

The Inria logo is written in a red, elegant cursive script.

IN PARTNERSHIP WITH:
CNRS

**Institut national des sciences
appliquées de Rennes**

Université Rennes 1

**École normale supérieure de
Rennes**

Activity Report 2019

Project-Team Myriads

Design and Implementation of Autonomous Distributed Systems

IN COLLABORATION WITH: Institut de recherche en informatique et systèmes aléatoires (IRISA)

RESEARCH CENTER
Rennes - Bretagne-Atlantique

THEME
Distributed Systems and middleware

Table of contents

1. Team, Visitors, External Collaborators	2
2. Overall Objectives	3
2.1. General Objectives	3
2.2. Context	4
2.3. Challenges	4
3. Research Program	5
3.1. Introduction	5
3.2. Scaling fogs and clouds	5
3.2.1. Resource management in hierarchical clouds	5
3.2.2. Resource management in fog computing architectures	5
3.2.3. Self-optimizing applications in multi-cloud environments	6
3.3. Greening clouds	6
3.3.1. Smart grids and clouds	6
3.3.2. Energy cost models	7
3.3.3. Energy-aware users	7
3.4. Securing clouds	8
3.4.1. Security monitoring SLO	8
3.4.2. Data protection in Cloud-based IoT services	8
3.5. Experimenting with Clouds	9
3.5.1. Experimentation methodologies for clouds	9
3.5.2. Use cases	9
4. Application Domains	9
5. Highlights of the Year	10
6. New Software and Platforms	10
6.1. PaaSAdapter	10
6.2. SAIDS	10
6.3. SimGrid	11
6.4. DiFFuSE	11
6.5. GinFlow	12
6.6. libcvss	12
7. New Results	12
7.1. Scaling Clouds	12
7.1.1. Efficient Docker container deployment in fog environments	12
7.1.2. Fog computing platform design	13
7.1.3. Edgification of micro-service applications	13
7.1.4. Community Clouds	13
7.1.5. Geo-distributed data stream processing	14
7.1.6. QoS-aware and energy-efficient resource management for Function-as-a-Service	14
7.2. Greening Clouds	14
7.2.1. Energy Models	14
7.2.2. End-to-end energy models for the Internet of Things	15
7.2.3. Exploiting renewable energy in distributed clouds	15
7.2.4. Smart Grids	16
7.3. Securing Clouds	16
7.3.1. Security monitoring in Cloud computing platforms	16
7.3.2. Privacy monitoring in Fog computing platforms	17
7.4. Experimenting with Clouds	17
7.4.1. Simulating distributed IT systems	17
7.4.2. Formal methods for IT systems	17

7.4.3.	Executing epidemic simulation applications in the Cloud	18
7.4.4.	Tools for experimentation	18
8.	Partnerships and Cooperations	18
8.1.	Regional Initiatives	18
8.2.	National Initiatives	19
8.2.1.	ADEME RennesGrid (2017-2020)	19
8.2.2.	Inria ADT Mc SimGrid (2019-2021)	19
8.2.3.	Inria IPL Discovery (2015-2019)	19
8.2.4.	Inria IPL Hac Specis (2016-2020)	19
8.2.5.	SESAME ASTRID project (2016-2019)	20
8.2.6.	CNRS GDS EcoInfo	20
8.3.	European Initiatives	20
8.3.1.	H2020 Projects	20
8.3.2.	Collaborations in European Programs, Except FP7 & H2020	21
8.3.3.	Inria Associate Teams Not Involved in an Inria International Labs	21
8.3.3.1.	FogCity	21
8.3.3.2.	FogRein	22
8.3.4.	Inria International Partners	22
8.3.5.	Participation in Other International Programs	22
8.4.	International Research Visitors	23
9.	Dissemination	23
9.1.	Promoting Scientific Activities	23
9.1.1.	Member of the Organizing Committees	23
9.1.2.	Scientific Events: Selection	23
9.1.2.1.	Chair of Conference Program Committees	23
9.1.2.2.	Member of the Conference Program Committees	23
9.1.3.	Journal	24
9.1.3.1.	Member of the Editorial Boards	24
9.1.3.2.	Reviewer - Reviewing Activities	24
9.1.4.	Invited Talks	24
9.1.5.	Leadership within the Scientific Community	24
9.1.6.	Scientific Expertise	24
9.1.7.	Research Administration	25
9.2.	Teaching - Supervision - Juries	25
9.2.1.	Teaching	25
9.2.2.	Supervision	26
9.2.3.	Juries	27
9.3.	Popularization	27
9.3.1.	Internal or external Inria responsibilities	27
9.3.2.	Education	27
9.3.3.	Internal action	27
9.3.4.	Creation of media or tools for science outreach	28
10.	Bibliography	28

Project-Team Myriads

Creation of the Team: 2010 January 01, updated into Project-Team: 2012 January 01

Keywords:

Computer Science and Digital Science:

- A1.1.9. - Fault tolerant systems
- A1.1.13. - Virtualization
- A1.2. - Networks
- A1.2.4. - QoS, performance evaluation
- A1.2.5. - Internet of things
- A1.3. - Distributed Systems
- A1.3.2. - Mobile distributed systems
- A1.3.4. - Peer to peer
- A1.3.5. - Cloud
- A1.3.6. - Fog, Edge
- A1.6. - Green Computing
- A2.1.7. - Distributed programming
- A2.2.5. - Run-time systems
- A2.3.2. - Cyber-physical systems
- A2.4.2. - Model-checking
- A2.6. - Infrastructure software
- A2.6.1. - Operating systems
- A2.6.2. - Middleware
- A2.6.3. - Virtual machines
- A2.6.4. - Ressource management
- A3.1.2. - Data management, quering and storage
- A3.1.3. - Distributed data
- A4.9. - Security supervision
- A4.9.1. - Intrusion detection
- A4.9.3. - Reaction to attacks
- A5.6. - Virtual reality, augmented reality
- A6.1.3. - Discrete Modeling (multi-agent, people centered)
- A7.1. - Algorithms
- A8.2. - Optimization

Other Research Topics and Application Domains:

- B2.3. - Epidemiology
- B3.1. - Sustainable development
- B3.2. - Climate and meteorology
- B4.3. - Renewable energy production
- B4.4. - Energy delivery
- B4.4.1. - Smart grids
- B4.5. - Energy consumption

- B4.5.1. - Green computing
- B5.1. - Factory of the future
- B5.8. - Learning and training
- B6.1. - Software industry
 - B6.1.1. - Software engineering
- B6.3. - Network functions
 - B6.3.3. - Network Management
- B6.4. - Internet of things
- B6.5. - Information systems
- B6.6. - Embedded systems
- B8.1. - Smart building/home
- B8.2. - Connected city
- B8.3. - Urbanism and urban planning
- B8.5. - Smart society
- B9.1. - Education
 - B9.1.1. - E-learning, MOOC
 - B9.1.2. - Serious games
- B9.5.1. - Computer science
- B9.7. - Knowledge dissemination
 - B9.7.1. - Open access
 - B9.7.2. - Open data
- B9.8. - Reproducibility
- B9.9. - Ethics
- B9.10. - Privacy

1. Team, Visitors, External Collaborators

Research Scientist

Anne-Cécile Orgerie [CNRS, Researcher]

Faculty Members

Guillaume Pierre [Team leader, Univ de Rennes I, Professor]
Martin Quinson [Ecole normale supérieure de Rennes, Professor, HDR]
Cédric Tedeschi [Univ de Rennes I, Associate Professor, HDR]
Nikos Parlavantzas [INSA Rennes, Associate Professor]
Jean-Louis Pazat [INSA Rennes, Professor, HDR]
Marin Bertier [INSA Rennes, Associate Professor]

Post-Doctoral Fellows

Kameswar Rao Vaddina [Inria, since Nov 2019]
Benjamin Camus [CNRS, until Sep 2019]
Aymen Jlassi [Univ de Rennes I]
Millian Poquet [Inria, until Jun 2019]

PhD Students

Adrien Gougeon [Ecole Normale Supérieure Rennes, since Sep [E19]]
Ali Fahs [Univ de Rennes I]
Amir Teshome Wonjiga [Inria, until Feb 2019]

Arif Ahmed [Univ de Rennes I]
Clement El Baz [Inria]
Davaadorj Battulga [U Hopper]
Dorra Boughzala [Inria, since Jul 2019]
Genc Tato [Inria, until Feb 2019]
Hamidreza Arkian [Univ de Rennes I]
Loic Guegan [Ecole normale supérieure de Rennes]
Mehdi Belkhiria [Univ de Rennes I]
Mozhdeh Farhadi [U Hopper]
Mulugeta Tamiru [Elastisys]
Paulo Rodrigues de Souza Junior [Univ de Rennes I]
The Anh Pham [Inria]
Yasmina Bouizem [Rennes Métropole]

Technical staff

Ehsan Azimi [Inria, since Dec 2019]
Sébastien Philippot [Inria]

Interns and Apprentices

Lorenzo Civolani [Inria, until Feb 2019]
Adrien Gougeon [Inria, from Feb 2019 until Jun 2019]
Romain Olivo [CNRS, from Jun 2019 until Aug 2019]
Joseph Sayegh [Univ de Rennes I, since Jun 2019 until Aug 2019]
Archana Wale [Univ de Rennes I, until Jul 2019]

Administrative Assistant

Caroline Tanguy [Univ de Rennes I, since Dec 2019]

Visiting Scientists

Gene Cooperman [Khoury College of Computer, May 2019]
Ayan Mondal [Indian Institute of Technology Kharagpur, since Jun 2019 until Sep 2019]
Deborah Agarwal [Berkeley Lab]

External Collaborators

Anne Blavette [CNRS]
Julie Montégu [Univ de Rennes I]
Louis Rilling [DGA]

2. Overall Objectives

2.1. General Objectives

MYRIADS is a joint team with INRIA, CNRS, UNIVERSITY RENNES 1, INSA RENNES and ENS RENNES. It is part of IRISA (DI department on large scale systems) and INRIA RENNES – BRETAGNE ATLANTIQUE.

The objective of MYRIADS is to design and implement systems for autonomous service and resource management in interconnected and distributed clouds. The team tackles the challenges of dependable application execution and efficient resource management in highly distributed clouds.

2.2. Context

The MYRIADS team research activities are conducted in the context of the future of Internet.

Internet of Services. Myriads of applications are provided to more than one billion users¹ all over the world. Over time, these applications are becoming more and more sophisticated, a given application being a composition of services likely to be executed on various sites located in different geographical locations. The Internet of Services is spreading all domains: home, administration, business, industry and science. Everyone is involved in the Internet of Services: citizens, enterprises, scientists are application, service and resource consumers and/or providers over the Internet.

Outsourcing. Software is provided as a service over the Internet. Myriads of applications are available on-line to billions of users as, for instance, *GoogleApps* (Gmail). After decades in which companies used to host their entire IT infrastructures in-house, a major shift is occurring where these infrastructures are outsourced to external operators such as Data Centers and Computing Clouds. In the Internet of Services, not only software but also infrastructure are delivered as a service. Clouds turned computing and storage into a utility. Just like water or electricity, they are available in virtually infinite amounts and their consumption can be adapted within seconds like opening or closing a water tap. The main transition, however, is the change in business models. Companies or scientists do not need to buy and operate their own data centers anymore. Instead, the compute and storage resources are offered by companies on a “pay-as-you-go” basis. There is no more need for large hardware investments before starting a business. Even more, the new model allows users to adapt their resources within minutes, e.g., scale up to handle peak loads or rent large numbers of computers for a short experiment. The risk of wasting money by either under-utilization or undersized data centers is shifted from the user to the provider.

Sharing and Cooperation. Sharing information and cooperating over the Internet are also important user needs both in the private and the professional spheres. This is exemplified by various services that have been developed in the last decade. Peer-to-peer networks are extensively used by citizens in order to share musics and movies. A service like *Flickr* allowing individuals to share pictures is also very popular. Social networks such as *FaceBook* or *LinkedIn* link millions of users who share various kinds of information within communities. Virtual organizations tightly connected to Grids allow scientists to share computing resources aggregated from different institutions (universities, computing centers...). The EGEE European Grid is an example of production Grid shared by thousands of scientists all over Europe.

