Activity Report 2017

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

Project-Team SPIRALS

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

Keywords:

Computer Science and Digital Science:

A1.1.6. - Cloud
A1.3. - Distributed Systems
A1.3.1. - Blockchain
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
B9.4.1. - Computer science
B9.4.5. - Data science
B9.8. - Privacy

1. Personnel

Research Scientists

Pierre Bourhis [CNRS, Researcher, from Sep 2017]
Philippe Merle [Inria, Researcher, HDR]
Simon Bliudze [Inria, Researcher, from Dec 2017]

Faculty Members

Lionel Seinturier [Team leader, Univ des sciences et technologies de Lille, Professor, HDR]
Laurence Duchien [Univ des sciences et technologies de Lille, Professor, HDR]
Martin Monperrus [Univ des sciences et technologies de Lille, Associate Professor, until Aug 2017, HDR]

Romain Rouvoy [Univ des sciences et technologies de Lille, Professor, HDR]
Walter Rudametkin Ivey [Univ des sciences et technologies de Lille, Associate Professor]
Clément Quinton [Univ des sciences et technologies de Lille, Associate Professor, from Sep 2017]

Post-Doctoral Fellows

Zhongxing Yu [Inria]
Faiez Zalila [Inria, from Feb 2017]
Fawaz Paraiso [Inria, until Mar 2017]

PhD Students

Zeinab Abou Khalil [Univ des sciences et technologies de Lille, from Nov 2017]
Yahya Al-Dhuraibi [Scalair, granted by CIFRE]
Stéphanie Challita [Inria]
Maxime Colmant [Univ des sciences et technologies de Lille, until Jun 2017]
Benjamin Danglot [Inria]
Thomas Durieux [Inria]
Guillaume Fieni [Univ des sciences et technologies de Lille, from Sep 2017]
Maria Gomez Lacruz [Inria, until Feb 2017]
Miguel Alejandro Gonzalez [CONACYT, until Aug 2017]
Sarra Habchi [Inria]
Geoffrey Hecht [Univ des sciences et technologies de Lille, until Aug 2017]
Lakhdar Meftah [Inria]
Gustavo Sousa [Erasmus Mundus]
Antoine Vastel [Univ des sciences et technologies de Lille]
Bo Zhang [Univ des sciences et technologies de Lille, until Apr 2017]
Amal Tahri [Orange, granted by CIFRE]

Technical staff

Aurélien Bourdon [Inria, from Jul 2017]
Julien Duribreux [Inria, until Jan 2017]
Christophe Gourdin [Inria]
Romain Sommerard [Inria, from Feb 2017]
Simon Urli [Inria]
Antoine Veuiller [Inria]

Interns

Moncef Aoudia [Inria, from Jun 2017 until Aug 2017]
Guillaume Fieni [Inria, from Mar 2017 until Aug 2017]
Bacem Hagui [Univ des sciences et technologies de Lille, from May 2017 until Jul 2017]
Mohammad Naseri [Inria, from Nov 2017]
Geoffrey Preudhomme [Inria, from Jun 2017 until Sep 2017]
Martin Rohmer [Univ des sciences et technologies de Lille, from May 2017 until Aug 2017]
Thomas Stievenard [Inria, from May 2017 until Aug 2017]
Rodolphe Toin [Univ des sciences et technologies de Lille, from May 2017 until Aug 2017]
Guillaume Willefert [Inria, from Apr 2017 until Jun 2017]

Administrative Assistants

Aurore Hermant [Inria, until Aug 2017]
Karine Lewandowski [Inria, Sep 2017]
Nathalie Bonte [Inria, from Oct 2017]

Visiting Scientists

Fernanda Madeiral Delfim [Univ. Uberlândia, PhD Student, until May 2017 & Sep 2017]
Nicolas Aguilera Miranda [Fundacion Inria Chile, from Feb 2017 until Jun 2017]
Fatima Belkouch Ameline [Univ du droit et de la santé de Lille, from Sep 2017]

External Collaborators

Christophe Gransart [Ifsttar, Researcher, External collaborator]
Martin Monperrus [KTH, Professor, from Sep 2017]

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 [71], the NESSI Research Priorities for the next Framework Programme for Research and Technological Development FP8 [74], the Roadmap for Advanced Cloud Technologies under H2020 [72], and research roadmaps, such as [79], [70], [61].

3. Research Program

3.1. Introduction

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

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

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

---

[1] http://cordis.europa.eu/fp7/ict/ssai

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 [76].

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** [77]. 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 [59].

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 [55] and [75]. 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" [69] 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 [60]. 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 homeostasis [2] 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 [78]. 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,

---

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

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 [68] 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 [57], 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 [53], 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 [56], but also the adoption of publish/subscribe abstractions [63] 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 [73] 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 [64], 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 compétivité Industries du Commerce*) cluster. We have established strong collaborations with local companies in the context of former funded projects, such as Cappucino and Macchiato, which focused on the development of a new generation of mobile computing platforms for e-commerce. We are also involved in the Datalyse and OCCIware funded projects that define cloud computing environments with applications for the retail industry. Finally, our activities in terms of crowd-sensing and data gathering on mobile devices with the APISENSE® platform share also applications for the retail industry.

4.3. Distributed software services for the digital home

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

5. Highlights of the Year

5.1. Highlights of the Year

5.1.1. *Awards*

María Gómez Lacruz defended her PhD thesis [58] in 2016 in the Spirals project-team. She is now a post-doctoral researcher at Saarland University, Germany. During her thesis, she worked in the domain of crowdsensed data. She proposed algorithms to mine traces of mobile applications in order to detect and reproduce application crashes. This research led to techniques that improve the robustness of mobile applications deployed in the wild. For these results, she was awarded an accessit for the thesis prize of the CNRS GDR *Génie de la programmation et du logiciel* (GPL). See http://gdr-gpl.cnrs.fr/node/284.

