



IN PARTNERSHIP WITH:
**Université des sciences et
technologies de Lille (Lille 1)**

Activity Report 2018

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

RESEARCH CENTER
Lille - Nord Europe

THEME
Distributed Systems and middleware

Table of contents

1. Team, Visitors, External Collaborators	1
2. Overall Objectives	3
2.1. Introduction	3
2.2. Scientific Foundations	3
3. Research Program	3
3.1. Introduction	3
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	4
3.2.2. Run-time self-healing	5
3.3. Objective #2: Self-optimization - Sharing runtime behaviors to continuously adapt software	6
3.3.1. Monitoring software in the wild	6
3.3.2. Collaborative decision-making approaches	7
3.3.3. Smart reconfigurations in the large	7
4. Application Domains	8
4.1. Introduction	8
4.2. Distributed software services for the retail industry	8
4.3. Distributed software services for the digital home	8
5. Highlights of the Year	9
6. New Software and Platforms	9
6.1. APISENSE	9
6.2. PowerAPI	10
6.3. Saloon	10
6.4. SPOON	10
7. New Results	11
7.1. Software Product Lines for Setup and Adaptation of Multi-Cloud Computing Systems	11
7.2. Automated Software Repair with Patch Generation in Production	11
7.3. Flexible Framework for Elasticity in Cloud Computing	11
7.4. Semantic Interoperability in Multi-Cloud Computing Systems	12
8. Bilateral Contracts and Grants with Industry	12
8.1. ip-label	12
8.2. Scalair	12
8.3. Davidson	12
8.4. Orange Labs #1	13
8.5. Orange Labs #2	13
8.6. Amaris	13
9. Partnerships and Cooperations	13
9.1. Regional Initiatives	13
9.1.1. Région Hauts-de-France	13
9.1.1.1. CIRBUS	13
9.1.1.2. Alloy@Scale	14
9.1.2. Inria Lille - Nord Europe	14
9.1.2.1. ADT LibRepair	14
9.1.2.2. ADT FingerKit	14
9.1.2.3. ADT e-Lens	14
9.2. National Initiatives	14
9.2.1. ANR	14
9.2.1.1. ANR BottleNet	14
9.2.1.2. ANR SATAS	15
9.2.1.3. ANR Headwork	15

9.2.1.4. ANR Delta	15
9.2.2. Competitivity Clusters	15
9.2.3. Inria National Initiatives	15
9.2.4. CNRS Momentum	16
9.3. European Initiatives	16
9.3.1. FP7 & H2020 Projects	16
9.3.2. Collaborations in European Programs, Except FP7 & H2020	17
9.4. International Initiatives	17
9.4.1. Inria Associate Teams Not Involved in an Inria International Labs	17
9.4.2. Participation in Other International Programs	18
9.5. International Research Visitors	18
9.5.1. Visits of International Scientists	18
9.5.2. Visits to International Teams	18
10. Dissemination	18
10.1. Promoting Scientific Activities	18
10.1.1. Scientific Events Organisation	18
10.1.2. Scientific Events Selection	19
10.1.3. Journal	20
10.1.3.1. Member of the Editorial Boards	20
10.1.3.2. Reviewer - Reviewing Activities	20
10.1.4. Invited Talks	20
10.1.5. Leadership within the Scientific Community	21
10.1.6. Scientific Expertise	21
10.1.7. Research Administration	21
10.2. Teaching - Supervision - Juries	21
10.2.1. Teaching	21
10.2.2. Supervision	23
10.2.3. Juries	23
10.3. Popularization	24
11. Bibliography	24

Project-Team SPIRALS

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

Keywords:

Computer Science and Digital Science:

- 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
- A3.1.9. - Database
- A3.2.1. - Knowledge bases
- A3.2.4. - Semantic Web
- A7.2. - Logic in Computer Science

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, Visitors, External Collaborators

Research Scientists

- Simon Bliudze [Inria, Researcher]
- Pierre Bourhis [CNRS, Researcher]
- Philippe Merle [Inria, Researcher, HDR]

Faculty Members

- Lionel Seinturier [Team leader, Université de Lille, Professor, HDR]
- Clément Quinton [Université de Lille, Associate Professor]
- Laurence Duchien [Université de Lille, Professor, HDR]

Romain Rouvoy [Université de Lille, Professor, HDR]
Walter Rudametkin Ivey [Université de Lille, Associate Professor]

Post-Doctoral Fellows

Alexandre Garnier [Inria, from Sep 2018]
Louis Jachiet [CNRS, from Oct 2018]
Hoang Gia Nguyen [Inria, from Nov 2018]
Zhongxing Yu [Inria, until Feb 2018]
Faiez Zalila [Inria]

PhD Students

Zeinab Abou Khalil [Université de Lille]
Yahya Al-Dhuraibi [Scalair]
Mohammed Chakib Belgaid [Inria]
Sacha Brisset [Amaris, from Nov 2018]
Stéphanie Challita [Inria]
Benjamin Danglot [Inria]
Antonin Durey [Université de Lille, from Oct 2018]
Thomas Durieux [Inria, until Aug 2018]
Guillaume Fieni [Université de Lille]
Sarrah Habchi [Inria]
Lakhdar Meftah [Inria]
Vikas Mishra [Inria, from Oct 2018]
Zakaria Ournani [Orange, from Nov 2018]
Gustavo Sousa [Université de Lille, until Feb 2018]
Marion Tommasi [Inria, from Oct 2018]
Antoine Vastel [Université de Lille]

Technical staff

Jordan Bouchoucha [Université de Lille, from Oct 2018]
Aurélien Bourdon [Inria, until Sep 2018]
Antoine Canda [Inria, from Oct 2018]
Arthur d'Azemar [Inria, from Oct 2018]
Julien Decaudin [Inria, from Jun 2018]
Bénédicte Delcourt [Inria, from Sep 2018]
Christophe Gourdin [Inria]
Abderrahman Lahiaoui [Inria, from Dec 2018]
Zakaria Ournani [Inria, from Apr 2018 until Oct 2018]
Romain Sommerard [Inria, until Feb 2018]
Simon Urli [Inria, until Aug 2018]
Antoine Veuilleur [Inria, until Mar 2018]

Interns

Waiss Azizian [Ecole Normale Supérieure Paris, from Jun 2018 until Aug 2018]
Jamal Boudjaj [Inria, from Jun 2018 until Oct 2018]
Chaima Chekhaba [Inria, from Mar 2018 until Aug 2018]
Rida Darmal [Inria, from May 2018 until Oct 2018]
Theo Delemazure [Ecole Normale Supérieure Paris, from Jun 2018 until Aug 2018]
Matthieu Delobelle Toussaint [Inria, from Jun 2018 until Aug 2018]
Anthony Durot [Inria, from May 2018 until Aug 2018]
Marine Gosset [Inria, from Apr 2018 until Jun 2018]
Manel Halima [Université de Lille, from Jun 2018 until Jul 2018]
Abderrahman Lahiaoui [Inria, from May 2018 until Oct 2018]
Mohammad Naseri [Inria, until Jan 2018]

Petra Rahme [Inria, from May 2018 until Aug 2018]
Paul Ribeiro [Inria, from May 2018 until Jul 2018]
Adam Taboulet [Inria, Feb 2018]
Marion Tommasi [Inria, from Mar 2018 until Aug 2018]

Administrative Assistants

Nathalie Bonte [Inria, until Apr 2018]
Karine Lewandowski [Inria, from Apr 2018]

Visiting Scientist

Fatima Belkouch Ameline [Université de Lille]

External Collaborator

Martin Monperrus [KTH (on leave from Université de Lille)]

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 declined for 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.*”¹ This pervasiveness continuously leads to new usages that in turn foster the emergence of novel requirements and concepts for new software services. Hence, it is necessary to establish new paradigms to design and execute software programs in these highly interconnected and heterogeneous environments, and it is necessary to ensure not only that these software systems can be adapted to new usages, new infrastructures, and new execution environments in the long term, but also that after the adaptation process the services still perform as expected.

This research project focuses on defining *self-adaptive* software services and middleware. From the perspective of the Internet of Services, this project fits in the vision sketched by *e.g.*, the FP8 Expert Group Services in the Future Internet [64], the NESSI Research Priorities for the next Framework Programme for Research and Technological Development FP8 [67], the Roadmap for Advanced Cloud Technologies under H2020 [65], and research roadmaps, such as [72], [63], [53].

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*.

¹<http://cordis.europa.eu/fp7/ict/ssai>

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 16 Dec 2011, the Eclipse bug repository contains 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 [69].

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** [70]. 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 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 [51].