2.3. Challenges

The term cloud was coined over 12 years ago. Today cloud computing is widely adopted for a wide range of usage: information systems outsourcing, web service hosting, scientific computing, data analytics, back-end of mobile and IoT applications. There is a wide variety of cloud service providers (IaaS, PaaS, SaaS) resulting in difficulties for customers to select the services fitting their needs. Production clouds are powered by huge data centers that customers reach through the Internet. This current model raises a number of issues. Cloud computing generates a lot of traffic resulting in ISP providers needing to increase the network capacity. An increasing amount of always larger data centers consumes a lot of energy. Cloud customers experience poor quality of experience for highly interactive mobile applications as their requests are dealt with in data centers that are several hops away. The centralization of data in clouds also raises (i) security issues as clouds are a target of choice for attackers and (ii) privacy issues with data aggregation.

Recently new cloud architectures have been proposed to overcome the scalability, latency, and energy issues of traditional centralized data centers. Various flavors of distributed cloud computing are emerging depending on the resources exploited: resources in the core network (distributed cloud), resources at the edge of the network (edge clouds) and even resources in the people swarms of devices (fog computing) enabling scalable cloud computing. These distributed clouds raise new challenges for resource and application management.

¹According to World Stats, there are 3.67 billion Internet users i.e. more than half of the total world population in June 2016 <http://www.internetworldstats.com/stats.htm>.

The ultimate goal of Myriads team is making highly distributed clouds sustainable. By sustainability we mean green, efficient and secure clouds. We plan to study highly distributed clouds including edge clouds and fog computing. In this context, we will investigate novel techniques for greening clouds including the optimization of energy consumption in distributed clouds in the context of smart grids. As more and more critical information system are outsourced in the cloud and personal data captured by sensors embedded in smart objects and smartphones are stored in the cloud, we will investigate security and privacy issues in two directions: cloud security monitoring and personal data protection in cloud-based IoT applications.

System research requires experimental validation based on simulation and/or prototyping. Reproducible experimentation is essential. We will contribute to the design and implementation of simulators well suited to the study of distributed clouds (architecture, energy consumption) and of large scale experimentation platforms for distributed systems enabling reproducible experiments.

3. Research Program

3.1. Introduction

In this section, we present our research challenges along four work directions: resource and application management in distributed cloud and fog computing architectures for scaling clouds in Section 3.2, energy management strategies for greening clouds in Section 3.3, security and data protection aspects for securing cloud-based information systems and applications in Section 3.4, and methods for experimenting with clouds in Section 3.5.

3.2. Scaling fogs and clouds

3.2.1. Resource management in hierarchical clouds

The next generation of utility computing appears to be an evolution from highly centralized clouds towards more decentralized platforms. Today, cloud computing platforms mostly rely on large data centers servicing a multitude of clients from the edge of the Internet. Servicing cloud clients in this manner suggests that locality patterns are ignored: wherever the client issues his/her request from, the request will have to go through the backbone of the Internet provider to the other side of the network where the data center relies. Besides this extra network traffic and this latency overhead that could be avoided, other common centralization drawbacks in this context stand in limitations in terms of security/legal issues and resilience.

At the same time, it appears that network backbones are over-provisioned for most of their usage. This advocates for placing computing resources directly within the backbone network. The general challenge of resource management for such clouds stands in trying to be locality-aware: for the needs of an application, several virtual machines may exchange data. Placing them *close* to each others can significantly improve the performance of the application they compose. More generally, building an overlay network which takes the hierarchical aspects of the platform without being a hierarchical overlay – which comes with load balancing and resilience issues is a challenge by itself.

We expect to integrate the results of these works in the Discovery initiative [33] which aims at revisiting OpenStack to offer a cloud stack able to manage utility computing platforms where computing resources are located in small computing centers in the backbone's PoPs (Point of Presence) and interconnected through the backbone's internal links.

3.2.2. Resource management in fog computing architectures

Fog computing infrastructures are composed of compute, storage and networking resources located at the edge of wide-area networks, in immediate proximity to the end users. Instead of treating the mobile operator's network as a high-latency dumb pipe between the end users and the external service providers, fog platforms aim at deploying cloud functionalities *within* the mobile phone network, inside or close to the mobile access points. Doing so is expected to deliver added value to the content providers and the end users by enabling

new types of applications ranging from Internet-of-Things applications to extremely interactive systems (e.g., augmented reality). Simultaneously, it will generate extra revenue streams for the mobile network operators, by allowing them to position themselves as cloud computing operators and to rent their already-deployed infrastructure to content and application providers.

Fog computing platforms have very different geographical distribution compared to traditional clouds. While traditional clouds are composed of many reliable and powerful machines located in a very small number of data centers and interconnected by very high-speed networks, mobile edge cloud are composed of a very large number of points-of-presence with a couple of weak and potentially unreliable servers, interconnected with each other by commodity long-distance networks. This creates new demands for the organization of a scalable mobile edge computing infrastructure, and opens new directions for research.

The main challenges that we plan to address are:

- How should an edge cloud infrastructure be designed such that it remains scalable, fault-tolerant, controllable, energy-efficient, etc.?
- How should applications making use of edge clouds be organized? One promising direction is to explore the extent to which stream-data processing platforms such as Apache Spark and Apache Flink can be adapted to become one of the main application programming paradigms in such environments.

3.2.3. *Self-optimizing applications in multi-cloud environments*

As the use of cloud computing becomes pervasive, the ability to deploy an application on a multi-cloud infrastructure becomes increasingly important. Potential benefits include avoiding dependence on a single vendor, taking advantage of lower resource prices or resource proximity, and enhancing application availability. Supporting multi-cloud application management involves two tasks. First, it involves selecting an initial multi-cloud application deployment that best satisfies application objectives and optimizes performance and cost. Second, it involves dynamically adapting the application deployment in order to react to changes in execution conditions, application objectives, cloud provider offerings, or resource prices. Handling price changes in particular is becoming increasingly complex. The reason is the growing trend of providers offering sophisticated, dynamic pricing models that allow buying and selling resources of finer granularities for shorter time durations with varying prices.

Although multi-cloud platforms are starting to emerge, these platforms impose a considerable amount of effort on developers and operations engineers, provide no support for dynamic pricing, and lack the responsiveness and scalability necessary for handling highly-distributed, dynamic applications with strict quality requirements. The goal of this work is to develop techniques and mechanisms for automating application management, enabling applications to cope with and take advantage of the dynamic, diverse, multi-cloud environment in which they operate.

The main challenges arising in this context are:

- selecting effective decision-making approaches for application adaptation,
- supporting scalable monitoring and adaptation across multiple clouds,
- performing adaptation actions in a cost-efficient and safe manner.

3.3. Greening clouds

The ICT (Information and Communications Technologies) ecosystem now approaches 5% of world electricity consumption and this ICT energy use will continue to grow fast because of the information appetite of Big Data, large networks and large infrastructures as Clouds that unavoidably leads to large power.

3.3.1. *Smart grids and clouds*

We propose exploiting Smart Grid technologies to come to the rescue of energy-hungry Clouds. Unlike in traditional electrical distribution networks, where power can only be moved and scheduled in very limited ways, Smart Grids dynamically and effectively adapt supply to demand and limit electricity losses (currently 10% of produced energy is lost during transmission and distribution).

For instance, when a user submits a Cloud request (such as a Google search for instance), it is routed to a data center that processes it, computes the answer and sends it back to the user. Google owns several data centers spread across the world and for performance reasons, the center answering the user's request is more likely to be the one closest to the user. However, this data center may be less energy efficient. This request may have consumed less energy, or a different kind of energy (renewable or not), if it had been sent to this further data center. In this case, the response time would have been increased but maybe not noticeably: a different trade-off between quality of service (QoS) and energy-efficiency could have been adopted.

While Clouds come naturally to the rescue of Smart Grids for dealing with this big data issue, little attention has been paid to the benefits that Smart Grids could bring to distributed Clouds. To our knowledge, no previous work has exploited the Smart Grids potential to obtain and control the energy consumption of entire Cloud infrastructures from underlying facilities such as air conditioning equipment (which accounts for 30% to 50% of a data center's electricity bill) to network resources (which are often operated by several actors) and to computing resources (with their heterogeneity and distribution across multiple data centers). We aim at taking advantage of the opportunity brought by the Smart Grids to exploit renewable energy availability and to optimize energy management in distributed Clouds.

3.3.2. Energy cost models

Cloud computing allows users to outsource the computer resources required for their applications instead of using a local installation. It offers on-demand access to the resources through the Internet with a pay-as-you-go pricing model. However, this model hides the electricity cost of running these infrastructures.

The costs of current data centers are mostly driven by their energy consumption (specifically by the air conditioning, computing and networking infrastructures). Yet, current pricing models are usually static and rarely consider the facilities' energy consumption per user. The challenge is to provide a fair and predictable model to attribute the overall energy costs per virtual machine and to increase energy-awareness of users.

Another goal consists in better understanding the energy consumption of computing and networking resources of Clouds in order to provide energy cost models for the entire infrastructure including incentivizing cost models for both Cloud providers and energy suppliers. These models will be based on experimental measurement campaigns on heterogeneous devices. Inferring a cost model from energy measurements is an arduous task since simple models are not convincing, as shown in our previous work. We aim at proposing and validating energy cost models for the heterogeneous Cloud infrastructures in one hand, and the energy distribution grid on the other hand. These models will be integrated into simulation frameworks in order to validate our energy-efficient algorithms at larger scale.

3.3.3. Energy-aware users

In a moderately loaded Cloud, some servers may be turned off when not used for energy saving purpose. Cloud providers can apply resource management strategies to favor idle servers. Some of the existing solutions propose mechanisms to optimize VM scheduling in the Cloud. A common solution is to consolidate the mapping of the VMs in the Cloud by grouping them in a fewer number of servers. The unused servers can then be turned off in order to lower the global electricity consumption.

Indeed, current work focuses on possible levers at the virtual machine suppliers and/or services. However, users are not involved in the choice of using these levers while significant energy savings could be achieved with their help. For example, they might agree to delay slightly the calculation of the response to their applications on the Cloud or accept that it is supported by a remote data center, to save energy or wait for the availability of renewable energy. The VMs are black boxes from the Cloud provider point of view. So, the user is the only one to know the applications running on her VMs.

We plan to explore possible collaborations between virtual machine suppliers, service providers and users of Clouds in order to provide users with ways of participating in the reduction of the Clouds energy consumption. This work will follow two directions: 1) to investigate compromises between power and performance/service quality that cloud providers can offer to their users and to propose them a variety of options adapted to their

workload; and 2) to develop mechanisms for each layer of the Cloud software stack to provide users with a quantification of the energy consumed by each of their options as an incentive to become greener.

3.4. Securing clouds

3.4.1. Security monitoring SLO

While the trend for companies to outsource their information system in clouds is confirmed, the problem of securing an information system becomes more difficult. Indeed, in the case of infrastructure clouds, physical resources are shared between companies (also called tenants) but each tenant controls only parts of the shared resources, and, thanks to virtualization, the information system can be dynamically and automatically reconfigured with added or removed resources (for example starting or stopping virtual machines), or even moved between physical resources (for example using virtual machine migration). Partial control of shared resources brings new classes of attacks between tenants, and security monitoring mechanisms to detect such attacks are better placed out of the tenant-controlled virtual information systems, that is under control of the cloud provider. Dynamic and automatic reconfigurations of the information system make it unfeasible for a tenant's security administrator to setup the security monitoring components to detect attacks, and thus an automated self-adaptable security monitoring service is required.