Benjamin Danglot and his colleagues—Thomas Durieux, Martin Monperrus, Simon Urli, contributing to the development of the Spoon software library—received the 2017 OW2 Community award for Spoon specific and dynamic community, and the use of Spoon in other projects. OW2 is an independent, global, open-source community that promotes the development of open-source middleware, generic business applications, and cloud computing platforms. Spoon is an OW2 project and a software building block that is used in many of our research activities and projects on self-adaptation. See https://www.ow2con.org/bin/view/2017/Awards_Results?year=2017&event=OW2con17.

Philippe Merle, Christophe Gourdin, and Nathalie Mitton (Inria Fun) received a best paper award in the 2nd IEEE International Congress on Internet of Things (ICIOT 2017) for their work on mobile cloud robotics. The paper proposes a novel interface to bridge the gap between cloud systems and mobile robot systems. See http://iciot.org/2018/news.html.

BEST PAPER AWARD:

[35]
P. MERLE, C. GOURDIN, N. MITTON. *Mobile Cloud Robotics as a Service with OCCIware*, in "2nd IEEE International Congress on Internet of Things, IEEE ICIOT 2017", Honolulu, Hawaii, United States, Proceedings of the 2nd IEEE International Congress on Internet of Things, IEEE ICIOT 2017, June 2017, pp. 50 - 57, Best Paper Award [*DOI :* 10.1109/IEEE.ICIOT.2017.15], https://hal.archives-ouvertes.fr/hal-01522684

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 Veuiller, Christophe Ribeiro, Julien Duribreux, Nicolas Haderer and Romain Rouvoy
- Partner: Université Lille 1
- Contact: Romain Rouvoy
- URL: http://apisense.io

6.2. Nopol

KEYWORD: Automatic software repair

FUNCTIONAL DESCRIPTION: Nopol is an automatic software repair tool for buggy conditional statements (i.e., if-then-else statements) in Java programs. Nopol takes a buggy program as well as a test suite as input and generates a patch with a conditional expression as output. The test suite is required to contain passing test cases to model the expected behavior of the program and at least one failing test case that reveals the bug to be repaired. The process of Nopol consists of three major phases. First, Nopol employs angelic fix localization to identify expected values of a condition during the test execution. Second, runtime trace collection is used to collect variables and their actual values, including primitive data types and objected-oriented features (e.g., nullness checks), to serve as building blocks for patch generation. Third, Nopol encodes these collected data into an instance of a Satisfiability Modulo Theory (SMT) problem, then a feasible solution to the SMT instance is translated back into a code patch.

- Contact: Martin Monperrus
- URL: https://github.com/SpoonLabs/nopol/

6.3. 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

6.4. 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

6.5. 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. A Domain-specific Language for The Control of Self-adaptive Component-based Architecture

In [12], together with Frederico Alvares (Inria Ascola) and Eric Rutten (Inria Ctrl-A), we have proposed Ctrl-F, a new domain-specific language for specifying reconfiguration policies in self-adaptable component-based software systems. Self-adaptive behaviors in the context of component-based architecture are generally designed based on past monitoring events, configurations (component assemblies) as well as behavioral programs defining the adaptation logics and invariant properties. The novelty of the proposed Ctrl-F language is to enable taking decisions on predictions on the possible futures of the system in order to avoid going into branches of the behavioral program leading to bad configurations. Ctrl-F is formally defined by a translation into *Finite State Automata* models. We use *Discrete Controller Synthesis* to automatically generate a controller to enforce correct self-adaptive behaviors. Ctrl-F is integrated with our FraSCAti middleware platform for distributed service and component oriented systems.

7.2. A New Interface for Mobile Cloud Robotics

In [35], together with Nathalie Mitton (Inria Fun), we have proposed OMCRI, a new interface for mobile cloud robotics. This interface enables to abstract from the heterogeneity of robotic platforms and to bring some resource management facilities to fleets of robots. This result is based on the expertise that we have developed in the management of resources for cloud computing environments, especially around the OCCI standard. To the best of our knowledge, OMCRI is the first interface that enables to concretize the vision of robotics as a service. This result has obtained the best award at the 2nd IEEE International Congress on Internet of Things (ICIOT 2017).

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 ongoing PhD thesis of Yahya Al-Dhuraibi.

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 will start in January 2018 in the context of this collaboration.

8.4. Orange Labs

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 [62]. This collaboration takes advantage of the expertise that we are developing since several years on reconfigurable component-based software systems [73], on cloud systems [67], and on the Alloy specification language [66].

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

9. Partnerships and Cooperations

9.1. Regional Initiatives

9.1.1. Région Hauts-de-France

9.1.1.1. Citizen Awareness and Contribution to Air Quality Monitoring
Participants: Moncef Ouadia, Romain Rouvoy [correspondant], Lionel Seinturier, Antoine Veuiller.

This is a 3-year project (2015–17) in the context of the so-called "*Chercheur citoyen*" program. The partners are LISIC/Université Côte d'Opale (leader), ATMO Nord-Pas De Calais, Association Bâtisseurs d'Economie Solidaire. This project targets the distributed monitoring of air quality with crowd-sensing solutions obtained via sensors connected to smart devices. We aim at inciting citizens to perform their own measures, and to obtain thanks to GPS geo-localization a large-scale database and a dynamic fine-grained cartography of air quality. This project takes advantage of the APISENSE® crowdsensing platform.

9.1.1.2. CIRRUS
Participants: Yahya Al-Dhuraibi, Stéphanie Challita, Guillaume Fieni, 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.

9.1.2. 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 [65] and Benoit Cornu [54].

9.1.2.2. North European Lab LLEX
Participants: Martin Monperrus, Lionel Seinturier [correspondant].

North European Lab LLEX (2015–17) is an international initiative supported by the Inria Lille - Nord Europe Center that takes place in the context of a collaboration between Inria and University College London. LLEX deals with research on automatic diagnosis and repair of software bugs. Automatic software repair is the process of fixing software bugs automatically. An automatic software repair system fixes software bugs with no human intervention. The goal of automatic software repair is to save maintenance costs and to enable systems to be more resilient to bugs and unexpected situations. This research may dramatically improve the quality of software systems. The objective of the partnership is to work on the automated diagnosis of exceptions with a focus on null pointer exceptions.

