

RESEARCH CENTRE

**Inria Centre
at Rennes University**

IN PARTNERSHIP WITH:

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

2023

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

Inria

Contents

Project-Team MYRIADS	1
1 Team members, visitors, external collaborators	3
2 Overall objectives	4
2.1 General Objectives	4
2.2 Context	4
2.3 Challenges	5
3 Research program	5
3.1 Introduction	5
3.2 Scaling fogs and clouds	6
3.2.1 Resource management in hierarchical clouds	6
3.2.2 Resource management in fog computing architectures	6
3.2.3 Self-optimizing applications in multi-cloud environments	6
3.3 Greening clouds	7
3.3.1 Smart grids and clouds	7
3.3.2 Energy cost models	7
3.3.3 Energy-aware users	8
3.4 Securing clouds	8
3.4.1 Security monitoring service level objectives	8
3.4.2 Data protection in Cloud-based IoT services	9
3.5 Experimenting with Clouds	9
3.5.1 Experimentation methodologies for clouds	9
3.5.2 Use cases	10
4 Application domains	10
4.1 Main application domains	10
5 Social and environmental responsibility	11
5.1 Footprint of research activities	11
5.2 Impact of research results	11
6 Highlights of the year	11
6.1 Awards	11
7 New software, platforms, open data	11
7.1 New software	11
7.1.1 SimGrid	11
7.1.2 Tansiv	12
7.1.3 EnOSlib	13
8 New results	13
8.1 Scaling Clouds	13
8.1.1 Fog computing platform design	13
8.1.2 Advanced data management in shared environments	14
8.1.3 Geo-distributed graph data processing	14
8.1.4 Resource Allocation for Serverless ML Workflows	15
8.1.5 Geo-distributed data stream processing	15
8.1.6 Fault tolerance for Function-as-a-Service environments	15
8.1.7 Flexible function placement for Function-as-a-Service in the fog	16
8.1.8 Balancing performance and sustainability for Function-as-a-Service in the fog	16
8.1.9 Reliable fog platforms in adverse natural environments	16
8.1.10 Modeling cloud infrastructures	17
8.2 Greening Clouds	17

8.2.1	Energy Models for Cloud infrastructures	17
8.2.2	End-to-end energy models for the Internet of Things	17
8.2.3	Exploiting renewable energy in distributed clouds	18
8.2.4	Smart Grids	18
8.2.5	End-to-end ecodesign of cloud platforms	18
8.3	Experimenting with Clouds	19
8.3.1	Simulating distributed IT systems	19
8.3.2	Toward stealth analysis of distributed applications	19
8.3.3	Tools for experimentation	20
9	Bilateral contracts and grants with industry	20
9.1	Bilateral contracts with industry	20
10	Partnerships and cooperations	21
10.1	National initiatives	21
10.2	Regional initiatives	23
11	Dissemination	23
11.1	Promoting scientific activities	23
11.1.1	Scientific events: organisation	24
11.1.2	Scientific events: selection	24
11.1.3	Journal	25
11.1.4	Invited talks	25
11.1.5	Leadership within the scientific community	26
11.1.6	Scientific expertise	26
11.1.7	Research administration	26
11.2	Teaching - Supervision - Juries	26
11.2.1	Teaching	26
11.2.2	Supervision	28
11.2.3	Juries	29
11.3	Popularization	29
11.3.1	Internal or external Inria responsibilities	29
11.3.2	Articles and contents	29
11.3.3	Education	29
12	Scientific production	29
12.1	Major publications	29
12.2	Publications of the year	29
12.3	Cited publications	31

Project-Team MYRIADS

Creation of the Project-Team: 2012 January 01

Keywords

Computer sciences and digital sciences

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

Other research topics and application domains

- B2.3. – Epidemiology
- B3.1. – Sustainable development
- B3.2. – Climate and meteorology
- B4.3. – Renewable energy production
- B4.4. – Energy delivery

- B4.4.1. – Smart grids
- B4.5. – Energy consumption
 - B4.5.1. – Green computing
- B5.1. – Factory of the future
- B5.8. – Learning and training
- B6.1. – Software industry
 - B6.1.1. – Software engineering
- B6.3. – Network functions
 - B6.3.3. – Network Management
- B6.4. – Internet of things
- B6.5. – Information systems
- B6.6. – Embedded systems
- B8.1. – Smart building/home
- B8.2. – Connected city
- B8.3. – Urbanism and urban planning
- B8.5. – Smart society
- B9.1. – Education
 - B9.1.1. – E-learning, MOOC
 - B9.1.2. – Serious games
- B9.5.1. – Computer science
- B9.7. – Knowledge dissemination
 - B9.7.1. – Open access
 - B9.7.2. – Open data
- B9.8. – Reproducibility
- B9.9. – Ethics
- B9.10. – Privacy

1 Team members, visitors, external collaborators

Research Scientists

- Shadi Ibrahim [INRIA, Researcher]
- Anne-Cécile Orgerie [CNRS, Senior Researcher, HDR]