Combining the two previous statements, there is a need for a dependable, automatic security monitoring service provided to tenants by the cloud provider. Our goal is to address the following challenges to design such a security monitoring service:

1. to define relevant Service-Level Objectives (SLOs) of a security monitoring service, that can figure in the Service-Level Agreement (SLA) signed between a cloud provider and a tenant;
2. to design heuristics to automatically configure provider-controlled security monitoring software components and devices so that SLOs are reached, even during automatic reconfigurations of tenants' information systems;
3. to design evaluation methods for tenants to check that SLOs are reached.

Moreover in challenges 2 and 3 the following sub-challenges must be addressed:

- although SLAs are bi-lateral contracts between the provider and each tenant, the implementation of the contracts is based on shared resources, and thus we must study methods to combine the SLOs;
- the designed methods should have a minimal impact on performance.

3.4.2. Data protection in Cloud-based IoT services

The Internet of Things is becoming a reality. Individuals have their own swarm of connected devices (e.g. smartphone, wearables, and home connected objects) continually collecting personal data. A novel generation of services is emerging exploiting data streams produced by the devices' sensors. People are deprived of control of their personal data as they don't know precisely what data are collected by service providers operating on Internet (oISP), for which purpose they could be used, for how long they are stored, and to whom they are disclosed. In response to privacy concerns the European Union has introduced, with the Global Data Protection Regulation (GDPR), new rules aimed at enforcing the people's rights to personal data protection. The GDPR also gives strong incentives to oISPs to comply. However, today, oISPs can't make their systems GDPR-compliant since they don't have the required technologies. We argue that a new generation of system is mandatory for enabling oISPs to conform to the GDPR. We plan to design an open source distributed operating system for native implementation of new GDPR rules and ease the programming of compliant cloud-based IoT services. Among the new rules, transparency, right of erasure, and accountability are the most challenging ones to be implemented in IoT environments but could fundamentally increase people's confidence in oISPs. Deployed on individuals' swarms of devices and oISPs' cloud-hosted servers, it will enforce detailed data protection agreements and accountability of oISPs' data processing activities. Ultimately we will show to what extent the new GDPR rules can be implemented for cloud-based IoT services.

3.5. Experimenting with Clouds

Cloud platforms are challenging to evaluate and study with a sound scientific methodology. As with any distributed platform, it is very difficult to gather a global and precise view of the system state. Experiments are not reproducible by default since these systems are shared between several stakeholders. This is even worsened by the fact that microscopic differences in the experimental conditions can lead to drastic changes since typical Cloud applications continuously adapt their behavior to the system conditions.

3.5.1. Experimentation methodologies for clouds

We propose to combine two complementary experimental approaches: direct execution on testbeds such as Grid'5000, that are eminently convincing but rather labor intensive, and simulations (using *e.g.*, SimGrid) that are much more light-weighted, but requires careful assessment. One specificity of the Myriads team is that we are working on these experimental methodologies *per se*, raising the standards of *good experiments* in our community.

We plan to make SimGrid widely usable beyond research laboratories, in order to evaluate industrial systems and to teach the future generations of cloud practitioners. This requires to frame the specific concepts of Cloud systems and platforms in actionable interfaces. The challenge is to make the framework both easy to use for simple studies in educational settings while modular and extensible to suit the specific needs of every advanced industrial-class users.

We aim at leveraging the convergence opportunities between methodologies by further bridging simulation and real testbeds. The predictions obtained from the simulator should be validated against some real-world experiments obtained on the target production platform, or on a similar platform. This (in)validation of the predicted results often improves the understanding of the modeled system. On the other side, it may even happen that the measured discrepancies are due to some mis-configuration of the real platform that would have been undetected without this (in)validation study. In that sense, the simulator constitutes a precious tool for the quality assurance of real testbeds such as Grid'5000.

Scientists need more help to make their Cloud experiments fully reproducible, in the spirit of Open Science exemplified by the HAL Open Archive, actively backed by Inria. Users still need practical solutions to archive, share and compare the whole experimental settings, including the raw data production (particularly in the case of real testbeds) and their statistical analysis. This is a long lasting task to which we plan to collaborate through the research communities gathered around the Grid'5000 and SimGrid scientific instruments.

Finally, since correction and performance can constitute contradictory goals, it is particularly important to study them jointly. To that extend, we want to bridge the performance studies, that constitute our main scientific heritage, to correction studies leveraging formal techniques. SimGrid already includes support to exhaustively explore the possible executions. We plan to continue this work to ease the use of the relevant formal methods to the experimenter studying Cloud systems.

3.5.2. Use cases

In system research it is important to work on real-world use cases from which we extract requirements inspiring new research directions and with which we can validate the system services and mechanisms we propose. In the framework of our close collaboration with the Data Science Technology department of the LBNL, we will investigate cloud usage for scientific data management. Next-generation scientific discoveries are at the boundaries of datasets, *e.g.*, across multiple science disciplines, institutions and spatial and temporal scales. Today, data integration processes and methods are largely adhoc or manual. A generalized resource infrastructure that integrates knowledge of the data and the processing tasks being performed by the user in the context of the data and resource lifecycle is needed. Clouds provide an important infrastructure platform that can be leveraged by including knowledge for distributed data integration.

4. Application Domains

4.1. Main application domains

The Myriads team investigates the design and implementation of system services. Thus its research activities address a broad range of application domains. We validate our research results with selected use cases in the following application domains:

- Smart city services,
- Smart grids,
- Energy and sustainable development,
- Home IoT applications,
- Bio-informatics applications,
- Data science applications,
- Computational science applications,
- Numerical simulations.

5. Highlights of the Year

5.1. Highlights of the Year

- SimGrid was named as one of ten “French scientific successes in year 2018” in the French government report “Vers une loi de programmation pluriannuelle de la Recherche”²
- The FogGuru European project has started a real-life experimentation in València (Spain) of Fog computing technologies applied to smart water supply management, in collaboration with Emivasa, the public-private company in charge of water supply.
- The RI/RE project (funded by the CNRS Momentum call) has started in 2019. This project will strengthen our exploration of Smart Grids relations with computing systems.

6. New Software and Platforms

6.1. PaaSage Adapter

KEYWORDS: Cloud computing - Dynamic adaptation - Cloud applications management

FUNCTIONAL DESCRIPTION: The purpose of the Adapter is to transform the current configuration of a cloud application into a target configuration in an efficient and safe way. The Adapter is part of PaaSage, an open-source platform for modeling, deploying and executing applications on different clouds in an optimal manner. The Adapter has the following responsibilities: (1) validating reconfiguration plans, (2) applying the plans to the running system, and (3) maintaining an up-to-date representation of the current system state.

- Contact: Nikolaos Parlavantzas
- URL: <https://team.inria.fr/myriads/software-and-platforms/paasage-adapter/>

6.2. SAIDS

self-adaptable intrusion detection system

KEYWORDS: Cloud - Security

²“Vers une loi de programmation pluriannuelle de la Recherche.” French government’s press release, Feb 2019, page 6. https://cache.media.enseignementsup-recherche.gouv.fr/file/Recherche/91/7/dp-loi_programmation_1069917.pdf

FUNCTIONAL DESCRIPTION: SAIDS is a self-adaptable intrusion detection system for IaaS clouds. To maintain an effective level of intrusion detection, SAIDS monitors changes in the virtual infrastructure of a Cloud environment and reconfigures its components (security probes) accordingly. SAIDS can also reconfigure probes in the case of a change in the list of running services.

- Authors: Anna Giannakou and Jean-Léon Cusinato
- Contact: Christine Morin

6.3. SimGrid

KEYWORDS: Large-scale Emulators - Grid Computing - Distributed Applications

SCIENTIFIC DESCRIPTION: SimGrid is a toolkit that provides core functionalities for the simulation of distributed applications in heterogeneous distributed environments. The simulation engine uses algorithmic and implementation techniques toward the fast simulation of large systems on a single machine. The models are theoretically grounded and experimentally validated. The results are reproducible, enabling better scientific practices.

Its models of networks, cpus and disks are adapted to (Data)Grids, P2P, Clouds, Clusters and HPC, allowing multi-domain studies. It can be used either to simulate algorithms and prototypes of applications, or to emulate real MPI applications through the virtualization of their communication, or to formally assess algorithms and applications that can run in the framework.

The formal verification module explores all possible message interleavings in the application, searching for states violating the provided properties. We recently added the ability to assess liveness properties over arbitrary and legacy codes, thanks to a system-level introspection tool that provides a finely detailed view of the running application to the model checker. This can for example be leveraged to verify both safety or liveness properties, on arbitrary MPI code written in C/C++/Fortran.

NEWS OF THE YEAR: There were 3 major releases in 2019: Python bindings were introduced, SMPI now partially supports some of the MPI/IO functions, a new model for Wifi networks was proposed, and the API for the simulation of storage resources was completely revisited. We also pursued our efforts to improve the documentation of the software, simplified the web site, and made a lot of bug fixing and code refactoring.

- Participants: Adrien Lèbre, Arnaud Legrand, Augustin Degomme, Florence Perronin, Frédéric Suter, Jean-Marc Vincent, Jonathan Pastor, Luka Stanisic and Martin Quinson
- Partners: CNRS - ENS Rennes
- Contact: Martin Quinson
- URL: <https://simgrid.org/>

6.4. DiFFuSE

Distributed framework for cloud-based epidemic simulations

KEYWORDS: Simulation - Cloud

FUNCTIONAL DESCRIPTION: The DiFFuSE framework enables simulations of epidemics to take full advantage of cloud environments. The framework provides design support, reusable code, and tools for building and executing epidemic simulations. Notably, the framework automatically handles failures and supports elastic allocation of resources from multiple clouds.

- Authors: Yvon Jégou, Manh Linh Pham, Nikolaos Parlavantzas and Christine Morin
- Contact: Nikolaos Parlavantzas
- Publication: [hal-01612979/](https://hal.archives-ouvertes.fr/hal-01612979/)
- URL: <https://team.inria.fr/myriads/software-and-platforms/diffuse/>

6.5. GinFlow

KEYWORDS: Workflow - Distributed computing - Distributed - Distributed Applications - Dynamic adaptation - Framework

FUNCTIONAL DESCRIPTION: GinFlow decentralizes the coordination of the execution of workflow-based applications. GinFlow relies on an architecture where multiple service agents (SA) coordinate each others through a shared space containing the workflow description and current status. GinFlow allows the user to define several variants of a workflow and to switch from one to the other during run time.

- Participants: Cédric Tedeschi, Hector Fernandez, Javier Rojas Balderrama, Matthieu Simonin and Thierry Priol
- Partner: Université de Rennes 1
- Contact: Cédric Tedeschi
- URL: <http://ginflow.inria.fr>

6.6. libcvss

KEYWORD: Cybersecurity

FUNCTIONAL DESCRIPTION: libcvss is a Rust implementation of the CVSS specification. The supported versions of CVSS are 2.0, 3.0 and 3.1.

The official CVSS website describes CVSS this way: "The Common Vulnerability Scoring System (CVSS) provides a way to capture the principal characteristics of a vulnerability and produce a numerical score reflecting its severity. The numerical score can then be translated into a qualitative representation (such as low, medium, high, and critical) to help organizations properly assess and prioritize their vulnerability management processes."

libcvss provides Rust users with a native way to manipulate CVSS-formatted vulnerability data. Rust is leveraged to provide a CVSS implementation focused on both performance and correctness.

- Participant: Clement El Baz
- Contact: Clement El Baz
- URL: <https://crates.io/crates/libcvss>

7. New Results

7.1. Scaling Clouds

7.1.1. Efficient Docker container deployment in fog environments

Participants: Arif Ahmed, Lorenzo Civolani, Guillaume Pierre, Paulo Rodrigues de Souza Junior.

Fog computing aims to extend datacenter-based cloud platforms with additional computing, networking and storage resources located in the immediate vicinity of the end users. By bringing computation where the input data was produced and the resulting output data will be consumed, fog computing is expected to support new types of applications which either require very low network latency (e.g., augmented reality applications) or which produce large data volumes which are relevant only locally (e.g., IoT-based data analytics).

