974077>.
- [46] P. Merle, S. Ben Rayana, L. Seinturier, R. Pissard-Gibollet, J.-B. Stefani and A. Ndeye Sylla. *Towards a formal reference computational model for cloud configuration management*. INRIA, 7th Jan. 2020. URL: <https://hal.inria.fr/hal-02940938>.

## 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] M. Benedikt, P. Bourhis, L. Jachiet and M. Thomazo. ‘Reasoning about disclosure in data integration in the presence of source constraints’. In: *IJCAI 2019 - 28th International Joint Conference on Artificial Intelligence*. Macao, China, Aug. 2019. URL: <https://hal.inria.fr/hal-02145369>.
- [52] 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>.

- [53] S. Bliudze, S. Furic, J. Sifakis and A. Viel. ‘Rigorous Design of Cyber-Physical Systems’. In: *Software and Systems Modeling*. Software and System Modeling 18.3 (2019), pp. 1613–1636. DOI: [10.1007/s10270-017-0642-5](https://doi.org/10.1007/s10270-017-0642-5). URL: <https://hal.inria.fr/hal-01636392>.
- [54] S. Bliudze, L. Henrio and E. Madelaine. ‘Verification of concurrent design patterns with data’. In: *COORDINATION 2019 - 21st International Conference on Coordination Models and Languages*. Ed. by H. R. Nielson and E. Tuosto. Vol. LNCS-11533. Coordination Models and Languages. t Part 4: Coordination Patterns. Kongens Lyngby, Denmark: Springer International Publishing, June 2019, pp. 161–181. DOI: [10.1007/978-3-030-22397-7\\_10](https://doi.org/10.1007/978-3-030-22397-7_10). URL: <https://hal.archives-ouvertes.fr/hal-02143782>.
- [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] L. Capra, W. Emmerich and C. Mascolo. ‘A micro-economic approach to conflict resolution in mobile computing’. In: *SIGSOFT FSE*. 2002, pp. 31–40.
- [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] V. Dallmeier, A. Zeller and B. Meyer. ‘Generating Fixes from Object Behavior Anomalies’. In: *Proceedings of the 2009 IEEE/ACM International Conference on Automated Software Engineering*. ASE ’09. Washington, DC, USA: IEEE Computer Society, 2009, pp. 550–554. DOI: [10.1109/ASE.2009.15](https://doi.org/10.1109/ASE.2009.15). URL: <http://dx.doi.org/10.1109/ASE.2009.15>.
- [60] F. Fouquet, E. Daubert, N. Plouzeau, O. Barais, J. Bourcier and J.-M. Jézéquel. ‘Dissemination of Reconfiguration Policies on Mesh Networks’. In: *Proceedings of 12th IFIP International Conference on Distributed Applications and Interoperable Systems (DAIS’12)*. Ed. by K. M. Göschka and S. Haridi. Vol. 7272. Lecture Notes in Computer Science. Stockholm, Sweden: Springer, June 2012, pp. 16–30.
- [61] M. Goldmann and G. Kreitz. ‘Measurements on the spotify peer-assisted music-on-demand streaming system’. In: *Peer-to-Peer Computing*. Ed. by T. Asami and T. Higashino. IEEE, 2011, pp. 206–211.
- [62] S. Habchi. ‘Understanding Mobile-Specific Code Smells’. Theses. Université de Lille, Dec. 2019. URL: <https://tel.archives-ouvertes.fr/tel-02414928>.
- [63] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann and I. H. Witten. ‘The WEKA data mining software: an update’. In: *SIGKDD Explorer Newsletter* 11.1 (2009), pp. 10–18.
- [64] M. Harman. ‘Automated patching techniques: the fix is in’. In: *Communications of the ACM* 53.5 (May 2010), pp. 108–108. DOI: [10.1145/1735223.1735248](https://doi.org/10.1145/1735223.1735248). URL: <http://doi.acm.org/10.1145/1735223.1735248>.
- [65] 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>.
- [66] 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>.



