RESEARCH CENTRE

Rennes - Bretagne Atlantique

IN PARTNERSHIP WITH:

CNRS, Institut national des sciences appliquées de Rennes, Université Rennes 1, École normale supérieure de Rennes

2020

ACTIVITY REPORT

Project-Team

MYRIADS

Design and Implementation of Autonomous Distributed Systems

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

DOMAIN

Networks, Systems and Services, Distributed Computing

THEME

Distributed Systems and middleware

Contents

Project-Team MYRIADS

Creation of the Team: 2010 January 01, updated into 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, from Oct 2020]
- Anne-Cécile Orgerie [CNRS, Researcher, HDR]

Faculty Members

- Guillaume Pierre [Team leader, Univ de Rennes I, Professor]
- Marin Bertier [INSA Rennes, Associate Professor]
- Nikolaos 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 Fellows

- Aymen Jlassi [Univ de Rennes I, until Jun 2020]
- Francois Lemercier [CNRS, from Mar 2020]
- Kameswar Rao Vaddina [Inria, until Oct 2020]

PhD Students

- Hamidreza Arkian [Univ de Rennes I]
- Davaadorj Battulga [U-Hopper - Italy]
- Mehdi Belkhiria [Univ de Rennes I]
- Dorra Boughzala [Inria, until Nov 2020]
- Yasmina Bouizem [COLLECTIVITE LOCALE]
- Clement Courageux-Sudan [Inria, from Oct 2020]
- Clement El Baz [Inria]
- Ali Fahs [Univ de Rennes I]
- Mozhdeh Farhadi [U-Hopper - Italy]
- Adrien Gougeon [École normale supérieure de Rennes]
- Loic Guegan [École normale supérieure de Rennes, until Feb 2021]
- Paulo Rodrigues De Souza Junior [Univ de Rennes I]
- Mulugeta Tamiru [Elastisys - Sweden]

Technical Staff

- Ehsan Azimi [Inria, Engineer]
- Thomas Lambert [Inria, Engineer, from Oct 2020]
- Nena Markovic [Univ de Rennes I, Engineer, from Dec 2020]
- Sébastien Philippot [Inria, Engineer, until Aug 2020]
- Matthieu Simonin [Inria, Engineer, SED RENNES]

Interns and Apprentices

- Clement Courageux-Sudan [Inria, from Feb 2020 until Jul 2020]
- Eloise Dagorn [Inria, from Jun 2020 until Aug 2020]

Administrative Assistants

- Cecile Bouton [Inria, from May 2020 until Nov 2020]
- Angélique Jarnoux [Inria, from Apr 2020]
- Caroline Tanguy [Inria, until May 2020]

Visiting Scientist

- Klervie Tocze [Université de Linköping - Suède, from Jun 2020 until Jul 2020]

External Collaborators

- Anne Blavette [CNRS]
- Julie Montegu [Univ de Rennes I]

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 users [1] 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.

[1] According to World Stats, there are 4.94 billion Internet users i.e. more than 60% of the total world population in December 2020 http://www.internetworldstats.com/stats.htm.

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 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 *Linkedln* 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 acccount 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 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 extend 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 sprit 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

1. Anne-Cécile Orgerie successfully defended her HDR on November 20th 2020;
2. Nikos Parlavantzas successfully defended his HDR on June 15th 2020;
3. Shadi Ibrahim (previously a member of the STACK team) joined the Myriads team.

6.1 Awards

1. Anne-Cécile Orgerie received the Bronze Medal of CNRS for her outstanding contributions to energy-efficient computing infrastructures;
2. Deb Agarwal was awarded a Doctorate Honoris Causa by the University of Rennes 1.
3. Shadi Ibrahim received the prestigious 2020 IEEE TCSC Award for Excellence in Scalable Computing (Middle Career Researcher) for his contributions on scalable and efficient data management in large-scale distributed systems.
4. Shadi Ibrahim received IEEE Outstanding Leadership Award as General Chair of the 6th IEEE International Conference on Smart Data as a part of the 2020 IEEE Cybermatics Congress (Blockchain/ CPSCom/ GreenCom/ Things/ SmartData-2020), 4-6 November 2020.

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/

Contacts: Arnaud Legrand, Martin Quinson, Frédéric Suter

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

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/

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

Contacts: Nikolaos Parlavantzas, Christine Morin

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 relyies 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

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

Contacts: Matthieu Simonin, Cédric Tedeschi, Christine Morin

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

Author: Clement El Baz

Contacts: Clement El Baz, Christine Morin, Louis Rilling

Participant: Clement El Baz

8 New results

8.1 Scaling Clouds

8.1.1 Efficient Docker container deployment in fog environments

Participants Arif Ahmed, Guillaume Pierre.

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

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

We finalized this line of research with the latest research contributions from Arif Ahmed's PhD thesis: in particular we showed how the usage of Checkpoint/Restart technologies can help to speed up the boot time of Docker containers after beeing started [6]. This technique effectively reduces the startup phase with speedups between 1x (no speedup) and 60x (in extreme cases). This contribution, and previous ones from the previous years, are also presented in Arif Ahmed's PhD thesis [26].

This line of research is now completed and we do not expect further work on it in the coming years.

8.1.2 Fog computing platform design

Participants Ali Fahs, Guillaume Pierre, Paulo Souza, Mulugeta Tamiru.

