



Activity Report 2021

Team MYRIADS

Design and Implementation of Autonomous Distributed Systems

Joint team with Inria Rennes – Bretagne Atlantique

D1 – Large Scale Systems



Contents

Project-Team MYRIADS	1
1 Team members, visitors, external collaborators	3
2 Overall objectives	4
2.1 General Objectives	4
2.2 Context	4
2.3 Challenges	5
3 Research program	5
3.1 Introduction	5
3.2 Scaling fogs and clouds	6
3.2.1 Resource management in hierarchical clouds	6
3.2.2 Resource management in fog computing architectures	6
3.2.3 Self-optimizing applications in multi-cloud environments	6
3.3 Greening clouds	7
3.3.1 Smart grids and clouds	7
3.3.2 Energy cost models	7
3.3.3 Energy-aware users	8
3.4 Securing clouds	8
3.4.1 Security monitoring SLO	8
3.4.2 Data protection in Cloud-based IoT services	9
3.5 Experimenting with Clouds	9
3.5.1 Experimentation methodologies for clouds	9
3.5.2 Use cases	10
4 Application domains	10
4.1 Main application domains	10
5 Social and environmental responsibility	11
5.1 Footprint of research activities	11
5.2 Impact of research results	11
6 Highlights of the year	11
6.1 Awards	11
7 New software and platforms	11
7.1 New software	11
7.1.1 PaaSage Adapter	11
7.1.2 SAIDS	11
7.1.3 SimGrid	12
7.1.4 DiFFuSE	13
7.1.5 GinFlow	13
7.1.6 libcvss	13
7.1.7 Tansiv	14
7.2 New platforms	14
8 New results	14
8.1 Scaling Clouds	14
8.1.1 Fog computing platform design	14
8.1.2 Advanced data management for fast deployment in distributed Clouds	15
8.1.3 Geo-distributed data stream processing	15
8.1.4 Streaming Graph Partitioning	16
8.1.5 QoS-aware resource management for Function-as-a-Service	16
8.2 Greening Clouds	17

8.2.1	Energy Models	17
8.2.2	End-to-end energy models for the Internet of Things	17
8.2.3	Exploiting renewable energy in distributed clouds	17
8.2.4	Smart Grids	18
8.3	Securing Clouds	18
8.3.1	Security monitoring in Cloud computing platforms	18
8.3.2	Privacy monitoring in Fog computing platforms	19
8.4	Experimenting with Clouds	19
8.4.1	Simulating distributed IT systems	19
8.4.2	Formal methods for IT systems	20
8.4.3	Toward stealth analysis of distributed applications	20
8.4.4	A systematic analysis of data locality in MapReduce	20
8.4.5	Tools for experimentation	21
9	Bilateral contracts and grants with industry	21
9.1	Bilateral contracts with industry	21
9.2	Bilateral Grants with Industry	21
10	Partnerships and cooperations	21
10.1	International initiatives	21
10.1.1	Associate Teams in the framework of an Inria International Lab or in the framework of an Inria International Program	21
10.1.2	Inria associate team not involved in an IIL or an international program	22
10.1.3	Participation in other International Programs	23
10.2	International research visitors	23
10.2.1	Visits of international scientists	23
10.2.2	Visits to international teams	23
10.3	European initiatives	23
10.3.1	Horizon Europe	23
10.3.2	FP7 & H2020 projects	23
10.3.3	Other european programs/initiatives	24
10.4	National initiatives	24
11	Dissemination	26
11.1	Promoting scientific activities	26
11.1.1	Scientific events: organisation	26
11.1.2	Scientific events: selection	26
11.1.3	Journal	27
11.1.4	Invited talks	27
11.1.5	Leadership within the scientific community	28
11.1.6	Scientific expertise	28
11.1.7	Research administration	28
11.2	Teaching - Supervision - Juries	28
11.2.1	Teaching	28
11.2.2	Supervision	30
11.2.3	Juries	30
11.3	Popularization	31
11.3.1	Internal or external Inria responsibilities	31
11.3.2	Articles and contents	31
11.3.3	Education	31
12	Scientific production	31
12.1	Major publications	31
12.2	Publications of the year	31
12.3	Other	33
12.4	Cited publications	34

Project-Team MYRIADS

Creation of the Project-Team: 2012 January 01

Keywords

Computer sciences and digital sciences

- 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
- A3.1.3. – Distributed data
- A4.9. – Security supervision
- A4.9.1. – Intrusion detection
- A4.9.3. – Reaction to attacks
- 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 members, visitors, external collaborators

Research Scientists

- Deborah Agarwal [Ernest Orlando Lawrence Berkeley National Laboratory, Chair]
- Shadi Ibrahim [Inria, Researcher]
- Anne-Cécile Orgerie [CNRS, Researcher, HDR]

Faculty Members

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

Post-Doctoral Fellow

- Francois Lemerrier [CNRS]

PhD Students

- Hamidreza Arkian [Univ de Rennes I, until Sep 2021]
- Davaadorj Battulga [U-Hopper - Italy, until Aug 2021]
- Mehdi Belkhiria [Univ de Rennes I, until Feb 2021]
- Yasmina Bouizem [Inria, until May 2021]
- Clement Courageux-Sudan [Inria]
- Mozhdeh Farhadi [U-Hopper - Italie, until Aug 2021]
- Wedan Emmanuel Gnibga [CNRS, from Oct 2021]
- Adrien Gougeon [École normale supérieure de Rennes]
- Chih Kai Huang [Univ de Rennes I, from Nov 2021]
- Paulo Rodrigues De Souza Junior [Univ de Rennes I]
- Mulugeta Tamiru [Elastisys - Suède, until Sep 2021]

Technical Staff

- Ehsan Azimi [Inria, Engineer, until Nov 2021]
- Thomas Lambert [Inria, Engineer, until Jul 2021]
- Nena Markovic [Univ de Rennes I, Engineer, until Nov 2021]
- Ayan Mondal [Univ de Rennes I, Engineer, until Oct 2021]
- Alessio Pagliari [Univ de Rennes I, Engineer, from Oct 2021]
- Matthieu Simonin [Inria, Engineer]

Interns and Apprentices

- Maryam Ahmadi [Univ de Rennes I, from Feb 2021 until Aug 2021]
- Jad Alhaji [Inria, from Feb 2021 until Aug 2021]
- Julien Hong [Inria, from Mar 2021 until Aug 2021]
- Damien Meranthe [Univ de Rennes I, from May 2021 until Jul 2021]
- Manon Sourisseau [Univ de Rennes I, from May 2021 until Jul 2021]
- Amisha Srivastava [École normale supérieure de Rennes, from May 2021 until Jul 2021]

Administrative Assistant

- Angélique Jarnoux [Inria]

External Collaborators

- Anne Blavette [CNRS]
- 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 (D1 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 online 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

¹According to World Stats, there are 4.94 billion Internet users i.e., more than 60% of the total world population in December 2020.

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 15 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 swarms of people's devices (fog computing) enabling scalable cloud computing. These distributed clouds raise new challenges for resource and application management.

The ultimate goal of the 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 systems 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 include 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 other can significantly improve the performance of the application they compose. More generally, building an overlay network that takes into account the hierarchical aspects of the platform without being a hierarchical overlay – which comes with load balancing and resilience issues – is a challenge by itself.

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 a 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.
- How data should be stored and managed to facilitate the deployment of Fog infrastructures and IoT applications while taking into account the limited storage capacity.