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
Participant: Pierre Bourhis [correspondant].

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. Competitivity Clusters

9.2.2.1. FUI StoreConnect
Participants: Aurélien Bourdon, Julien Duribreux, Romain Rouvoy, Lionel Seinturier [correspondant], Antoine Veuiller.

StoreConnect is a 24-month project (2016–18) funded by FUI and labelled by the PICOM (Pôle des Industries du COMmerce) competitivity 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. Programme Investissement d'Avenir (PIA)

9.2.3.1. PIA OCCIware
Participants: Stéphanie Challita, Christophe Gourdin, Romain Rouvoy, Philippe Merle [correspondant], Lionel Seinturier, Faiez Zalila.

OCCIware is a 36-month project (2014–17) of the Programme Investissement d'Avenir Cloud Computing and Big Data 4th call for projects. The partners are Smile (leader), ActiveEon SA, Scalair, Institut Mines-Télécom/Télécom SudParis, Inria, Linagora GSO, Obeo, OW2 Consortium, and Université Grenoble Alpes. The project aims at defining a formal framework for managing every digital resources in the clouds, based on *Open Cloud Computing Interface* (OCCI) recommendations from *Open Grid Forum* (OGF).

9.2.4. Inria National Initiatives

9.2.4.1. Inria IPL BetterNet
Participants: Lakhdar Meftah, Romain Rouvoy [correspondant], Romain Sommerard, Antoine Veuiller.

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.

9.2.5. Others

9.2.5.1. CNRS MASTODONS 2017 DoMaSQ'Air
Participants: Moncef Ouadia, Romain Rouvoy [correspondant].

DoMaSQ'Air is a 1-year project funded by the CNRS INS2I MASTODONS program on research in big data. This project gathers a pluridisciplinary team on the measurement and the continuous analysis of indoor and outdoor air quality. This project takes advantage of crowds of cheap and miniaturized sensors in relation with the Internet of Things and smart cities. In addition to the challenges raised by the massive amount of data generated by these cyber-physical systems, the project tackles questions related to the quality and privacy of data. DoMaSQ'Air is led by Romain Rouvoy with the participation of the PC2A laboratory on PhysicoChemistry of Combustion of the Atmosphere (CNRS/U. Lille) and the LISIC laboratory on Computer Science, Signal and Image (U. Côte d'Opale).

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].

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 a 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-years (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

Fernanda Madeiral Delfim, PhD Student from the Federal University of Uberlândia, Brazil, visited us from January to May 2017, and again in September 2017.

9.5.1.1. Internships

Mohammad Naseri, MSc. Student in Computer Science from Saarland University, Germany, is visiting us for 3 months, starting in November 2017.

Chaima Chakhava, MSc. Student in Computer Science from ESI Alger, Algeria, is visiting us for 6 months, starting in December 2017.

9.5.2. *Visits to International Teams*

9.5.2.1. Research Stays Abroad

Thomas Durieux, PhD Student, spent 4 months from September to December 2017 in KTH, Sweden.

10. Dissemination

10.1. Promoting Scientific Activities

10.1.1. *Scientific Events Selection*

10.1.1.1. Chair of Conference Program Committees

Laurence Duchien

Doctoral Symposium of European Conference on Software Architecture (ECSA)

10.1.1.2. Member of the Conference Program Committees

Laurence Duchien

International Conference on Web Engineering (ICWE)

European Conference on Software Architecture (ECSA)

International Systems and Software Product Line Conference (SPLC)

BElgian-NEtherlands software eVOLution symposium (BENEVOL)

International Conference on Software Engineering education and Training (CSEE&T)

Joint 5th ICSE International Workshop on Software Engineering for Systems-of-Systems and 11th Workshop on Distributed Software Development, Software Ecosystems and Systems-of-Systems (SESoS/WDES)

Philippe Merle

International Conference on Cooperative Information Systems (CoopIS)

International Symposium on Theoretical Aspects of Software Engineering (TASE)

International Conference on Web Engineering (ICWE)

International Conference on Cloud Computing, GRIDs, and Virtualization (CLOUD COMPUT-ING)

International Conference on Advanced Service Computing (SERVICE COMPUTATION)

Symposium on Machine Learning and Metaheuristics Techniques and Applications for Dependable Distributed Systems (MADS)

International Symposium on Security in Computing and Communications (SSCC)

International Workshop on Adaptive and Reflective Middleware (ARM)

Workshop on CrossCloud Infrastructures & Platforms (CrossCloud)

Martin Monperrus

International Conference on Software Engineering (ICSE)

Genetic Improvement Workshop (GI)

Clément Quinton

International Workshop on Dynamic Software Product Lines (DSPL)

Romain Rouvoy

    International Conference on Pervasive Computing and Communications (PerCom)

    IEEE/ACM International Conference on Mobile Software Engineering and Systems (MOBILESoft)

    International Symposium on Applied Computing (SAC), track on Dependable, Adaptive, and Trustworthy Distributed Systems (DADS)

    International Conference on Service Oriented Computing (ICSOC)

    International Conference on Principles of Distributed Systems (OPODIS)

    International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC)

    International Conference on Ambient Systems, Networks and Technologies (ANT)

    International Conference on Information Systems Development (ISD)

    International Conference on Algorithms and Architectures for Parallel Processing (ICA3PP)

    International Symposium on Computers and Communication (ISCC)

    International Workshop on Scalable Computing For Real-Time Big Data Applications (SCRAMBL)

Lionel Seinturier

    IEEE/ACM International Conference on Mobile Software Engineering and Systems (MOBILESoft)

    International Conference on Service Oriented Computing (ICSOC)

    ACM Symposium on Applied Computing (SAC), tracks Software Architecture Theory Technology and Applications (SA-TTA), Operating Systems (OS)

    Euromicro Conference on Software Engineering and Advanced Applications (SEAA), track Model-based development Components and Services (MOCS)

    International Workshop on Models@run.time @ MODELS

    International Workshop on Models@run.time @ ICAC

    Workshop on Context-Oriented Programming (COP) @ ECOOP

10.1.2. Journal

10.1.2.1. Member of the Editorial Boards

Martin Monperrus is member of the editorial board of the international journal Springer Empirical Software Engineering (IF-2016: 3.275).