There does not yet exist any reference platform for fog computing platforms. We therefore investigated how Kubernetes could be adapted to support the specific needs of fog computing platforms. In particular we focused on the problem of redirecting end-user traffic to a nearby instance of the application. When different users impose various load on the system, any traffic routing system must necessarily implement a tradeoff between proximity and fair load-balancing between the application instances. We demonstrated how 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.

First, we designed and implemented location-aware pod scheduling strategies within Kubernetes, such that every user of the platform has access to at least one nearby replica reachable within pre-defined latency bounds [16]. We further extended this work by integrating an improved version of these pod placement algorithms with auto-scaling algorithms such that the system maintains the user-perceived performance within its pre-defined bounds while avoiding to overload resources and while using no more resources than strictly necessary [15]. These contributions, and other ones from previous years, are also presented in Ali Fahs's PhD thesis [27].

We continued the work on Kubernetes by performing an extensive experimental study of the different auto-scalers that are broadly available in this platform. Specifically we compared the performance of the Kubernetes Cluster Autoscaler (which selects resources from a single pool) and the new CA-NAP node auto-provisioning capability. We showed that, overall, CA-NAP outperforms CA and that autoscaling performance depends mainly on the composition of the workload [24].

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 however showed that, in the current implementation of KubeFed, delays and transient network failures coupled with static configuration, including the default configuration parameter values, can lead to instability of application deployments in Kubernetes Federation, making applications unavailable for long periods of time. Leveraging on the benefits of configuration tuning, we proposed a feedback controller to dynamically adjust the concerned configuration parameter to improve the stability of application deployments without slowing down the detection of hard failures [23].

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 [21]. 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.

Some of these contributions, and others developed by our partners in the FogGuru H2020 project, have been integrated in a unified fog computing platform which has been deployed in València (Spain) as part of the FogGuru "Living Lab" [2]. This fog platform based on the LoRaWAN wireless networking technology has first been applied to the management of water consumption data produced by smart water meters[2] and they are now being extended at a larger scale to manage IoT data produced within the Marina in Valencia. The FogGuru project organizes a Hackathon in March 2021 where external participants will be able to experiment first-hand with this platform.

[2] http://www.fogguru.eu/living-lab/

8.1.3 Community Clouds

Participants Jean-Louis Pazat, Bruno Stevant.

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

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

This year we consolidated the results of the optimization algorithm using our prototype (5 "boxes" installed in different home locations connected by fiber or ADSL). Results showed that the algorithm sucessfuy handles the variation of the network available bandwidth by modifying the deployment of applications [22] As it was not possible to find a predictive model of this variation during a day, we used a regular monitoring of the response time of applications.

8.1.4 Geo-distributed data stream processing

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

We investigated a decentralized scaling mechanism for stream processing applications where the different operators composing the processing topology are able to take their own scaling decisions independently, based on local information. While these works were validated by several publications in 2019, our recent work extended this work with a more generic approach to the inherent synchronization problems induced by decentralized scaling. More precisely, a group mutual exclusion problem was formulated and a new algorithm was proposed [9].

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 such a model to represent the performance of a single stream processing operator [7]. This model is very accurate with predictions $\pm 2\%$ of the actual values even in the presence of heterogeneous network latencies. Individual operator models can be composed together and, after the initial calibration of a first operator, a reasonably accurate model for other operators can be derived from a single measurement only. We continued this work with the design of 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. A publication on this topic is currently in preparation.

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. This work started in November 2020 (with several months of delay due to the Covid-19 crisis), and we expect to obtain the first results during the coming year.

In the context of the ANR KerStream project, we have been working on implementing and evaluating a performance-aware task scheduling for stream data applications which aims at placing operators considering the network heterogeneity and node capacities. The initial design of the scheduler was presented in [31]. In addition, we present an empirical study to evaluate the impact of two widely used preemption techniques (i.e., kill and pause-resume) on fairness in multi-tenant data-intensive clusters [29].

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 [8].

In the context of the ANR Sesame project, we invistigated strategies for scheduling multiple stream processing applications. Based on the assumption that applications, dynamically submitted to a platform share some operational similarities, we devised algorithms based on the reuse of the data produced by running applications [18]. We also helped scaling a maching learning application developed at IMT Atlantique that detects dynamically abnormal vessel trajectories [19].

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 plaforms 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 integrated an alternative fault-tolerance mechanism, namely active-standby failover, into an open-source FaaS framework based on Kubernetes, namely the Fission framework. In 2020, we performed extensive experiments to compare this mechanism to the typical retry-based approach followed by current FaaS frameworks. The experiments showed that our approach outperforms the typical approach in terms of response time and availability while incurring a limited overhead in resource consumption [11].

8.2 Greening Clouds

8.2.1 Energy Models

Participants Dorra Boughzala, Loic Guegan, Anne-Cécile Orgerie, Martin Quinson.

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

To tackle the energy consumption challenge while continuing to provide tremendous performance, the HPC community have rapidly adopted GPU-based systems. Today, GPUs have became the most prevailing components in the massively parallel HPC landscape thanks to their high computational power and energy efficiency. Modeling the energy consumption of applications running on GPUs has gained a lot of attention for the last years. However, the HPC community lacks simple yet accurate simulators to predict the energy consumption of general purpose GPU applications. We proposed a simple and lightweight energy model that we implemented using the open-source framework SimGrid. Our proposed model is validated across a diverse set of CUDA kernels and on two different NVIDIA GPUs (Tesla M2075 and Kepler K20Xm). This work was published in [10].