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), this 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. The request may have consumed less energy, or a different kind of energy (renewable or not), if it had been sent to a more distant 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 (oISPs), 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. In addition, we are also working on new approaches to allow the running of graph applications in geo-distributed Clouds while respecting the data protection regulations in different locations.

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 is eminently convincing but rather labor intensive, and simulations (using e.g., SimGrid) that are much more light-weight, but require careful assessments. 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 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 ad hoc 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 Social and environmental responsibility

5.1 Footprint of research activities

Anne-Cécile Orgerie is involved in the CNRS GDS EcoInfo that 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 reduce the environmental impact of ICT equipment in use in public research organizations.

5.2 Impact of research results

One of the research axes of the team consists in measuring and decreasing the energy consumption of Cloud computing infrastructures. The work associated to this axis contributes to increasing the energy efficiency of distributed infrastructures. This work has been conducted in particular in the IPL Hac Specis.

In the context of the ADEME RennesGrid and CNRS RI/RE projects, work is also conducted on the current challenges of the energy sector and more specifically on the smart digitization of power grid management through the joint optimization of electricity generation, distribution and consumption. This work aims to optimize the computing infrastructure in charge of managing the electricity grids: guaranteeing their performance while minimizing their energy consumption.

6 Highlights of the year

- The FogGuru European project coordinated by the Myriads team successfully completed on November 30th 2021.
- The ANR FACTO project coordinated by the Myriads team started on November 1st 2021.

6.1 Awards

- Shadi Ibrahim was elevated to the grade of IEEE Senior Member on December 2021.

7 New software and platforms

7.1 New software

7.1.1 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.

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

Contact: Nikolaos Parlavantzas

7.1.2 SAIDS

Name: self-adaptable intrusion detection system

Keywords: Cloud, Security

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, Jean-Léon Cusinato

Contact: Christine Morin

7.1.3 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.

Functional 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 2 major releases in 2020. SMPI is now regularly tested on medium scale benchmarks of the exascale suite. The Wifi support was improved, through more example and documentation, and an energy model of wifi links was proposed. Many bugs were fixed in the bindings to the ns-3 packet-level network simulator, which now allows to simulate Wifi links using ns-3 too. We enriched the API expressiveness to allow the construction of activity tasks. 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.

URL: <https://simgrid.org/>

Contact: Martin Quinson

Participants: Adrien Lebre, Anne-Cécile Orgerie, Arnaud Legrand, Augustin Degomme, Emmanuelle Saillard, Frédéric Suter, Jean-Marc Vincent, Jonathan Pastor, Luka Stanisic, Martin Quinson, Samuel Thibault

Partners: CNRS, ENS Rennes

7.1.4 DiFFuSE

Name: 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.

URL: <https://team.inria.fr/myriads/software-and-platforms/diffuse/>

Publication: [hal-01612979/](https://hal.archives-ouvertes.fr/hal-01612979/)

Authors: Yvon Jégou, Manh Linh Pham, Nikolaos Parlavantzas, Christine Morin

Contact: Nikolaos Parlavantzas

7.1.5 GinFlow

Name: 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.

URL: <http://ginflow.inria.fr>

Contact: Cédric Tedeschi

Participants: Cédric Tedeschi, Hector Fernandez, Javier Rojas Balderrama, Matthieu Simonin, Thierry Priol

Partner: Université de Rennes 1

7.1.6 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.

URL: <https://crates.io/crates/libcvss>

Contact: Clement El Baz

Participant: Clement El Baz

7.1.7 Tansiv

Name: Time-Accurate Network Simulation Interconnecting Vms

Keywords: Operating system, Virtualization, Cloud, Simulation, Cybersecurity

Functional Description: Tansiv: Time-Accurate Network Simulation Interconnecting Virtual machines (VMs). Tansiv is a novel way to run an unmodified distributed application on top of a simulated network in a time accurate and stealth way. To this aim, the VMs execution is coordinated (interrupted and restarted) in order to guarantee accurate arrival and transfer of network packets while ensuring realistic time flow within the VMs. The initial prototype uses Simgrid for simulating the data transfer and control the execution of the VMs. Also, Qemu is used to encapsulate the application, intercept the network traffic and enforce the interruption decision. Tansiv can be used in various situations: malware analysis (e.g. to defeat malware evasion technique based on network timing measures) or analysis of an application on a geo-distributed context.

Contact: Louis Rilling

Partner: DGA-MI

7.2 New platforms

n/a

8 New results

8.1 Scaling Clouds

8.1.1 Fog computing platform design

Participants Ali Fahs, Chih-Kai Huang, Guillaume Pierre, Paulo Souza, Mulegeta 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 Kubernetes may be adapted to handle a range of difficult issues to make it suitable for becoming a major reference platform for future fog computing platforms.

An interesting technology for extending Kubernetes to large-scale geo-distributed scenarios is Kubernetes Federations (KubeFed), which allow one to aggregate resources provided by multiple independent Kubernetes clusters in various locations. We continued our previous work on Kubernetes by proposing mck8s, an extension to the KubeFed Kubernetes federation which gives additional control to the users regarding the placement of applications within the federation. mck8s also integrates extensive functionalities for application and platform autoscaling [13]. This work constitutes the final contribution of Mulegeta Tamiru's PhD thesis [24].

To allow Kubernetes to become a major platform for geo-distributed fog computing, an important and currently missing technique is the migration of Kubernetes pods across different servers located in different locations. In this context, one difficult sub-problem is the migration of disk state across the

system: although in a single data center the usage of network attached storage allows administrators to migrate VMs without having to migrate their disk state, in geo-distributed migration migrating this state becomes necessary. We showed how the usage of layered file systems for keeping this disk state can significantly reduce the migration time of such disk state [33]. We further extended this work toward a complete implementation of seamless geo-distributed Kubernetes pod migration, and a publication on this topic is currently in preparation.

8.1.2 Advanced data management for fast deployment in distributed Clouds

Participants Thomas Lambert, Shadi Ibrahim.

Containers have been widely used in various cloud and fog platforms as they enable agile and elastic application deployment through their process-based virtualization and layered image system. Different layers of a container image may contain substantial duplicate and unnecessary data, which slows down its deployment due to long image downloading time and increased burden on the image registry. To accelerate the deployment and reduce the size of the registry, we introduce a new image format that consists of two parts: a Gear index describing the structure of the image's file system and a set of files that are required when running an application. The Gear index is represented as a single-layer image compatible with the existing deployment framework. Containers can be launched by pulling a Gear index and on demand retrieving files pointed to by the index. Furthermore, the Gear image enables a file-level sharing mechanism, which helps remove duplicate data in the registry and avoid repeated downloading of identical files by a client [8].

Another approach to provide fast and predictable service provisioning time in distributed Clouds and Fog platforms is distributing container images across Fog nodes. In our previous work [30], we introduced two placement algorithms that distribute container images (layers) across Fog nodes considering the layer sizes, the storage capacity of Fog servers and network heterogeneity in the Fog (we refer to as the initial placement). In 2021, we extend our work by refining the initial placement based the popularity of the containers. A scientific paper on this topic is currently in preparation.

8.1.3 Geo-distributed data stream processing

Participants Hamidreza Arkian, Davaadorj Battulga, Mehdi Belkhiria, Ayan Mondal, Thomas Lambert, Guillaume Pierre, Cédric Tedeschi, Shadi Ibrahim.

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 models that capture stream processing in geo-distributed environments are theoretical works only, and have not been validated against empirical measurements. We developed and experimentally validated an auto-scaling mechanism which can dynamically add or remove resources such that a geo-distributed stream processing application can dynamically adjust its processing capacity to wide variations in the intensity of the workload imposed by the arrival of incoming data in the system [5]. This work constitutes part of Hamidreza Arkian's PhD thesis [20].