Romain Rouvoy is member of the editorial board of the journal Lavoisier Technique et Science Informatiques (TSI).

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

10.1.2.2. Reviewer - Reviewing Activities

Laurence Duchien: IET Software, IEEE Software, Wiley Software: Practice and Experience (SPE).

Philippe Merle: ACM Transactions on Internet Technology (TOIT), Wiley Journal of Software Testing, Verification and Reliability (STVR), International Journal of Web Services Research (IJWSR), Springer Journal of Cloud Computing, IEEE Cloud Computing.

Martin Monperrus: IEEE Transactions on Software Engineering (TSE), ACM Transactions on Software Engineering and Methodology (TOSEM), Elsevier Journal of Systems and Software (JSS), IEEE Transactions on Reliability (TR).

Clément Quinton: Elsevier Journal of System and Software (JSS), IEEE Software, IET Software, IEEE Transactions on Services Computing (TSC).

Romain Rouvoy: Elsevier Sustainable Computing, Informatics and Systems (SUSCOM), IEEE Transactions on Software Engineering (TSE), Springer Journal of Internet Services and Applications (JISA), Wiley Journal of Software: Evolution and Process (JSEP).

Lionel Seinturier: Elsevier Journal of System and Software (JSS), ACM Transactions on Internet Technology (TOIT), Elsevier Computer Languages, Systems and Structures (COMLAN), Journal of Universal Computer Science (JUCS), Ingénierie des systèmes d'information (ISI), Elsevier Future Generation Computer Systems (FGCS).

10.1.3. Invited Talks

Philippe Merle, Walid Gaaloul and **Faiez Zalila** were invited to give a tutorial on "Seamless Lifecycle Management of Cloud Resources Using OCCI" at Services Conference Federation (SCF) 2017.

Martin Monperrus was invited to give a talk on "When and how to automatically repair bugs?" at the 4th Workshop on Design Automation for Understanding Hardware Designs (DUHDe 2017).

10.1.4. Leadership within the Scientific Community

Martin Monperrus is the co-head of the "Groupe de Travail Génie Logiciel Empirique" of the GDR GPL.

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

10.1.5. Scientific Expertise

Laurence Duchien did some scientific expertise for Direction des Relations Internationales of the French Ministry of Research.

Philippe Merle was member of the recruitment committee for assistant professors at INSA Toulouse. He is member of the Inria Scientific Board.

Lionel Seinturier was member of the recruitment committee for junior researchers at Inria Rennes. He was scientific expert for ANRT, ECOS Nord, COFECUB, Belgium FWO. He was member of the Agence Nationale de la Recherche (ANR) Scientific Evaluation Committee for Software and Network (CES25).

10.1.6. Research Administration

Laurence Duchien is member of the CNRS CoCNRS section 6 committee, and of the "bureau" of this committee. She was member of Hcéres committees for DIENS Laboratory (Paris) and LIS Laboratory (Marseille).

Philippe Merle is president of the CUMI (Comité des Utilisateurs des Moyens Informatiques), project manager for Irill at Lille, 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 leading the CIRRUS research joint team between Scalair company and Spirals Inria project-team. He was the scientific and technical leader of the OCCIware PIA funded project. He is member of the steering committee of CIEL (Conférence en IngénieriE du Logiciel).

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 heads the committee (so-called "vivier 27 rang A") that selects members of recruitment committees in Computer Science at the University of Lille 1. He is Scientific Advisor for the evaluation of ICT research laboratories at the Hcéres.

10.2. Teaching - Supervision - Juries

10.2.1. Teaching

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).

> Bases de données, 18h, Cycle Polytechnique Info553

Laurence Duchien teaches at the University of Lille 1 in the FST faculty. She heads the Carrières et Emplois service and is referent for the professional insertion in the PhD program in Computer Science at ComUE University Lille Nord de France. She is Director of Doctoral Studies for Computer Science in the Doctoral School Engineering Science (SPI) - ComUE Lille Nord de France.

> Software Project Management, 50h, Level M2, Master MIAGE
>
> Design of distributed applications, 42h, Level M1, Master of Computer Science
>
> Software Product Lines, 8h, Level M2, 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

Martin Monperrus teaches at the University of Lille 1 in the FST faculty. Until August 2017, he has headed the IAGL specialty of the Master of Computer Science at the University of Lille 1.

> Introduction to programming, 48h, Level L1, Licence of Computer Science
>
> Object-oriented design, 39h, Level L3, Licence of Computer Science
>
> Automated software engineering, 40h, Level M2 IAGL, Master of Computer Science

Clément Quinton teaches at the University of Lille 1 in the FST faculty.

> Design of distributed applications, 42h, Level M1, Master of Computer Science
>
> Software engineering, 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 Science

Romain Rouvoy teaches at the University of Lille 1 in the FST faculty. He heads the Master of Computer Science program at the University of Lille 1.

> 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 1 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 1 in the FST faculty. He heads the Computer Science Department at the Faculty of Science and Technology of the University of Lille 1.

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 in progress: Zeinab Abou Khalil, November 2017, Laurence Duchien, co-supervision with Tom Mens (University of Mons, Belgium).

PhD in progress: Guillaume Fieni, GreenData : Vers un traitement efficient et éco-responsable des grandes masses de données numériques, October 2017, Romain Rouvoy & Lionel Seinturier.

PhD in progress: 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).

In progress PhD: Sarra Habchi, Une supervision de contexte sensible à la confidentialité pour les développements logiciels en crowdsource, 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.

PhD in progress: Yahya Al Dhuraibi, Un cadre flexible pour l'élasticité dans les nuages, October 2015, Philippe Merle.