Fog computing architectures are fundamentally different from traditional clouds: to provide computing resources in the physical proximity of any end user, fog computing platforms must necessarily rely on very large numbers of small Points-of-Presence connected to each other with commodity networks whereas clouds are typically organized with a handful of extremely powerful data centers connected by dedicated ultra-high-speed networks. This geographical spread also implies that the machines used in any Point-of-Presence may not be datacenter-grade servers but much weaker commodity machines.

We investigated the challenges of efficiently deploying Docker containers in fog platforms composed of tiny single-board computers such as Raspberry Pis. Significant improvements in the Docker image cache hit rate can be obtained by sharing the caches of multiple co-located servers rather than letting them operate independently [9]. In the case when an image must be downloaded and locally installed, large performance gains can be obtained with relatively simple modifications in the way Docker imports container images [3]. Finally, we showed (in collaboration with Prof. Paolo Bellavista from the University of Bologna) that it is possible to let a container start producing useful work even before its image has been fully downloaded [14]. Another paper in this direction of work is in preparation about the way to speedup the boot phase of Docker containers. We are also exploring innovative techniques to improve the performance of live container migration in fog computing environments.

7.1.2. Fog computing platform design

Participants: Ali Fahs, Ayan Mondal, Nikos Parlavantzas, Guillaume Pierre, Mulugeta Tamiru.

There does not yet exist any reference platform for fog computing platforms. We therefore investigated how Kubernetes could be adapted to support the specific needs of fog computing platforms. In particular we focused on the problem of redirecting end-user traffic to a nearby instance of the application. When different users impose various load on the system, any traffic routing system must necessarily implement a tradeoff between proximity and fair load-balancing between the application instances. We demonstrated how such customizable traffic routing policies can be integrated in Kubernetes to help transform it in a suitable platform for fog computing [15]. We extended this work to let the platform automatically choose (and maintain over time) the best locations where application replicas should be deployed. A paper on this topic is currently under submission. We finally started addressing the topic of application autoscaling such that the system can enforce performance guarantees despite traffic variations. We expect one or two publications on this topic next year.

In collaboration with Prof. Misra from IIT Kharagpur (India), and thanks to the collaboration established by the FogCity associate team, we developed mechanisms based on game theory to assign resources to competing applications in a fog computing platform. The objective of those mechanisms is to satisfy user preferences while maximizing resource utilisation. We evaluated the mechanisms using an emulated fog platform built on Kubernetes and Grid'5000, and showed that they significantly outperform baseline algorithms. A paper on this topic is in preparation.

7.1.3. Edgification of micro-service applications

Participants: Genc Tato, Cédric Tedeschi, Marin Bertier.

Last year, we investigated in collaboration with Etienne Riviere from UC Louvain the feasibility and possible benefits brought about by the *edgification* of a legacy micro-service-based application [35]. In other words, we devised a method to classify services composing the application as *edgifiable* or not, based on several criteria. We applied this method to the particular case of the ShareLatex application which enables the collaborative edition of LaTeX documents. Recently, we continue this work by automate the localization and the migration of microservices. Our middleware, based on Koala [36], a lightweight Distributed Hash Table, allows adapting compatible legacy microservices applications for hybrid core/edge deployments [21].

7.1.4. Community Clouds

Participants: Jean-Louis Pazat, Bruno Stevant.

Small communities of people who need to share data and applications can now buy inexpensive devices in order to use only "on premise" resources instead of public Clouds. This "self-hosting-and-sharing" solution provides a better privacy and does not need people to pay any monthly fee to a resource provider. We have implemented a prototype based on micro-services in order to be able to distribute the load of applications among devices.

However, such a distributed platform needs to rely on a very good distribution of the computing and communication load over the devices. Using an emulator of the system, we have shown that, thanks to well known optimization techniques (Particle Swarm Optimization), it is possible to quickly find a service placement resulting in a response time close to the optimal one.

This year we evaluated the results of the optimization algorithm on a prototype (5 "boxes" installed in different home locations connected by fiber or ADSL). Results shown that due to the variation of the network available bandwidth it is necessary to dynamically modify the deployment of applications. This was not a big surprise, but we were not able to find any predictive model of this variation during a day. So, we developed and experimented a dynamic adaptation of the placement of micro-services based applications based on a regular monitoring of the response time of applications. We plan to submit a paper on this topic in early 2020.

7.1.5. *Geo-distributed data stream processing*

Participants: Hamidreza Arkian, Davaadorj Battulga, Mehdi Belkhiria, Guillaume Pierre, Cédric Tedeschi.

We investigated a decentralized scaling mechanism for stream processing applications where the different operators composing the processing topology are able to take their own scaling decisions independently, based on local information. We built a simulation tool to validate the ability of our algorithm to react to load variation. Then, we started the development of a software prototype of a decentralized Stream Processing Engine including this autoscaling mechanism, and deployed it over the Grid'5000 platform. Two papers have been accepted in 2019 about this work [11], [12].

Although data stream processing platforms such as Apache Flink are widely recognized as an interesting paradigm to process IoT data in fog computing platforms, the existing performance model to capture of stream processing in geo-distributed environments are theoretical works only, and have not been validated against empirical measurements. We developed and experimentally validated such a model to represent the performance of a single stream processing operator [10]. This model is very accurate with predictions $\pm 2\%$ of the actual values even in the presence of heterogeneous network latencies. Individual operator models can be composed together and, after the initial calibration of a first operator, a reasonably accurate model for other operators can be derived from a single measurement only.

7.1.6. *QoS-aware and energy-efficient resource management for Function-as-a-Service*

Participants: Yasmina Bouizem, Christine Morin, Nikos Parlavantzas.

Recent years have seen the widespread adoption of serverless computing, and in particular, Function-as-a-Service (FaaS) systems. These systems enable users to execute arbitrary functions without managing underlying servers. However, existing FaaS frameworks provide no quality of service guarantees to FaaS users in terms of performance and availability. Moreover, they provide no support for FaaS providers to reduce energy consumption. The goal of this work is to develop an automated resource management solution for FaaS platforms that takes into account performance, availability, and energy efficiency in a coordinated manner. This work is performed in the context of the thesis of Yasmina Bouizem. In 2019, we integrated a fault-tolerance mechanism into Fission, an open-source FaaS framework based on Kubernetes, and are currently evaluating its impact on performance, availability, and energy consumption.

7.2. Greening Clouds

7.2.1. *Energy Models*

Participants: Loic Guegan, Anne-Cécile Orgerie, Martin Quinson.

Cloud computing allows users to outsource the computer resources required for their applications instead of using a local installation. It offers on-demand access to the resources through the Internet with a pay-as-you-go pricing model. However, this model hides the electricity cost of running these infrastructures.

The costs of current data centers are mostly driven by their energy consumption (specifically by the air conditioning, computing and networking infrastructures). Yet, current pricing models are usually static and rarely consider the facilities' energy consumption per user. The challenge is to provide a fair and predictable model to attribute the overall energy costs per virtual machine and to increase energy-awareness of users. We aim at proposing such energy cost models without heavily relying on physical wattmeters that may be costly to install and operate. These results have been published in [24].

Another goal consists in better understanding the energy consumption of computing and networking resources of Clouds in order to provide energy cost models for the entire infrastructure including incentivizing cost models for both Cloud providers and energy suppliers. These models should be based on experimental measurement campaigns on heterogeneous devices. As hardware architectures become more complex, measurement campaigns are required to better understand their energy consumption and to identify potential sources of energy waste. These results, conducted with Amina Guermouche (IMT Telecom SudParis), have been presented in [30].

Similarly, software stacks add complexity in the identification of energy inefficiencies. For HPC applications, precise measurements are required to determine the most efficient options for the runtime, the resolution algorithm and the mapping on physical resources. An example of such a study has been published in collaboration with HiePACS (Bordeaux) and NACHOS (Sophia) teams in [8].

The fine-grain measurements lead us to propose models that have been used to compare different Cloud architectures (from fog and edge to centralized clouds) in terms of energy consumption on a given scenario. These results have been published in [4].

Inferring a cost model from energy measurements is an arduous task since simple models are not convincing, as shown in our previous work. We aim at proposing and validating energy cost models for the heterogeneous Cloud infrastructures in one hand, and the energy distribution grid on the other hand. These models will be integrated into simulation frameworks in order to validate our energy-efficient algorithms at larger scale. In particular, this year we implemented in SimGrid a flow-based energy model for wired network devices [17].

7.2.2. *End-to-end energy models for the Internet of Things*

Participants: Anne-Cécile Orgerie, Loic Guegan.

The development of IoT (Internet of Things) equipment, the popularization of mobile devices, and emerging wearable devices bring new opportunities for context-aware applications in cloud computing environments. The disruptive potential impact of IoT relies on its pervasiveness: it should constitute an integrated heterogeneous system connecting an unprecedented number of physical objects to the Internet. Among the many challenges raised by IoT, one is currently getting particular attention: making computing resources easily accessible from the connected objects to process the huge amount of data streaming out of them.

While computation offloading to edge cloud infrastructures can be beneficial from a Quality of Service (QoS) point of view, from an energy perspective, it is relying on less energy-efficient resources than centralized Cloud data centers. On the other hand, with the increasing number of applications moving on to the cloud, it may become untenable to meet the increasing energy demand which is already reaching worrying levels. Edge nodes could help to alleviate slightly this energy consumption as they could offload data centers from their overwhelming power load and reduce data movement and network traffic. In particular, as edge cloud infrastructures are smaller in size than centralized data center, they can make a better use of renewable energy.

We investigate the end-to-end energy consumption of IoT platforms. Our aim is to evaluate, on concrete use-cases, the benefits of edge computing platforms for IoT regarding energy consumption. We aim at proposing end-to-end energy models for estimating the consumption when offloading computation from the objects to the Cloud, depending on the number of devices and the desired application QoS. This work has been published in [18].

7.2.3. *Exploiting renewable energy in distributed clouds*

Participants: Benjamin Camus, Anne-Cécile Orgerie.

The growing appetite of Internet services for Cloud resources leads to a consequent increase in data center (DC) facilities worldwide. This increase directly impacts the electricity bill of Cloud providers. Indeed, electricity is currently the largest part of the operation cost of a DC. Resource over-provisioning, energy non-proportional behavior of today's servers, and inefficient cooling systems have been identified as major contributors to the high energy consumption in DCs.

In a distributed Cloud environment, on-site renewable energy production and geographical energy-aware load balancing of virtual machines allocation can be associated to lower the brown (i.e. not renewable) energy consumption of DCs. Yet, combining these two approaches remains challenging in current distributed Clouds. Indeed, the variable and/or intermittent behavior of most renewable sources – like solar power for instance – is not correlated with the Cloud energy consumption, that depends on physical infrastructure characteristics and fluctuating unpredictable workloads.

7.2.4. *Smart Grids*

Participants: Anne Blavette, Benjamin Camus, Anne-Cécile Orgerie, Martin Quinson.

Smart grids allow to efficiently perform demand-side management in electrical grids in order to increase the integration of fluctuating and/or intermittent renewable energy sources in the energy mix. In this work, we consider the computing infrastructure that controls the smart grid. This infrastructure comprises communication and computing resources to allow for a smart management of the electrical grid. In particular, we study the influence of communication latency over a shedding scenario on a small-scale electrical network. We show that depending on the latency some shedding strategies are not feasible [13].

7.3. Securing Clouds

7.3.1. *Security monitoring in Cloud computing platforms*

Participants: Clément Elbaz, Christine Morin, Louis Rilling, Amir Teshome Wonjiga.

In the INDIC project we aim at making security monitoring a dependable service for IaaS cloud customers. To this end, we study three topics:

- defining relevant SLA terms for security monitoring,
- enforcing and verifying SLA terms,
- making the SLA terms enforcement mechanisms self-adaptable to cope with the dynamic nature of clouds.

The considered enforcement and verification mechanisms should have a minimal impact on performance.