As part of the new DiPET project we started exploring how transprecision techniques may be used in the context of geo-distributed data stream processing systems to implement a tradeoff between execution performance and the precision of the obtained results. For example, when a short-lived load surge is detected, it might be preferable to temporarily reduce the precision rather than to trigger expensive auto-scaling operations. In other scenarios such as the continuous monitoring of network traffic for intrusion detection, it may be useful to operate with low precision and low resource usage most of the time, only to increase the precision level when a suspicious event is detected.

In the context of the ANR KerStream project, we extend our preliminary work on resource provisioning of data processing workflows [34] by adding one more use case involving two types of system dynamics including the cloud performance dynamics and dynamic system failures. We demonstrate the effectiveness

and also generality of our proposed algorithm using real-world workflows on two real cloud platforms and a large-scale shared cluster [4]. In addition, we focus on stragglers (slow tasks) mitigation in data analytic frameworks. In particular, we perform extensive experiments to study and analyze speculative execution (launching other copies of stragglers) implementation in Apache Spark. Moreover, we study how heartbeat arrivals and using outdated information when detecting stragglers impact the accuracy of stragglers detection. Our findings demonstrate that considering heartbeat time-stamps can improve the accuracy of stragglers detection. A scientific paper on this topic is submitted for publication.

In the context of Davaadorj Battulga's PhD thesis, we are exploring mechanisms to bring stream processing applications in a geo-distributed environment. We base our approach on the collaboration of multiple, geo-distributed stream processing engines. Preliminary experiments over a small commodity cluster were obtained [28] in 2020. In 2021, a more complete software prototype was developed and deployed over multiple sites on the Grid'5000 platform. This helped us validate the idea of having stream processing pipelines deployed over multiple autonomous computing sites and offer a playground for future experiments with features to be designed regarding self-adaptation of such a platform [6].

8.1.4 Streaming Graph Partitioning

Participants Anne-Cécile Orgerie.

Graph partitioning, a preliminary step of distributed graph processing, has been attracting increasing attention in the last decade. A high quality graph partitioning algorithm should facilitate graph processing by minimizing the communication overhead and maintaining the load balancing among distributed computing units. Offline partitioning algorithms usually require the knowledge of a complete graph, and therefore, are not adaptive to handle massive graph-structured data. On the contrary, streaming partitioning algorithms take edges or vertices as a stream and make partitioning decisions on the fly. However, the streaming manner faces dilemmas from time to time because of a lack of knowledge. Furthermore, an unmindful partitioning decision in such a dilemma could significantly decrease the partition quality. We proposed a novel window-based streaming graph partitioning algorithm (WSGP). WSGP leverages a greedy-based heuristic to perform edge partitioning. When facing a decision dilemma, WSGP utilizes a size-bounded window to buffer the edges. When the window is fully filled, an edge is popped and assigned to a partition. The assignment is decided by knowledge obtained from both the edges already settled and the ones still cached in the buffer window. Our experiments take into account various real-world benchmark graphs. The experimental results demonstrate that WSGP consistently has a smaller replication factor than the state-of-the-art algorithms by up to 23%, at a limited cost in terms of memory and comprehensive running time [18].

8.1.5 QoS-aware 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. The goal of this work is to develop an automated resource management solution for FaaS platforms that takes into account performance and availability in a coordinated manner. This work is performed in the context of the thesis of Yasmina Bouizem.

We have proposed the integration of additional fault-tolerance mechanisms beyond the typical retry mechanism into FaaS frameworks. In 2021, we integrated an active replication mechanism into an open source, Kubernetes-based FaaS framework, namely Fission. We then performed a detailed experimental comparison of this mechanism with the retry mechanism and an Active-Standby mechanism, offering insights on how to select fault-tolerance approaches based on the application type and performance and availability requirements [25].

8.2 Greening Clouds

8.2.1 Energy Models

Participants 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.

Vectorized instructions were introduced to improve the performance of applications. However, they come at the cost of an increase in the power consumption. As a consequence, processors are designed to limit their frequency when such instructions are used in order to respect the thermal design power limit. We studied and compare the impact of thermal design power and SIMD instructions on performance, power and energy consumption of processors and memory. The study shows that, because of processor frequency, performance and power consumption are strongly related to thermal design power [2].

Making distributed systems more energy and thermal efficient will require understanding of individual power dissipation and temperature contributions of multiple hardware system components and their accompanying software. We proposed an experimental workflow for energy and temperature profiling on systems running distributed applications. It allows full and dynamic control over the execution of applications for the entire frequency range. Through its use, we show that the energy response to frequency scaling is highly dependent on the workload characteristics and it is convex in nature with an optimal frequency point [19].

8.2.2 End-to-end energy models for the Internet of Things

Participants Clément Courageux-Sudan, Loic Guegan, Anne-Cécile Orgerie, Martin Quinson.

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. In particular, in 2021, we investigated the end-devices' side and their Wi-Fi interfaces. We proposed scalable accurate flow-level Wi-Fi models that have been integrated within SimGrid [23].

8.2.3 Exploiting renewable energy in distributed clouds

Participants Anne Blavette, Emmanuel Gnibga, 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.

8.2.4 Smart Grids

Participants Anne Blavette, Adrien Gougeon, François Lemerrier, 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. Our results show that latency may have a significant impact on energy management strategies in terms of performance for the smart grid, and that it should be considered when designing such strategies [10].

8.3 Securing Clouds

8.3.1 Security monitoring in Cloud computing platforms

Participants Clément Elbaz, Christine Morin, Louis Rilling.

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.

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 exists 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. A major

challenge in automation during this critical time window is that newly published vulnerabilities do not contain machine-readable data and this data only appears up to several weeks later.

In the past years we have worked on a first step of mitigation, which consists in deciding if a newly published vulnerability impacts a given information system and how urgent it is to assess this impact before, if necessary, mitigating it. More precisely, we have designed an active learning process to build a knowledge base of the information system, which aims at representing the set of software products used in the information system and for which vulnerabilities may be important to mitigate. This knowledge base is composed of past public vulnerabilities processed using the keyword extraction process that we previously designed [31] and annotated about their potential criticality with respect to the information system to defend by a security expert of the information system. Using this knowledge base, newly published vulnerabilities are processed using the same keyword extraction process and compared with similar vectors in the knowledge base to infer if and how fast it should be handed to a security analyst. In 2021 we completed a preliminary experimental evaluation of this active learning-based reaction process and published the results [16].

All the work on automated mitigation was part of the PhD thesis of Clément Elbaz who defended his thesis in 2021 [21].

8.3.2 Privacy monitoring in Fog computing platforms

Participants Mozhdeh 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. We demonstrated how machine-learning techniques may be used to analyze the (encrypted) network traffic produced by fog applications to identify the types of data they communicate, and compare this type with the claims made in their privacy policy [9]. This work constitutes part of Mozhdeh Farhadi's PhD thesis [22].

8.4 Experimenting with Clouds

8.4.1 Simulating distributed IT systems

Participants Anne-Cécile Orgerie, 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 (60 PhD thesis, 150 journal articles, and 300 conference articles), it was used as a scientific instrument by thousands of 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, in preparation to ongoing works on the simulation of the computing infrastructure behind the SKA scientific infrastructure [26].