- [67] D. Jackson. ‘Alloy: A Lightweight Object Modelling Notation’. In: *ACM Transactions on Software Engineering and Methodology* 11.2 (Apr. 2002), pp. 256–290.
- [68] M. Léger, T. Ledoux and T. Coupaye. ‘Reliable Dynamic Reconfigurations in a Reflective Component Model’. In: *CBSE*. Ed. by L. Grunske, R. Reussner and F. Plasil. Vol. 6092. Lecture Notes in Computer Science. Springer, 2010, pp. 74–92.
- [69] 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>.
- [70] W. Li, S. Hu, J. Li and H.-A. Jacobsen. ‘Community Clustering for Distributed Publish/Subscribe Systems’. In: *CLUSTER*. IEEE, 2012, pp. 81–89.
- [71] P. Merle, A. Ndeye Sylla, M. Ouzzif, F. Klamm and K. Guillouard. ‘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>.
- [72] 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>.
- [73] 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>.
- [74] 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>.
- [75] L. Provensi, F. Eliassen, R. Vitenberg and R. Rouvoy. ‘Improving Context Interpretation by Using Fuzzy Policies: The Case of Adaptive Video Streaming’. In: *28th ACM Symposium on Applied Computing (SAC) - 8th Track on Dependable and Adaptive Distributed Systems (DADS)*. Ed. by K. M. Göschka, R. Oliveira, P. Pietzuch and G. Russello. Vol. 1. Best paper award. Coimbra, Portugal: ACM, Mar. 2013, pp. 415–422. URL: <http://hal.inria.fr/hal-00799136>.
- [76] M. Rinard, C. Cadar, D. Dumitran, D. Roy, T. Leu and W. Beebe Jr. ‘Enhancing server availability and security through failure-oblivious computing’. In: *Proceedings of the 6th conference on Symposium on Operating Systems Design & Implementation (OSDI’06)*. USENIX Association. 2004, pp. 21–21.
- [77] 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>.
- [78] 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](http://www.eurosfair.prdd.fr/7pc/doc/1303227305_fp8_future_internet_28_02_2011.pdf).
- [79] 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](http://ec.europa.eu/information_society/newsroom/cf/dae/document.cfm?doc_id=2165).
- [80] 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>.

- [81] 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>.
- [82] 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](http://www.nessi-europe.eu/files/Docs/NESSI%5C%20SRA_update_May_2011_V1-0.pdf).
- [83] A. Vastel. 'Tracking Versus Security: Investigating the Two Facets of Browser Fingerprinting'. Theses. Université de Lille Nord de France, Oct. 2019. URL: <https://tel.archives-ouvertes.fr/tel-02343930>.
- [84] Y. Wei, Y. Pei, C. A. Furia, L. S. Silva, S. Buchholz, B. Meyer and A. Zeller. 'Automated fixing of programs with contracts'. In: *Proceedings of the 19th international symposium on Software testing and analysis*. ISSTA '10. Trento, Italy: ACM, 2010, pp. 61–72. DOI: [10.1145/1831708.1831716](https://doi.org/10.1145/1831708.1831716). URL: <http://doi.acm.org/10.1145/1831708.1831716>.
- [85] C. Weiss, R. Premraj, T. Zimmermann and A. Zeller. 'How long will it take to fix this bug?' In: *Proceedings of the Fourth International Workshop on Mining Software Repositories*. IEEE Computer Society, 2007, p. 1.
- [86] T. Xie, S. Thummalapenta, D. Lo and C. Liu. 'Data Mining for Software Engineering'. In: *IEEE Computer* 42.8 (Aug. 2009), pp. 55–62. URL: <http://doi.ieeecomputersociety.org/10.1109/MC.2009.256>.
- [87] W. Xu, L. Huang, A. Fox, D. A. Patterson and M. I. Jordan. 'Detecting large-scale system problems by mining console logs'. In: *Proceedings of the 22nd ACM Symposium on Operating Systems Principles (SOSP'09)*. Ed. by J. N. Matthews and T. E. Anderson. Big Sky, Montana, USA: ACM, Oct. 2009, pp. 117–132.