Faculty Members

- Guillaume Pierre [Team leader, UNIV RENNES, Professor]
- Marin Bertier [INSA RENNES, Associate Professor]
- François Lemerrier [UNIV RENNES, Associate Professor]
- Nikolaos Parlavantzas [INSA RENNES, Associate Professor, HDR]
- Jean-Louis Pazat [INSA RENNES, Professor, HDR]
- Martin Quinson [ENS RENNES, Professor, HDR]
- Cédric Tedeschi [UNIV RENNES, Professor, HDR]

Post-Doctoral Fellow

- Alessio Pagliari [UNIV RENNES]

PhD Students

- Khaled Arsalane [UNIV RENNES]
- Leo Cosseron [ENS RENNES]
- Clément Courageux-Sudan [ENS RENNES]
- Wedan Gnibga [CNRS]
- Chih-Kai Huang [UNIV RENNES]
- Ammar Kazem [UNIV RENNES, from Oct 2023]
- Govind Kovilkkatt Panickerveetil [INRIA, from Mar 2023]
- Mohamed Cherif Zouaoui Latreche [INSA RENNES, from Oct 2023]
- Mathieu Laurent [ENS RENNES, from Sep 2023]
- Volodia Parol–Guarino [INRIA]
- Mohammad Rizk [INRIA, from Nov 2023]
- Matthieu Silard [CNRS, from Feb 2023]

Technical Staff

- Adrien Gougeon [ENS RENNES]
- Mathieu Simonin [INRIA, Engineer]

Interns and Apprentices

- Hugo Depuydt [ENS Rennes, Intern]
- Weissem Trabelsi [UNIV RENNES, Intern]

Administrative Assistants

- Stephanie Gosselin Lemaile [INRIA]
- Angélique Jarnoux [INRIA]

Visiting Scientist

- Maxwell Pirtle [Northeastern University]

External Collaborators

- Anne Blavette [CNRS]
- Louis Rilling [DGA, until Nov 2023]

2 Overall objectives

2.1 General Objectives

MYRIADS is a joint team with INRIA, CNRS, UNIVERSITY RENNES 1, INSA RENNES and ENS RENNES. It is part of IRISA (D1 department on large scale systems) and INRIA RENNES – BRETAGNE ATLANTIQUE.

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

2.2 Context

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

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

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

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

of computers for a short experiment. The risk of wasting money by either under-utilization or undersized data centers is shifted from the user to the provider.

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

2.3 Challenges

The term cloud was coined 18 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 may advocate for placing computing resources directly within the backbone network. The general challenge of resource management for such clouds stands in trying to be locality-aware: for the needs of an application, several virtual machines may exchange data. Placing them *close* to each other can significantly improve the performance of the application they compose. More generally, building an overlay network that takes into account the hierarchical aspects of the platform without being a hierarchical overlay – which comes with load balancing and resilience issues – is a challenge by itself.

3.2.2 Resource management in fog computing architectures

Fog computing infrastructures are composed of compute, storage and networking resources located at the edge of wide-area networks, in immediate proximity to the end users. Instead of treating the mobile operator's network as a high-latency dumb pipe between the end users and the external service providers, fog platforms aim at deploying cloud functionalities *within* the mobile phone network, inside or close to the mobile access points. Doing so is expected to deliver added value to the content providers and the end users by enabling new types of applications ranging from Internet-of-Things applications to extremely interactive systems (e.g., augmented reality). Simultaneously, it will generate extra revenue streams for the mobile network operators, by allowing them to position themselves as cloud computing operators and to rent their already-deployed infrastructure to content and application providers.

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

The main challenges that we plan to address are:

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

3.2.3 Self-optimizing applications in multi-cloud environments

As the use of cloud computing becomes pervasive, the ability to deploy an application on a multi-cloud infrastructure becomes increasingly important. Potential benefits include avoiding dependence on a single vendor, taking advantage of lower resource prices or resource proximity, and enhancing application availability. Supporting multi-cloud application management involves two tasks. First, it involves

selecting an initial multi-cloud application deployment that best satisfies application objectives and optimizes performance and cost. Second, it involves dynamically adapting the application deployment in order to react to changes in execution conditions, application objectives, cloud provider offerings, or resource prices. Handling price changes in particular is becoming increasingly complex. The reason is the growing trend of providers offering sophisticated, dynamic pricing models that allow buying and selling resources of finer granularities for shorter time durations with varying prices.

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

The main challenges arising in this context are:

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

3.3 Greening clouds

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

3.3.1 Smart grids and clouds

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

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

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

3.3.2 Energy cost models

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

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

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

3.3.3 Energy-aware users

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

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

We plan to explore possible collaborations between virtual machine suppliers, service providers and users of Clouds in order to provide users with ways of participating in the reduction of the Clouds energy consumption. This work will follow two directions: 1) to investigate compromises between power and performance/service quality that cloud providers can offer to their users and to propose them a variety of options adapted to their workload; and 2) to develop mechanisms for each layer of the Cloud software stack to provide users with a quantification of the energy consumed by each of their options as an incentive to become greener.