PhD in progress: Stéphanie Challita, Un cadre formel et outillé pour la gestion de toute ressource en nuage, October 2015, Philippe Merle.

PhD in progress: Thomas Durieux, Search-based Monitoring and Root Cause Diagnosis in Production, September 2015, Lionel Seinturier & Martin Monperrus.

PhD in progress: Gustavo Sousa, Towards Dynamic Software Product Lines to Optimize Management and Reconfiguration of Cloud Applications, October 2012, Laurence Duchien & Walter Rudametkin Ivey.

PhD in progress: Amal Tahri, Evolution logicielle multi-vues, des réseaux domestiques au Cloud, March 2013, Laurence Duchien.

10.2.3. Juries

Laurence Duchien

Hanae Rateau, Université Lille, chair

Sylvain Defourny, Université Lille, chair

Thi-Kim-Dung Pham, Conservatoire National des Arts et Métiers, examiner

Xhevahire Ternava, Université Côte d'Azur, chair

Julien Perolat, Université Lille, chair

Philippe Merle

Mohamed Boussaa, Université Rennes 1, reviewer

Jean-François Weber, Université de Franche-Comté, reviewer

Martin Monperrus

Bogdan Marculescu (Blekinge Institute of Technology, Sweden), member

Romain Rouvoy

      Stéphane Delbruel (University of Rennes 1), reviewer

      Yaroslav Hayduk (University of Neuchâtel, Switzerland), reviewer

      Gustavo Sousa (University of Lille 1), president

      Sabbir Hassan (University of Rennes 1), examiner

      Yunbo Li (IMT Atlantique), examiner

      Cyril Cecchinel (University Côte d'Azur), reviewer

      Elyas Ben Hadj Yahia (University of Bordeaux), reviewer

      Marcelino Rodriguez Cancio (University of Rennes 1), reviewer

Lionel Seinturier

      HDR Abderrahmane Seriai (University of Montpellier), president

      HDR Bilel Derbel (University of Lille 1), president

      HDR Marc Frincu (University of Timisoara, Romania), examiner

      Akram Kamoun (University of Sfax, Tunisia), reviewer

      André Sales Fonteles (University of Grenoble), reviewer

      Adja Ndeye Sylla (University of Grenoble), reviewer

10.3. Popularization

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.

Antoine Vastel 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 browser fingerprinting and user privacy tracking.

11. Bibliography

Major publications by the team in recent years

[1] B. BAUDRY, M. MONPERRUS. *The Multiple Facets of Software Diversity: Recent Developments in Year 2000 and Beyond*, in "ACM Computing Surveys", July 2015, pp. 1-26, https://hal.inria.fr/hal-01182103

[2] 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

[3] B. CORNU, E. T. BARR, L. SEINTURIER, M. MONPERRUS. *Casper: Automatic Tracking of Null Dereferences to Inception with Causality Traces*, in "Journal of Systems and Software", 2016, vol. 122, pp. 52-62 [*DOI :* 10.1016/J.JSS.2016.08.062], https://hal.archives-ouvertes.fr/hal-01354090

[4] G. HECHT, B. OMAR, R. ROUVOY, N. MOHA, L. DUCHIEN. *Tracking the Software Quality of Android Applications along their Evolution*, in "30th IEEE/ACM International Conference on Automated Software Engineering", Lincoln, Nebraska, United States, L. GRUNSKE, M. WHALEN (editors), Proceedings of the 30th IEEE/ACM International Conference on Automated Software Engineering (ASE 2015), IEEE, November 2015, 12 p. , https://hal.inria.fr/hal-01178734

[5] P. LAPERDRIX, W. RUDAMETKIN, B. BAUDRY. *Beauty and the Beast: Diverting modern web browsers to build unique browser fingerprints*, in "37th IEEE Symposium on Security and Privacy (S&P 2016)", San Jose, United States, May 2016, https://hal.inria.fr/hal-01285470

[6] M. MARTINEZ, M. MONPERRUS. *Mining Software Repair Models for Reasoning on the Search Space of Automated Program Fixing*, in "Empirical Software Engineering", 2015, vol. 20, n^o 1, pp. 176–205 [*DOI :* 10.1007/s10664-013-9282-8], https://hal.inria.fr/hal-00903808

[7] V. MUSCO, M. MONPERRUS, P. PREUX. *A Large-scale Study of Call Graph-based Impact Prediction using Mutation Testing*, in "Software Quality Journal", 2016 [*DOI :* 10.1007/s11219-016-9332-8], https://hal.inria.fr/hal-01346046

[8] 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

[9] 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

[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

Articles in International Peer-Reviewed Journals

[11] 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)", September 2017, 18 p. [*DOI :* 10.1109/TSC.2017.2711009], https://hal.inria.fr/hal-01529654

[12] F. ALVARES, E. RUTTEN, L. SEINTURIER. *A Domain-specific Language for The Control of Self-adaptive Component-based Architecture*, in "Journal of Systems and Software", January 2017, https://hal.archives-ouvertes.fr/hal-01450517

[13] S. BLIUDZE, S. FURIC, J. SIFAKIS, A. VIEL. *Rigorous Design of Cyber-Physical Systems: Linking Physicality and Computation*, in "Software and Systems Modeling", 2017, pp. 1-24, forthcoming [*DOI :* 10.1007/s10270-017-0642-5], https://hal.inria.fr/hal-01636392

[14] B. DANGLOT, P. PREUX, B. BAUDRY, M. MONPERRUS. *Correctness Attraction: A Study of Stability of Software Behavior Under Runtime Perturbation*, in "Empirical Software Engineering", 2017, https://arxiv.org/abs/1611.09187 [*DOI :* 10.1007/s10664-017-9571-8], https://hal.archives-ouvertes.fr/hal-01378523

[15] M. GÓMEZ, B. ADAMS, W. MAALEJ, M. MONPERRUS, R. ROUVOY. *App Store 2.0: From Crowd Information to Actionable Feedback in Mobile Ecosystems*, in "IEEE Software", March 2017, vol. 34, pp. 81-89 [*DOI :* 10.1109/MS.2017.46], https://hal.inria.fr/hal-01406422