While distributed computing infrastructures can provide infrastructure-level techniques for managing energy consumption, application-level energy consumption models have also been developed to support energy-efficient scheduling and resource provisioning algorithms. We analyzed the accuracy of a widely-used application-level model that has been developed and used in the context of scientific workflow executions. To this end, we profiled two production scientific workflows on a distributed platform instrumented with power meters. Our analysis showed that power consumption is not linearly related to CPU utilization and that I/O operations significantly impact power, and thus energy, consumption. We also proposed a power consumption model that accounts for I/O operations, including the impact of waiting for these operations to complete, and for concurrent task executions on multi-socket, multi-core compute nodes. We implemented our proposed model as part of a simulator that allows us to draw direct comparisons between real-world and modeled power and energy consumption. Our model improves accuracy by about two orders of magnitude when compared to the traditional models used in the energy-efficient workflow scheduling literature. This work has been done in collaboration with the University of Southern California and the University of Hawai [5].

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 2020, 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.

8.2.3 Exploiting renewable energy in distributed clouds

Participants Adrien Gougeon, 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. The Fog computing paradigm represents a distributed architecture closer to the end-user. We explored energy-efficient Fog architectures considering the integration of renewable energy. Our simulation results, based on real traces, showed that the intrinsic low computing capability of the nodes in a Fog context makes it harder to exploit renewable energy. In addition, the share of the consumption from the communication network between the computing resources increases in this context [17].

8.2.4 Smart Grids

Participants Anne Blavette, Adrien Gougeon, François Lemercier, Anne-Cécile Orgerie, Martin Quinson.

Smart grids allow to efficiently perform demand-side management in electrical grids in order to increase the integration of fluctuating and/or intermittent renewable energy sources in the energy mix. In this work, we consider the computing infrastructure that controls the smart grid. This infrastructure comprises communication and computing resources to allow for a smart management of the electrical grid. In particular, we study the influence of communication latency over a shedding scenario on a small-scale electrical network. We show that depending on the latency some shedding strategies are not feasible [12].

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.

In the past years we proposed a verification method for security monitoring SLOs [34] and we have then studied a methodology to define security monitoring SLOs that are at the same time relevant for the tenant, achievable for the provider, and verifiable. In 2020, the methodology to define security monitoring SLOs was submitted for publication and accepted at the forthcoming SAC 2021 conference [25].

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 winwow 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. Our previous work on a first step of mitigation, which consists in extracting keywords characterizing the products impacted by a newly published vulnerability despite the sole availability of a free-form text description, was published at the NOMS 2020 international conference [13]. Assuming that an automatic vulnerability processing system has some knowledge of the considered information system, this first contribution is an important step to discriminate between pointless vulnerabilities, that do not impact products used in the information system, and vulnerabilities that require further processing before making a decision.

In 2020 we proposed another feature extraction process to further help to make a decision about a newly published vulnerability based on its potential severity. The proposed process predicts the base Common Vulnerability Severity Scoring (CVSS) vector, again using only the free-form text description of the vulnerability. The CVSS vector is standard metadata added to vulnerability publications after manual analysis and is composed of a set of properties of the vulnerability that characterize technical means required to exploit it, the required strength for an attacker to exploit it, and the impact of an attacker successfully exploiting it. The base vector is the part of the vector that only contains intrinsic properties of the vulnerability, that depend neither on time after the publication nor on the context of a specific information system. The prediction process is based on linear regression fits of the base CVSS vectors of the previously published vulnerabilities, which enables immediate prediction at the publication time of the vulnerability as well as a good explicability of the prediction made. Experimental results show that the accuracy is good enough and that retraining the predictor using updated data is fast enough for the prediction process to be used in production. This work was published at the ARES 2020 international conference [14]. In future work the extracted features should be combined with a knowledge base of the information system to automatically trigger an adequate reaction in the information system.

Our results were published in [13, 14].

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 published a first paper on this topic with a methodology to analyze the security and privacy of fog computing platforms [4]. We further extended this work by demonstrating 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. A paper on this topic has been submitted for publication.

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, further improved the documentation. The work on SimGrid is fully integrated into the other research efforts of the Myriads team. This year we had two major contributions now integrated to the official SimGrid distribution. First, we added energy models for the wired communications as well as a new wifi model [28] to enable the studies that we envision in the team. Secondly, we added the ability to co-simulate IT systems with SimGrid and physical systems modeled with equational systems [12]. This work was developed to study the co-evolution of thermal systems or of the electric grid with the IT system.

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 [33]. This year, we pursued our effort on this topic, in collaboration with Thierry Jéron (EPI SUMO).

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.

Even if Mc SimGrid was initially written to verify asynchronous distributed applications such as MPI-based ones, we proposed this year an adaptation of this work to multithreaded applications written with the pthread interface [3]. This is particularly interesting to teach multithread programming, as the learners can experiment by themselves to detect and understand the flaws in their implementation. This work was conducted in collaboration with Northeastern University, USA, as part of the *Fog Rein* Inria associated team.

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 2020 we have built a first prototype of the interconnection between SimGrid and a VMM, based on the Qemu PC emulator running in full emulation mode. An internship topic was published to further evaluate the approach. In future work we will extend the approach to hardware-assisted virtualization and study the integration of TANSIV with the time-related tools available in malware analysis sandboxes.