3.4 Securing clouds

3.4.1 Security monitoring service level objectives

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

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

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

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

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

3.4.2 Data protection in Cloud-based IoT services

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

3.5 Experimenting with Clouds

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

3.5.1 Experimentation methodologies for clouds

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

We plan to make SimGrid widely usable beyond research laboratories, in order to evaluate industrial systems and to teach the future generations of cloud practitioners. This requires to frame the specific concepts of Cloud systems and platforms in actionable interfaces. The challenge is to make the framework

both easy to use for simple studies in educational settings while modular and extensible to suit the specific needs of advanced industrial-class users.

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

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

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

3.5.2 Use cases

In system research, it is important to work on real-world use cases from which we extract requirements inspiring new research directions and with which we can validate the system services and mechanisms we propose. In the framework of our close collaboration with the Data Science Technology department of the Lawrence Berkeley National Lab (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 CNRS 80Prime DECORUS project.

In the context of the 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.

In the CNRS FACTO project, the energy aspect of a sustainable smart home is studied by reducing the oversized number of wireless technologies that is actually connecting all the devices. This work aims to propose a versatile network based on only one optimized and energy efficient technology (Wi-Fi 7) that could meet all connected devices requirements.

The Myriads team also engaged in the Inria FrugalCloud challenge, in collaboration with the OVH-cloud company. The objective is to participate in the end-to-end eco-design of Cloud platforms in order to reduce their environmental impact.

6 Highlights of the year

- Deborah Agarwal from the Lawrence Berkeley National Laboratory, who held an Inria International Chair position in the Myriads team from 2014 to 2022, received the **Doctorate Honoris Causa from Université de Rennes** on April 4th 2023.
- The Myriads research team reached its end of life on December 31st, 2023. A new team called **Magellan** will pursue its research activities starting from January 1st, 2024.

6.1 Awards

- The paper entitled “Latency, Energy and Carbon Aware Collaborative Resource Allocation with Consolidation and QoS Degradation Strategies in Edge Computing” authored by Emmanuel Gribga, Anne Blavette and Anne-Cécile Orgerie received an *Outstanding Paper Award* at the ICPADS 2023 conference (IEEE International Conference on Parallel and Distributed Systems) [10].

7 New software, platforms, open data

7.1 New software

7.1.1 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 simulation toolkit that provides core functionalities for the simulation of distributed applications in large scale heterogeneous distributed environments.

News of the Year: There were 3 major releases in 2022. The SimDag API for the simulation of the scheduling of Directed Acyclic Graphs has been dropped and replaced by the SimDag++ API which provides the different features of SimDag directly on top of the S4U API. We also dropped the old and clumsy Lua bindings to create platforms in a programmatic way. It can be done in C++ in a much cleaner way now, which motivates this suppression. The C++ platform description has been improved to reject forbidden topologies, improve exporting for visualization, and allow users to dynamically change injected costs for MPI_* operations. The Python API to S4U has been extended. A new solver for parallel task (BMF) has been introduced and provides with more realistic sharing of heterogeneous resources compared to the fair bottleneck solver used by ptask_L07. Although this is still ongoing work, this paves the way to efficient macroscopic modeling of streaming activities and parallel applications. The internals of the Model Checker have been heavily reworked and new test suites from the MPI Bugs Initiative (MBI) are now used. The documentation was thoroughly overhauled to ease the use of the framework. We also pursued our efforts to improve the overall framework, through bug fixes, code refactoring and other software quality improvements.

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

Contact: Martin Quinson

Participants: Adrien Lebre, Anne-Cécile Orgerie, Arnaud Legrand, Augustin Degomme, Arnaud Giersch, Emmanuelle Saillard, Frédéric Suter, Jonathan Pastor, Martin Quinson, Samuel Thibault

Partners: CNRS, ENS Rennes

7.1.2 Tansiv

Name: Time-Accurate Network Simulation Interconnecting Vms

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

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

Authors: Louis Rilling, Matthieu Simonin, Martin Quinson

Contact: Louis Rilling

Partner: DGA-MI

7.1.3 EnOSlib

Name: EnOSlib is a library to help you with your experiments

Keywords: Distributed Applications, Distributed systems, Evaluation, Grid Computing, Cloud computing, Experimentation, Reproducibility, Linux, Virtualization

Functional Description: EnOSlib is a library to help you with your distributed application experiments. The main parts of your experiment logic is made reusable by the following EnOSlib building blocks:

- Reusable infrastructure configuration: The provider abstraction allows you to run your experiment on different environments (locally with Vagrant, Grid'5000, Chameleon and more)
- Reusable software provisioning: In order to configure your nodes, EnOSlib exposes different APIs with different level of expressivity
- Reusable experiment facilities: Tasks help you to organize your experimentation workflow.

EnOSlib is designed for experimentation purpose: benchmark in a controlled environment, academic validation ...

URL: <https://discovery.gitlabpages.inria.fr/enoslib/>

Publications: [hal-01664515](#), [hal-01689726](#)

Contact: Mathieu Simonin

Participants: Mathieu Simonin, Baptiste Jonglez, Marie Delavergne, Alexis Bitailou

8 New results

8.1 Scaling Clouds

8.1.1 Fog computing platform design

Participants: Chih-Kai Huang, Guillaume Pierre.

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.