[16] M. MARTINEZ, T. DURIEUX, R. SOMMERARD, J. XUAN, M. MONPERRUS. *Automatic Repair of Real Bugs in Java: A Large-Scale Experiment on the Defects4J Dataset*, in "Empirical Software Engineering", 2017, vol. 22, n⁰ 4, pp. 1936-1964 [*DOI : 10.1007/s10664-016-9470-4*], https://hal.archives-ouvertes.fr/hal-01387556

[17] M. MONPERRUS. *Automatic Software Repair: a Bibliography*, in "ACM Computing Surveys", 2017, https://hal.archives-ouvertes.fr/hal-01206501

[18] V. MUSCO, M. MONPERRUS, P. PREUX. *A Large-scale Study of Call Graph-based Impact Prediction using Mutation Testing*, in "Software Quality Journal", September 2017, vol. 25, n⁰ 3, pp. 921–950 [*DOI : 10.1007/s11219-016-9332-8*], https://hal.inria.fr/hal-01346046

[19] 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", 2017, vol. 43, n⁰ 1, pp. 34-55 [*DOI : 10.1109/TSE.2016.2560811*], https://hal.archives-ouvertes.fr/hal-01285008

International Conferences with Proceedings

[20] Y. AL-DHURAIBI, F. PARAISO, N. DJARALLAH, P. MERLE. *Autonomic Vertical Elasticity of Docker Containers with ElasticDocker*, in "10th IEEE International Conference on Cloud Computing, IEEE CLOUD 2017", Honolulu, Hawaii, United States, Proceedings of the 10th IEEE International Conference on Cloud Computing, IEEE CLOUD 2017, June 2017, pp. 472 - 479, https://hal.archives-ouvertes.fr/hal-01522940

[21] F. ALVARES, G. DELAVAL, E. RUTTEN, L. SEINTURIER. *Language Support for Modular Autonomic Managers in Reconfigurable Software Components*, in "SeAC 2017 - 2nd Workshop on Self-Aware Computing @ ICAC'17", Colombus, United States, July 2017, pp. 271 - 278 [*DOI : 10.1109/ICAC.2017.48*], https://hal.inria.fr/hal-01575353

[22] A. CARETTE, M. A. AIT YOUNES, G. HECHT, N. MOHA, R. ROUVOY. *Investigating the Energy Impact of Android Smells*, in "24th International IEEE Conference on Software Analysis, Evolution and Reengineering (SANER)", Klagenfurt, Austria, A. MARCUS, G. BAVOTA (editors), Proceedings of the 24th International IEEE Conference on Software Analysis, Evolution and Reengineering (SANER), IEEE, February 2017, 10 p. , https://hal.inria.fr/hal-01403485

[23] S. CHALLITA, F. PARAISO, P. MERLE. *A Study of Virtual Machine Placement Optimization in Data Centers*, in "7th International Conference on Cloud Computing and Services Science, CLOSER 2017", Porto, Portugal, SCITEPRESS (editor), INSTICC, April 2017, pp. 343–350 [*DOI : 10.5220/0006236503430350*], https://hal.inria.fr/hal-01481631

[24] S. CHALLITA, F. PARAISO, P. MERLE. *Towards Formal-based Semantic Interoperability in Multi-Clouds: The FCLOUDS Framework*, in "10th IEEE International Conference on Cloud Computing (CLOUD)", Honolulu, Hawaii, United States, June 2017, pp. 710 - 713, https://hal.inria.fr/hal-01519831

[25] M. COLMANT, P. FELBER, R. ROUVOY, L. SEINTURIER. *WattsKit: Software-Defined Power Monitoring of Distributed Systems*, in "17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)", Madrid, Spain, F. CAPELLO, G. FOX, J. GARCIA-BLAS (editors), Proceedings of the 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid), IEEE, May 2017, 10 p. , https://hal.inria.fr/hal-01439889

[26] M. A. GONZÁLEZ, W. RUDAMETKIN, M. MONPERRUS, R. ROUVOY. *Challenging Anti-fragile Blockchain Applications*, in "11th EuroSys Doctoral Workshop EuroDW'17", Belgrade, Serbia, April 2017, https://hal.inria.fr/hal-01653986

[27] S. HABCHI, G. HECHT, R. ROUVOY, N. MOHA. *Code Smells in iOS Apps: How do they compare to Android?*, in "MOBILESoft'17 - 4th IEEE/ACM International Conference on Mobile Software Engineering and Systems", Buenos Aires, Argentina, Proceedings of the 4th IEEE/ACM International Conference on Mobile Software Engineering and Systems, May 2017, https://hal.inria.fr/hal-01471294

[28] T. HARTMANN, F. FOUQUET, M. JIMENEZ, R. ROUVOY, Y. LE TRAON. *Analyzing Complex Data in Motion at Scale with Temporal Graphs*, in "The 29th International Conference on Software Engineering & Knowledge Engineering (SEKE'17)", Pittsburgh, United States, X. HE, O. PEREIRA, A. PERKUSICH (editors), Proceedings of the 29th International Conference on Software Engineering & Knowledge Engineering (SEKE'17), KSI Research, July 2017, 6 p. , https://hal.inria.fr/hal-01511636

[29] A. HAVET, R. PIRES, P. FELBER, M. PASIN, R. ROUVOY, V. SCHIAVONI. *SecureStreams: A Reactive Middleware Framework for Secure Data Stream Processing*, in "DEBS'17 - The 11th ACM International Conference on Distributed and Event-Based Systems", Barcelona, Spain, M. CHANDY, B. KOLDEHOFE (editors), Proceedings of the 11th ACM International Conference on Distributed and Event-Based Systems, ACM, June 2017, pp. 124-133 [*DOI :* 10.1145/3093742.3093927], https://hal.inria.fr/hal-01510699