8.4.4 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 [30] a possible solution to the above desiderata.

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

The tool is used in several articles (see https://discovery.gitlabpages.inria.fr/enoslib/theyuseit.html). In particular, in [20] the tool is used to build an ad hoc framework for studying FOG applications.

9 Partnerships and cooperations

9.1 International initiatives

9.1.1 Inria International Labs

Inria@SiliconValley

Associate Teams involved in the Inria International Lab:

Hermes

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

Duration: 2019 - 2021

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.

9.1.2 Inria Associate Teams not involved in an IIL

FogCity

Title: QoS-aware Resource Management for Smart Cities

Duration: 2018 - 2020

Coordinator: Guillaume Pierre

Partners:

- IIT Kharagpur (India)

- IIT Kanpur (India)
- Inria (France)

Inria contact: Guillaume Pierre

Participants: Ayan Mondal, Nikos Parlavantzas, Guillaume Pierre

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

Unfortunately the sanitary situation in 2020 did not allow us to achieve the expected student and staff mobility scheme. Nevertheless, thanks to the collaboration established within the team, we developed mechanisms based on game theory to assign resources to competing applications in a fog computing platform. The objective of those mechanisms is to satisfy user preferences while maximizing resource utilisation. A paper on this topic is under preparation.

FogRein

Title: Steering Efficiency for Distributed Applications

Duration: 2019 – 2021

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.

9.1.3 Inria international partners

Informal international partners

Huazhong university of Science and Technology (HUST) We collaborate on container storage management.

National University of Singapore (NUS) We collaborate on resource optimisation for data processing workflows.

Shenzhen University We collaborate on resource optimization for data processing workflows and optimizing graph processing in geo-distributed data-centers.

University of California, Santa Barbara We collaborate on optimizing graph processing in geo-distributed data-centers.

University of Tlemcen, Algeria We collaborate on energy-efficient, fault-tolerant resource and application management in containerized clouds.

9.1.4 Participation in other international programs

Managing epidemic simulation applications in the cloud

Title: Managing epidemic simulation applications in cloud environments

Duration: 2020 - 2022

Project type: Scientific cooperation project, French Embassy in Vietnam

Coordinator: Nikos Parlavantzas

Partners:

- Vietnam National University (Vietnam)

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 improve the DiFFuSE framework and apply it to epidemic simulation applications developed in the Vietnam National University. DiFFuSE is a service-based framework [32] developed by Myriads that enables simulation applications to fully exploit cloud platforms. Unfortunately the sanitary situation in 2020 has not allowed us to achieve the planned mobility scheme so far.

Optimization Framework for Future (5G and beyond) Green Wireless Networks

Title: Optimization Framework for Future (5G and beyond) Green Wireless Networks

Duration: 2020 - 2022

Project type: PHC Peridot

Coordinator: Anne-Cécile Orgerie

Partners:

- University of Engineering and Technology, Taxila, Pakistan
- University of Lyon 1

Summary: The project aims to optimize the 5G networks by using energy harvesting capable base stations and on demand resource allocation while meeting quality of service requirements.

9.2 International research visitors

9.2.1 Visits of international scientists

- Subhas Chandra Misra and Abhishek Shukla
  - Date: 13-20 January 2020
  - Institution: IIT Kanpur (India)
  - Participants: Guillaume Pierre, Nikos Parlavantzas

- Gene Cooperman
  - Date: 19-25 January 2020
  - Institution: College of Computer and Information Science, Northeastern University (United States)
  - Participants: Martin Quinson, Anne-Cécile Orgerie

9.3 European initiatives

9.3.1 FP7 & H2020 Projects

H2020 MSCA FogGuru

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

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

Program: H2020 MSCA ITN EID

Duration: September 2017 - November 2021

Coordinator: Guillaume Pierre (UR1)

Participants:

- University of Rennes 1, France (coordinator)
- Technische Universität Berlin, Germany
- Elastisys AB, Sweden
- U-Hopper srl, Italy
- EIT Digital Rennes, France
- Las Naves, Spain

Abstract: FogGuru is a doctoral training project which aims to 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.

9.3.2 Collaborations in European programs, except FP7 and H2020

CHIST-ERA DiPET (2020-2023)

Participants Ayan Mondal, Guillaume Pierre.

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

Program: CHIST-ERA

Duration: 2020 - 2023

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

Partners:

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

Abstract: 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.

9.4 National initiatives

ADEME RennesGrid (2017-2020)

Participants Anne Blavette, Anne-Cécile Orgerie, Martin Quinson.

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

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.

INRIA IPL Discovery (2015-2019)

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

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

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

INRIA IPL Hac Specis (2016-2020)

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

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

We are pursuing a collaboration with E. Saillard (CR Inria Bordeaux), initiated during this project. Our goal is to evaluate the practical usability of several tools that allow to assess the correction of MPI programs.

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

SESAME ASTRID project (2016-2019)

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

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

CNRS Momentum RI/RE (2019-2022)

Participants Anne-Cécile Orgerie.

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.

KerStream (ANR)

Participants Shadi Ibrahim.

The KerStream project (Big Data Processing: Beyond Hadoop!) is an ANR JCJC (Young Researcher) project (ANR-16-CE25-0014-1) running for 48 months (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.