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 [47] and [68]. 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” [62] 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 [52]. We think 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 homeostasis² 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

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

[71]. 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 [61] 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 [50], can be used to build a decentralized decision-making algorithm 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 [44], 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 [49], but also the adoption of publish/subscribe abstractions [55] 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 [66] 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 [56], 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 Cappuccino and **Macchiato**, which focused on the development of a new generation of mobile computing platforms for e-commerce. We are also involved in the Datalyse and OCCIware funded projects that define cloud computing environments with applications for the retail industry. Finally, our activities in terms of crowd-sensing and data gathering on mobile devices with the APISENSE[®] platform share also applications for the retail industry.

4.3. Distributed software services for the digital home

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

5. Highlights of the Year

5.1. Highlights of the Year

5.1.1. Awards

Thomas Durieux, who defended his PhD thesis [11] in September 2018 on automated software repair, has obtained two best paper awards in 2018: at the 29th IEEE International Symposium on Software Reliability Engineering (ISSRE) [26] for his work on self-healing web applications based on HTML and JavaScript rewriting, and at the 6th Workshop on Software Visualization, Evolution and Maintenance (VEM) [38] for his contribution to the automatic detection of bug fixing patterns, which is a joint work with Fernanda Madeiral and colleagues from the Federal University of Uberlândia in Brazil.

Browser fingerprinting [48] has emerged as a technique to track users without their consent. Studying the implication of this technique on user privacy, and proposing software systems that can self-protect against such techniques, have become an important topic in the research activities of Spirals on self-adaptation. In 2018, two papers authored by Spirals members in the context of the ongoing PhD thesis of Antoine Vastel have been published in A* conferences of the domain of security: [36] at the 39th IEEE Symposium on Security and Privacy, and [35] at 27th USENIX Security Symposium.

In 2018, we published at the 40th International Conference on Software Engineering, Track Software Engineering in Practice, a paper on Repairator [34], which this is the very first bot for automated software repair. This work has been conducted by Simon Urli and his coworkers in Spirals in the context of Inria ADT Librepare. An [article on Repairator](#) has also been published on the Inria web site.

Stéphanie Challita won the [Prix L'Oréal-UNESCO For Women in Science](#). She is among the 30 laureates rewarded out of 900 candidates. Stéphanie Challita has defended her PhD thesis in December 2018 on FLOUDS, a framework to achieve semantic interoperability in multi-cloud computing systems.

[Xscalibur](#) is a transfer project led by Christophe Gourdin and Philippe Merle that aims at creating a startup company with innovative solutions for the management of multi-cloud systems. The project has been selected by Serre Numérique Valenciennes. The research activities that led to this project have been, in part, conducted in the context of the joint lab between Spirals and the [Scalair](#) cloud hosting company. Several papers in relation with these research activities have also been published in 2018: [30], [25], [24], [13], [23].

BEST PAPERS AWARDS:

[26]

T. DURIEUX, Y. HAMADI, M. MONPERRUS. *Fully Automated HTML and Javascript Rewriting for Constructing a Self-healing Web Proxy*, in "Proceedings of the 29th IEEE International Symposium on Software Reliability Engineering (ISSRE 2018)", Memphis, United States, October 2018, <https://arxiv.org/abs/1803.08725> [DOI : 10.1109/ISSRE.2018.00012], <https://hal.inria.fr/hal-01746141>

[38]

F. MADEIRAL, T. DURIEUX, V. SOBREIRA, M. MAIA. *Towards an automated approach for bug fix pattern detection*, in "VEM '18 - Proceedings of the VI Workshop on Software Visualization, Evolution and Maintenance", São Carlos, Brazil, September 2018, <https://arxiv.org/abs/1807.11286>, <https://hal.archives-ouvertes.fr/hal-01851813>

6. New Software and Platforms

6.1. 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.

- Participants: Antoine Veuille, Christophe Ribeiro, Julien Duribreux, Nicolas Haderer, Romain Rouvoy, Romain Sommerard and Lakhdar Meftah
- Partner: Université de Lille
- Contact: Romain Rouvoy
- URL: <https://apisense.io>

6.2. 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 tool in its software orientation, with a fully customizable and modular solution that let the user to 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.).

- Participants: Adel Noureddine, Loïc Huertas, Maxime Colmant and Romain Rouvoy
- Contact: Romain Rouvoy
- URL: <http://powerapi.org>

6.3. Saloon

KEYWORDS: Feature Model - Software Product Line - Cloud computing - Model-driven engineering - Ontologies

FUNCTIONAL DESCRIPTION: Saloon is a framework for the selection and configuration of Cloud providers according to application requirements. The framework enables the specification of such requirements by defining ontologies. Each ontology provides a unified vision of provider offers in terms of frameworks, databases, languages, application servers and computational resources (i.e., memory, storage and CPU frequency). Furthermore, each provider is related to a Feature Model (FM) with attributes and cardinalities, which captures its capabilities. By combining the ontology and FMs, the framework is able to match application requirements with provider capabilities and select a suitable one. Specific scripts to the selected provider are generated in order to enable its configuration.

- Participants: Clément Quinton, Daniel Romero Acero, Laurence Duchien, Lionel Seinturier and Romain Rouvoy
- Partner: Université Lille 1
- Contact: Lionel Seinturier
- URL: <https://gitlab.irisa.fr/drome00A/saloon>

6.4. SPOON

KEYWORDS: Java - Code analysis

FUNCTIONAL DESCRIPTION: Spoon is an open-source library that enables you to transform (see below) and analyze Java source code (see example) . Spoon provides a complete and fine-grained Java metamodel where any program element (classes, methods, fields, statements, expressions...) can be accessed both for reading and modification. Spoon takes as input source code and produces transformed source code ready to be compiled.

- Participants: Gérard Paligot, Lionel Seinturier, Martin Monperrus and Nicolas Petitprez
- Contact: Martin Monperrus
- URL: <http://spoon.gforge.inria.fr>

7. New Results

7.1. Software Product Lines for Setup and Adaptation of Multi-Cloud Computing Systems

In 2018, in the domain of cloud computing, we proposed a new software product line-based approach for managing the variability in order to automate the setup and adaptation of multi-cloud environments. Building such systems is still very challenging and time consuming due to the heterogeneity across cloud providers' offerings and the high-variability in the configuration of cloud providers. This variability is expressed by the large number of available services and the many different ways in which they can be combined and configured. In order to ensure correct setup of a multi-cloud environment, developers must be aware of service offerings and configuration options from multiple cloud providers. Our results enable to automatically generate a configuration or reconfiguration plan for a multi-cloud environment from a description of its requirements. The conducted experiments aim to assess the impact of the approach on the automated analysis of feature models and the feasibility of the approach to automate the setup and adaptation of multi-cloud environments. These results have been obtained in the context of the PhD thesis of Gustavo Sousa [12] defended in June 2018.

7.2. Automated Software Repair with Patch Generation in Production

In 2018, in the domain of automated software repair, we proposed new patch generation techniques. Patch creation is one of the most important actions in the life cycle of an application. Creating patches is a time-consuming task. Not only because it is difficult to create a sound and valid patch, but also because it requires the intervention of humans. Our work proposes new patch generation techniques that remove the human intervention. Our idea is to put as close as possible the patch generation in the production environment. We adopt this approach because the production environment contains all the data and human interactions that lead to the bug. We show how to exploit this data to detect bugs, generate and validate patches. We evaluate this approach on seven different benchmarks of real bugs collected from open-source projects. During the evaluation, we are particularly attentive to the number of generated patches, to their correctness, readability and to the time required for generating them. Our evaluation shows the applicability and feasibility of our approach to generate patches in the production environment without the intervention of a developer. These results have been obtained in the context of the PhD thesis of Thomas Durieux [11] defended in September 2018.

7.3. Flexible Framework for Elasticity in Cloud Computing

In 2018, in the domain of cloud computing, we proposed a new framework for managing elasticity. The main factor motivating the use of cloud is its ability to provide resources according to the customer needs or what is referred to as elasticity. Adapting cloud applications during their execution according to demand variation is nevertheless a challenging task. In addition, cloud elasticity is diverse and heterogeneous because it encompasses different approaches, policies, purposes, etc. In this work, three contributions are proposed: (1)

an up-to-date state-of-the-art of the cloud elasticity for both virtual machines and containers, (2) ELASTIC-DOCKER, an approach to manage container elasticity including vertical elasticity, live migration, and elasticity combination between different virtualization techniques, and (3) MODEMO, a new unified standard-based, model-driven, highly extensible and reconfigurable framework that supports multiple elasticity policies, vertical and horizontal elasticity, different virtualization techniques and multiple cloud providers. These results have been obtained in the context of the PhD thesis of Yahya Al-Dhuraibi defended in December 2018.