Our work on SimGrid is fully integrated into the other research efforts of the Myriads team. This year we had two major contributions on this topic. First, we proposed a methodology to semi-automatically model micro-services applications [7], enabling the study of their performance within the simulator. In addition, we built upon our previous contributions toward the co-simulation of cyberphysical systems coupling SimGrid to other domain-specific simulators [10].

8.4.2 Formal methods for IT systems

Participants Ehsan Azimi, Martin Quinson.

The SimGrid framework also provides a state of the art Model-Checker for MPI applications called Mc SimGrid. This can be used to formally verify whether the application entails synchronization issues such as deadlocks or livelocks.

This year, Ehsan Azimi worked as an engineer to integrate the results of The Anh Pham, who defended his thesis last year, into the SimGrid framework. We worked to reorganize the existing prototype to ease the full integration of the modern research works in SimGrid. This work was conducted in collaboration with Thierry Jéron (EPI SUMO).

In addition, we proposed a benchmark to evaluate bug finding tools dedicated to distributed MPI programs [11]. This benchmark was used to compare Mc SimGrid to its direct competitors, to better understand their main strengths and drawbacks. We hope to initiate a yearly tool competition, to make the state of the art more transparent and foster the adoption of such tools by their potential users.

8.4.3 Toward stealth analysis of distributed applications

Participants Martin Quinson, Louis Rilling, Matthieu Simonin.

In the TANSIV project we aim at extending the usability of SimGrid to software using arbitrary network communication APIs and paradigms. For instance this enables SimGrid to run distributed services implemented in operating systems kernels, such as distributed file systems, and high performance network applications based on poll-mode network interface card drivers like in the DPDK framework. To this end we proposed to interconnect SimGrid with Virtual Machine Monitors (VMM) and let all the network packets output by the virtual machines (VM) flow through SimGrid. This proposal also enhances SimGrid with applications to security, as the interconnected VMMs can be malware analysis sandboxes. Thus SimGrid enables malware analysis sandboxes to feature scalable performance-realistic network environments in order to defeat anti-analysis techniques developed by malware authors.

In 2021 we have started experimentally evaluating the soundness of the network simulation provided by SimGrid in the TANSIV prototype. To this end network micro-benchmarks as well distributed applications are run on top of TANSIV. These experiments will be continued in 2022.

Two internships were started in 2021 and master thesis should be defended in 2022 about a first approach to use TANSIV in hardware-virtualization-based VMMs. A PhD thesis topic was also published about a complete work on hardware-assisted virtualization in TANSIV and its integration with the time-related tools available in malware analysis sandboxes

8.4.4 A systematic analysis of data locality in MapReduce

Participants Jad Alhajj, Thomas Lambert, Shadi Ibrahim, Martin Quinson.

Big Data analytic frameworks leverage data locality to improve the performance of data-intensive applications: launching tasks on the nodes where data resides to avoid data transfer. Much work have focused on optimizing data locality of data-intensive applications by targeting task scheduling, job scheduling, or data distribution. However, less efforts have been dedicated to compare and to study the complementary of these approaches.

To this end, we perform extensive experiments to evaluate the performance of three state-of-the-art data locality approaches in a contrast and complementary approach. In particular, we implement these approaches in iHadoop (Hadoop simulator build on the top of SimGrid) and observe the achieved data locality and the performance when using individual approaches and when combined. We find that the

impact of data locality is less important in heterogeneous environments. Interestingly, we show that combining heterogeneity-aware data distribution and locality-aware task and job scheduling yields better performance since this increases the probability of executing local tasks on fast nodes.

8.4.5 Tools for experimentation

Participants 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 [29] a possible solution to the above desiderata.

The outcome is a library (EnOSlib) target reusability in experiment driven research in distributed systems.

The tool is used in several articles (see [here](#)). In particular, in [32] the tool is used to build an ad hoc framework for studying FOG applications.

9 Bilateral contracts and grants with industry

9.1 Bilateral contracts with industry

Défi Inria OVHcloud (2021-2025)

Participants Anne-Cécile Orgerie, Shadi Ibrahim, Guillaume Pierre.

The goal of this collaborative framework between the OVHcloud and Inria is to explore new solutions for the design of cloud computing services that are more energy-efficient and environment friendly. Five Inria project-teams are involved in this challenge including Avalon, Inocs, Myriads, Spirals, Stack.

Members of the Myraids team will contribute to four sub-challenges including (1) Software codesign of a data stream processing service; (2) energy-efficient data management; (3) observation of bare metal co-location platforms and proposal of energy reduction catalogues and models; and (4) modelling and designing a framework and its environmental Gantt Chart to manage physical and logical levers. The actual work is expected to start in 2022.

9.2 Bilateral Grants with Industry

n/a

10 Partnerships and cooperations

10.1 International initiatives

10.1.1 Associate Teams in the framework of an Inria International Lab or in the framework of an Inria International Program

Hermes

Title: Accelerating the Performance of Multi-Site Scientific applications through Coordinated Data management

Duration: 2019 - 2022

Coordinator: Shadi Ibrahim

Partners:

- Scientific Data Management Group, Lawrence Berkeley National Laboratory (United States)

Inria contact: Shadi Ibrahim

Summary: Advances in computing, experimental, and observational facilities are enabling scientists to generate and analyze unprecedented volumes of data. A critical challenge facing scientists in this era of data deluge is storing, moving, sharing, retrieving, and gaining insight from massive collections of data efficiently. Existing data management and I/O solutions on high-performance computing (HPC) systems require significant enhancements to handle the three V's of Big Data (volume, velocity, and variety) in order to improve productivity of scientists. Even more challenging, many scientific Big Data and machine learning applications require data to be shared, exchanged, and transferred among multiple HPC sites. Towards overcoming these challenges, in this project, we aim at accelerating scientific Big Data application performance through coordinated data management that addresses performance limitations of managing data across multiple sites. In particular, we focus on challenges related to the management of data and metadata across sites, distributed burst buffers, and online data analysis across sites.

10.1.2 Inria associate team not involved in an IIL or an international program

FogRein

Title: Steering Efficiency for Distributed Applications

Duration: 2019 - 2022

Coordinator: Martin Quinson

Partners:

- College of Computer and Information Science, Northeastern University (United States)

Inria contact: Martin Quinson

Summary: 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. Most planned visits were canceled in 2020 because of the sanitary situation. Prof. Cooperman came to Rennes in January, and one of his PhD student spent a 4 months internship in our team during spring.

10.1.3 Participation in other International Programs

Managing epidemic simulation applications in cloud environments

Title: Managing epidemic simulation applications in cloud environments

Partner Institution: Vietnam National University (Vietnam)

Duration: 2020-2022

Project type: Scientific cooperation project, French Embassy in Vietnam

Coordinator: Nikos Parlavantzas

Summary: The project concerns a collaboration between the Myriads team and the FIMO center, VNU University of Engineering and Technology, Vietnam National University. The main aim is to extend the DiFFuSE framework and apply it to epidemic simulation applications developed in the Vietnam National University. DiFFuSE is a service-based framework developed by Myriads that enables simulation applications to fully exploit cloud platforms. In 2021, we produced a joint publication on a framework supporting the simulation of the spread of African swine fever in Vietnam [3].

10.2 International research visitors

10.2.1 Visits of international scientists

Inria International Chair

Other international visits to the team

10.2.2 Visits to international teams

Sabbatical programme

Research stays abroad

10.3 European initiatives

10.3.1 Horizon Europe

n/a

10.3.2 FP7 & H2020 projects

H2020 MSCA FogGuru

Title: Training the next Generation of European Fog Computing Experts

Duration: 2017-2021

Coordinator: Guillaume Pierre (Université de Rennes 1)

Partners:

- Université de Rennes 1 (France)
- Elastisys AB (Sweden)
- Technische Universität Berlin (Germany)
- U-Hopper srl (Italy)
- EIT Digital Rennes (France)
- Las Naves (Spain)

Summary: 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.