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 demonstrated that these federations suffer from scalability issues due in particular to the amount of monitoring information they need to collect from the member clusters to make proper scheduling decisions. We therefore proposed a new technique based on metrics aggregation and deduplication to reduce the volume of monitoring traffic up to 99% without impacting the quality of scheduling decisions [11]. We also showed how dynamically adapting the metrics collection frequency allows one to control the tradeoff between communication costs and monitoring accuracy [12].

We proposed algorithms to allow a fog computing platform based on Kubernetes to dynamically add and exploit new computing resources in case an overload is detected in some part of the system [7].

8.1.2 Advanced data management in shared environments

Participants: Shadi Ibrahim.

Solid State Drives (SSDs) are widely used in data-intensive scenarios due to their high performance and decreasing cost. However, in shared environments, concurrent workloads can interfere with each other, leading to a violation of Quality of Service (QoS). While QoS mechanisms like fairness guarantees and latency constraints have been integrated into SSDs, existing transaction processing frameworks offer limited QoS guarantees and can significantly degrade overall performance in a shared environment. The reason is that the internal components of an SSD, originally designed to exploit parallelism, struggle to coordinate effectively when QoS mechanisms are applied to them. In this collaborative work, we propose a novel QoS-enhanced transaction processing framework, called QoS-pro, which enhances QoS guarantees for concurrent workloads while maintaining high parallelism for SSDs [4]. QoS-pro achieves this by redesigning transaction processing procedures to fully exploit the parallelism of shared SSDs and enhancing QoS-oriented transaction translation and scheduling with parallelism features in mind.

The continuous growth in data volume increases the interest in using peer-to-peer (P2P) systems not only to store static data (i.e., immutable data) but also to store and share mutable data – data that are updated and modified by multiple users. Unfortunately, current P2P systems are mainly optimized to manage immutable data. Thus, each modification creates a new copy of the file, which leads to a high "useless" network usage. Conflict-free Replicated Data Types (CRDTs) are specific data types built in a way that mutable data can be managed without the need for consensus-based concurrency control. A few studies have demonstrated the potential benefits of integrating CRDTs in the InterPlanetary File System (IPFS), an open-source widely used P2P content sharing system. However, they have not been implemented and evaluated in a real IPFS deployment. This work tries to fill the gap between theory and practice and provides a quantitative measurement of the performance of CRDTs in IPFS. Accordingly, in collaboration with the COAST team, we introduce IM-CRDT [8], an implementation of CRDTs in IPFS that focuses on the simple data type (i.e., Set); and carry out extensive experiments to verify whether CRDTs can efficiently be utilized in IPFS to handle mutable data. Experiments on Grid'5000 show that IM-CRDT can sustain low convergence time under concurrent updates.

The high-performance computing I/O stack has been complex due to multiple software layers, the inter-dependencies among these layers, and the different performance tuning options for each layer. In this complex stack, the definition of an "I/O access pattern" has been reappropriated to describe what an application is doing to write or read data from the perspective of different layers of the stack, often comprising a different set of features. It has become common to have to redefine what is meant when discussing a pattern in every new study, as no assumption can be made. In the framework of the **Hermes Associate team** and the context of the ongoing collaboration between the Myriads project team at Inria and Lawrence Berkeley National Laboratory, we conduct a comprehensive study and propose a baseline taxonomy [2], harnessing the I/O community's knowledge over the past 20 years. This definition can serve as a common ground for high-performance computing I/O researchers and developers to apply known I/O tuning strategies and design new strategies for improving I/O performance. We seek to summarize and bring a consensus to the multiple ways to describe a pattern based on common features already used by the community over the years.

8.1.3 Geo-distributed graph data processing

Participants: Shadi Ibrahim.

Graph processing is a popular computing model for big data analytics. Emerging big data applications are often maintained in multiple geographically distributed (geo-distributed) data centers (DCs) to provide low-latency services to global users. Graph processing in geo-distributed DCs poses a challenge when personal data are involved, such as in social network graphs. Indeed, conducting graph processing algorithms in such a geo-distributed system requires coordination among multiple DCs, while

simultaneously addressing privacy concerns and complying with different laws and regulations across countries and regions. With Tristan Allard, Cédric Eichler, Benjamin Nguyen, and Haoying Zhang (a master intern co-supervised by Myriads, Spicy, and SDS@LIFO), we address these challenges and propose an innovative framework that enables privacy-preserving geo-distributed graph processing through the use of synthetic graphs. With this framework, each DC can generate a differentially private view of the global graph, allowing the application of various graph processing algorithms under the guarantee of differential privacy.

8.1.4 Resource Allocation for Serverless ML Workflows

Participants: Shadi Ibrahim.