7.4. Semantic Interoperability in Multi-Cloud Computing Systems

In 2018, in the domain of cloud computing, we proposed two major results related to semantic interoperability. First, an approach based on reverse-engineering to extract knowledge from the ambiguous textual documentation of cloud APIs and to enhance its representation using MDE techniques has been proposed. This approach is applied to Google Cloud Platform (GCP), where we provide GCP Model, a precise model-driven specification for GCP. GCP Model is automatically inferred from GCP textual documentation, conforms to the OCCIWARE METAMODEL and is implemented within OCCIWARE STUDIO. It allows one to perform qualitative and quantitative analysis of the GCP documentation. Second, we have proposed the FLOUDS framework to achieve semantic interoperability in multi-clouds, i.e., to identify the common concepts between cloud APIs and to reason over them. The FLOUDS language is a formalization of OCCI concepts and operational semantics in Alloy formal specification language. To demonstrate the effectiveness of the FLOUDS language, we formally specify thirteen case studies and verify their properties. Then, thanks to formal transformation rules and equivalence properties, we draw a precise alignment between my case studies, which promotes semantic interoperability in multi-clouds. These results have been obtained in the context of the PhD thesis of Stéphanie Challita defended in December 2018.

8. Bilateral Contracts and Grants with Industry

8.1. ip-label

Participant: Romain Rouvoy [correspondant].

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

8.2. Scalair

Participants: Yahya Al-Dhuraibi, Philippe Merle [correspondant].

This collaboration (2015–18) aims at proposing a framework to deal with elasticity in cloud computing environments. This framework must cover all kinds of resources, IaaS, PaaS, SaaS, must provide a solution for interoperability between different clouds and virtualization technologies, and must enable the specification and composition of reactive and predictive strategies.

This collaboration is conducted in the context of the PhD thesis of Yahya Al-Dhuraibi defended in December 2018.

8.3. Davidson

Participants: Romain Rouvoy [correspondant], Lionel Seinturier.

This collaboration (2017–20) 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.

8.4. Orange Labs #1

Participants: Philippe Merle [correspondant], Lionel Seinturier.

This collaboration (2017–18) aims at defining a computational model for software infrastructures layered on top of virtualized and interconnected cloud resources. This computational model will provide application programming and management facilities to distributed applications and services. This computational model will define a pivot model that will enable the interoperability of various existing and future standards for cloud systems such as OCCI and TOSCA. This pivot model will be defined with the Alloy specification language [54]. This collaboration takes advantage of the expertise that we are developing since several years on reconfigurable component-based software systems [66], on cloud systems [60], and on the Alloy specification language [58].

This collaboration with Orange Labs is a joint project with Jean-Bernard Stefani from the [Spades](#) Inria project-team.

8.5. Orange Labs #2

Participants: Zakaria Ournani, Romain Rouvoy [correspondant], 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.

8.6. Amaris

Participants: Sacha Brisset, Romain Rouvoy [correspondant], 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.

9. Partnerships and Cooperations

9.1. Regional Initiatives

9.1.1. Région Hauts-de-France

9.1.1.1. CIRRUS

Participants: Yahya Al-Dhuraibi, Stéphanie Challita, Guillaume Fieni, Alexandre Garnier, Christophe Gourdin, Philippe Merle [correspondant], Romain Rouvoy, Lionel Seinturier, Faiez Zalila.

CIRRUS is an 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.

9.1.1.2. Alloy@Scale

Participants: Philippe Merle [correspondant], Romain Rouvoy, Lionel Seinturier, Zakaria Ournani.

Alloy@Scale is a 12-month project funded in the context of CPER Data program. Alloy@Scale aims at overcoming the limits of the formal verification of large software systems specified with the Alloy formal specification language. For that, the program combines the Grid’5000 infrastructure and the Docker container technology.

9.1.2. Inria Lille - Nord Europe

9.1.2.1. ADT LibRepair

Participants: Benjamin Danglot, Martin Monperrus, Lionel Seinturier [correspondant], Simon Urli.

ADT LibRepair (2016–18) is a technology development initiative supported by the Inria Lille - Nord Europe Center that aims at supporting the development of an integrated library of automated software repair algorithms and techniques. This ADT builds on our results about with the Astor, Nopol and NpeFix that have been obtained in the context of the defended PhD theses of Matias Martinez [57] and Benoit Cornu [46].

9.1.2.2. ADT FingerKit

Participants: Antoine Canda, Walter Rudametkin Ivey [correspondant], Antoine Vastel.

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 ongoing PhD thesis of Antoine Vastel.

9.1.2.3. ADT e-Lens

Participants: Arthur d’Azémar, Guillaume Fieni, Romain Rouvoy [correspondant].

ADT e-Lens (2018–20) 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 [59] and Maxime Colmant [45], and with the ongoing PhD thesis of Guillaume Fieni.

9.2. National Initiatives

9.2.1. ANR

9.2.1.1. ANR BottleNet

Participants: Romain Rouvoy [correspondant], Walter Rudametkin Ivey, Lionel Seinturier.

BottleNet is a 48-month project (2015–19) funded by ANR. The objective of BottleNet is to deliver methods, algorithms, and software systems to measure Internet *Quality of Experience* (QoE) and diagnose the root cause of poor Internet QoE. Our goal calls for tools that run directly at users’ devices. We plan to collect network and application performance metrics directly at users’ devices and correlate it with user perception to model Internet QoE, and to correlate measurements across users and devices to diagnose poor Internet QoE. This data-driven approach is essential to address the challenging problem of modeling user perception and of diagnosing sources of bottlenecks in complex Internet services. BottleNet will lead to new solutions to assist users, network and service operators as well as regulators in understanding Internet QoE and the sources of performance bottleneck.

9.2.1.2. ANR SATAS

Participants: Philippe Merle [correspondant], Romain Rouvoy, Lionel Seinturier.

SATAS is a 48-month project (2015–19) 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.

9.2.1.3. ANR Headwork

Participants: Pierre Bourhis [correspondant], 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 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.

9.2.1.4. ANR Delta

Participant: Pierre Bourhis [correspondant].

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.

9.2.2. Competitiveness Clusters

9.2.2.1. FUI StoreConnect

Participants: Aurélien Bourdon, Julien Duribreux, Romain Rouvoy, Lionel Seinturier [correspondant], Antoine Veuille.

StoreConnect is a 36-month project (2016–19) funded by FUI and labelled by the PICOM (**Pôle des Industries du COMmerce**) competitiveness cluster. The partners are Tevolys, UbuDu (leader), Smile, STIME, Leroy Merlin, Insiteo, Inria Spirals, **Inria Fun**, **Inria Stars**. The goal of the project is to define a modular multi-sensors middleware platform for indoor geolocation.

9.2.3. Inria National Initiatives

9.2.3.1. Inria IPL BetterNet

Participants: Lakhdar Meftah, Romain Rouvoy [correspondant], Romain Sommerard, Antoine Veuille.

BetterNet (2016–19) aims at building and delivering a scientific and technical collaborative observatory to measure and improve the Internet service access as perceived by users. In this Inria Project Lab, we will propose new original user-centered measurement methods, which will associate social sciences to better understand Internet usage and the quality of services and networks. Our observatory can be defined as a vantage point, where: (1) tools, models and algorithms/heuristics will be provided to collect data, (2) acquired data will be analyzed, and shared appropriately with scientists, stakeholders and civil society, and (3) new value-added services will be proposed to end-users. IPL BetterNet is led by Isabelle Chrisment (Inria Madynes), with the participation of the [Diana](#), [Dionysos](#), [Inria Chile](#), [Muse](#), and Spirals Inria project-teams, as well as the ARCEP French agency and the ip-label company.

9.2.4. CNRS Momentum

9.2.4.1. Manage Your Data Without Information Leakage

Participants: Pierre Bourhis [correspondant], Louis Jachiet.

“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.

9.3. European Initiatives

9.3.1. FP7 & H2020 Projects

Program: H2020 ICT-10-2016.

Project acronym: STAMP.

Project title: Software Testing Amplification.

Duration: 36 months (2016–19).

Coordinator: Inria.

Other partners: ActiveEon (France), Atos (Spain), Engineering (Italy), OW2 (France), SINTEF (Norway), TellU (Norway), TU Delft (The Netherlands), XWiki (France).

Abstract: By leveraging advanced research in automatic test generation, STAMP aims at pushing automation in DevOps one step further through innovative methods of test amplification. It will reuse existing assets (test cases, API descriptions, dependency models), in order to generate more test cases and test configurations each time the application is updated. Acting at all steps of development cycle, STAMP techniques aim at reducing the number and cost of regression bugs at unit level, configuration level and production stage.