10 Dissemination

10.1 Promoting scientific activities

10.1.1 Scientific events: organisation

General chair, scientific chair

- Shadi Ibrahim was general co-chair for the 6th IEEE International Conference on Smart Data (SmartData-2020), November 02 - 06, 2020.

Member of the organizing committees

- Guillaume Pierre is a member of the steering committee of the ACM Middleware conference.
- Anne-Cécile Orgerie was co-organizer of the GDR RSD and ASF Winter School on Distributed Systems and Networks in Feburary 2020.
- Anne-Cécile Orgerie was member of the scientific committee of the annual congress of SIF (Société Informatique de France) in February 2020.

10.1.2 Scientific events: selection

Chair of conference program committees

- Guillaume Pierre was co-chair of the "Intelligent Systems and Infrastructure" track of TheWebConf 2020.
- Guillaume Pierre was co-chair of the ALGOCLOUD 2020 conference.
- Shadi Ibrahim was program chair of the 5th International Parallel Data Systems Workshop (PDSW@SC2020), November 12, 2020.
- Anne-Cécile Orgerie was Doctoral Symposium co-chair for CCGrid 2020: IEEE/ACM International Symposium in Cluster, Cloud, and Grid Computing.
- Anne-Cécile Orgerie was Demo and Poster co-chair for ICDCS 2020: IEEE International Conference on Distributed Computing Systems.

Member of the conference program committees

- Guillaume Pierre was a member of the program committees of IEEE IPDPS 2020 and IEEE/ACM CCGrid 2020.
- Shadi Ibrahim was a member of the program committees of IEEE ISPA 2020.
- Nikos Parlavantzas was a member of the program committees of IEEE/ACM UCC 2020, IEEE CloudCom 2020, and IEEE ISPDC 2020
- Cédric Tedeschi was a member of the program committees of ICPP 2020.
- Anne-Cécile Orgerie was a member of the program committees of CCGrid 2020, IC2E 2020, IPDPS 2020, ICPP 2020, Cluster 2020, AlgoTel 2020, Compas 2020, and JCAD 2020.

10.1.3 Journals

Member of the editorial boards

- Shadi Ibrahim is a guest co-editor of the Journal of Grid Computing - From Grids to Cloud Federations. Special Issue on Orchestration of computing resources in the Cloud-to-Things continuum.
- Shadi Ibrahim is a young Associate Editor of the Springer Frontiers of Computer Science journal.
- Anne-Cécile Orgerie is member of the editorial board of IJDSN: International Journal of Distributed Sensor Networks.
- Anne-Cécile Orgerie is member of the editorial board of TPDS: IEEE Transactions on Parallel and Distributed Systems.

Reviewer - reviewing activities

- Shadi Ibrahim was a reviewer of the IEEE Transactions on Parallel and Distributed Systems, ACM Transactions on Internet Technology.
- Marin Bertier was a reviewer of the Journal of Supercomputing

10.1.4 Invited talks

- "The FogGuru Living Lab: A LoRa-enabled fog computing testbed." Invited presentation at the French Journées Cloud 2020 [3], 24 November 2020.
- Anne-Cécile Orgerie gave a talk on Green computing at Big Day, the annual scientific day of the bioinformatics Master of University of Rennes 1, November 2020.

[3] https://journeescloud20.sciencesconf.org/resource/page/id/11

10.1.5 Leadership within the scientific community

- Anne-Cécile Orgerie was vice-chair of ASF, the French Chapter of ACM SIGOPS until end of 2020.
- Anne-Cécile Orgerie is director of the CNRS service group on ICT environmental impact (GDS EcoInfo) since Feburary 2020.
- Anne-Cécile Orgerie is chief scientist for the Rennes site of Grid'5000.

10.1.6 Scientific expertise

- Guillaume Pierre expertised two PhD funding requests submitted to Region Grand Est.
- Guillaume Pierre was a member of the selection committee for a full professor at TU Wien (Austria).
- Anne-Cécile Orgerie expertised projects submitted to the i-Nov concours organized by BPI France.
- Anne-Cécile Orgerie was a member of the selection committees for an assistant professor position at ENSIMAG and Inria research scientist positions (CRCN and ISFP) at the Nancy site in 2020.
- Jean-Louis Pazat is the coordinator of the expertises of bi-national cooperations in the domain of numerical sciences for the Ministry of Education, Research and Innovation.

10.1.7 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.

10.2 Teaching - Supervision - Juries

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 Mathematiques L3, Insa Rennes.
- Bachelor: Guillaume Pierre, Systèmes Informatiques, L3 MIAGE, Univ. Rennes 1.
- Bachelor: Guillaume Pierre, Systèmes d'exploitation, L3 Informatique, Univ. Rennes 1.

- Bachelor: Martin Quinson, Architecture et Systèmes, 60 hETD, L3 Informatique, ENS Rennes.
- Bachelor: Martin Quinson, Pedagogy, 15 hETD, L3 Informatique, ENS Rennes.
- Master: Marin Bertier, Operating Systems, Département Informatique M1, INSA de Rennes
- Master: Marin Bertier, Distributed systems, Département Informatique M2, INSA de Rennes
- Master: Anne-Cécile Orgerie, Cloud & Big Data, 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 de Science Industrielle (Programming and Software Engineering, 20h ETD; Operating Systems and C programming, 20 hETD; Networking, 20h ETD), ENS Rennes.
- Master: Martin Quinson, Scientific Outreach, M2, 30 hEDT, ENS Rennes.
- Master: Cédric Tedeschi, Concurrency in Systems and Networks, M1, Univ. Rennes 1.
- Master: Cédric Tedeschi, Service Technology, M1, Univ. Rennes 1.
- Master: Cédric Tedeschi, Parallel Programming, M1, Univ. Rennes 1.
- Master: 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.