Machine Learning (ML) workflows are increasingly deployed on serverless computing platforms to benefit from their elasticity and fine-grain pricing. Proper resource allocation is crucial to achieve fast and cost-efficient execution of serverless ML workflows (specially for hyperparameter tuning and model training). Unfortunately, existing resource allocation methods are static, treat functions equally, and rely on offline prediction, which limit their efficiency. In this collaborative work, we introduce CE-scaling – a Cost-Efficient autoscaling framework for serverless ML work-flows [20]. During the hyperparameter tuning, CE-scaling partitions resources across stages according to their exact usage to minimize resource waste. Moreover, it incorporates an online prediction method to dynamically adjust resources during model training. In this collaborative work, we implement and evaluate CE-scaling on AWS Lambda using various ML models. Evaluation results show that compared to state-of-the-art static resource allocation methods, CE-scaling can reduce the job completion time and the monetary cost for hyperparameter tuning and model training.

8.1.5 Geo-distributed data stream processing

Participants: Khaled Aرسالane, Davaadorj Battulga, Alessio Pagliari, Guillaume Pierre, Cédric Tedeschi.

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. In previous work we developed an auto-scaling mechanism which can dynamically add or remove resources to adjust the processing capacity to wide variations in the intensity of the workload. However, in a fog computing environment additional resources may not always be available. Similarly some streaming applications cannot scale horizontally. In such situations it becomes useful to exploit transprecision where the precision as well as the compute intensity can be reduced when necessary. We proposed a dual-controller autoscaler which exploits both horizontal scalability and transprecision to maintain an application's processing capacity [15].

In the context of Khaled Aرسالane's PhD thesis, we started exploring the limitations of stream processing autoscaling systems when dealing with stateful data processing operators in geo-distributed environments. A publication on this topic is in preparation.

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. We worked on the autonomous decentralized adaptation of such a collaborative platform. In particular, we designed a fully-decentralized adaptation algorithm for the migration of jobs composing a stream processing pipeline. A software prototype of it was built and experimented over an emulated geo-distributed platform based on Grid'5000.

8.1.6 Fault tolerance for Function-as-a-Service environments

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. One of the main challenges of FaaS systems is providing fault tolerance for the deployed applications. The basic fault tolerance mechanism in current FaaS platforms is automatically retrying function invocations. Although the retry mechanism is well-suited for transient faults, it incurs delays in recovering from other types of faults, such as node crashes. To address this limitation, we proposed the integration of a Request Replication mechanism in FaaS platforms and described how this integration was implemented in a well-known, open-source platform. We also provided a detailed experimental comparison of the proposed mechanism with retry and Active-Standby mechanisms in terms of performance, availability, and resource consumption under different failure scenarios [3].

8.1.7 Flexible function placement for Function-as-a-Service in the fog

Participants: Volodia Parol-Guarino, Nikos Parlavantzas.

Function-as-a-Service (FaaS) is a compelling programming model for developing applications that run on fog infrastructures. FaaS applications are composed of ephemeral, event-triggered functions, which can be flexibly deployed along the Cloud-to-thing continuum. However, deciding where to place those functions in the fog poses many challenges, including the heterogeneity and resource constraints of fog nodes combined with the stringent latency requirements of FaaS applications. Another challenge is that fog nodes are typically owned by different entities that should be incentivized to share their resources.

To address these challenges, we propose a market-based approach for placing FaaS applications in the fog. Clients submit function placement requests including latency and resource requirements. The approach then organizes auctions to determine the nodes that will host the functions and the associated client payments. We developed an open-source implementation of the framework, called GIRAFFE, and evaluated this on the Grid'5000 testbed. Experiments showed that our framework can reduce client spending by up to three times while delivering service quality that matches or exceeds that of baseline methods [16]. This research is part of the thesis project of Volodia Parol-Guarino, started in October 2022.

8.1.8 Balancing performance and sustainability for Function-as-a-Service in the fog

Participants: Mohamed Cherif Zouaoui Latreche, Nikos Parlavantzas.

A main challenge faced by FaaS platforms is effectively balancing application QoS requirements with the imperative to reduce energy consumption and carbon footprint of fog nodes. This challenge is exacerbated by the dynamic nature and high deployment density of FaaS workloads, characterized by short function durations and small function sizes. These factors lead to interferences between workloads, making it difficult to predict the performance and energy impact of management actions. This research explores QoS-driven, energy-aware management of FaaS applications in fog environments. To simplify decision-making for managing FaaS workloads, we are exploring machine learning methods, specifically, reinforcement learning. This research is part of the thesis project of Mohamed Cherif Zouaoui Latreche, started in October 2023, in collaboration with Hector Duran-Limon of the University of Guadalajara, Mexico.

8.1.9 Reliable fog platforms in adverse natural environments

Participants: Ammar Kazem, Guillaume Pierre.

An interesting use-case for fog computing platforms is environmental monitoring to help Earth Sciences researchers (biologists, hydrologists, etc.) measure and understand specific natural environments. A fog platform in this context needs to support the needs of a wide range of applications which potentially compete for limited available computing resources, and exhibit self-management capabilities to remain operational with minimal human intervention despite potentially adverse execution conditions such as limited energy resources. In the context of Ammar Kazem's PhD thesis we started surveying the state of the art, and organized a survey for better understanding the current and potential future practice of data management in this domain.

8.1.10 Modeling cloud infrastructures

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

Wi-Fi networks are extensively used to provide Internet access to end-users and to deploy applications at the edge. By playing a major role in modern networking, Wi-Fi networks are getting bigger and denser. However, studying their performance at large-scale and in a reproducible manner remains a challenging task.