Participants: Benjamin Danglot, Martin Monperrus [correspondant].

Program: H2020 JU Shift2Rail.

Project acronym: X2Rail-1.

Project title: Start-up activities for Advanced Signalling and Automation System.

Duration: 36 months (2016–19).

Coordinator: Siemens.

Other partners: 19 partners, among others Bombardier, Siemens, Thales, IRT Railenium.

Abstract: Our contribution to the project is focused on adaptive communication middleware for cyber-physical railway systems.

Participants: Lionel Seinturier [correspondant].

Program: EIT Digital.

Project acronym: MCS.

Project title: Multi-Cloud Studio.

Duration: 12 months (2018).

Coordinator: Inria.

Other partners: Santer Reply (Italy) and Scalair (France).

Abstract: The goal of this EIT Digital activity is to create the Xscalibur start-up and develop its first Multi-Cloud Studio product. This product is a model-driven graphical interface to design, deploy and administrate multi-cloud systems based on Amazon Web Service, OpenStack, and VMware.

Participants: Jamal Boudjaj, Rida Darmal, Julien Decaudin, Bénédicte Delcourt, Christophe Gourdin, Philippe Merle [correspondant].

Program: EIT Digital.

Project acronym: Bankable.

Project title: Bankable, Deliver the best mobile banking services and customer experience

Duration: 12 months (2018).

Coordinator: Alfstore.

Other partners: Inria, Cefriel

Abstract: BANKABLE helps consumer banks continuously deliver the best mobile banking experience. The platform will anonymously analyze data gathered from mobile sensors, user context and interactions. Enriched by online customers' feedbacks, the bank will find deeper insights & drive real-time actions.

Participants: Romain Rouvroy [correspondant].

9.3.2. Collaborations in European Programs, Except FP7 & H2020

Program: EUREKA Celtic-Plus.

Project acronym: SENDATE.

Project title: SEcure Networking for a DATa Center Cloud in Europe.

Duration: 36 months (2016–19).

Coordinator: Nokia.

Other partners: 50+ partners in Finland, France, Germany, Norway, and Sweden. Selected partners involved: Nokia, Orange.

Abstract: The project addresses the convergence of telecommunication networks and IT in the context of distributed data centers. We are involved in the TANDEM subproject that targets the infrastructure of such a distributed system. More specifically, we are studying new approaches in terms of software engineering and component-based solutions for enabling this convergence of network and IT.

Participants: Lionel Seinturier [correspondant].

9.4. International Initiatives

9.4.1. Inria Associate Teams Not Involved in an Inria International Labs

9.4.1.1. SOMCA

Title: Self-Optimization of Service Oriented Architectures for Mobile and Cloud Applications

International Partner (Institution - Laboratory - Researcher):

Université du Québec À Montréal (Canada) - LATECE - Naouel MOHA

Start year: 2017

See also: <http://sofa.uqam.ca/somca.php>

The long-term goal of this research program is to propose a novel and innovative methodology embodied in an software platform, to support the runtime detection and correction of anti-patterns in large-scale service-oriented distributed systems in order to continuously optimize their quality of service. One originality of this program lies in the dynamic nature of the service-oriented environments and the application on emerging frameworks for embedded and distributed systems (e.g., Android/iOS for mobile devices, PaaS/SaaS for Cloud environments), and in particular mobile systems interacting with remote services hosted on the Cloud.

9.4.2. Participation in Other International Programs

9.4.2.1. PHC Zenon Cyprus - Project RRI-MobDev

Participants: Sarra Habchi, Lakhdar Meftah, Mohammad Naseri, Romain Rouvoy [correspondant], Walter Rudametkin Ivey, Romain Sommerard, Antoine Vastel.

RRI-MobDev (*Responsible Research and Innovation for Mobile Application Development*) is a 2-year (2017–2018) bilateral collaboration with UCLan Cyprus, an overseas campus of the University of Central Lancashire. Mobile applications are part of a complex ecosystem involving various stakeholders (developers, users, app stores, etc.) exposed to various threats, including not only malware, but also potential information leaks through the continuous interactions with remote servers. This project aims to study and alleviate this problem by intervening both with the users and the developers of mobile apps, with an aim of enabling a cleaner, safer and more responsible mobile app ecosystem.

9.5. International Research Visitors

9.5.1. Visits of International Scientists

Rajkumar Buyya, Redmond Barry Distinguished Professor and Director of the Cloud Computing and Distributed Systems (CLOUDS) Laboratory at the University of Melbourne, visited us in February 2018.

9.5.1.1. Internships

Mohammad Naseri, MSc. Student in Computer Science from Saarland University, Germany, visited us for 3 months, from November 2017 to January 2018.

Chaima Chakhaba, MSc. Student in Computer Science from ESI Alger, Algeria, visited us for 9 months, from December 2017 to August 2018.

9.5.2. Visits to International Teams

9.5.2.1. Research Stays Abroad

Thomas Durieux, PhD Student, spent 3 months from April to June 2018 in KTH, Sweden.

10. Dissemination

10.1. Promoting Scientific Activities

10.1.1. Scientific Events Organisation

10.1.1.1. General Chair, Scientific Chair

Simon Bliudze

- co-chair International Conference on Verification and Evaluation of Computer and Communication Systems (VECoS)
- co-chair International Workshop on Methods and Tools for Rigorous System Design (MeTRiD)

Laurence Duchien

- IEEE TCSE Distinguished Women in Science and Engineering Leadership Award Committee
- co-chair ECSA Track for Women in Software Architecture

10.1.2. Scientific Events Selection*10.1.2.1. Member of the Conference Program Committees***Simon Bludze**

- International Conference on Formal Aspects of Component Software (FACS)
- International Symposium on Formal Approaches to Parallel and Distributed Systems (4PAD)
- Interaction and Concurrency Experience International Workshop (ICE)
- Euromicro Conference on Software Engineering and Advanced Applications (SEAA)
- International Conference on Coordination Models and Languages (COORDINATION)
- FME Workshop on Formal Methods in Software Engineering (FormaliSE)
- International Workshop on Foundations of Coordination Languages and Self-Adaptive Systems (Foclasa)

Pierre Bourhis

- International Joint Conferences on Artificial Intelligence (IJCAI)
- AAAI Conference on Artificial Intelligence

Laurence Duchien

- International Workshop on Variability Modelling of Software-Intensive Systems (VAMOS)
- Live Adaptation of Software Systems (LASSY)
- International Workshop on Software Qualities and their Dependencies (SQUADE)
- Workshop on Distributed Software Development, Software Ecosystems and Systems-of-Systems (WDES)
- Doctoral Symposium ECSA
- ICSE Poster Session
- International Workshop on Software Engineering Research and Industrial Practise (ICSE/SER&IP)
- Conference on Software Engineering Education and Training (CSEE&T)
- European Conference on Software Architecture (ECSA)

Philippe Merle

- International Conference on Cooperative Information Systems (CoopIS)
- International Conference on Cloud Computing Work-in-Progress Track (IEEE CLOUD)
- International Conference on Web Engineering (ICWE)
- International Conference on Mobile Ubiquitous Computing, Systems, Services and Technologies (UBICOMM)
- International Conference on Cloud Computing, GRIDs, and Virtualization (CLOUD COMPUTING)
- International Conference on Advanced Service Computing (SERVICE COMPUTATION)
- Workshop on CrossCloud Infrastructures & Platforms (CrossCloud)

Clément Quinton

- International Workshop on Software Engineering for Systems-of-Systems (SESOS)
- Euromicro Conference on Software Engineering and Advanced Applications (SEAA)
- International Systems and Software Product Line Conference (SPLC)

Romain Rouvoy

- IEEE/ACM International Conference on Automated Software Engineering (ASE), demonstration track
- ACM/USENIX European Conference on Computer Systems (Eurosys)
- IEEE International Conference on Cloud Engineering (IC2E)
- IEEE/ACM International Conference on Mobile Software Engineering and Systems (MOBILESoft)
- ACM Symposium on Applied Computing (SAC), DADS track

Lionel Seinturier

- International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS)
- IEEE/ACM International Conference on Mobile Software Engineering and Systems (MOBILESoft)
- ACM Symposium on Applied Computing (SAC) SA-TTA & OS Tracks
- Australasian Software Engineering Conference (ASWEC)
- International Conference on Software Technologies (ICSOFT)
- International Conference on Service Oriented Computing (ICSOC)
- ACM International Conference on Management of Emergent Digital EcoSystems (MEDES)

Faiez Zalila

- International Conference on Verification and Evaluation of Computer and Communication Systems (VECoS)

10.1.3. Journal*10.1.3.1. Member of the Editorial Boards*

Lionel Seinturier is editor for software engineering of the ISTE-Wiley Computer Science and Information Technology book collection.