In the past years we proposed a verification method for security monitoring SLOs [37] and we have then studied a methodology to define security monitoring SLOs that are at the same time relevant for the tenant, achievable for the provider, and verifiable. The methodology is based on metrics benchmarks that a cloud service provider runs on a set of basic setups of an NIDS (Network Intrusion Detection), the basic setups covering together the variety of NIDS rules that may interest tenants. In order to make it achievable for a cloud service provider to run such benchmarks despite thousands of rules that could be chosen individually by tenants, we proposed a rule clustering strategy to lower the number of sets of rules that should be benchmarked and thus the number of benchmarks run. Finally we proposed extensions to an existing cloud SLA language to define security monitoring SLOs. These results were published in a technical report [27] as well as in Amir Teshome Wonjiga's thesis (to appear) and were submitted for publication in an international conference.

In a side project with Dr Sean Peisert at LBNL, the work on security SLO verification was extended to the use case of data integrity, where tenants outsource data to a cloud storage provider. This work allowed us to tackle a challenge in SLO verification because, in this use case as well as in the security monitoring use case, tenants cannot verify SLOs without a minimal trust in providers involvement in the verification process. We proposed a strategy based on blockchains that allows tenants as well as providers to do SLO verification without having to trust any individual entity. This work was published in the CIFS security workshop [22].

To make security monitoring SLOs adaptable to context changes like the evolution of threats and updates to the tenants' software, we have worked on automating the mitigation of new threats during the time window in which no intrusion detection rule exist and no security patch is applied yet (if available). This time window is critical because newly published vulnerabilities get exploited up to five orders of magnitude right after they are published and the time window may last several days or weeks. We have worked on a first step of mitigation, which consists in deciding if a newly published vulnerability impacts a given information system. A major challenge in automating this step is that newly published vulnerabilities do not contain machine-readable data and this data only appears up to several weeks later. For this reason we designed and evaluated a keyword extraction process from the free-form text description of a vulnerability to map a given vulnerability to product names. This keyword extraction process was first published at the RESSI French security conference [23] and will appear in the NOMS 2020 international conference. In future work this mapping should be combined with a knowledge base of the information system to automatically score the impact of a new vulnerability on the information system.

Our results were published in [27], [28], [22], [23], [25].

7.3.2. *Privacy monitoring in Fog computing platforms*

Participants: Mozhddeh Farhadi, Guillaume Pierre.

IoT devices are integrated in our daily lives, and as a result they often have access to lots of private information. For example many digital assistants (Alexa, Amazon Echo...) were shown to have violated the privacy policy they had established themselves. To increase the level of confidence that end users may have in these devices and the applications which process their data, we started designing monitoring mechanisms such that the fog or the cloud platform can certify whether an application actually follows its own privacy policy or not. A survey paper on security of fog computing platforms is under submission, and we expect another paper on privacy monitoring in 2020.

7.4. Experimenting with Clouds

7.4.1. *Simulating distributed IT systems*

Participants: Toufik Boubehziz, Benjamin Camus, Anne-Cécile Orgerie, Millian Poquet, Martin Quinson.

Our team plays a major role in the advance of the SimGrid simulator of IT systems. This framework has a major impact on the community. Cited by over 900 papers, it was used as a scientific instrument by more than 300 publications over the years.

This year, we pursued our effort to ensure that SimGrid becomes a *de facto* standard for the simulation of distributed IT platforms. We further polished the new interface to ensure that it correctly captures the concepts needed by the experimenters, and provided a Python binding to smooth the learning curve. To that extend, we also continued our rewriting of the documentation.

The work on SimGrid is fully integrated to the other research efforts of the Myriads team. This year, we added the ability to co-simulate IT systems with SimGrid and physical systems modeled with equational systems [13]. This work, developed to study the co-evolution of thermal systems or of the electric grid with the IT system, is now distributed as an official plugin of the SimGrid framework.

7.4.2. *Formal methods for IT systems*

Participants: Ehsan Azimi, The Anh Pham, Martin Quinson.

The SimGrid framework also provide a state of the art Model-Checker for MPI applications. This can be used to formally verify whether the application entails synchronization issues such as deadlocks or livelocks [32]. This year, we pursued our effort on this topic, in collaboration with Thierry Jérón (EPI SUMO).

The Anh Pham defended his thesis this year on techniques to mitigate the state space explosion while verifying asynchronous distributed applications. He adapted an algorithm leveraging event folding structures to this context. This allows to efficiently compute how to not explore equivalent execution traces more than once. This work was published this year [19]. This work, co-advised by Martin Quinson with Thierry Jéron (team SUMO, formal methods), was important to bridge the gap between the involved communities.

Ehsan Azimi joined the Myriads team as an engineer in December to integrate the results of this thesis into the SimGrid framework.

7.4.3. Executing epidemic simulation applications in the Cloud

Participants: Christine Morin, Nikos Parlavantzas, Manh Linh Pham.

In the context of the DiFFuSE ADT and in collaboration with INRA researchers, we transformed a legacy application for simulating the spread of Mycobacterium avium subsp. paratuberculosis (MAP) to a cloud-enabled application based on the DiFFuSE framework (Distributed framework for cloud-based epidemic simulations). This is the second application to which the DiFFuSE framework is applied. The first application was a simulator of the spread of the bovine viral diarrhea virus, developed within the MIHMES project (2012-2017). Using both the MAP and BVDV applications, we performed extensive experiments showing the advantages of the DiFFuSE framework. Specifically, we showed that DiFFuSE enhances application performance and allows exploring different cost-performance trade-offs while supporting automatic failure handling and elastic resource acquisition from multiple clouds [7].

7.4.4. Tools for experimentation

Participant: Matthieu Simonin.

In collaboration with the STACK team and in the context of the Discovery IPL, novel experimentation tools have been developed. In this context experimenting with large software stacks (OpenStack, Kubernetes) was required. These stacks are often tedious to handle. However, practitioners need a right abstraction level to express the moving nature of experimental targets. This includes being able to easily change the experimental conditions (e.g underlying hardware and network) but also the software configuration of the targeted system (e.g service placement, fined-grained configuration tuning) and the scale of the experiment (e.g migrate the experiment from one small testbed to another bigger testbed).

In this spirit we discuss in [31] a possible solution to the above desiderata. We illustrate its use in a real world use case study which has been completed in [34]. We show that an experimenter can express their experimental workflow and execute it in a safe manner (side effects are controlled) which increases the repeatability of the experiments.

The outcome is a library (EnOSlib) target reusability in experiment driven research in distributed systems. The library can be found in <https://bil.inria.fr/fr/software/view/3589/tab>.

8. Partnerships and Cooperations

8.1. Regional Initiatives

8.1.1. INDIC - Cybersecurity Pole of Excellence (2014-2020)

Participants: Clément Elbaz, Christine Morin, Louis Rilling, Amir Teshome Wonjiga.

Our study carried out in the framework of a collaboration with DGA-MI aims at defining and enforcing SLA for security monitoring of virtualized information systems. To this aim we study three topics:

- defining relevant SLA terms for security monitoring,
- enforcing and evaluating SLA terms,
- making the SLA terms enforcement mechanisms self-adaptable to cope with the dynamic nature of clouds.

The considered enforcement and evaluation mechanisms should have a minimal impact on performance. The funding from DGA funded the PhD of Anna Giannakou (defended in 2017) and Amir Teshome Wonjiga (defended in 2019). Clément Elbaz is partially funded by the Brittany Regional Council in the PEC framework.

8.2. National Initiatives

8.2.1. ADEME RennesGrid (2017-2020)

Participants: Anne Blavette, Benjamin Camus, Anne-Cécile Orgerie, Martin Quinson.

The aim of the RennesGrid project is to design and implement a large-scale preindustrial microgrid demonstrator in the territory of Rennes Metropole to organize the shared self-consumption of a group of photovoltaic panels coupled to stationary storage devices. Traditional approaches to power grid management tend to overlook the costs, both energy and economic, of using computers to ensure optimal electricity network management. However, these costs can be significant. It is therefore necessary to take them into account along with the design of IT tools during studies of optimal energy management of smart grids. In addition, telecommunication networks are generally considered to have an ideal functioning, that is to say they can not negatively affect the performance of the electricity network. However, this is not realistic and it is necessary to analyze the impact of phenomena such as congestion, latency, failures related to computer equipment or impact on the batteries of sensors, etc. on strategies for optimal management of the electricity network. In this project, we closely collaborate with Anne Blavette (CR CNRS in electrical engineering, SATIE, Rennes) and co-supervise the post-doc of Benjamin Camus who started in April 2018 on evaluating the impact of the IT infrastructure in the management of smart grids.

8.2.2. Inria ADT Mc SimGrid (2019-2021)

Participants: Ehsan Azimi, Martin Quinson.

The Mc SimGrid technological development action funded by INRIA targets the refactoring of model checker that is integrated to the SimGrid simulation framework. Its software quality should be improved to be on par with the rest of the SimGrid framework. Our ultimate goal is to make this model-checker usable in production, both to assess real-size applications and as a workbench for the researchers designing new techniques and algorithms for the verification of distributed asynchronous applications and algorithms.

The technical actions envisioned for this ADT are the complete re-factoring of this software module, and the exposure of a sensible python interface to experiment with new exploration algorithms. This work is lead by Ehsan Azimi, in collaboration with Thierry Jérón from the Sumo team.

8.2.3. Inria IPL Discovery (2015-2019)

Participants: Anne-Cécile Orgerie, Matthieu Simonin, Genc Tato, Cédric Tedeschi.

The Inria IPL Discovery officially started in September 2015. It targets the design, development and deployment of a distributed Cloud infrastructure within the network's backbone. It will be based upon a set of building blocks whose design will take locality as a primary constraint, so as to minimize distant communications and consequently achieve better network traffic, partition management and improved availability.

Its developments are planned to get integrated within the OpenStack framework. Myriads is involved in the design of new overlay networks for such environments so as to support efficient messaging and routing. Myriads is also involved in the energy/cost benefit analysis of distributed edge-cloud architectures.

8.2.4. Inria IPL Hac Specis (2016-2020)

Participants: Dorra Boughzala, Anne-Cécile Orgerie, The Anh Pham, Martin Quinson.

The goal of the HAC SPECIS (High-performance Application and Computers: Studying PErformance and Correctness In Simulation) project (<http://hacspecis.gforge.inria.fr/>) is to answer methodological needs of HPC application and runtime developers and to allow to study real HPC systems both from the correctness and performance point of view. To this end, we gather experts from the HPC, formal verification and performance evaluation community.

The Anh Pham defended his thesis on December 6., on techniques to mitigate the state space explosion while verifying asynchronous distributed applications. He proposed a new algorithm to mitigate the state space explosion problem (published this year [19]), using event folding structures to efficiently compute how to not explore equivalent execution traces more than once. This work, co-advised by Martin Quinson with Thierry Jérón (team SUMO, formal methods), was important to bridge the gap between the involved communities.

During her PhD thesis, Dorra Boughzala studied the energy consumption of GPU and the simulation tools of the literature related to this aspect. Her work is co-advised by Laurent Lefèvre (Avalon team, Lyon), Martin Quinson and Anne-Cécile Orgerie.

8.2.5. *SESAME ASTRID project (2016-2019)*

Participants: Mehdi Belkhiria, Pascal Morillon, Christine Morin, Matthieu Simonin, Cédric Tedeschi.

The Sesame project (<http://www.agence-nationale-recherche.fr/Project-ANR-16-ASTR-0026>) led by IMT Atlantique aims at develop efficient infrastructures and tools for the maritime traffic surveillance. The role of Myriads is to define a robust and scalable infrastructure for the real-time and batch processing of vessel tracking information. In 2019, we focused on autoscaling and placement for Stream Processing applications.

In 2019, we investigated the dynamic, decentralized scaling of stream processing applications. Also, we collaborated with the Inria OBELIX team to scale and deploy a machine learning application they developed to build a model of a *normal* vessel trajectory.

8.2.6. *CNRS GDS EcoInfo*

Participant: Anne-Cécile Orgerie.

The EcoInfo group deals with reducing environmental and societal impacts of Information and Communications Technologies from hardware to software aspects. This group aims at providing critical studies, lifecycle analyses and best practices in order to improve the energy efficiency of printers, servers, data centers, and any ICT equipment in use in public research organizations.