Supervision

- Defended PhD: Ali Jawad Fahs, “Decentralized Fog Computing Infrastructure Control”, defended on December 16th 2020, supervised by Guillaume Pierre.
- Defended PhD: Loic Guegan, “Scalable end-to-end models for the time and energy performance of Fog infrastructures”, defended on January 29th 2021, supervised by Martin Quinson and Anne-Cécile Orgerie.
- 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: Paulo Rodrigues De Souza Junior, “fog computing service roaming techniques”, started in December 2018, supervised by Guillaume Pierre and Daniele Miorandi (U-Hopper srl, Italy).
- PhD in progress: Davaadorj Battulga, “Scalable data pipelines for fog computing applications”, started in September 2018, supervised by Cédric Tedeschi and Daniele Miorandi (U-Hopper srl, Italy).
- PhD in progress: Mulugeta Tamiru, “Automatic optimization of autonomous management systems”, started in September 2018, supervised by Guillaume Pierre and Erik Elmroth (Elastisys AB, Sweden).
- PhD in progress: Mozhdeh Farhadi, “Fog computing-enabled IoT situation-aware services”, started in June 2018, supervised by Guillaume Pierre and Daniele Miorandi (U-Hopper srl, Italy).
- PhD in progress: Hamidreza Arkian, “Stream processing operator placement”, started in June 2018, supervised by Guillaume Pierre and Erik Elmroth (Elastisys AB, Sweden).
- PhD in progress: Mehdi Belkhiria, “Dynamic Stream Processing for Maritime Traffic Surveillance,” started in December 2017, supervised by Cédric Tedeschi.
- PhD in progress (*co-tutelle*): Yasmina Bouizem, “Energy-efficient, fault-tolerance mechanisms for containerized cloud applications,” started in December 2017, supervised by Didi Fedoua (Tlemcen University, Algeria), Djawida Dib (Tlemcen University, Algeria), Christine Morin and Nikos Parlavantzas.
- PhD in progress: Clément El Baz, “Reactive security monitoring in clouds,” started in October 2017, supervised by Louis Rilling and Christine Morin.
- PhD in progress: Dorra Boughzala, “Simulating Energy Consumption of Continuum Computing between Heterogeneous Numerical Infrastructures in HPC”, started in December 2017, ended in November 2020, supervised by Laurent Lefèvre (Avalon team in Lyon), Anne-Cécile Orgerie and Martin Quinson.
- PhD in progress: Baptiste Goupille-Lescar, “Designing agile, distributed cyber-physical systems with advanced collaboration capabilities”, started in January 2016, ended in September 2020, supervised by Eric Lenormand (Thales), Christine Morin, Nikos Parlavantzas.
- PhD in progress: Bruno Stevant, “Resource allocation strategies for service distribution at the Internet edge to optimize end-to-end latency,” started in December 2014 (part-time), supervised by Jean-Louis Pazat.

Juries

- Guillaume Pierre was a member of the PhD committee of Cédric Morin, IMT Atlantique, November 18th 2020.
- Martin Quinson chaired the HDR committee of Anne-Cécile Orgerie, ENS Rennes, November 20th 2020.
- Anne-Cécile Orgerie is a member of the jury for *Agrégation externe de Sciences Industrielles de l'Ingénieur option ingénierie informatique*.
- Martin Quinson is a member of the jury for *Capes d'informatique*.
- Anne-Cécile Orgerie was a member of the PhD committee of Ioanna Stypsanelli, INSA de Toulouse, December 7th 2020.
- Anne-Cécile Orgerie was a member of the PhD committee of Yewan Wang, IMT-Atlantique, March 9th 2020.
- Shadi Ibrahim was a member of the PhD committee of Ali Fahs, Universite de Rennes 1, December 17th 2020.
- Cédric Tedeschi was a member of the PhD committee of Arif Ahmed, Université de Rennes 1, January 13, 2020.
- Jean-Louis Pazat was a member of the PhD committee of Johan Pelay, January 2020.

10.3 Popularization

10.3.1 Articles and contents

- “Le vrai coût énergétique du numérique” Anne-Cécile Orgerie and Laurent Lefèvre, Pour la Science, December 2020, pages 48-59.

10.3.2 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 occured in April 2019, and a new session has started for 2020-2021. The program is currently supported by: Fondation Blaise Pascal, ED MathSTIC, ENS de Rennes, Université Rennes 1 and Fondation Rennes 1.

11 Scientific production

11.1 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.

11.2 Publications of the year

International journals

[2] H. Arkian, D. Giouroukis, P. Souza Junior and G. Pierre. ‘Potable Water Management with integrated Fog computing and LoRaWAN technologies’. In: *IEEE IoT Newsletter* (11th Mar. 2020), pp. 1–3. URL: https://hal.inria.fr/hal-02513467.