10.1.3.2. Reviewer - Reviewing Activities

Yahya Al-Dhuraibi: IEEE Transactions on Services Computing (TSC), The Computer Journal of Oxford University Press.

Simon Bludze: Elsevier Science of Computer Programming (SCICO).

Stéphanie Challita: Wiley Software: Practice & Experience (SPE), Wiley Concurrency and Computation: Practice and Experience (CCPE).

Laurence Duchien: Elsevier Journal of Systems and Software (JSS), IEEE Transaction on Cloud Computing (TCC).

Philippe Merle: Elsevier Future Generation Computer Systems (FGCS), Wiley Software: Practice and Experience (SPE), Elsevier Sustainable Computing, Informatics and Systems (SUSCOM), Wiley Concurrency and Computation: Practice and Experience (CCPE), IGI Global International Journal of Web Services Research (IJWSR).

Romain Rouvoy: Elsevier Information and Software Technology (IST).

Lionel Seinturier: Springer Software Quality Journal (SQJ), Elsevier Sustainable Computing, Informatics and Systems (SUSCOM), Springer Empirical Software Engineering (EMSE), Elsevier Future Generation Computer Systems (FGCS).

Faiez Zalila: Elsevier Future Generation Computer Systems (FGCS), IEEE ACCESS, Springer Journal of Internet Services and Applications (JISA).

10.1.4. Invited Talks

Philippe Merle: International Conference on Cloud and Robotics (ICCR), invited talk.

Faiez Zalila: Cloudification of the Internet of Things Conference (CIoT), tutorial.

10.1.5. Leadership within the Scientific Community

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

10.1.6. Scientific Expertise

Laurence Duchien was member of the recruitment committee for Professor in Computer Science at Université du Littoral Côte d’Opale. She was scientific expert for DGRI PHC Ulysse, NSERC-Canada, PHC Van Gogh Netherlands, PHC Amadeus Austria, Programme Emergence Ville de Paris Vinci, Prix de thèse GDR GPL, Prix de valorisation de la MESHS.

Philippe Merle was member of the recruitment committee for junior researcher (CRCN) at Inria Lille - Nord Europe, and Associate Professor at Télécom SudParis. He was scientific reviewer for ANR.

Romain Rouvoy was member of the recruitment committees for associate professor in Computer Science at ENSEIRB MATMECA (Bordeaux), Université de Nantes and Université de Rennes. He was also a member of the scientific evaluation committee of ANR (CES25). He was scientific expert for DRRT IDF, ADEME PERFECTO and FET-OPEN - Novel Ideas for Radically New Technologies - *Research and Innovation Actions* (RIA).

Lionel Seinturier was member of the recruitment committee for Professor in Computer Science at CNAM Paris, and Associate Professor at University of Lille. He was member of the Promotion Committee to Full Professor for Charles University, Prague. He was scientific expert for ANRT, ECOS Nord, DRRT IDF, PICOM/Silab Retail Booster.

10.1.7. Research Administration

Laurence Duchien is member of the CNRS CoCNRS section 6 committee, and of the “bureau” of this committee. She is elected member-at-large for the **IEEE Technical Council of Software Engineering (TCSE)** 2016–18.

Philippe Merle is member of the Inria scientific board, 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 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 the steering committee of CIEL (Conférence en Ingénierie du Logiciel).

Romain Rouvoy is an elected member of the “bureau” of the ACM SIGOPS French chapter (ASF) and elected member of the administrative council of SpecifCampus.

Lionel Seinturier is president of the CDT (Commission Développement Technologique), and member of the BCEP (Bureau du Comité des Équipes-Projets), for the Inria Lille - Nord Europe research center. He is Scientific Advisor for the evaluation of ICT research laboratories at the Hcéres.

10.2. Teaching - Supervision - Juries

10.2.1. Teaching

Simon Bliudze is, in addition to his tenure junior research position at Inria, chargé d’enseignement at École Polytechnique, Palaiseau, France, in the Department of Computer Sciences (DIX).

- INF411: Les bases de la programmation et de l’algorithmique, 40h, 2e année du Cycle ingénieur
- INF442: Traitement des données massives, 40h, 2e année du Cycle ingénieur

Pierre Bourhis is, in addition to his tenure junior research position at CNRS, chargé d'enseignement spécialité Sciences des données at École Polytechnique, Palaiseau, France, in the Department of Computer Sciences (DIX).

- École Polytechnique, Bases de données, 18h, Cycle Polytechnique Info553
- École Polytechnique, Modal Graphe Géant, 36h
- École Centrale de Lille, Master DAD, Bases de données, 12h

Laurence Duchien teaches at the University of Lille in the FST faculty. She headed the Carrières et Emplois service. She was Director of Doctoral Studies for Computer Science in the Doctoral School Engineering Science (SPI) - ComUE Lille Nord de France. She is now project leader for the doctoral studies at University of Lille.

- Software Project Management, 50h, Level M2, Master MIAGE
- Design of distributed applications, 42h, Level M1, Master of Computer Science
- Research and Innovation Initiation, 22h, Level M2 IAGL, Master of Computer Science
- Tutoring Internship, 16h, Level M2, Master of Computer Science

Clément Quinton teaches at the University of 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 University of Lille in the FST faculty. He heads the **Master of Computer Science** program at the University of 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 University of Lille in the Polytech 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 University of Lille in the FST faculty. He heads the Computer Science Department at the Faculty of Science and Technology of the University of 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

10.2.2. Supervision

- PhD: Yahya Al-Dhuraibi, Flexible Framework for Elasticity in Cloud Computing, University of Lille, December 2018, Philippe Merle.
- PhD: Stéphanie Challita, Inferring Models from Cloud APIs and Reasoning over Them: a Tooled and Formal Approach, University of Lille, December 2018, Philippe Merle.
- PhD: Thomas Durieux, From Runtime Failures to Patches: Study of Patch Generation in Production, University of Lille, September 2018, Lionel Seinturier & Martin Monperrus. [11]
- PhD: Gustavo Sousa, A Software Product Lines-Based Approach for the Setup and Adaptation of Multi-Cloud Environments, University of Lille, June 2018, Laurence Duchien & Walter Rudametkin Ivey. [12]
- 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: Vikas Mishra, Collaborative Strategies to Protect Against Browser Fingerprinting, October 2018, Romain Rouvoy & Walter Rudametkin & Lionel Seinturier.
- PhD in progress: Marion Tommasi, Collaborative Data-centric Workflows: Towards Knowledge Centric Workflows and Integrating Uncertain Data, October 2018, Pierre Bourhis & Lionel Seinturier.
- 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: 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: 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: Zeinab Abou Khalil, November 2017, Laurence Duchien & Clément Quinton, co-supervision with Tom Mens (University of Mons, Belgium).
- PhD in progress: Guillaume Fieni, GreenData : Vers un traitement efficace et éco-responsable des grandes masses de données numériques, October 2017, Romain Rouvoy & Lionel Seinturier.
- PhD in progress: Benjamin Danglot, Software Testing Amplification, December 2016, Martin Monperrus & Lionel Seinturier.
- PhD in progress: Lakhdar Meftah, Cartography of the Quality of Experience for Mobile Internet Access, November 2016, Romain Rouvoy, co-supervision with Isabelle Chrisment (Inria Madynes).
- PhD in progress: Sarra Habchi, Une supervision de contexte sensible à la confidentialité pour les développements logiciels en crowdsourcing, October 2016, Romain Rouvoy.
- PhD in progress: Antoine Vastel, Cartographie de la qualité d'expérience pour l'accès à l'internet mobile, October 2016, Romain Rouvoy & Walter Rudametkin.