This year, we introduced a new Wi-Fi energy model for large-scale simulations. This model, based on flow-level simulation and integrated in SimGrid, requires fewer computations than state-of-the-art energy models to estimate bandwidth sharing over a wireless medium, leading to better scalability. The study shows that our approach yields to performance evaluations that are close to the ones of the classical ns-3 simulator while improving the runtime of simulations by several orders of magnitude [9].

8.2 Greening Clouds

8.2.1 Energy Models for Cloud infrastructures

Participants: Anne-Cécile Orgerie, Martin Quinson.

The global energy demand for digital activities is constantly growing. Computing nodes and cloud services are at the heart of these activities. Understanding their energy consumption is an important step towards reducing it. On one hand, physical power meters are very accurate in measuring energy but they are expensive, difficult to deploy on a large scale, and are not able to provide measurements at the service level. On the other hand, power models and vendor-specific internal interfaces are already available or can be implemented on existing systems. Plenty of tools, called software-based power meters, have been developed around the concepts of power models and internal interfaces, in order to report the power consumption at levels ranging from the whole computing node to applications and services. However, we have found that it can be difficult to choose the right tool for a specific need. In this work in collaboration with Avalon and Datamove teams, we qualitatively and experimentally compare several software-based power meters able to deal with CPU or GPU-based infrastructures [13]. For this purpose, we evaluate them against high-precision physical power meters while executing various intensive workloads. We extend this empirical study to highlight the strengths and limitations of each software-based power meter.

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

Participants: Clément Courageux-Sudan, 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.

8.2.3 Exploiting renewable energy in distributed clouds

Participants: Emmanuel Gnibga, Anne-Cécile Orgerie.

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

In a distributed Cloud environment, on-site renewable energy production and geographical energy-aware load balancing of virtual machines allocation can be associated to lower the brown (i.e. not renewable) energy consumption of DCs. Yet, combining these two approaches remains challenging in current distributed Clouds. Indeed, the variable and/or intermittent behavior of most renewable sources – like solar power for instance – is not correlated with the Cloud energy consumption, that depends on physical infrastructure characteristics and fluctuating unpredictable workloads [5, 10].

8.2.4 Smart Grids

Participants: Anne-Cécile Orgerie, Matthieu Silard.

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. In particular, in the context of the EDEN4SG ANR project, we consider the wide-scale deployment of electrical vehicles. The energy management of power systems closer-and-closer to real-time will require the intensive use of pervasive ICT that may suffer from an imperfect quality of service (e.g. delays) which may greatly decrease the performance of smart grids. They also have an energy and environmental impact. Hence, the project will consider the “cost of information”.

8.2.5 End-to-end ecodesign of cloud platforms

Participants: Anne-Cécile Orgerie, Guillaume Pierre, Govind Kovilkkatt Panickerveetil.

In the context of Inria's challenge on End-to-end ecodesign of cloud platforms, in collaboration with OVHCloud, we started building a system to automatically deploy data stream processing systems with their full software ecosystem in a cloud environment. The system uses software energy probes which can attribute the measured energy consumption to fine-grained elements such as the underlying cloud infrastructure, the different software elements, etc [14]. We then proceeded to using this tool to study the main factors which determine a stream processing system's energy consumption.

Numerous energy, power, and environmental leverages exist and can help cloud providers and data center managers to reduce some of these impacts. But dealing with such heterogeneous leverages can be a challenging task that requires some support from a dedicated framework. In collaboration with OVHCloud, we propose a new approach for modeling, evaluating, and orchestrating a large set of technological and logistical leverages. Our framework is based on leverages modeling and Gantt chart leverages mapping. First experimental results based on selected scenarios show the pertinence of the proposed approach in terms of management facilities and potential impacts reduction [6].

8.3 Experimenting with Clouds

8.3.1 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, the work on the framework did not lead to any new publication but instead we pursued our effort on the technical framework to prepare future scientific endeavors. We worked to ensure that it correctly captures the concepts needed by the experimenters, in preparation to a post-doctoral work on the simulation of the computing infrastructure behind the SKA (Square Kilometer Array telescope) scientific infrastructure starting in January 2023.

Our work on SimGrid is fully integrated into the other research efforts of the Myriads team. This year our main contribution on this topic was a new energy model of the Wi-Fi networks that is now integrated to the framework [9]. This enables the study of these networks within the simulator. Along the same line, the work on the TANSIV project described in the following section also required some adaptation to the simulator, to further increase the prediction accuracy when exchanging small data packets on local area networks.

8.3.2 Toward stealth analysis of distributed applications

Participants: Léo Cosseron, 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 2023 two directions were studied. First, two soundness issues of TANSIV were identified regarding the simulated network performance. The first issue is a buffer bloat effect caused by the performance difference between an emulated network interface card and the network link performance simulated by

the network simulator. The emulated network interface card is indeed only bounded by the CPU speed and the memory bandwidth, which, without specific precautions, makes the emulated hardware output packets at much higher speed than the one which should be enforced by the simulation. Thus, network packets are accumulated between the emulator and the simulator in a large buffer, each packet in the buffer adding an apparent latency between the source and destination node. Several attempts to solve this issue while respecting the constraints of not changing the guest software and minimizing changes in Qemu have been tried during the year and new ones will be considered in 2024.