10.3.3 Other european programs/initiatives

CHIST-ERA DiPET

Title: Distributed Stream Processing on Fog and Edge Systems via Transprecise Computing

Duration: 2020-2023

Coordinator: Hans van Dierendonck (Queen's University Belfast)

Partners:

- Queen's university Belfast (QUB, United Kingdom)
- Foundation for research and Technolohy - Hellas (FORTH, Greece)
- Université de Rennes 1 (UR1, France)
- Universitat Politècnica de Catalunya (UPC, Spain)
- West University of Timisoara (UVT, Romania)

Summary: The DiPET project investigates models and techniques that enable distributed stream processing applications to seamlessly span and redistribute across fog and edge computing systems. The goal is to utilize devices dispersed through the network that are geographically closer to users to reduce network latency and to increase the available network bandwidth. However, the network that user devices are connected to is dynamic. For example, mobile devices connect to different base stations as they roam, and fog devices may be intermittently unavailable for computing. In order to maximally leverage the heterogeneous compute and network resources present in these dynamic networks, the DiPET project pursues a bold approach based on transprecise computing. Transprecise computing states that computation need not always be exact and proposes a disciplined trade-off of precision against accuracy, which impacts on computational effort, energy efficiency, memory usage and communication bandwidth and latency. Transprecise computing allows to dynamically adapt the precision of computation depending on the context and available resources. This creates new dimensions to the problem of scheduling distributed stream applications in fog and edge computing environments and will lead to schedules with superior performance, energy efficiency and user experience. The DiPET project will demonstrate the feasibility of this unique approach by developing a transprecise stream processing application framework and transprecision-aware middleware. Use cases in video analytics and network intrusion detection will guide the research and underpin technology demonstrators.

10.4 National initiatives

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.

This year, we refactored the code that was written during the thesis of The Anh Pham to make it easier to integrate it in the future to the SimGrid framework.

SESAME ASTRID project (2016-2021)

Participants Mehdi Belkhiria, Pascal Morillon, Matthieu Simonin, Cédric Tedeschi.

The **Sesame project** 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.

CNRS Momentum RI/RE (2019-2021)

Participants Anne-Cécile Orgerie, François Lemercier.

This project focuses on the current challenges of the energy sector and more specifically on the smart digitization of power grid management through the joint optimization of electricity generation, distribution and consumption.

The project aims to optimize the computing infrastructure in charge of managing the electricity grids: guaranteeing their performance while minimizing their energy consumption.

ANR KerStream (2017-2022)

Participants Shadi Ibrahim.

The KerStream project (Big Data Processing: Beyond Hadoop!) is an ANR JCJC (Young Researcher) project (ANR-16-CE25-0014-1) starting in January 2017 with an allocated budget of 238 k€.

The goal of the KerStream project is to address the limitations of Hadoop when running Big Data stream applications on large-scale clouds and do a step beyond Hadoop by proposing a new approach, called KerStream, for scalable and resilient Big Data stream processing on clouds. The KerStream project can be seen as the first step towards developing the first French middleware that handles Stream Data processing at Scale.

ANR FACTO (2021-2024)

Participants Anne-Cécile Orgerie, Martin Quinson, François Lemercier.

The number of smart homes is rapidly expanding worldwide with an increasing amount of wireless IT devices. The diversity of these devices is accompanied by the development of multiple wireless protocols and technologies that aim to connect them. However, these technologies offer overlapping capabilities. This overprovisioning is highly suboptimal from an energy point of view and can be viewed as a first barrier towards sustainable smart homes. Therefore, in the FACTO project, we propose to design a multi-purpose network based on a single optimized technology (namely Wi-Fi), in order to offer an energy-efficient, adaptable and integrated connectivity to all smart home's devices.

Smartobs (2021)

Participants Guillaume Pierre, Shadi Ibrahim, Maryam Ahmadi.

SmartObs is a collaborative project funded by the University of Rennes 1 which brings together the Myriads and Taran teams of Inria as well as OSUR, the university's department of geosciences. The objective of Smartobs was to initiate a collaboration between the three teams in the domain of natural environment monitoring. In this context, we implement a web application that leverages Fog technologies to collect data about the natural environment including temperature, humidity and barometric pressure.

11 Dissemination

11.1 Promoting scientific activities

11.1.1 Scientific events: organisation

General chair, scientific chair

- Shadi Ibrahim was the general chair for the 6th International Parallel Data Systems Workshop (PDSW@SC2021), November 15, 2021.

Member of organizing committees

- Anne-Cécile Orgerie was member of the program committee of the Inria Scientific Days in 2021.
- Shadi Ibrahim is a member of the steering committee of the International Parallel Data Systems Workshop.
- Shadi Ibrahim was publicity co-chair for the 50th International Conference on Parallel Processing (ICPP 2021).

11.1.2 Scientific events: selection

Chair of conference program committees

- Anne-Cécile Orgerie was Doctoral Symposium co-chair for CCGrid 2021, Demo and Poster co-chair for ICDCS 2021 and chair of the track on Performance Evaluation for SBAC-PAD 2021.
- Shadi Ibrahim was a program co-chair for the The 23rd IEEE International Conference on High Performance Computing and Communications (HPCC-2021).

Member of the conference program committees

- Guillaume Pierre was a member of the program committees of TheWebConf 2021, IEEE ICFC 2021, ACM Middleware 2021, IEEE JCC 2021, WEEE 2021, VNI 2021, IEEE AINet 2021, SIGNIS 2021.
- Cédric Tedeschi was a member of the program committee of IEEE/ACM CCGrid 2021.
- Anne-Cécile Orgerie was a member of the program committees of IEEE/ACM CCGrid 2021, IEEE IPDPS 2021, IEEE Cluster 2021, ACM Student Research Competition hosted at IEEE/ACM Supercomputing 2021, AlgoTel 2021 and JCAD 2021.
- Shadi Ibrahim was a member of the program committees of IEEE IPDPS 2021, IEEE ICDCS 2021, IEEE/ACM CCGrid 2021, IEEE Cluster 2021, IEEE ICPADS 2021, and IEEE Cloud Summit 2021.
- Nikos Parlavantzas was a member of the program committees of IEEE/ACM UCC 2021, IEEE ISPD 2020, VNI 2021, SIGNIS 2021

Reviewer

11.1.3 Journal

Member of the editorial boards

- Anne-Cécile Orgerie is a member of the editorial board of IEEE Transactions on Parallel and Distributed Systems.
- Shadi Ibrahim is an associate editor of IEEE Internet Computing Magazine.
- Shadi Ibrahim is a young associate editor of the Springer Frontiers of Computer Science journal.

Reviewer - reviewing activities

- Shadi Ibrahim was a reviewer for the IEEE Transactions on Big Data, IEEE Transactions on Network and Service Management, and ACM Transactions on Embedded Computing Systems.
- Nikos Parlavantzas was a reviewer for the IEEE Computer and IEEE Transactions on Cloud Computing.