10.2.3. Juries

Simon Bliudze

- Maaz Mashood Mohiuddin (EPFL, Switzerland), examiner
- Emmanouela Stachtari (Aristotle U. of Thessaloniki, Greece), examiner

Laurence Duchien

- HDR Chouki Tibermacine (U. Montpellier), reviewer
- HDR Thomas Ledoux (U. Nantes), reviewer
- HDR Nikolaos Georgantas (U. Pierre et Marie Curie), examiner
- HDR Matthieu Roy (U. Toulouse), president
- HDR Raul Mazo (U. Paris 1 Sorbonne), reviewer
- Anne-Lucie Vion (U. Grenoble), reviewer
- Van Cam Pham, CEA, examiner
- Benjamin Allaert (U. Lille), president
- Farzan Kalantari (U. Lille), president
- Eddy Habach (U. Nice Côte d'Azur), examiner
- Mathieu Allon (U. Lille), president
- Juan Carlos Munoz (U. Paris 1 Sorbonne)
- Paola Andreas Gomez Barreto (U. Grenoble), examiner

Philippe Merle

- HDR Sylvain Vauttier (U. Montpellier), reviewer
- Nabila Belhaj (Télécom SudParis), examiner

Lionel Seinturier

- Mohamed Mosli (Telecom SudParis), president
- William Excoffon (U. Toulouse), president
- Alex Palesandro (U. Lyon 3), president
- Bassirou Ngom (Sorbonne U.), reviewer
- Gabriel Campeanu, (Märladalen U., Sweden), examiner
- Isabelle Astic (CNAM Paris), examiner

10.3. Popularization

10.3.1. Interventions

Lionel Seinturier participated in October to the Chercheur itinérant event organized by the Inria Lille Nord Europe research center in the context of La fête de la science. Two classrooms (2nde and terminale) have been visited. The theme of the visit was crowdsensing and data gathering from mobile devices.

11. Bibliography

Major publications by the team in recent years

- [1] S. ABITEBOUL, P. BOURHIS, 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, <https://hal.inria.fr/hal-01744978>
- [2] Y. AL-DHURAIBI, F. PARAISSO, N. DJARALLAH, P. MERLE. *Elasticity in Cloud Computing: State of the Art and Research Challenges*, in "IEEE Transactions on Services Computing (TSC)", March 2018, vol. 11, n^o 2, pp. 430-447 [DOI : 10.1109/TSC.2017.2711009], <https://hal.inria.fr/hal-01529654>

- [3] M. COLMANT, M. KURPICZ, P. FELBER, L. HUERTAS, R. ROUVOY, A. SOBE. *Process-level Power Estimation in VM-based Systems*, in "European Conference on Computer Systems (EuroSys)", Bordeaux, France, T. HARRIS, M. HERLIHY (editors), EuroSys'15: Proceedings of the Tenth European Conference on Computer Systems, ACM, April 2015, 14 p. [DOI : 10.1145/2741948.2741971], <https://hal.inria.fr/hal-01130030>
- [4] T. DURIEUX, Y. HAMADI, M. MONPERRUS. *Fully Automated HTML and Javascript Rewriting for Constructing a Self-healing Web Proxy*, in "Proceedings of the 29th IEEE International Symposium on Software Reliability Engineering (ISSRE 2018)", Memphis, United States, October 2018, <https://arxiv.org/abs/1803.08725>, <https://hal.inria.fr/hal-01746141>
- [5] S. HABCHI, X. BLANC, 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", Montpellier, France, ACM Press, September 2018, vol. 11 [DOI : 10.1145/3238147.3238197], <https://hal.inria.fr/hal-01829135>
- [6] C. QUINTON, D. ROMERO, L. DUCHIEN. *SALOON: a platform for selecting and configuring cloud environments*, in "Software: Practice and Experience", January 2016, vol. 46, pp. 55-78 [DOI : 10.1002/SPE.2311], <https://hal.inria.fr/hal-01103560>
- [7] E. STACHTIARI, A. MAVRIDOU, P. KATSAROS, S. BLIUDZE, J. SIFAKIS. *Early validation of system requirements and design through correctness-by-construction*, in "Journal of Systems and Software", 2018, vol. 145, pp. 52-78 [DOI : 10.1016/J.JSS.2018.07.053], <https://hal.archives-ouvertes.fr/hal-01873999>
- [8] S. URLI, Z. YU, L. SEINTURIER, M. MONPERRUS. *How to Design a Program Repair Bot? Insights from the Repairnator Project*, in "40th International Conference on Software Engineering, Track Software Engineering in Practice", Gothenburg, Sweden, May 2018, pp. 95-104, update for oadoi on Nov 02 2018 [DOI : 10.1145/3183519.3183540], <https://hal.inria.fr/hal-01691496>
- [9] A. VASTEL, P. LAPERDRIX, W. RUDAMETKIN, R. ROUVOY. *FP-STALKER: Tracking Browser Fingerprint Evolutions*, in "IEEE S&P 2018 - 39th IEEE Symposium on Security and Privacy", San Francisco, United States, B. PARNO, C. KRUEGEL (editors), Proceedings of the 39th IEEE Symposium on Security and Privacy (S&P), IEEE, May 2018, pp. 728-741 [DOI : 10.1109/SP.2018.00008], <https://hal.inria.fr/hal-01652021>
- [10] J. XUAN, M. MARTINEZ, F. DEMARCO, M. CLÉMENT, S. LAMELAS, T. DURIEUX, D. LE BERRE, M. MONPERRUS. *Nopol: Automatic Repair of Conditional Statement Bugs in Java Programs*, in "IEEE Transactions on Software Engineering", 2016 [DOI : 10.1109/TSE.2016.2560811], <https://hal.archives-ouvertes.fr/hal-01285008>

Publications of the year

Doctoral Dissertations and Habilitation Theses

- [11] T. DURIEUX. *From Runtime Failures to Patches: Study of Patch Generation in Production*, University of Lille, September 2018, <https://tel.archives-ouvertes.fr/tel-01883676>
- [12] G. SOUSA. *A Software Product Lines-Based Approach for the Setup and Adaptation of Multi-Cloud Environments*, Université de Lille, June 2018, <https://hal.inria.fr/tel-01808779>

Articles in International Peer-Reviewed Journals

- [13] Y. AL-DHURAIBI, F. PARAISO, N. DJARALLAH, P. MERLE. *Elasticity in Cloud Computing: State of the Art and Research Challenges*, in "IEEE Transactions on Services Computing (TSC)", March 2018, vol. 11, n^o 2, pp. 430-447 [DOI : 10.1109/TSC.2017.2711009], <https://hal.inria.fr/hal-01529654>
- [14] M. COLMANT, R. ROUYOY, M. KURPICZ, A. SOBE, P. FELBER, L. SEINTURIER. *The Next 700 CPU Power Models*, in "Journal of Systems and Software", July 2018, <https://hal.inria.fr/hal-01827132>
- [15] B. DANGLLOT, P. PREUX, B. BAUDRY, M. MONPERRUS. *Correctness Attraction: A Study of Stability of Software Behavior Under Runtime Perturbation*, in "Empirical Software Engineering", August 2018, vol. 23, n^o 4, pp. 2086–2119, <https://arxiv.org/abs/1611.09187> [DOI : 10.1007/s10664-017-9571-8], <https://hal.archives-ouvertes.fr/hal-01378523>
- [16] B. DANGLLOT, O. L. VERA-PÉREZ, B. BAUDRY, M. MONPERRUS. *Automatic Test Improvement with DSpot: a Study with Ten Mature Open-Source Projects*, in "Empirical Software Engineering", 2018, pp. 1-35, <https://hal.inria.fr/hal-01923575>
- [17] H. GUESMI, B. BEN HEDIA, S. BLIUDZE, S. BENSALÉM, B. LENABEC. *TT-BIP: using correct-by-design BIP approach for modelling real-time system with time-triggered paradigm*, in "Innovations in Systems and Software Engineering", June 2018, vol. 14, n^o 2, pp. 117 - 142 [DOI : 10.1007/s11334-018-0312-Y], <https://hal.inria.fr/hal-01823749>
- [18] M. MOHIUDDIN, W. SAAB, S. BLIUDZE, J.-Y. LE BOUDEC. *Axo: Detection and Recovery for Delay and Crash Faults in Real-Time Control Systems*, in "IEEE Transactions on Industrial Informatics", July 2018, vol. 14, n^o 7, pp. 3065 - 3075 [DOI : 10.1109/TII.2017.2772219], <https://hal.archives-ouvertes.fr/hal-01846124>
- [19] E. STACHTIARI, A. MAVRIDOU, P. KATSAROS, S. BLIUDZE, J. SIFAKIS. *Early validation of system requirements and design through correctness-by-construction*, in "Journal of Systems and Software", 2018, vol. 145, pp. 52-78 [DOI : 10.1016/J.JSS.2018.07.053], <https://hal.archives-ouvertes.fr/hal-01873999>
- [20] O. L. VERA-PÉREZ, B. DANGLLOT, M. MONPERRUS, B. BAUDRY. *A Comprehensive Study of Pseudo-tested Methods*, in "Empirical Software Engineering", 2018, pp. 1-33 [DOI : 10.1007/s10664-018-9653-2], <https://hal.inria.fr/hal-01867423>
- [21] Z. YU, M. MARTINEZ, B. DANGLLOT, T. DURIEUX, M. MONPERRUS. *Alleviating Patch Overfitting with Automatic Test Generation: A Study of Feasibility and Effectiveness for the Nopol Repair System*, in "Empirical Software Engineering", 2018 [DOI : 10.1007/s10664-018-9619-4], <https://hal.inria.fr/hal-01774223>