[3] G. Cooperman and M. Quinson. 'Sthread: In-Vivo Model Checking of Multithreaded Programs'. In: *The Art, Science, and Engineering of Programming* (2020). DOI: 10.22152/programming-journal.org/2020/4/13. URL: https://hal.inria.fr/hal-02449080.

[4] M. Farhadi, J.-L. Lanet, G. Pierre and D. Miorandi. 'A systematic approach towards security in Fog computing: assets, vulnerabilities, possible countermeasures'. In: *Software: Practice and Experience* 50.6 (2020), pp. 973–997. DOI: 10.1002/spe.2804. URL: https://hal.inria.fr/hal-02441639.

[5] R. Ferreira Da Silva, H. Casanova, A.-C. Orgerie, R. Tanaka, E. Deelman and F. Suter. 'Characterizing, Modeling, and Accurately Simulating Power and Energy Consumption of I/O-intensive Scientific Workflows'. In: *Journal of computational science* 44 (June 2020), p. 101157. DOI: 10.1016/j.jocs.2020.101157. URL: https://hal.archives-ouvertes.fr/hal-02876736.

International peer-reviewed conferences

[6] A. Ahmed, A. Mohan, G. Cooperman and G. Pierre. 'Docker Container Deployment in Distributed Fog Infrastructures with Checkpoint/Restart'. In: IEEE International Conference on Mobile Cloud Computing. Oxford, United Kingdom, 13th Apr. 2020. URL: https://hal.inria.fr/hal-02473333.

[7] H. Arkian, G. Pierre, J. Tordsson and E. Elmroth. 'An Experiment-Driven Performance Model of Stream Processing Operators in Fog Computing Environments'. In: SAC 2020 - ACM/SIGAPP Symposium On Applied Computing. Brno, Czech Republic: https://www.sigapp.org/sac/sac2020/, 30th Mar. 2020, pp. 1–9. URL: https://hal.inria.fr/hal-02394396.

[8] 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, 24th Feb. 2020. URL: https://hal.inria.fr/hal-02463206.

[9] M. Belkhiria, M. Bertier and C. Tedeschi. 'Group Mutual Exclusion to Scale Distributed Stream Processing Pipelines'. In: UCC 2020 - 13th IEEE/ACM International Conference on Utility and Cloud Computing. Leicester / Virtual, United Kingdom, 7th Dec. 2020. URL: https://hal.inria.fr/hal-02993099.

[10] D. Boughzala, L. Lefèvre and A.-C. Orgerie. 'Predicting the energy consumption of CUDA kernels using SimGrid'. In: SBAC-PAD 2020 - 32nd IEEE International Symposium on Computer Architecture and High Performance Computing. SBAC-PAD 2020 - 32nd IEEE International Symposium on Computer Architecture and High Performance Computing. Porto, Portugal, Sept. 2020, pp. 191–198. URL: https://hal.archives-ouvertes.fr/hal-02924028.

[11] Y. Bouizem, D. Dib, N. Parlavantzas and C. Morin. 'Active-Standby for High-Availability in FaaS'. In: Sixth International Workshop on Serverless Computing (WoSC6) 2020. Delft, Netherlands, 7th Dec. 2020. DOI: 10.1145/3429880.3430097. URL: https://hal.inria.fr/hal-03043479.

[12] B. Camus, A. Blavette, A.-C. Orgerie and J.-B. Blanc-Rouchossé. 'Co-simulation of an electrical distribution network and its supervision communication network'. In: CCNC 2020 - IEEE Consumer Communications & Networking Conference. CCNC: IEEE Consumer Communications & Networking Conference. Las Vegas, United States, Jan. 2020, pp. 1–6. URL: https://hal.archives-ouvertes.fr/hal-02352832.

[13] 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, 20th Apr. 2020, pp. 1–9. URL: https://hal.inria.fr/hal-02506364.

[14] C. Elbaz, L. Rilling and C. Morin. 'Fighting N-Day Vulnerabilities with Automated CVSS Vector Prediction at Disclosure'. In: ARES 2020 - International Conference on Availability, Reliability and Security. Virtual Event, Ireland, 25th Aug. 2020, pp. 1–10. URL: https://hal.inria.fr/hal-02895913.

[15] A. Fahs, G. Pierre and E. Elmroth. 'Voilà: Tail-Latency-Aware Fog Application Replicas Autoscaler'. In: MASCOTS 2020 - 27th IEEE Symposium on Modelling, Analysis, and Simulation of Computer and Telecommunication Systems. Nice, France, 17th Nov. 2020. URL: https://hal.inria.fr/hal-02932484.

[16] A. J. Fahs and G. Pierre. 'Tail-Latency-Aware Fog Application Replica Placement'. In: ICSOC 2020 - 18th International Conference on Service Oriented Computing. Dubai, United Arab Emirates, 14th Dec. 2020, pp. 1–16. URL: https://hal.inria.fr/hal-02917191.

[17] A. Gougeon, B. Camus and A.-C. Orgerie. 'Optimizing Green Energy Consumption of Fog Computing Architectures'. In: SBAC-PAD 2020 - 32nd IEEE International Symposium on Computer Architecture and High Performance Computing. SBAC-PAD 2020 - 32nd IEEE International Symposium on Computer Architecture and High Performance Computing. Porto, Portugal, 8th Sept. 2020, pp. 75–82. URL: https://hal.archives-ouvertes.fr/hal-02924022.