11.1.4 Invited talks

- Guillaume Pierre: "La plateforme FogGuru d'expérimentation fog computing pour les villes intelligentes", GDR meeting "Cluster d'architectures à base de SoC pour des calculs hautes performances", January 20th 2021.
- Guillaume Pierre: "Fog Computing Research in the Inria/IRISA Myriads team". Journées 5G, Inria Rennes, October 14th 2021.
- Guillaume Pierre: "The LivingFog platform, from water consumption measurement to natural environment monitoring". Journées Environnement et Numérique, Inria Rennes, October 20th 2021.
- Guillaume Pierre: "mck8s: An Kubernetes-based orchestration platform for geo-distributed multi-cluster environments". Journées Cloud, ENSTA Brest, November 26th 2021.
- Guillaume Pierre: "Stateful Container Migration in Geo-Distributed Environments". Huawei Annual Storage Technology Summit, December 13th 2021.
- Anne-Cécile Orgerie: "Impacts sur les dépenses énergétiques de la 5G", seminar Penser la science - La 5G en question of the Université libre de Bruxelles, Belgium, February 5, 2021.
- Anne-Cécile Orgerie: "Comprendre et réduire la consommation énergétique des systèmes distribués", seminar IDCHP (on distributed systems and HPC) at the Fédération informatique de Lyon, Lyon, July 1, 2021.
- Anne-Cécile Orgerie: "Energy efficiency of computing and communication infrastructures for HPC", seminar at IBENS (Institut de Biologie de l'École Normale Supérieure), Paris, September 9, 2021.
- Anne-Cécile Orgerie: "Understanding and improving the energy efficiency of distributed systems", seminar at LISN (Laboratoire Interdisciplinaire des Sciences du Numérique), Orsay, October 5, 2021.
- Anne-Cécile Orgerie: "Measuring and modeling the energy consumption of Cloud infrastructures", Inria Workshop On Systems, Rennes, October 12, 2021.
- Anne-Cécile Orgerie: "Modeling and optimizing Smart Grids", Journées Environnement et Numérique, Inria Rennes, October 20th 2021..

- Anne-Cécile Orgerie: “Measuring and modeling the energy consumption of ICT distributed systems”, Power efficient deep learning Workshop co-located with Asian Conference on Machine Learning (ACML), virtually, November 17, 2021.
- Anne-Cécile Orgerie: “Measuring the energy consumption of HPC systems”, forum on frugal HPC of ORAP (French Organization of Parallelism), virtually, December 9th, 2021.

11.1.5 Leadership within the scientific community

- Anne-Cécile Orgerie is director of the CNRS service group on ICT environmental impact (GDS EcoInfo).
- Anne-Cécile Orgerie is chief scientist for the Rennes site of Grid’5000.

11.1.6 Scientific expertise

- Cédric Tedeschi was an expert for the Natural Sciences and Engineering Research Council of Canada (NSERC).
- Anne-Cécile Orgerie was a member of the selection committees for an assistant professor position at Université de Lorraine and Inria research scientist positions (CRCN and ISFP) at the Nancy site and Grenoble site in 2021.
- Nikos Parlavantzas was an expert for ANR (French National Research Agency)

11.1.7 Research administration

- Anne-Cécile Orgerie is an officer (chargée de mission) for the IRISA cross-cutting axis on Green IT.
- Anne-Cécile Orgerie is member of the Inria Evaluation Committee.
- Anne-Cécile Orgerie is co-coordinator of the working group Ancre-Allistene on computer science and energy.
- Anne-Cécile Orgerie is member of the steering committee of the CNRS GDR RSD.

11.2 Teaching - Supervision - Juries

11.2.1 Teaching

- Bachelor: Marin Bertier, Networks, Département Informatique L3, Insa Rennes.
- Bachelor: Marin Bertier, C Language Département Informatique L3, Insa Rennes.
- Bachelor: Marin Bertier, C Language, Département Mathématique L3, Insa Rennes.
- Bachelor: Nikos Parlavantzas, Theoretical and practical study, Département Informatique L3, Insa Rennes.
- Bachelor: Nikos Parlavantzas, Networks, Département Informatique L3, Insa Rennes.
- Bachelor: Nikos Parlavantzas, Multi-core architectures, Département Informatique L3, Insa Rennes.
- Bachelor: Jean-Louis Pazat, Introduction to programming L1, Département STPI, INSA de Rennes.
- Bachelor: Jean-Louis Pazat, High Performance Computing, Département Informatique L3, Insa Rennes.
- Bachelor: Jean-Louis Pazat, High Performance Computing, 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.
- Bachelor: Cédric Tedeschi, Cloud and networks, L3, Univ. Rennes 1.
- 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, 25 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, Operating Systems, M1, INSA Rennes.
- Master: Nikos Parlavantzas, Parallel programming, M1, INSA Rennes.
- Master: Nikos Parlavantzas, Big Data Storage and Processing, M2, INSA Rennes.
- Master: Nikos Parlavantzas, NoSQL, M2, Statistics for Smart Data, ENSAI, Bruz.
- Master: Nikos Parlavantzas, 4th-year Project, 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, Advanced Cloud Infrastructures, M2, Univ. Rennes 1.
- Master: Martin Quinson, Préparation à l'Agrégation d'Informatique (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, Advanced Cloud Infrastructures, M2, Univ. Rennes 1.
- Master: Shadi Ibrahim, Cloud Computing and Hadoop Technologies, 36hETD, M2 : Statistics for Smart Data, ENSAI, Bruz.
- Master: Shadi Ibrahim, Smart City Services: From applications to Infrastructures (SCS), 18hETD, M2 , Univ. Rennes 1.
- Master: Shadi Ibrahim, Cloud and Big data, 25hETD, M1, ENS Rennes, Bruz.
- Master: Shadi Ibrahim, Cloud1 (Map-Reduce), 17.5hETD, M2 , IMT-Atlantique, Nantes.
- Master: Shadi Ibrahim, Distributed Big Data, 36hETD, M2, ENSAI, Bruz.

11.2.2 Supervision

- PhD in progress: Clément Courageux-Sudan, “Reducing the energy consumption of Internet of Things”, started in October 2020, supervised by Anne-Cécile Orgerie and Martin Quinson.
- 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: Emmanuel Gnibga, “Modeling and optimizing edge computing infrastructures and their electrical system”, started in November 2021, supervised by Anne-Cécile Orgerie and Anne Blavette.
- PhD in progress (co-tutorship): Yasmina Bouzenn, “Energy-efficient, fault-tolerance mechanisms for containerized cloud applications”, started in January 2018, supervised by Djawida Dib (Tlemcen University, Algeria), Fedoua Lahfa (Tlemcen University, Algeria), Christine Morin, and Nikos Parlavantzas.
- PhD in progress: Davaadorj Battulga, “Creating scalable data pipelines for processing IoT data over fog computing infrastructures”, started in September 2018, supervised by Cédric Tedeschi and Daniele Miorandi (U-Hopper).
- PhD in progress: Paulo Souza, "stateful container migration in geo-distributed environments", started in December 2018, supervised by Guillaume Pierre and Daniele Miorandi (U-Hopper srl, Italy).
- PhD in progress: Chih-Kai Huang: "Decentralized control in large-scale fog computing infrastructures", started in November 2021, supervised by Guillaume Pierre.