[30] A. HAVET, V. SCHIAVONI, P. FELBER, M. COLMANT, R. ROUVOY, C. FETZER. *GenPack: A Generational Scheduler for Cloud Data Centers*, in "5th IEEE International Conference on Cloud Engineering (IC2E)", Vancouver, Canada, I. GUPTA, J. LIU (editors), Proceedings of the 5th IEEE International Conference on Cloud Engineering (IC2E), IEEE, April 2017, 10 p. , https://hal.inria.fr/hal-01403486

[31] F. KORTE, S. CHALLITA, F. ZALILA, 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, https://hal.inria.fr/hal-01678945

[32] A. KOYUNCU, T. F. BISSYANDÉ, D. KIM, J. KLEIN, M. MONPERRUS, Y. LE TRAON. *Impact of Tool Support in Patch Construction*, in "26th ACM SIGSOFT International Symposium on Software Testing and Analysis", Santa Barbara, United States, ISSTA 2017, ACM, July 2017, pp. 237–248 [*DOI :* 10.1145/3092703.3092713], https://hal.archives-ouvertes.fr/hal-01575214

[33] M. MARTINEZ, T. DURIEUX, R. SOMMERARD, J. XUAN, M. MONPERRUS. *Automatic repair of real bugs in java: a large-scale experiment on the defects4j dataset*, in "ESEC/FSE 2017 - 11th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering", PADERBORN, Germany, September 2017 [*DOI :* 10.1007/S10664-016-9470-4], https://hal.inria.fr/hal-01574671

[34] L. MEFTAH, M. GOMEZ, R. ROUVOY, I. CHRISMENT. *AndroFleet: Testing WiFi Peer-to-Peer Mobile Apps in the Large*, in "ASE 2017 - 32nd IEEE/ACM International Conference on Automated Software Engineering", Urbana-Champaign, Illinois, United States, ASE 2017 - The 32nd IEEE/ACM International Conference on Automated Software Engineering - Tool demonstration, October 2017, https://hal.inria.fr/hal-01574466

[35] *Best Paper*
P. MERLE, C. GOURDIN, N. MITTON. *Mobile Cloud Robotics as a Service with OCCIware*, in "2nd IEEE International Congress on Internet of Things, IEEE ICIOT 2017", Honolulu, Hawaii, United States, Proceedings of the 2nd IEEE International Congress on Internet of Things, IEEE ICIOT 2017, June 2017, pp. 50 - 57, Best Paper Award [*DOI :* 10.1109/IEEE.ICIOT.2017.15], https://hal.archives-ouvertes.fr/hal-01522684.

[36] M. MONPERRUS. *Principles of Antifragile Software*, in "Proceedings of the Salon des Refusés 2017", Brussels, Belgium, 2017, https://arxiv.org/abs/1404.3056 [*DOI :* 10.1145/3079368.3079412], https://hal.archives-ouvertes.fr/hal-00978310

[37] F. PETRILLO, P. MERLE, N. MOHA, Y.-G. GUÉHÉNEUC. *Towards a REST Cloud Computing Lexicon*, in "7th International Conference on Cloud Computing and Services Science, CLOSER 2017", Porto, Portugal, SCITEPRESS (editor), INSTICC, April 2017, pp. 376–383 [*DOI :* 10.5220/0006281203760383], https://hal.archives-ouvertes.fr/hal-01480593

[38] G. SOUSA, W. RUDAMETKIN, L. DUCHIEN. *Extending Dynamic Software Product Lines with Temporal Constraints*, in "12th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS 2017)", Buenos Aires, Argentina, May 2017, https://hal.inria.fr/hal-01482014

[39] 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. 1-14, https://hal.inria.fr/hal-01652021

[40] F. ZALILA, S. CHALLITA, P. MERLE. *A Model-Driven Tool Chain for OCCI*, in "25th International Conference on Cooperative Information Systems (CoopIS)", Rhodes, Greece, October 2017, https://hal.inria.fr/hal-01583165

[41] B. ZHANG, Y. AL-DHURAIBI, R. ROUVOY, F. PARAISO, L. SEINTURIER. *CloudGC: Recycling Idle Virtual Machines in the Cloud*, in "5th IEEE International Conference on Cloud Engineering (IC2E)", Vancouver, Canada, I. GUPTA, J. LIU (editors), Proceedings of the 5th IEEE International Conference on Cloud Engineering (IC2E), IEEE, April 2017, 10 p. , https://hal.inria.fr/hal-01403488

[42] B. ZHANG, F. KRIKAVA, R. ROUVOY, L. SEINTURIER. *Hadoop-Benchmark: Rapid Prototyping and Evaluation of Self-Adaptive Behaviors in Hadoop Clusters*, in "12th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS'17)", Buenos Aires, Argentina, J. CAMARA, B. NUSEIBEH, D. GARLAN (editors), Proceedings of the 12th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS), May 2017, 6 p. , https://hal.inria.fr/hal-01475635

Conferences without Proceedings

[43] C. DORFFER, M. PUIGT, G. DELMAIRE, G. ROUSSEL, R. ROUVOY, I. SAGNIER. *Matrix Factorisation-based Calibration For Air Quality Crowd-sensing*, in "European Geosciences Union General Assembly 2017", Vienna, Austria, April 2017, vol. 19, pp. EGU2017-16389, https://hal.archives-ouvertes.fr/hal-01495581

[44] T. DURIEUX, B. CORNU, L. SEINTURIER, M. MONPERRUS. *Dynamic Patch Generation for Null Pointer Exceptions Using Metaprogramming*, in "IEEE International Conference on Software Analysis, Evolution and Reengineering", Klagenfurt, Austria, February 2017, pp. 349-358 [*DOI :* 10.1109/SANER.2017.7884635], https://hal.archives-ouvertes.fr/hal-01419861

[45] T. DURIEUX, Y. HAMADI, M. MONPERRUS. *Production-Driven Patch Generation*, in "Proceedings of the 39th International Conference on Software Engineering: New Ideas and Emerging Results Track", Buenos Aires, Argentina, 2017, pp. 23-26 [*DOI :* 10.1109/ICSE-NIER.2017.8], https://hal.archives-ouvertes.fr/hal-01463689