The second issue is to simulate the effects of network congestion using SimGrid at the packet level whereas it is a flow-based simulator. In its current models SimGrid considers packets as flows and thus just slows them down during network congestion. However the real effect of network congestion on individual packets is to slow them down within the size limits of the queues in the different routers in the network, and to drop packets once the queues are full. Two approaches are followed to solve this issue. First TANSIV was adapted to use the ns-3 network simulator instead of SimGrid, which allowed to do experiments with network congestion. Second a student project was proposed and started to design SimGrid models which are able to drop packets on congestion.

The second direction studied is using TANSIV with actual malware analysis sandboxes and their tooling. Sandbox tools especially rely on Virtual Machine Introspection (VMI), which consists in intercepting guest events of interest for the analysis and inspecting the guest state at arbitrary times as well. LibVMI is a popular library to do VMI and can be used with Qemu-KVM as well as with Xen in the Drakvuf sandbox. We adapted TANSIV to work with LibVMI on Qemu-KVM and could show that similar mechanisms allow TANSIV to synchronize VMs with a discrete-event network simulator and hide the guest pauses induced by VMI activity. This work was submitted for publication (acceptance notification is due in 2024) and an early version was published as a research report [22]. In parallel we have started porting TANSIV to the Xen hypervisor, which should allow us to more extensively study and validate the compatibility of the TANSIV approach and the sandbox tooling that manipulates time in the guest. Porting to the Xen hypervisor should also allow us to validate the portability of the TANSIV approach between hypervisors. A paper about the principles introduced for TANSIV was submitted for publication and should benefit from results of the Xen port.

In addition to the research report published in 2023 [22] a communication of Léo Cosseron was accepted at the [first SOSP Doctoral Workshop](#) of the SOSP conference.

8.3.3 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 [23] a possible solution to the above desiderata.

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

The tool is used in several articles (see [here](#)). In particular, in [24] the tool is used to build an ad hoc framework for studying FOG applications.

9 Bilateral contracts and grants with industry

9.1 Bilateral contracts with industry

Défi Inria OVHcloud (2021-2025)

Participants: Anne-Cécile Orgerie, Shadi Ibrahim, Govind Kovilkatt Panickerveetil, Guillaume Pierre.

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

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

Défi Inria Hive (2022-2026)

Participants: Shadi Ibrahim, Mohammad Rizk, Guillaume Pierre.

The goal of this collaborative framework between Hive and Inria is to explore new solutions for the design and realization of large scale secure and reliable Peer-to-Peer Cloud storage. Four Inria project-teams are involved in this challenge including COAST, Myriads, WIDE, COATI.

Members of the Myriads team will contribute to two axes. Specifically, the Myriads team will coordinate the axis on reliable and cost-efficient data placement and repair in P2P storage over immutable data; and contribute to the axis on the management of mutable data over P2P storage.

10 Partnerships and cooperations

10.1 National initiatives

ANR FACTO (2021-2024)

Participants: Anne-Cécile Orgerie, Martin Quinson, François Lemerancier.

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

ANR EDEN4SG (2023-2027)

Participants: Anne-Cécile Orgerie.

Climate change as well as geopolitical tensions have led a large number of countries to target a massive integration of renewables in their energy mix. This will be achieved among others by increasing the electrification rate of several sectors such as transport. In this context, the wide-scale deployment of electrical vehicles (EVs) represents a challenge as well as an opportunity to render more efficient and affordable the transformation of the current power system into a smarter grid. The project targets to develop methods for the intelligent coordination of large-scale EV fleets and as well to determine the associated cost of information for piloting the required smart grid.

ANR Dark-Era (2021-2025)

Participants: Martin Quinson.

The future Square Kilometer Array (SKA) radio telescope poses unprecedented challenges to the underlying computational system. The instrument is expected to produce a sustained rate of Terabytes of data per second, mandating on-site pre-processing to reduce the size of data to be transferred. However, the electromagnetic noise of a traditional computing center would hinder the quality of the measurements if located near to the instrument. As a result, the Science Data Processor (SDP) pipeline will only have an energy budget of only 1 MWatt to execute a complex algorithm chain estimated at 250 Petaops/s. Because of these requirements, the SDP must be an innovative data-oriented infrastructure running on a disaggregated architecture combining standard HPC systems with dedicated accelerators such as GPU or FPGA.

The goal of the DarkEra project is to contribute to the performance assessment both in time and energy of new complex scientific algorithms on not-yet-existing complex computing infrastructures. To that extend, a prototyping tool will be developed for the prospective profiling of data-oriented applications during their development.

CARECloud (PEPR Cloud) (2023-2030)

Participants: Anne-Cécile Orgerie.