11.2.3 Juries

- Guillaume Pierre was chairperson of the HDR defense of Guillaume Doyen (UTC Compiègne), January 12th 2021
- Guillaume Pierre was reviewer of the PhD defense of Alessio Pagliari (université Côte d’Azur), February 17th 2021
- Guillaume Pierre was a reviewer of the PhD defense of Li Wu (TU Berlin), October 18th 2021
- Guillaume Pierre was a member of the PhD committee of Mehdi Belkhiria (université de Rennes 1), October 25th 2021
- Guillaume Pierre was chairperson of the PhD defense of Alex Auvolat (Université de Rennes 1), November 2nd 2021
- Anne-Cécile Orgerie was a reviewer of the PhD defense of François Bouchaud (Université de Lille), January 30 th 2021.
- Anne-Cécile Orgerie was a member of the PhD committee of Alexandre Rio (Université Rennes 1), February 26th 2021.
- Anne-Cécile Orgerie was a member of the PhD committee of Youssef Ait El Mahjoub (Université Paris-Saclay), March 18th 2021.
- Anne-Cécile Orgerie was a reviewer of the PhD committee of Zakaria Ournani (Université de Lille), November 8th 2021.
- Anne-Cécile Orgerie was a member of the PhD defense of Andrea Segalini (Université Côte d’Azur), November 29th 2021.
- Nikos Parlavantzas was a member of the PhD committee of Antonios Papaioannou, University of Crete, October 22nd 2021

- Cédric Tedeschi was a member of the PhD committee of Hamidreza Arkian (université de Rennes 1), December 7th 2021.
- Cédric Tedeschi was a member of the PhD committee of Jonathan Sid-Othmane (Sorbonne université), December 13th 2021.
- Cédric Tedeschi was a member of the PhD committee of Zeina Houmani (Ecole normale supérieure de Lyon), December 16th 2021.

11.3 Popularization

11.3.1 Internal or external Inria responsibilities

n/a

11.3.2 Articles and contents

- Anne-Cécile Orgerie was interviewed for the radio program La méthode scientifique by Antoine Beauchamp with Guillaume Pitron, France Culture, 22nd September 2021.

11.3.3 Education

- “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 2021. The program is currently supported by: Fondation Blaise Pascal, ED MathSTIC, ENS de Rennes, Université Rennes 1, Fondation Rennes 1 and Alstom.
- Cédric Tedeschi is responsible for the Cloud and Network Infrastructures (CNI) Master program, hosted by University of Rennes 1 and part of the European Institute of Innovation & Technology (EIT) Digital Master program. He is also responsible for the first year of the Cloud et Réseaux master program, also hosted by University of Rennes 1.

12 Scientific production

12.1 Major publications

- [1] H. Casanova, A. Legrand, M. Quinson and 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, Nov. 2018, pp. 1–10. URL: <https://hal.inria.fr/hal-01891513>.

12.2 Publications of the year

International journals

- [2] A. Guermouche and A.-C. Orgerie. ‘Thermal design power and vectorized instructions behavior’. In: *Concurrency and Computation: Practice and Experience* 34.2 (Jan. 2022), e6261. DOI: [10.1002/cpe.6261](https://doi.org/10.1002/cpe.6261). URL: <https://hal.archives-ouvertes.fr/hal-03185821>.
- [3] L. M. Pham, N. Parlavantzis, H.-H. Le and Q. H. Bui. ‘Towards a Framework for High-Performance Simulation of Livestock Disease Outbreak: A Case Study of Spread of African Swine Fever in Vietnam’. In: *Animals* 11.9 (Sept. 2021), pp. 1–19. DOI: [10.3390/ani11092743](https://doi.org/10.3390/ani11092743). URL: <https://hal.inria.fr/hal-03360693>.

- [4] A. C. Zhou, W. Xue, Y. Xiao, B. He, S. Ibrahim and R. Cheng. ‘Taming System Dynamics on Resource Optimization for Data Processing Workflows: A Probabilistic Approach’. In: *IEEE Transactions on Parallel and Distributed Systems* 33.1 (22nd June 2021), pp. 231–248. DOI: [10.1109/TPDS.2021.3091400](https://doi.org/10.1109/TPDS.2021.3091400). URL: <https://hal.inria.fr/hal-03363176>.

International peer-reviewed conferences

- [5] H. Arkian, G. Pierre, J. Tordsson and E. Elmroth. ‘Model-based Stream Processing Auto-scaling in Geo-Distributed Environments’. In: *ICCCN 2021 - 30th International Conference on Computer Communications and Networks*. Athens, Greece, 19th July 2021, pp. 1–11. URL: <https://hal.inria.fr/hal-03206689>.
- [6] D. Battulga, D. Miorandi and C. Tedeschi. ‘SpecK: Composition of Stream Processing Applications over Fog Environments’. In: *DAIS 2021 - 21st International Conference on Distributed Applications and Interoperable Systems*. Valetta, Malta, 9th June 2021, pp. 38–54. DOI: [10.1007/978-3-030-78198-9_3](https://doi.org/10.1007/978-3-030-78198-9_3). URL: <https://hal.inria.fr/hal-03259975>.
- [7] C. Courageux-Sudan, A.-C. Orgerie and M. Quinson. ‘Automated performance prediction of microservice applications using simulation’. In: *MASCOTS 2021 - International Symposium on the Modeling, Analysis, and Simulation of Computer and Telecommunication Systems*. virtually, France, 3rd Nov. 2021, pp. 1–8. URL: <https://hal.archives-ouvertes.fr/hal-03389508>.
- [8] H. Fan, S. Bian, S. Wu, S. Jiang, S. Ibrahim and H. Jin. ‘Gear: Enable Efficient Container Storage and Deployment with a New Image Format’. In: *ICDCS’21 - 41st IEEE International Conference on Distributed Computing Systems*. Washington DC, United States, 7th July 2021, pp. 1–11. URL: <https://hal.inria.fr/hal-03363179>.
- [9] M. Farhadi, G. Pierre and D. Miorandi. ‘Towards automated privacy compliance checking of applications in Cloud and Fog environments’. In: *FiCloud 2021 - 8th International Conference on Future Internet of Things and Cloud*. Rome / Virtual, Italy, 23rd Aug. 2021, pp. 1–8. URL: <https://hal.inria.fr/hal-03257252>.
- [10] A. Gougeon, B. Camus, F. Lemercier, M. Quinson, A. Blavette and A.-C. Orgerie. ‘Co-Simulation of Power Systems and Computing Systems using the FMI Standard’. In: *IM 2021 - IFIP/IEEE International Symposium on Integrated Network Management*. Bordeaux / Virtual, France, 17th May 2021, pp. 130–131. URL: <https://hal.archives-ouvertes.fr/hal-03217565>.
- [11] M. Laurent, E. Saillard and M. Quinson. ‘The MPI BUGS INITIATIVE: a Framework for MPI Verification Tools Evaluation’. In: *Correctness 2021: Fifth International Workshop on Software Correctness for HPC Applications*. St. Louis, United States, 19th Nov. 2021, pp. 1–9. URL: <https://hal.inria.fr/hal-03474762>.
- [12] D. Rosendo, A. Costan, G. Antoniu, M. Simonin, J.-C. Lombardo, A. Joly and P. Valduriez. ‘Reproducible Performance Optimization of Complex Applications on the Edge-to-Cloud Continuum’. In: *Cluster 2021 - IEEE International Conference on Cluster Computing*. Portland, OR, United States, 7th Sept. 2021. URL: <https://hal.archives-ouvertes.fr/hal-03310540>.
- [13] M. A. Tamiru, G. Pierre, J. Tordsson and E. Elmroth. ‘mck8s: An orchestration platform for geo-distributed multi-cluster environments’. In: *ICCCN 2021 - 30th International Conference on Computer Communications and Networks*. Athens, Greece, 19th July 2021, pp. 1–12. URL: <https://hal.inria.fr/hal-03205743>.
- [14] S. Tavonatti, D. Battulga, M. Farhadi, C. Caprini and D. Miorandi. ‘An experimental evaluation of the scalability of permissioned blockchains’. In: *FiCloud 2021 - 8th International Conference on Future Internet of Things and Cloud*. Rome / Virtual, Italy, 23rd Aug. 2021, pp. 1–8. URL: <https://hal.inria.fr/hal-03263551>.
- [15] A. T. Wonjiga, L. Rilling and C. Morin. ‘SLA Definition for Network Intrusion Detection Systems in IaaS Clouds’. In: *SAC 2021 - 36th ACM/SIGAPP Symposium on Applied Computing*. Virtual Event, Republic of Korea., South Korea: ACM, 22nd Mar. 2021, pp. 1–10. DOI: [10.1145/3412841.3441885](https://doi.org/10.1145/3412841.3441885). URL: <https://hal.inria.fr/hal-03085554>.