8.3. European Initiatives

8.3.1. *H2020 Projects*

8.3.1.1. *H2020 MSCA FogGuru*

Participants: Hamidreza Arkian, Davaadorj Battulga, Mozhdeh Farhadi, Julie Montégu, Guillaume Pierre, Mulugeta Ayalew Tamiru, Cédric Tedeschi, Paulo Rodrigues de Souza Junior.

Title: FogGuru – Training the Next Generation of European Fog Computing Experts

Program: H2020 MSCA ITN EID

Duration: September 2017 - August 2021

Coordinator: Guillaume Pierre

Participants:

University of Rennes 1, France (coordinator)

Technische Universität Berlin, Germany

Elastisys AB, Sweden

U-Hopper srl, Italy

EIT Digital Rennes, France

Las Naves, Spain

Abstract: FogGuru is a doctoral training project which aims to train eight talented PhD students with an innovative and inter-sectoral research program to constitute the next generation of European Cloud and Fog computing experts. Besides their scientific and technical education, FogGuru's PhD students will receive extensive training in technological innovation and entrepreneurship as well as soft skills. These combined skills will enable them to fully master the innovation process stemming from fundamental research towards invention and development of innovative products and services, and to real-life deployment, experimentation and engagement with beta-testers.

8.3.2. Collaborations in European Programs, Except FP7 & H2020

8.3.2.1. EIT Digital DriveTrust

Participant: Guillaume Pierre.

Program: EIT Digital

Project acronym: DriveTrust

Project title: AI-Powered Driving Evaluation

Duration: January 2019 - December 2019

Coordinator: University of Rennes 1, France

Other partners:

Eurapco, Switzerland

Achmea, the Netherlands

Imec, Belgium

Abstract: This project aims to develop and commercialize an AI-powered dash cam with short range V2X and LTE communication capabilities. The product uses the newest AI capable hardware for real-time object detection. The device can detect street signs, traffic lights, other cars, and pedestrians. Combined with sensor data from the accelerometer, GPS and weather data from the cloud we use the data to calculate different dimensions of driving profiles. In addition the V2X and object detection capabilities allow us to warn the driver in real-time about dangers on the road.

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

8.3.3.1. FogCity

Participants: Ayan Mondal, Nikos Parlavatzas, Guillaume Pierre.

Title: QoS-aware Resource Management for Smart Cities

International Partner (Institution - Laboratory - Researcher):

IIT Kharagpur (India) - Department of Computer Science and Engineering - Sudip Misra

IIT Kanpur (India) - Department of Industrial and Management Engineering - Subhas Chandra Misra

Start year: 2018

See also: <https://team.inria.fr/myriads/projects/fogcity/>

Abstract: The FogCity associate team proposal concerns a collaboration between the Myriads project-team, and two research teams at Indian Institute of Technology Kharagpur and Indian Institute of Technology Kanpur. The proposal focuses on a smart city scenario in which data from static and mobile sensors is routed to appropriate fog data centres based on application QoS requirements. The main goal of the research is to select suitable nodes within the fog data centers to optimize the QoS of the applications in terms of latency. The different teams have complementary expertise in theoretical research (Indian partners) and system research (Inria Myriads project-team) and share a strong research interest in IoT and Fog Computing.

8.3.3.2. *FogRein*

Participants: Anne-Cécile Orgerie, Martin Quinson.

Title: Steering Efficiency for Distributed Applications

International Partner: Gene Cooperman, College of Computer and Information Science, Northeastern University (USA).

Start year: 2019

In Fog Computing, the Internet of Things (IoT), and Intermittent Computing, low-power devices migrate geographically, and are required to rapidly assimilate new data in a power-efficient manner. This is a key component of any Smart Interfaces solution as devices migrate from the IT infrastructure to the Edge of the Cloud in order to provide Function-as-a-Service, High-availability mobility, and IT infrastructure malleability. A three-tier strategy is proposed toward steering Fog applications in order to optimize the energy efficiency and sustainability. The strategy will leverage the backgrounds of the participants in Fog Computing, checkpointing, scheduling, Green Levers within the IT infrastructure, and a simulation infrastructure for predicting and efficiently steering such distributed applications. The Inria team and the Northeastern team are uniquely positioned to make rapid progress due to their long history of collaborative research based on visits by both permanent members and PhD students in the two directions.

8.3.4. *Inria International Partners*

8.3.4.1. *Informal International Partners*

UC Louvain (Belgium): We collaborate with Prof. Etienne Riviere on legacy application edgification. Genc Tato spent six month at UCL.

Tlemcen University (Algeria): We collaborate with Dr. Djawida Dib on energy-efficient and fault-tolerant resource management in containerized clouds. Christine Morin and Nikos Parlavantzis are co-advising Yasmina Bouizem, who is enrolled in both Tlemcen University and University of Rennes 1.

University of Bologna (Italy): We collaborate with Prof. Paolo Bellavista on the design of performance-efficient fog computing platforms. Lorenzo Civolani from University of Bologna spent 6 months in the Myriads team to complete his master thesis internship. A paper on his work has been accepted for publication [14].

Umeå University (Sweden): We collaborate with Prof. Erik Elmroth on the control of large-scale cloud and fog computing platforms. Ali Fahs spent 6 months at Umeå University where he worked on autoscaling techniques for future fog computing platforms.

Rutgers University (USA): We collaborate with Prof. Manish Parashar on improving the energy efficiency of distributed data-intensive applications. This collaboration is outlined by a joint publication in [20].

University of Hawaii (USA): We collaborate with Prof. Henri Casanova on simulating the energy consumption of scientific workflows. This collaboration is outlined by a joint publication in [16].

University of Southern California (USA): We collaborate with Dr. Rafael Ferreira da Silva and Prof. Ewa Deelman on simulating the energy consumption of scientific workflows. This collaboration is outlined by a joint publication in [16].

8.3.5. *Participation in Other International Programs*

8.3.5.1. *Inria International Chairs*

Deborah AGARWAL

Title: Workflow, user centered design, and data management as well as mobile applications for data science

International Partner (Institution - Laboratory - Researcher):

Université californienne de Santa Barbara (United States) - Computational Research
Division - Deborah Agarwal

Duration: 2015 - 2019

Start year: 2015

8.4. International Research Visitors

8.4.1. Visits of International Scientists

8.4.1.1. Internships

Lorenzo CIVOLANI

Date: Sep 2018 - Feb 2019

Institution: University of Bologna (Italy)

Supervisor: Guillaume Pierre

Adrien GOUGEON

Date: Feb 2019 - Jun 2019

Institution: ENS Rennes

Supervisors: Anne-Cécile Orgerie and Benjamin Camus

Archana WALE

Date: Jan 2019 - Jun 2019

Institution: University of Rennes 1

Supervisor: Guillaume Pierre

Romain Olivo

Date: Juin 2019 - Aug 2019

Institution: Inria

Supervisor: Matthieu Simonin

9. Dissemination

9.1. Promoting Scientific Activities

9.1.1. Member of the Organizing Committees

- Guillaume Pierre is a member of the ACM/IFIP Middleware steering committee.

9.1.2. Scientific Events: Selection

9.1.2.1. Chair of Conference Program Committees

- Guillaume Pierre is co-chair of the Intelligent Systems and Infrastructure track of the TheWebConf 2020 conference.
- Anne-Cécile Orgerie is co-chair of the Architecture and Networking track of the IEEE/ACM CCGrid 2019 conference.
- Anne-Cécile Orgerie is global co-chair of the Cluster and Cloud Computing track of the EuroPar 2019 conference.

9.1.2.2. Member of the Conference Program Committees

- Guillaume Pierre was a PC member of the TheWebConf 2019 conference.
- Guillaume Pierre was a PC member of the ACM/IEEE CCGRID 2019 conference.

- Guillaume Pierre was a PC member of the ACM/IFIP Middleware 2019 conference.
- Guillaume Pierre was a PC member of the Conpas 2019 conference.
- Guillaume Pierre was a PC member of the InterCloud-HPC 2019 workshop.
- Mozhdeh Farhadi was a PC member of the CIFS 2019 workshop.
- Anne-Cécile Orgerie was a PC member of the IEEE HiPC 2019 conference.
- Anne-Cécile Orgerie was a PC member of the ICA3PP 2019 conference.
- Anne-Cécile Orgerie was a PC member of the IEEE/ACM CCGrid 2019 conference.
- Anne-Cécile Orgerie was a PC member of the ICPP 2019 conference.
- Anne-Cécile Orgerie was a PC member of the AlgoTel 2019 conference.
- Anne-Cécile Orgerie was a PC member of the IEEE IC2E 2019 conference.
- Cédric Tedeschi was a PC member of the ICPP 2019 conference.
- Cédric Tedeschi was a PC member of the IEEE ICWS 2019 conference.
- Cédric Tedeschi was a PC member of the Conpas 2019 conference.
- Nikos Parlavantzas was a PC member of the UCC 2019 conference.
- Nikos Parlavantzas was a PC member of the CloudCom 2019 conference.
- Nikos Parlavantzas was a PC member of the ISPDC 2019 conference.
- Nikos Parlavantzas was a PC member of the CrossCloud 2019 workshop.
- Nikos Parlavantzas was a PC member of the VHPC 2019 workshop.

9.1.3. Journal

9.1.3.1. Member of the Editorial Boards

- Anne-Cécile Orgerie is member of the editorial board of International Journal of Distributed Sensor Networks, SAGE Publishing.

9.1.3.2. Reviewer - Reviewing Activities

- Guillaume Pierre reviewed one article for the IEEE Networking Letters.
- Guillaume Pierre reviewed one article for the Software: Practice and Experience journal.
- Cédric Tedeschi reviewed one article for the TPDS journal.
- Nikos Parlavantzas reviewed one article for the Applied Computing and Informatics journal.
- Marin Bertier reviewed one article for the Journal of Supercomputing.
- Martin Quinson reviewed one article for the Journal of Computational Design and Engineering.

9.1.4. Invited Talks

- Guillaume Pierre gave a keynote speech at the IEEE CloudNet 2019 conference: "From Cloud to Fog: The Tao of IT Infrastructure Decentralization."
- Anne-Cécile Orgerie gave a seminar at LORIA lab in September 2019.
- Anne-Cécile Orgerie gave a seminar at the scientific council of INS2I in September 2019.
- Anne-Cécile Orgerie gave a talk at the Cloud Control Workshop in June 2019.
- Martin Quinson gave a talk at the "Algorithmes et programmation" workshop, at Centre International des Rencontres Mathématiques (CIRM) in May 2019.

9.1.5. Leadership within the Scientific Community

- Anne-Cécile Orgerie is vice-chair of ASF, the French Chapter of ACM SIGOPS.

9.1.6. Scientific Expertise

- Jean-Louis Pazat is the coordinator of experts in Information Technology for the evaluation of international bilateral collaborations at the ministry of research and education.

9.1.7. Research Administration

- Anne-Cécile Orgerie is an officer (chargée de mission) for the IRISA cross-cutting axis on Green IT.
- Martin Quinson is the leader of the “Large Scale Systems” department of IRISA.

9.2. Teaching - Supervision - Juries

9.2.1. Teaching

Bachelor: Marin Bertier, Networks, Département Informatique L3, Insa Rennes.

Bachelor: Marin Bertier, Language C, Département Informatique L3, Insa Rennes.

Bachelor: Marin Bertier, Language C, Département Mathématique L3, Insa Rennes.

Bachelor: Jean-Louis Pazat, Introduction à la programmation, L1, Département STPI, INSA de Rennes.

Bachelor: Jean-Louis Pazat, Calcul Hautes Performances, Département Informatique L3, Insa Rennes.

Bachelor: Jean-Louis Pazat, Calcul Hautes Performances, Département Mathématiques L3, Insa Rennes.

Bachelor: Guillaume Pierre, Systèmes Informatiques, L3 MIAGE, Univ. Rennes 1.