International Conferences with Proceedings

- [22] S. ABITEBOUL, P. BOURHIS, 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, <https://hal.inria.fr/hal-01744978>
- [23] Y. AL-DHURAIBI, F. ZALILA, N. B. DJARALLAH, P. MERLE. *Coordinating Vertical Elasticity of both Containers and Virtual Machines*, in "8th International Conference on Cloud Computing and Services Science - CLOSER 2018", Funchal, Madeira, Portugal, March 2018, <https://hal.archives-ouvertes.fr/hal-01683041>

- [24] S. CHALLITA, F. ZALILA, C. GOURDIN, P. MERLE. *A Precise Model for Google Cloud Platform*, in "6th IEEE International Conference on Cloud Engineering (IC2E)", Orlando, Florida, United States, <http://conferences.computer.org/IC2E/2018/>, April 2018, pp. 177-183, <https://hal.inria.fr/hal-01689659>
- [25] S. CHALLITA, F. ZALILA, P. MERLE. *Specifying Semantic Interoperability between Heterogeneous Cloud Resources with the FCLOUDS Formal Language*, in "CLOUD 2018 - 11th IEEE International Conference on Cloud Computing", San Francisco, United States, July 2018, <https://hal.inria.fr/hal-01790629>
- [26] *Best Paper*
T. DURIEUX, Y. HAMADI, M. MONPERRUS. *Fully Automated HTML and Javascript Rewriting for Constructing a Self-healing Web Proxy*, in "Proceedings of the 29th IEEE International Symposium on Software Reliability Engineering (ISSRE 2018)", Memphis, United States, October 2018, <https://arxiv.org/abs/1803.08725> [DOI : 10.1109/ISSRE.2018.00012], <https://hal.inria.fr/hal-01746141>.
- [27] T. DURIEUX, Y. HAMADI, Z. YU, B. BAUDRY, M. MONPERRUS. *Exhaustive Exploration of the Failure-oblivious Computing Search Space*, in "ICST 2018 - 11th IEEE Conference on Software Testing, Validation and Verification", Västerås, Sweden, April 2018, <https://arxiv.org/abs/1710.09722> - arXiv admin note: substantial text overlap with arXiv:1603.07631 [DOI : 10.1109/ICST.2018.00023], <https://hal.inria.fr/hal-01624988>
- [28] G. FIENI, R. ROUVOY, L. SEINTURIER. *Towards Sustainable Software Infrastructures for Data-Intensive Systems*, in "EuroDW 2018 - 12th EuroSys Doctoral Workshop", Porto, Portugal, April 2018, <https://hal.inria.fr/hal-01784225>
- [29] S. HABCHI, X. BLANC, 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", Montpellier, France, ACM Press, September 2018, vol. 11 [DOI : 10.1145/3238147.3238197], <https://hal.inria.fr/hal-01829135>
- [30] F. KORTE, S. CHALLITA, Z. FAIEZ, P. MERLE, J. GRABOWSKI. *Model-Driven Configuration Management of Cloud Applications with OCCI*, in "8th International Conference on Cloud Computing and Services Science (CLOSER)", Funchal, Madeira, Portugal, March 2018, pp. 100-111, <https://hal.inria.fr/hal-01678945>
- [31] M. NASERI, N. P. BORGES JR., A. ZELLER, 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, <https://hal.inria.fr/hal-01929049>
- [32] W. SAAB, M. MOHIUDDIN, S. BLIUDZE, J.-Y. LE BOUDEC. *Ordering Events Based on Intentionality in Cyber-Physical Systems*, in "ICCPS 2018 - 9th ACM/IEEE International Conference on Cyber-Physical Systems", Porto, Portugal, IEEE, April 2018 [DOI : 10.1109/ICCPS.2018.00019], <https://hal.archives-ouvertes.fr/hal-01889067>
- [33] V. SOBREIRA, T. DURIEUX, F. MADEIRAL, M. MONPERRUS, M. A. MAIA. *Dissection of a Bug Dataset: Anatomy of 395 Patches from Defects4J*, in "SANER 2018 - 25th edition of the IEEE International Conference on Software Analysis, Evolution and Reengineering", Campobasso, Italy, March 2018, <https://arxiv.org/abs/1801.06393> [DOI : 10.1109/SANER.2018.8330203], <https://hal.inria.fr/hal-01689684>

- [34] S. URLI, Z. YU, L. SEINTURIER, M. MONPERRUS. *How to Design a Program Repair Bot? Insights from the Repairator Project*, in "40th International Conference on Software Engineering, Track Software Engineering in Practice", Gothenburg, Sweden, May 2018, pp. 95-104, update for oadoi on Nov 02 2018 [DOI : 10.1145/3183519.3183540], <https://hal.inria.fr/hal-01691496>
- [35] A. VASTEL, P. LAPERDRIX, W. RUDAMETKIN, R. ROUVOY. *FP-Scanner: The Privacy Implications of Browser Fingerprint Inconsistencies*, in "Proceedings of the 27th USENIX Security Symposium", Baltimore, United States, August 2018, <https://hal.inria.fr/hal-01820197>
- [36] A. VASTEL, P. LAPERDRIX, W. RUDAMETKIN, R. ROUVOY. *FP-STALKER: Tracking Browser Fingerprint Evolutions*, in "IEEE S&P 2018 - 39th IEEE Symposium on Security and Privacy", San Francisco, United States, B. PARNO, C. KRUEGEL (editors), Proceedings of the 39th IEEE Symposium on Security and Privacy (S&P), IEEE, May 2018, pp. 728-741 [DOI : 10.1109/SP.2018.00008], <https://hal.inria.fr/hal-01652021>
- [37] A. VASTEL, W. RUDAMETKIN, R. ROUVOY. *FP-TESTER: Automated Testing of Browser Fingerprint Resilience*, in "IWPE 2018 - 4th International Workshop on Privacy Engineering", London, United Kingdom, Proceedings of the 4th International Workshop on Privacy Engineering (IWPE'18), April 2018, pp. 1-5, <https://hal.inria.fr/hal-01717158>

National Conferences with Proceedings

- [38] *Best Paper*
F. MADEIRAL, T. DURIEUX, V. SOBREIRA, M. MAIA. *Towards an automated approach for bug fix pattern detection*, in "VEM '18 - Proceedings of the VI Workshop on Software Visualization, Evolution and Maintenance", São Carlos, Brazil, September 2018, <https://arxiv.org/abs/1807.11286> , <https://hal.archives-ouvertes.fr/hal-01851813>.

Conferences without Proceedings

- [39] M. WHITE, M. TUFANO, M. MARTINEZ, M. MONPERRUS, D. POSHYVANYK. *Sorting and Transforming Program Repair Ingredients via Deep Learning Code Similarities*, in "IEEE International Conference on Software Analysis, Evolution and Reengineering", Hangzhou, China, 2019, <https://arxiv.org/abs/1707.04742> , <https://hal.archives-ouvertes.fr/hal-01581170>

Research Reports

- [40] C. BARILLOT, I. FIJALKOW, I. QUEINNEC, F. THEOLEYRE, H. TOUZET, M. BEAUDOUIN-LAFON, M.-P. CANI, F. CHAUMETTE, G. CONREUR, V. CORTIER, C. GAVOILLE, G. GOGNIAT, J. GOSSA, A. HERZIG, J. KRIVINE, P. LAMARRE, F. LAMNABHI-LAGARRIGUE, A.-C. LETOURNEL, L. SEINTURIER, I. TELLIER, S. TORRES, C. TRUCHET, A. TSOUKIAS, M. VERLEYSSEN. *Prospective Report of the scientific council of the Information Sciences and their interactions Institute of CNRS*, CNRS, December 2018, pp. 1-87, <https://hal.archives-ouvertes.fr/hal-01956087>
- [41] X. QIN, S. BLIUDZE, E. MADELAINE, M. ZHANG. *Using SMT engine to generate Symbolic Automata -Extended version*, Inria & Université Cote d'Azur, CNRS, I3S, Sophia Antipolis, France ; inria, June 2018, n° RR-9177, <https://hal.inria.fr/hal-01823507>
- [42] L. SEINTURIER. *Adlet: A Java-based Architecture Description Language*, Inria Lille - Nord Europe, December 2018, n° RR-9242, <https://hal.inria.fr/hal-01964792>