Scientific Books (or Scientific Book chapters)

[46] F. KRIKAVA, P. COLLET, R. ROUVOY, L. SEINTURIER. *Contracts-based Control Integration into Software Systems*, in "Software Engineering for Self-Adaptive Systems 3: Assurances", R. DE LEMOS, D. GARLAN, C. GHEZZI, H. GIESE (editors), LNCS, Springer, May 2017, vol. 9640, https://hal.inria.fr/hal-01257205

Research Reports

[47] T. DURIEUX, B. DANGLOT, Z. YU, M. MARTINEZ, S. URLI, M. MONPERRUS. *The Patches of the Nopol Automatic Repair System on the Bugs of Defects4J version 1.1.0*, Université Lille 1 - Sciences et Technologies, 2017, n^O hal-01480084, https://hal.archives-ouvertes.fr/hal-01480084

Other Publications

[48] T. DURIEUX, Y. HAMADI, Z. YU, M. MONPERRUS. *Exhaustive Exploration of the Failure-oblivious Computing Search Space*, October 2017, https://arxiv.org/abs/1710.09722 - arXiv admin note: substantial text overlap with arXiv:1603.07631, https://hal.inria.fr/hal-01624988

[49] M. MONPERRUS, B. DANGLOT, O. VERA-PEREZ, Z. YU, B. BAUDRY. *The Emerging Field of Test Amplification: A Survey*, November 2017, https://arxiv.org/abs/1705.10692 - working paper or preprint, https://hal.archives-ouvertes.fr/hal-01634288

[50] M. WHITE, M. TUFANO, M. MARTINEZ, M. MONPERRUS, D. POSHYVANYK. *Sorting and Transforming Program Repair Ingredients via Deep Learning Code Similarities*, September 2017, https://arxiv.org/abs/1707.04742 - working paper or preprint, https://hal.archives-ouvertes.fr/hal-01581170

[51] Z. YU, M. MARTINEZ, B. DANGLOT, T. DURIEUX, M. MONPERRUS. *Test Case Generation for Program Repair: A Study of Feasibility and Effectiveness*, August 2017, https://arxiv.org/abs/1703.00198 - working paper, https://hal.archives-ouvertes.fr/hal-01575219

[52] B. ZHANG, F. KŘIKAVA, R. ROUVOY, L. SEINTURIER. *Hadoop-Benchmark: Rapid Prototyping and Evaluation of Self-Adaptive Behaviors in Hadoop Clusters (Artifact)*, May 2017, 3 p. , Artifact [*DOI :* 10.4230/DARTS.3.1.1], https://hal.inria.fr/hal-01525304

References in notes

[53] L. CAPRA, W. EMMERICH, C. MASCOLO. *A micro-economic approach to conflict resolution in mobile computing*, in "SIGSOFT FSE", 2002, pp. 31-40

[54] B. CORNU. *Automatic Analysis and Repair of Exception Bugs for Java Programs*, Université de Lille, November 2015, https://tel.archives-ouvertes.fr/tel-01250092

[55] 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

[56] 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

[57] 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

[58] M. GOMEZ. *Towards Improving the Quality of Mobile Apps by Leveraging Crowdsourced Feedback*, Universite Lille 1 ; Inria Lille - Nord Europe, December 2016, https://tel.archives-ouvertes.fr/tel-01418298

[59] 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, n⁰ 1, pp. 10–18

[60] M. HARMAN. *Automated patching techniques: the fix is in*, in "Communications of the ACM", May 2010, vol. 53, n⁰ 5, pp. 108–108, http://doi.acm.org/10.1145/1735223.1735248

[61] 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, n⁰ 1, pp. 23-45, http://dx.doi.org/10.1007/s13174-011-0021-3

[62] D. JACKSON. *Alloy: A Lightweight Object Modelling Notation*, in "ACM Transactions on Software Engineering and Methodology", April 2002, vol. 11, n⁰ 2, pp. 256-290

[63] W. LI, S. HU, J. LI, H.-A. JACOBSEN. *Community Clustering for Distributed Publish/Subscribe Systems*, in "CLUSTER", IEEE, 2012, pp. 81-89

[64] M. LÉGER, T. LEDOUX, T. COUPAYE. *Reliable Dynamic Reconfigurations in a Reflective Component Model*, in "CBSE", L. GRUNSKE, R. REUSSNER, F. PLASIL (editors), Lecture Notes in Computer Science, Springer, 2010, vol. 6092, pp. 74-92

[65] M. MARTINEZ. *Extraction and Analysis of Knowledge for Automatic Software Repair*, Université Lille 1, October 2014, https://hal.archives-ouvertes.fr/tel-01078911

[66] P. MERLE, J.-B. STEFANI. *A formal specification of the Fractal component model in Alloy*, Inria, November 2008, n⁰ 6721, 44 p. , http://hal.inria.fr/inria-00338987

[67] 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 Comput-

ing", May 2016, vol. 98, n⁰ 5, pp. 539-565 [*DOI :* 10.1007/S00607-014-0421-X], https://hal.inria.fr/hal-01019385

[68] 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

[69] 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

[70] M. SALEHIE, L. TAHVILDARI. *Self-adaptive software: Landscape and research challenges*, in "ACM Transactions on Autonomous and Adaptive Systems", May 2009, vol. 4, n⁰ 2, pp. 14:1–14:42, http://doi.acm.org/10.1145/1516533.1516538

[71] 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

[72] 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

[73] 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, n⁰ 5, pp. 559-583 [*DOI :* 10.1002/SPE.1077], http://hal.inria.fr/inria-00567442

[74] 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

[75] 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

[76] 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.

[77] T. XIE, S. THUMMALAPENTA, D. LO, C. LIU. *Data Mining for Software Engineering*, in "IEEE Computer", August 2009, vol. 42, n⁰ 8, pp. 55-62, http://doi.ieeecomputersociety.org/10.1109/MC.2009.256

[78] 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

[79] 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