[18] A. Jlassi and C. Tedeschi. 'Merge, Split, and Cluster: Dynamic Deployment of Stream Processing Applications'. In: CCGRID 2020 - 20th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing. Melbourne, France, 2nd Nov. 2020, pp. 1–10. URL: https://hal.inria.fr/hal-02508987.

[19] D. Nguyen, M. Simonin, G. Hajduch, R. Vadaine, C. Tedeschi and R. Fablet. 'Detection of Abnormal Vessel Behaviors from AIS data using GeoTrackNet: from the Laboratory to the Ocean'. In: MBDW 2020 : 2nd Maritime Big Data Workshop part of MDM 2020 : 21st IEEE International Conference on Mobile Data Management. Versailles, France, 30th June 2020, pp. 264–268. DOI: 10.1109/MDM48529.2020.00061. URL: https://hal.archives-ouvertes.fr/hal-02523279.

[20] 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, 14th Sept. 2020, pp. 1–11. DOI: 10.1109/CLUSTER49012.2020.00028. URL: https://hal.archives-ouvertes.fr/hal-02916032.

[21] 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, 14th Dec. 2020, pp. 1–8. URL: https://hal.inria.fr/hal-02963913.

[22] B. Stévant, J.-L. Pazat and A. Blanc. 'QoS-aware autonomic adaptation of microservices placement on edge devices'. In: CLOSER 2020 : 10th International Conference on Cloud Computing and Services Science. Prague, Czech Republic, 7th May 2020, pp. 1–8. DOI: 10.5220/0009319902370244. URL: https://hal-imt-atlantique.archives-ouvertes.fr/hal-02870961.

[23] M. A. Tamiru, G. Pierre, J. Tordsson and E. Elmroth. 'Instability in Geo-Distributed Kubernetes Federation: Causes and Mitigation'. In: MASCOTS 2020 - 27th IEEE Symposium on Modelling, Analysis, and Simulation of Computer and Telecommunication Systems. Nice, France, 17th Nov. 2020. URL: https://hal.inria.fr/hal-02934475.

[24] M. A. Tamiru, J. Tordsson, E. Elmroth and G. Pierre. 'An Experimental Evaluation of the Kubernetes Cluster Autoscaler in the Cloud'. In: CloudCom 2020 - 12th IEEE International Conference on Cloud Computing Technology and Science. Bangkok, Thailand, 14th Dec. 2020, pp. 1–9. URL: https://hal.inria.fr/hal-02958916.

[25] A. T. Wonjiga, L. Rilling and C. Morin. 'SLA Definition for Network Intrusion Detection Systems in IaaS Clouds'. In: 36th ACM/SIGAPP Symposium on Applied Computing (SAC '21). Virtual Event, Republic of Korea., South Korea, 26th Mar. 2021. DOI: 10.1145/3412841.3441885. URL: https://hal.inria.fr/hal-03085554.

Doctoral dissertations and habilitation theses

[26] A. Ahmed. 'Déploiement efficace d'applications cloud dans les infrastructures fog distribuées'. Université Rennes 1, 20th Jan. 2020. URL: https://tel.archives-ouvertes.fr/tel-02932386.

[27] A. Fahs. 'Proximity-Aware Replicas Management in Geo-Distributed Fog Computing Platforms'. Université de Rennes 1, 16th Dec. 2020. URL: https://hal.inria.fr/tel-03080983.

[28] 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.

Reports & preprints

[29] O. Yildiz and S. Ibrahim. *Preserving Fairness in Shared Hadoop Cluster: A Study on the Impact of (Non-) Preemptive Approaches*. Inria Rennes - Bretagne Atlantique, 30th Dec. 2020. URL: https://hal.inria.fr/hal-03091371.

11.3 Cited publications

[30] 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. URL: https://hal.inria.fr/hal-01916484.

[31] T. Lambert, D. Guyon and S. Ibrahim. 'Rethinking Operators Placement of Stream Data Application in the Edge'. In: *CIKM 2020- 29th ACM International Conference on Information and Knowledge Management*. Dublin, Ireland, Oct. 2020. URL: https://hal.archives-ouvertes.fr/hal-02942759.

[32] N. Parlavantzas, L. M. Pham, C. Morin, S. Arnoux, G. Beaunée, L. Qi, P. Gontier and P. Ezanno. 'A Service-based Framework for Building and Executing Epidemic Simulation Applications in the Cloud'. In: *Concurrency and Computation: Practice and Experience* 32.5 (2019), e5554. DOI: 10.1002/cpe.5554. URL: https://hal.archives-ouvertes.fr/hal-02312313.

[33] T. A. Pham, T. Jéron and M. Quinson. 'Verifying MPI Applications with SimGridMC'. In: *Correctness 2017 - First International Workshop on Software Correctness for HPC Applications*. Denver, United States, Nov. 2017, pp. 28–33. DOI: 10.1145/3145344.3145345. URL: https://hal.inria.fr/hal-01632421.

[34] A. T. Wonjiga, L. Rilling and C. Morin. 'Verification for Security Monitoring SLAs in IaaS Clouds: the Example of a Network IDS'. In: *NOMS 2018 - Network Operations and Management Symposium*. Taipei, Taiwan: IEEE, Apr. 2018, pp. 1–7. DOI: 10.1109/NOMS.2018.8406157. URL: https://hal.inria.fr/hal-01663941.