Bachelor: Guillaume Pierre, Systèmes d’exploitation, L3 Informatique, Univ. Rennes 1.

Bachelor: Martin Quinson, Architecture et Systèmes, 60 hETD, L3 Informatique, ENS Rennes.

Bachelor: Martin Quinson, Pedagogy, 15 hETD, L3 Informatique, ENS Rennes.

Master: Marin Bertier, Operating Systems, Département Informatique M1, INSA de Rennes

Master: Marin Bertier, Distributed systems, Département Informatique M2, INSA de Rennes

Master: Anne-Cécile Orgerie, Cloud & Big Data, 30 hETD, M1, ENS Rennes.

Master: Anne-Cécile Orgerie, Green ICT, 4.5 hETD, M2, Telecom SudParis Evry.

Master: Anne-Cécile Orgerie, Green ICT, 3 hETD, M2, ENSSAT Lannion.

Master: Anne-Cécile Orgerie, Green IT, 6 hETD, M1, INSA Rennes.

Master: Nikos Parlavantzas, Clouds, M1, INSA Rennes.

Master: Nikos Parlavantzas, Performance Evaluation, M1, INSA Rennes.

Master: Nikos Parlavantzas, Project in Large-Scale Systems, M2, INSA Rennes.

Master: Nikos Parlavantzas, Big Data and Applications, M1, INSA Rennes.

Master: Jean-Louis Pazat, Parallel Computing, M1 Département Informatique Insa Rennes.

Master: Jean-Louis Pazat, Internet Of Things, M1 & M2 Département Informatique Insa Rennes.

Master: Guillaume Pierre, Distributed Systems, M1, Univ. Rennes 1.

Master: Guillaume Pierre, Service technology, M1, Univ. Rennes 1.

Master: Guillaume Pierre, Techniques de développement pour le cloud, M2, Univ. Rennes 1.

Master: Guillaume Pierre, Advanced Cloud Infrastructures, M2, Univ. Rennes 1.

Master: Martin Quinson, Préparation à l’Agrégation de Science Industrielle (Programming and Software Engineering, 20h ETD; Operating Systems and C programming, 20 hETD; Networking, 20h ETD), ENS Rennes.

Master: Martin Quinson, Scientific Outreach, M2, 30 hEDT, ENS Rennes.

Master: Cédric Tedeschi, Concurrency in Systems and Networks, M1, Univ. Rennes 1.

Master: Cédric Tedeschi, Service Technology, M1, Univ. Rennes 1.

Master: Cédric Tedeschi, Parallel Programming, M1, Univ. Rennes 1.

Master: Cédric Tedeschi, Systems, L3, Univ. Rennes 1.

9.2.2. Supervision

Defended PhD: Genc Tato, “Locality-aware Lazy Overlay Networks for WANs,” defended on December 19th 2019, supervised by Marin Bertier, Cédric Tedeschi.

Defended PhD: The Anh Pham, “Dynamic Formal Verification of High Performance Runtimes and Applications,” defended on December 6th 2019, supervised by Martin Quinson, Thierry Jéron.

Defended PhD: Amir Teshome Wonjiga, “User-Centric Security Monitoring in Cloud Environments”, defended on June 3rd 2019, supervised by Louis Rilling and Christine Morin.

PhD in progress: Adrien Gougeon, “Designing an energy-efficient communication network for the dynamic and distributed control of the electrical grid”, started in September 2019, supervised by Anne-Cécile Orgerie and Martin Quinson.

PhD in progress: Paulo Rodrigues De Souza Junior, “fog computing service roaming techniques”, started in December 2018, supervised by Guillaume Pierre and Daniele Miorandi (U-Hopper srl, Italy).

PhD in progress: Davaadorj Battulga, “Scalable data pipelines for fog computing applications”, started in September 2018, supervised by Cédric Tedeschi and Daniele Miorandi (U-Hopper srl, Italy).

PhD in progress: Mulugeta Tamiru, “Automatic optimization of autonomous management systems”, started in September 2018, supervised by Guillaume Pierre and Erik Elmroth (Elastisys AB, Sweden).

PhD in progress: Mozhdeh Farhadi, “Fog computing-enabled IoT situation-aware services”, started in June 2018, supervised by Guillaume Pierre and Daniele Miorandi (U-Hopper srl, Italy).

PhD in progress: Hamidreza Arkian, “Stream processing operator placement”, started in June 2018, supervised by Guillaume Pierre and Erik Elmroth (Elastisys AB, Sweden).

PhD in progress: Mehdi Belkhiria, “Dynamic Stream Processing for Maritime Traffic Surveillance,” started in December 2017, supervised by Cédric Tedeschi.

PhD in progress (*co-tutelle*): Yasmina Bouizem, “Energy-efficient, fault-tolerance mechanisms for containerized cloud applications,” started in December 2017, supervised by Didi Fedoua (Tlemcen University, Algeria), Djawida Dib (Tlemcen University, Algeria), Christine Morin and Nikos Parlavantzas.

PhD in progress: Ali Jawad Fahs, “Decentralized Fog Computing Infrastructure Control”, started in October 2017, supervised by Guillaume Pierre.

PhD in progress: Clément El Baz, “Reactive security monitoring in clouds,” started in October 2017, supervised by Louis Rilling and Christine Morin.

PhD in progress: Loic Guegan, “Simulating Internet of Things”, started in October 2017, supervised by Martin Quinson and Anne-Cécile Orgerie.

PhD in progress: Dorra Boughzala, “Simulating Energy Consumption of Continuum Computing between Heterogeneous Numerical Infrastructures in HPC”, started in December 2017, supervised by Laurent Lefèvre (Avalon team in Lyon), Anne-Cécile Orgerie and Martin Quinson.

PhD in progress: Arif Ahmed, “Scalable Decentralized Edge Cloud Infrastructures”, started in October 2016, supervised by Guillaume Pierre.

PhD in progress: Baptiste Goupille-Lescar, “Designing agile, distributed cyber-physical systems with advanced collaboration capabilities”, started in January 2016, supervised by Eric Lenormand (Thales), Christine Morin, Nikos Parlavantzas.

PhD in progress: Bruno Stevant, “Resource allocation strategies for service distribution at the Internet edge to optimize end-to-end latency,” started in December 2014 (part-time), supervised by Jean-Louis Pazat.

9.2.3. *Juries*

- Guillaume Pierre chaired the PhD committee of Alexandre Veith, ENS Lyon, September 23rd 2019.
- Guillaume Pierre chaired the PhD committee of Jad Darrous, ENS Lyon, December 17th 2019.
- Martin Quinson chaired the PhD committee of Pierre Huchant, Ecole Nationale Supérieure d’Électronique, Informatique, Télécommunications, Mathématique et Mécanique de Bordeaux (ENSEIRB-MATMECA), March 29th 2019.
- Anne-Cécile Orgerie is a member of the jury for *Agrégation externe de Sciences Industrielles de l’Ingénieur option ingénierie informatique*.
- Martin Quinson is a member of the jury for *Capes d’informatique*.
- Anne-Cécile Orgerie was a member of the PhD committee of Ayham Kassab, Université de Franche-Comté, November 14th 2019.
- Anne-Cécile Orgerie was a member of the PhD committee of Gustavo Rostirolla, Université Toulouse 3, November 25th 2019.
- Louis Rilling was a member of the PhD committee of Guillaume Averlant, INSA Toulouse, October 2nd 2019.
- Louis Rilling was a guest of the PhD committee of Mohammad Mahdi Bazm, Université de Nantes, July 8th 2019.

9.3. Popularization

9.3.1. *Internal or external Inria responsibilities*

- Martin Quinson is on the scientific board of the Blaise Pascal foundation, boosting the scientific outreach in the domains of Maths and Computer Science.

9.3.2. *Education*

- Guillaume Pierre shared his experience as the coordinator of a H2020 Maria Skłodowska Curie Action project with CRRDT, the Brittany Region’s consultative committee for research and technological development. January 14th 2019.
- “L codent L créent” is an outreach program to send PhD students to teach Python to middle school students in 8 sessions of 45 minutes. Tassadit Bouadi (Lacodam), Camille Maumet (Empenn) and Anne-Cécile Orgerie (Myriads) are coordinating the local version of this program, initiated in Lille. The first session in Rennes occurred in April 2019, and a new session has started for 2019-2020. The program is currently supported by: Fondation Blaise Pascal, ED MathSTIC, Université Rennes 1 and Fondation Rennes 1.
- “J’peux pas, j’ai informatique” is an outreach program to deconstruct IT stereotypes. Organized over one day, this program welcomes 110 college students (5ème) each year and is coordinated by Anne-Cécile Orgerie. The next event is planned in April 2020.
- Martin Quinson participated to the DIU EIL (Enseigner l’Informatique au Lycée), training existing maths and engineering teachers so that they can teach Informatics in their colleges.
- Martin Quinson is in charge of a M2 lecture on scientific vulgarization at ENS Rennes. The students are asked to come up with new activities (or improve existing ones) on their research field.
- Martin Quinson gave a doctoral teaching on December 13 on scientific outreach targeting either schools or science forums at IRISA laboratory.

9.3.3. *Internal action*

- Anne-Cécile Orgerie, Martin Quinson and Matthieu Simonin: organize experimenter meetings for the large-scale systems department of IRISA: <https://xug.gitlabpages.inria.fr/meetings/>

9.3.4. Creation of media or tools for science outreach

- Martin Quinson is working on a set of new outreach activities with his Master 2 class at ENS Rennes since several years. The most promising activities are then tested on field in schools during another lecture at ENS Rennes, at Licence 3 level. All available resources are freely available at <https://github.com/InfoSansOrdi/pedago-rennes/>.

10. Bibliography

Major publications by the team in recent years

- [1] H. CASANOVA, A. LEGRAND, M. QUINSON, F. SUTER. *SMPI Courseware: Teaching Distributed-Memory Computing with MPI in Simulation*, in "EduHPC-18 - Workshop on Education for High-Performance Computing", Dallas, United States, November 2018, pp. 1-10, <https://hal.inria.fr/hal-01891513>

Publications of the year

Doctoral Dissertations and Habilitation Theses

- [2] T. A. PHAM. *Efficient state-space exploration for asynchronous distributed programs: Adapting unfolding-based dynamic partial order reduction to MPI programs*, ENS Rennes, December 2019, <https://hal.archives-ouvertes.fr/tel-02420950>

Articles in International Peer-Reviewed Journals

- [3] A. AHMED, G. PIERRE. *Docker-pi: Docker Container Deployment in Fog Computing Infrastructures*, in "International Journal of Cloud Computing", 2019, pp. 1-20, forthcoming, <https://hal.inria.fr/hal-02271434>
- [4] E. AHVAR, A.-C. ORGERIE, A. LEBRE. *Estimating Energy Consumption of Cloud, Fog and Edge Computing Infrastructures*, in "IEEE Transactions on Sustainable Computing", April 2019, pp. 1-12 [DOI : 10.1109/TSUSC.2019.2905900], <https://hal.archives-ouvertes.fr/hal-02083080>
- [5] G. COOPERMAN, M. QUINSON. *Sthread: In-Vivo Model Checking of Multithreaded Programs*, in "The Art, Science, and Engineering of Programming", 2020, forthcoming, <https://hal.inria.fr/hal-02449080>
- [6] M. FARHADI, J.-L. LANET, G. PIERRE, D. MIORANDI. *A systematic approach towards security in Fog computing: assets, vulnerabilities, possible countermeasures*, in "Software: Practice and Experience", 2020, forthcoming, <https://hal.inria.fr/hal-02441639>
- [7] N. PARLAVANTZAS, L. M. PHAM, C. MORIN, S. ARNOUX, G. BEAUNÉE, L. QI, P. GONTIER, P. EZANNO. *A Service-based Framework for Building and Executing Epidemic Simulation Applications in the Cloud*, in "Concurrency and Computation: Practice and Experience", 2019, forthcoming, <https://hal.archives-ouvertes.fr/hal-02312313>