Conferences without proceedings

- [16] C. Elbaz, L. Rilling and C. Morin. ‘Automated Risk Analysis of a Vulnerability Disclosure Using Active Learning’. In: Proceedings of the 28th Computer & Electronics Security Application Rendezvous. Rennes, France, 16th Nov. 2021. URL: <https://hal.inria.fr/hal-03515662>.
- [17] A. Gougeon, B. Camus, A. Blavette and A.-C. Orgerie. ‘Impact of wired telecommunication network latency on demand-side management in smart grids’. In: IM 2021 - 17th IFIP/IEEE International Symposium on Integrated Network Management. Bordeaux / Virtual, France: IEEE, 17th May 2021, pp. 1–9. URL: <https://hal.archives-ouvertes.fr/hal-03217562>.
- [18] Y. Li, C. Li, A.-C. Orgerie and P. R. Parvédy. ‘WSGP: A Window-based Streaming Graph Partitioning Approach’. In: CCGrid 2021 - 21th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing. Melbourne, Australia: IEEE, 10th May 2021, pp. 1–10. URL: <https://hal.archives-ouvertes.fr/hal-03217558>.
- [19] K. Rao Vaddina, L. Lefèvre and A.-C. Orgerie. ‘Experimental Workflow for Energy and Temperature Profiling on HPC Systems’. In: ISCC 2021 - IEEE Symposium on Computers and Communications. Athens, Greece: IEEE, 5th Sept. 2021, pp. 1–7. URL: <https://hal.archives-ouvertes.fr/hal-03335184>.

Doctoral dissertations and habilitation theses

- [20] H. Arkian. ‘Resource Management for Data Stream Processing in Geo-Distributed Environments’. Université de Rennes 1, 7th Dec. 2021. URL: <https://hal.inria.fr/tel-03477454>.
- [21] C. Elbaz. ‘Reacting to “n-day” vulnerabilities in information systems’. Université Rennes 1, 30th Mar. 2021. URL: <https://tel.archives-ouvertes.fr/tel-03350455>.
- [22] M. Farhadi. ‘Automated application privacy compliance checking in distributed Fog environments’. Université de Rennes 1, 5th Nov. 2021. URL: <https://hal.inria.fr/tel-03420748>.
- [23] L. Guegan. ‘Scalable end-to-end models for the time and energy performance of Fog infrastructures’. École normale supérieure de Rennes, 29th Jan. 2021. URL: <https://tel.archives-ouvertes.fr/tel-03134046>.
- [24] M. A. Tamiru. ‘Automatic resource management in geo-distributed multi-cluster environments’. Université Rennes 1, 20th Sept. 2021. URL: <https://tel.archives-ouvertes.fr/tel-03435237>.

Reports & preprints

- [25] Y. Bouizem, D. Dib, N. Parlavantzas, C. Morin and F. Lahfa. *Request Replication for FaaS Fault Tolerance*. RR-9444. Inria, 4th Jan. 2022, 19 p. URL: <https://hal.inria.fr/hal-03510322>.

Other scientific publications

- [26] D. Charlet, K. Desnos, M. Dardaillon, A. Ferrari, C. Ferrari, N. Gac, J. F. Nezan, F. Orioux, S. Prunet, M. Quinson, F. Suter, C. Tasse and C. Dumez-Viou. *Dataflow Algorithm aRchitecture co-design of SKA pipeline for Exascale Radio Astronomy*. Virtual conference, France, 24th June 2021. URL: <https://hal.archives-ouvertes.fr/hal-03233205>.

12.3 Other

Scientific popularization

- [27] G. Pierre. ‘Le « fog computing » est l’avenir du cloud – en plus frugal et plus efficace’. In: *The Conversation* (30th Aug. 2021), pp. 1–4. URL: <https://hal.inria.fr/hal-03336292>.

12.4 Cited publications

- [28] D. Battulga, D. Miorandi and C. Tedeschi. ‘FogGuru: a Fog Computing Platform Based on Apache Flink’. In: *23rd Conference on Innovation in Clouds, Internet and Networks (ICIN 2020)*. Paris, France, Feb. 2020. URL: <https://hal.inria.fr/hal-02463206>.
- [29] R.-A. Cherrueau, M. Simonin and 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, Apr. 2018, pp. 336–341. DOI: [10.1109/INFCOMW.2018.8407024](https://doi.org/10.1109/INFCOMW.2018.8407024). URL: <https://hal.inria.fr/hal-01916484>.
- [30] J. Darrous, T. Lambert and S. Ibrahim. ‘On the Importance of Container Image Placement for Service Provisioning in the Edge’. In: *ICCCN 2019 - 28th International Conference on Computer Communication and Networks*. Valencia, Spain, July 2019. DOI: [10.1109/ICCCN.2019.8846920](https://doi.org/10.1109/ICCCN.2019.8846920). URL: <https://hal.inria.fr/hal-02134507>.
- [31] C. Elbaz, L. Rilling and C. Morin. ‘Automated Keyword Extraction from ”One-day” Vulnerabilities at Disclosure’. In: *NOMS 2020 - IEEE/IFIP Network Operations and Management Symposium*. Budapest, Hungary: IEEE, Apr. 2020, pp. 1–9. URL: <https://hal.inria.fr/hal-02506364>.
- [32] D. Rosendo, P. Silva, M. Simonin, A. Costan and G. Antoniu. ‘E2Clab: Exploring the Computing Continuum through Repeatable, Replicable and Reproducible Edge-to-Cloud Experiments’. In: *Cluster 2020 - IEEE International Conference on Cluster Computing*. Kobe, Japan, Sept. 2020, pp. 1–11. DOI: [10.1109/CLUSTER49012.2020.00028](https://doi.org/10.1109/CLUSTER49012.2020.00028). URL: <https://hal.archives-ouvertes.fr/hal-02916032>.
- [33] P. Souza Junior, D. Miorandi and G. Pierre. ‘Stateful Container Migration in Geo-Distributed Environments’. In: *CloudCom 2020 - 12th IEEE International Conference on Cloud Computing Technology and Science*. Bangkok, Thailand: IEEE, Dec. 2020, pp. 1–8. URL: <https://hal.inria.fr/hal-02963913>.
- [34] A. C. Zhou, Y. Xiao, B. He, S. Ibrahim and R. Cheng. ‘Incorporating Probabilistic Optimizations for Resource Provisioning of Data Processing Workflows’. In: *ICPP 2019 - 48th International Conference on Parallel Processing*. Kyoto, Japan: ACM Press, Aug. 2019, pp. 1–10. DOI: [10.1145/3337821.3337847](https://doi.org/10.1145/3337821.3337847). URL: <https://hal.inria.fr/hal-02389078>.