Other Publications

- [43] A. AMARILLI, P. BOURHIS, M. MONET, P. SENELLART. *Evaluating Datalog via Tree Automata and Cycluits*, October 2018, <https://arxiv.org/abs/1808.04663> - 53 pages, 61 references. Journal version of "Combined Tractability of Query Evaluation via Tree Automata and Cycluits (Extended Version)" at arXiv:1612.04203. To appear in *Theory of Computing Systems*, <https://hal.inria.fr/hal-01891814>

References in notes

- [44] L. CAPRA, W. EMMERICH, C. MASCOLO. *A micro-economic approach to conflict resolution in mobile computing*, in "SIGSOFT FSE", 2002, pp. 31-40
- [45] M. COLMANT. *Multi-Dimensional Analysis of Software Power Consumptions in Multi-Core Architectures*, Université Lille 1 - Sciences et Technologies, November 2016, <https://tel.archives-ouvertes.fr/tel-01403559>
- [46] B. CORNU. *Automatic Analysis and Repair of Exception Bugs for Java Programs*, Université de Lille, November 2015, <https://tel.archives-ouvertes.fr/tel-01250092>
- [47] V. DALLMEIER, A. ZELLER, B. MEYER. *Generating Fixes from Object Behavior Anomalies*, in "Proceedings of the 2009 IEEE/ACM International Conference on Automated Software Engineering", Washington, DC, USA, ASE '09, IEEE Computer Society, 2009, pp. 550–554, <http://dx.doi.org/10.1109/ASE.2009.15>
- [48] P. ECKERSLEY. *How Unique is Your Web Browser?*, in "Proceedings of the 10th International Conference on Privacy Enhancing Technologies", Berlin, Heidelberg, PETS'10, Springer-Verlag, 2010, pp. 1–18, <http://dl.acm.org/citation.cfm?id=1881151.1881152>
- [49] F. FOUQUET, E. DAUBERT, N. PLOUZEAU, O. BARAIS, J. BOURCIER, 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)", Stockholm, Sweden, K. M. GÖSCHKA, S. HARIDI (editors), Lecture Notes in Computer Science, Springer, June 2012, vol. 7272, pp. 16-30
- [50] M. GOLDMANN, G. KREITZ. *Measurements on the spotify peer-assisted music-on-demand streaming system*, in "Peer-to-Peer Computing", T. ASAMI, T. HIGASHINO (editors), IEEE, 2011, pp. 206-211
- [51] M. HALL, E. FRANK, G. HOLMES, B. PFAHRINGER, P. REUTEMANN, I. H. WITTEN. *The WEKA data mining software: an update*, in "SIGKDD Explorer Newsletter", 2009, vol. 11, n^o 1, pp. 10–18
- [52] M. HARMAN. *Automated patching techniques: the fix is in*, in "Communications of the ACM", May 2010, vol. 53, n^o 5, pp. 108–108, <http://doi.acm.org/10.1145/1735223.1735248>
- [53] V. ISSARNY, N. GEORGANTAS, S. HACHEM, A. ZARRAS, P. VASSILIADIS, M. AUTILI, M. A. GEROSA, A. BEN HAMIDA. *Service-Oriented Middleware for the Future Internet: State of the Art and Research Directions*, in "Journal of Internet Services and Applications", May 2011, vol. 2, n^o 1, pp. 23-45, <http://dx.doi.org/10.1007/s13174-011-0021-3>
- [54] D. JACKSON. *Alloy: A Lightweight Object Modelling Notation*, in "ACM Transactions on Software Engineering and Methodology", April 2002, vol. 11, n^o 2, pp. 256-290

- [55] W. LI, S. HU, J. LI, H.-A. JACOBSEN. *Community Clustering for Distributed Publish/Subscribe Systems*, in "CLUSTER", IEEE, 2012, pp. 81-89
- [56] M. LÉGER, T. LEDOUX, T. COUPAYE. *Reliable Dynamic Reconfigurations in a Reflective Component Model*, in "CBSE", L. GRUNSKÉ, R. REUSSNER, F. PLASIL (editors), Lecture Notes in Computer Science, Springer, 2010, vol. 6092, pp. 74-92
- [57] M. MARTINEZ. *Extraction and Analysis of Knowledge for Automatic Software Repair*, Université Lille 1, October 2014, <https://hal.archives-ouvertes.fr/tel-01078911>
- [58] P. MERLE, J.-B. STEFANI. *A formal specification of the Fractal component model in Alloy*, Inria, November 2008, n^o 6721, 44 p. , <http://hal.inria.fr/inria-00338987>
- [59] A. NOUREDDINE. *Towards a Better Understanding of the Energy Consumption of Software Systems*, Université des Sciences et Technologie de Lille - Lille I, March 2014, <https://tel.archives-ouvertes.fr/tel-00961346>
- [60] F. PARAISO, P. MERLE, L. SEINTURIER. *soCloud: A service-oriented component-based PaaS for managing portability, provisioning, elasticity, and high availability across multiple clouds*, in "Springer Computing", May 2016, vol. 98, n^o 5, pp. 539-565 [DOI : 10.1007/s00607-014-0421-x], <https://hal.inria.fr/hal-01019385>
- [61] L. PROVENSI, F. ELIASSEN, R. VITENBERG, 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)", Coimbra, Portugal, K. M. GÖSCHKA, R. OLIVEIRA, P. PIETZUCH, G. RUSSELLO (editors), ACM, March 2013, vol. 1, pp. 415-422, Best paper award, <http://hal.inria.fr/hal-00799136>
- [62] M. RINARD, C. CADAR, D. DUMITRAN, D. ROY, T. LEU, W. BEEBEE 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
- [63] M. SALEHIE, L. TAHVILDARI. *Self-adaptive software: Landscape and research challenges*, in "ACM Transactions on Autonomous and Adaptive Systems", May 2009, vol. 4, n^o 2, pp. 14:1-14:42, <http://doi.acm.org/10.1145/1516533.1516538>
- [64] H. SCHAFFERS, M. SHARPE. *Services in the Future Internet*, April 2011, FP8 Expert Group, EU Directorate-General Information Society & Media, <http://cordis.europa.eu/fp7/ict/ssai/docs/softwareconsultationreportfeb2011.pdf>
- [65] L. SCHUBERT, K. JEFFERY, B. NEIDECKER-LUTZ. *A Roadmap for Advanced Cloud Technologies under H2020*, Dec 2012, <http://cordis.europa.eu/fp7/ict/ssai/docs/cloud-expert-group/roadmap-dec2012-vfinal.pdf>
- [66] L. SEINTURIER, P. MERLE, R. ROUVOY, D. ROMERO, V. SCHIAVONI, J.-B. STEFANI. *A Component-Based Middleware Platform for Reconfigurable Service-Oriented Architectures*, in "Software: Practice and Experience", May 2012, vol. 42, n^o 5, pp. 559-583 [DOI : 10.1002/SPE.1077], <http://hal.inria.fr/inria-00567442>

- [67] J. URBAN. *NESSI Research Priorities for the next Framework Programme for Research and Technological Development FP8*, May 2011, http://www.nessi-europe.eu/files/Docs/NESSI%20SRA_update_May_2011_V1-0.pdf
- [68] Y. WEI, Y. PEI, C. A. FURIA, L. S. SILVA, S. BUCHHOLZ, B. MEYER, A. ZELLER. *Automated fixing of programs with contracts*, in "Proceedings of the 19th international symposium on Software testing and analysis", New York, NY, USA, ISSTA '10, ACM, 2010, pp. 61–72, <http://doi.acm.org/10.1145/1831708.1831716>
- [69] C. WEISS, R. PREMRAJ, T. ZIMMERMANN, 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, 1 p.
- [70] T. XIE, S. THUMMALAPENTA, D. LO, C. LIU. *Data Mining for Software Engineering*, in "IEEE Computer", August 2009, vol. 42, n^o 8, pp. 55-62, <http://doi.ieeecomputersociety.org/10.1109/MC.2009.256>
- [71] W. XU, L. HUANG, A. FOX, D. A. PATTERSON, 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)", Big Sky, Montana, USA, J. N. MATTHEWS, T. E. ANDERSON (editors), ACM, October 2009, pp. 117-132
- [72] 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. PREHOFE, 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, J. WUTTKE. *Software Engineering for Self-Adaptive Systems: A Second Research Roadmap*, in "Software Engineering for Self-Adaptive Systems", R. DE LEMOS, H. GIESE, H. MÜLLER, M. SHAW (editors), Dagstuhl Seminar Proceedings, Springer, 2013, vol. 7475, pp. 1-26, <http://hal.inria.fr/inria-00638157>