International Conferences with Proceedings

- [8] E. AGULLO, L. GIRAUD, S. LANTERI, G. MARAIT, A.-C. ORGERIE, L. POIREL. *Energy Analysis of a Solver Stack for Frequency-Domain Electromagnetics*, in "PDP 2019 - 27th Euromicro International Conference on Parallel, Distributed and Network-Based Processing", Pavia, Italy, 2019 27th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP), IEEE, February 2019, pp. 385-391 [DOI : 10.1109/EMPDP.2019.8671555], <https://hal.archives-ouvertes.fr/hal-02191331>
- [9] A. AHMED, G. PIERRE. *Docker Image Sharing in Distributed Fog Infrastructures*, in "CloudCom 2019 - 11th IEEE International Conference on Cloud Computing Technology and Science", Sydney, Australia, IEEE, December 2019, <https://hal.inria.fr/hal-02304285>
- [10] H. ARKIAN, G. PIERRE, J. TORDSSON, E. ELMROTH. *An Experiment-Driven Performance Model of Stream Processing Operators in Fog Computing Environments*, in "SAC 2020 - ACM/SIGAPP Symposium On Applied Computing", Brno, Czech Republic, March 2020, pp. 1-9, <https://hal.inria.fr/hal-02394396>
- [11] M. BELKHIRIA, C. TEDESCHI. *A Fully Decentralized Autoscaling Algorithm for Stream Processing Applications*, in "Auto-DaSP 2019 - Third International Workshop on Autonomic Solutions for Parallel and Distributed Data Stream Processing", Göttingen, Germany, August 2019, pp. 1-12, <https://hal.inria.fr/hal-02171172>
- [12] M. BELKHIRIA, C. TEDESCHI. *Design and Evaluation of Decentralized Scaling Mechanisms for Stream Processing*, in "CloudCom 2019 - 11th IEEE International Conference on Cloud Computing Technology and Science", Sydney, Australia, IEEE, December 2019, <https://hal.inria.fr/hal-02351108>
- [13] B. CAMUS, A. BLAVETTE, A.-C. ORGERIE, J.-B. BLANC-ROUCHOSSÉ. *Co-simulation of an electrical distribution network and its supervision communication network*, in "CCNC 2020 - IEEE Consumer Communications & Networking Conference", Las Vegas, United States, CCNC: IEEE Consumer Communications & Networking Conference, IEEE, January 2020, pp. 1-6, <https://hal.archives-ouvertes.fr/hal-02352832>
- [14] L. CIVOLANI, G. PIERRE, P. BELLAVISTA. *FogDocker: Start Container Now, Fetch Image Later*, in "UCC 2019 - 12th IEEE/ACM International Conference on Utility and Cloud Computing", Auckland, New Zealand, ACM, December 2019, pp. 51-59 [DOI : 10.1145/3344341.3368811], <https://hal.inria.fr/hal-02332679>
- [15] A. FAHS, G. PIERRE. *Proximity-Aware Traffic Routing in Distributed Fog Computing Platforms*, in "CCGrid 2019 - IEEE/ACM International Symposium in Cluster, Cloud, and Grid Computing", Larnaca, Cyprus, IEEE, May 2019, pp. 1-10 [DOI : 10.1109/CCGRID.2019.00062], <https://hal.inria.fr/hal-02048965>
- [16] R. FERREIRA DA SILVA, A.-C. ORGERIE, H. CASANOVA, R. TANAKA, E. DEELMAN, F. SUTER. *Accurately Simulating Energy Consumption of I/O-intensive Scientific Workflows*, in "ICCS 2019 - International Conference on Computational Science", Faro, Portugal, ICCS 2019 - International Conference on Computational Science, Springer, June 2019, pp. 138-152 [DOI : 10.1007/978-3-030-22734-0_11], <https://hal.archives-ouvertes.fr/hal-02112893>
- [17] L. GUEGAN, B. L. AMERSHO, A.-C. ORGERIE, M. QUINSON. *A Large-Scale Wired Network Energy Model for Flow-Level Simulations*, in "AINA 2019 - 33rd International Conference on Advanced Information Networking and Applications", Matsue, Japan, Advances in Intelligent Systems and Computing, Springer, March 2019, vol. 926, pp. 1047-1058 [DOI : 10.1007/978-3-030-15032-7_88], <https://hal.inria.fr/hal-02020045>

- [18] L. GUEGAN, A.-C. ORGERIE. *Estimating the end-to-end energy consumption of low-bandwidth IoT applications for WiFi devices*, in "CloudCom 2019 - 11th IEEE International Conference on Cloud Computing Technology and Science", Sydney, Australia, IEEE, December 2019, <https://hal.archives-ouvertes.fr/hal-02352637>
- [19] T. A. PHAM, T. JÉRON, M. QUINSON. *Unfolding-based Dynamic Partial Order Reduction of Asynchronous Distributed Programs*, in "FORTE 2019 - 39th International Conference on Formal Techniques for Distributed Objects, Components, and Systems", Copenhagen, Denmark, J. A. PÉREZ, N. YOSHIDA (editors), Formal Techniques for Distributed Objects, Components, and Systems, Springer International Publishing, 2019, vol. LNCS-11535, pp. 224-241, Part 1: Full Papers [DOI : 10.1007/978-3-030-21759-4_13], <https://hal.inria.fr/hal-02109769>
- [20] I. RAÏS, D. BALOUËK-THOMERT, A.-C. ORGERIE, L. LEFÈVRE, M. PARASHAR. *Leveraging energy-efficient non-lossy compression for data-intensive applications*, in "HPCS 2019 - 17th International Conference on High Performance Computing & Simulation", Dublin, Ireland, HPCS 2019 - 17th International Conference on High Performance Computing & Simulation, July 2019, pp. 1-7, <https://hal.archives-ouvertes.fr/hal-02179621>
- [21] G. TATO, M. BERTIER, E. RIVIÈRE, C. TEDESCHI. *Split and migrate: Resource-driven placement and discovery of microservices at the edge*, in "OPODIS 2019 : 23rd International Conference On Principles Of Distributed Systems", Neuchâtel, Switzerland, 2019, pp. 1-16 [DOI : 10.4230/LIPIcs.OPODIS.2019.29], <https://hal.inria.fr/hal-02401933>
- [22] A. T. WONJIGA, S. PEISERT, L. RILLING, C. MORIN. *Blockchain as a Trusted Component in Cloud SLA Verification*, in "UCC 2019 - 12th IEEE/ACM International Conference on Utility and Cloud Computing", Auckland, New Zealand, ACM, December 2019, pp. 93-100 [DOI : 10.1145/3368235.3368872], <https://hal.inria.fr/hal-02354647>

National Conferences with Proceedings

- [23] C. ELBAZ, L. RILLING, C. MORIN. *Towards Automated Risk Analysis of "One-day" Vulnerabilities*, in "RESSI 2019 - Rendez-vous de la Recherche et de l'Enseignement de la Sécurité des Systèmes d'Information", Erquy, France, August 2019, pp. 1-3, <https://hal.inria.fr/hal-02267192>

Scientific Books (or Scientific Book chapters)

- [24] A. OLEKSIK, L. LEFÈVRE, P. ALONSO, G. DA COSTA, V. DE MAIO, N. FRASHERI, V. GARCIA, J. GUERRERO, S. LAFOND, A. LASTOVETSKY, R. R. MANUMACHU, B. MUIITE, A.-C. ORGERIE, W. PIATEK, J.-M. PIERSON, R. PRODAN, P. STOLF, E. SHEME, S. VARRETTE. *Energy aware ultrascale systems*, in "Ultrascale Computing Systems", Institution of Engineering and Technology, January 2019, pp. 127-188 [DOI : 10.1049/PBPC024E_CH], <https://hal.inria.fr/hal-02163289>

Research Reports

- [25] C. ELBAZ, L. RILLING, C. MORIN. *Automated Keyword Extraction from "One-day" Vulnerabilities at Disclosure*, Inria Rennes - Bretagne Atlantique, November 2019, n° RR-9299, pp. 1-22, <https://hal.inria.fr/hal-02362062>
- [26] A. JLASSI, C. TEDESCHI. *Merge, Split, and Cluster: Dynamic Deployment of Stream Processing Applications*, Inria, December 2019, <https://hal.inria.fr/hal-02418771>

- [27] A. T. WONJIGA, L. RILLING, C. MORIN. *Defining Security Monitoring SLAs in IaaS Clouds: the Example of a Network IDS*, Inria Rennes Bretagne Atlantique, March 2019, n^o RR-9263, pp. 1-37, <https://hal.inria.fr/hal-02079860>
- [28] A. T. WONJIGA, L. RILLING, C. MORIN. *Security Monitoring SLA Verification in Clouds: the Case of Data Integrity*, Inria Rennes - Bretagne Atlantique, March 2019, n^o RR-9267, pp. 1-29, <https://hal.inria.fr/hal-02084186>

Other Publications

- [29] M. BELKHIRIA, C. TEDESCHI. *Decentralized Scaling for Stream Processing Engines*, May 2019, working paper or preprint, <https://hal.inria.fr/hal-02127609>
- [30] A. GUERMOUCHE, A.-C. ORGERIE. *Experimental analysis of vectorized instructions impact on energy and power consumption under thermal design power constraints*, June 2019, working paper or preprint, <https://hal.archives-ouvertes.fr/hal-02167083>

References in notes

- [31] R.-A. CHERRUEAU, M. SIMONIN, A. VAN KEMPEN. *EnosStack: A LAMP-like stack for the experimenter*, in "INFOCOM WKSHPS 2018 - IEEE International Conference on Computer Communications", Honolulu, United States, IEEE, April 2018, pp. 336-341 [DOI : 10.1109/INFCOMW.2018.8407024], <https://hal.inria.fr/hal-01916484>
- [32] M. GUTHMULLER, G. CORONA, M. QUINSON. *System-level state equality detection for the formal dynamic verification of legacy distributed applications*, in "Journal of Logical and Algebraic Methods in Programming", April 2018, vol. 96, pp. 1 - 11 [DOI : 10.1016/J.JLAMP.2017.12.004], <https://hal.inria.fr/hal-01900120>
- [33] A. LEBRE, J. PASTOR, M. BERTIER, F. DESPREZ, J. ROUZAUD-CORNABAS, C. TEDESCHI, P. ANEDDA, G. ZANETTI, R. NOU, T. CORTES, E. RIVIÈRE, T. ROPARS. *Beyond The Cloud, How Should Next Generation Utility Computing Infrastructures Be Designed?*, Inria, July 2013, n^o RR-8348, <https://hal.inria.fr/hal-00854204>
- [34] J. ROJAS BALDERRAMA, M. SIMONIN. *Scalability and Locality Awareness of Remote Procedure Calls: An Experimental Study in Edge Infrastructures*, in "CloudCom", Nicosia, Cyprus, December 2018, <https://hal.inria.fr/hal-01891567>
- [35] G. TATO, M. BERTIER, E. RIVIÈRE, C. TEDESCHI. *ShareLatex on the Edge: Evaluation of the Hybrid Core/Edge Deployment of a Microservices-based Application*, in "The 3rd workshop on Middleware for Edge Clouds & Cloudlets", Rennes, France, December 2018, <https://hal.inria.fr/hal-01942807>
- [36] G. TATO, M. BERTIER, C. TEDESCHI. *Koala: Towards Lazy and Locality-Aware Overlays for Decentralized Clouds*, in "ICFEC 2018 - 2nd IEEE International Conference on Fog and Edge Computing", Washington, United States, IEEE, May 2018, pp. 1-10 [DOI : 10.1109/CFEC.2018.8358728], <https://hal.inria.fr/hal-01938582>
- [37] A. T. WONJIGA, L. RILLING, C. MORIN. *Verification for Security Monitoring SLAs in IaaS Clouds: the Example of a Network IDS*, in "NOMS 2018 - Network Operations and Management Symposium", Taipei, Taiwan, IEEE, April 2018, pp. 1-7 [DOI : 10.1109/NOMS.2018.8406157], <https://hal.inria.fr/hal-01663941>