Cloud computing and its many variations offer users considerable computing and storage capacities. The maturity of virtualization techniques has enabled the emergence of complex virtualized infrastructures, capable of rapidly deploying and reconfiguring virtual and elastic resources in increasingly distributed infrastructures. This resource management, transparent to users, gives the illusion of access to flexible, unlimited and almost immaterial resources. However, the power consumption of these clouds is very real and worrying, as are their overall greenhouse gas (GHG) emissions and the consumption of critical raw materials used in their manufacture. In a context where climate change is becoming more visible and impressive every year, with serious consequences for people and the planet, all sectors (transport, building, agriculture, industry, etc.) must contribute to the effort to reduce GHG emissions. Despite their ability to optimize processes in other sectors (transport, energy, agriculture), clouds are not immune to this observation: the increasing slope of their greenhouse gas emissions must be reversed, otherwise their potential benefits in other sectors will be erased. This is why the CARECloud project (understanding, improving, reducing the environmental impacts of Cloud Computing) aims to drastically reduce the environmental impacts of cloud infrastructures.

NF-JEN (PEPR 5G and future networks) (2023-2030)

Participants: Anne-Cécile Orgerie.

Communication networks are often presented as a necessary means of reducing the impact on the environment of various sectors of industry. In practice, the roll-out of new generations of mobile broadband networks has required increased communication resources for wireless access networks. This has proved an effective approach in terms of performance but concerns remain about its energy cost and more generally its environmental impacts. Exposure to electromagnetic fields also remains a cause of concern despite existing protection limits. In the JEN (Just Enough Networks) project, we propose to develop just enough networks: network whose dimension, performance, resource usage and energy consumption are just enough to satisfy users needs. Along with designing energy-efficient and sober networks, we will provide multi-indicators models that could help policy-makers and inform the public debate.

Taranis (PEPR Cloud) (2023-2030)

Participants: Shadi Ibrahim, Nikos Parlavantzas, Guillaume Pierre.

New infrastructures, such as Edge Computing or the Cloud-Edge-IoT computing continuum, complicate the cloud landscape as they add new challenges related to resource diversity and heterogeneity (from small sensors to data centers/HPC, from low power networks to core networks), geographical distribution, as well as increased dynamicity and security needs, all under energy consumption and regulatory constraints. In order to efficiently exploit new infrastructures, we propose a strategy based on a significant abstraction of the application structure description to further automate application and infrastructure management. Thus, it will be possible to globally optimize used resources with respect to multi-criteria objectives (price, deadline, performance, energy, etc.) on both the user side (applications) and the provider side (infrastructures). This abstraction also includes the challenges related to facilitating application reconfiguration and to automatically adapting the use of resources. The Taranis project addresses these issues through four scientific work packages, each focusing on a phase of the application lifecycle: application and infrastructure description models, deployment and reconfiguration, orchestration, and optimization.

Members of the Myriads team will contribute to four sub-topics including (1) Decentralized, market-based application orchestration for Fog and IoT environments; (2) Realizing and optimizing Serverless Computing in the Edge-Cloud continuum; (3) Orchestrating multi-dimensional resources in the Edge-Cloud continuum; and (4) Resource Provisioning for stream data processing in the Fog.

10.2 Regional initiatives

ARED TDFDE (2022-2025)

Participants: Khaled Arsalane, Guillaume Pierre.

The PhD thesis of Khaled Arsalane is funded at 50% by the TDFDE project under the ARED program of Région Bretagne. This project investigates performance modeling and optimization of aggregation operators within geo-distributed data stream processing platforms.

ARED TANSIV (2022-2025)

Participants: Léo Cosseron, Martin Quinson, Louis Rilling, Matthieu Simonin.

The PhD thesis of Léo Cosseron is funded at 50% by the ARED program "Breizh Cyber Valley" of Région Bretagne and 50% by the Creach Labs. This project studies how to defeat malware evasion techniques based on network performance analysis by adding a virtual realistic networks to analysis environments.

11 Dissemination

11.1 Promoting scientific activities

Participants: Shadi Ibrahim, François Lemerrier, Anne-Cécile Orgerie, Nikos Parlavantzas, Guillaume Pierre.

11.1.1 Scientific events: organisation

General chair, scientific chair

- Anne-Cécile Orgerie was conference chair for ICT4S 2023: International Conference on Information and Communications Technology for Sustainability, Rennes, France.
- Guillaume Pierre is a member of the ACM/IFIP Middleware conference's Steering Committee.
- Shadi Ibrahim was the Tutorials deputy chair for ISC High Performance 2023, Hamburg, Germany.
- Shadi Ibrahim was a member of the steering committee for the 7th edition of the Workshop Performance and Scalability of Storage Systems, May 30 2023, Paris.
- Shadi Ibrahim was a member of the steering committee for ISC High Performance 2023, Hamburg, Germany.
- Shadi Ibrahim is a member of the steering committee of the International Parallel Data Systems Workshop.

Member of the organizing committees

- Shadi Ibrahim was publicity co-chair for the 53th International Conference on Parallel Processing (ICPP 2023), 2023, Utah, USA.

11.1.2 Scientific events: selection

Chair of conference program committees

- Guillaume Pierre was PC co-chair of the Future Compute Continuum track of IEEE/ACM CCGrid 2023, and PC co-chair of IEEE IC2E 2023.
- Shadi Ibrahim was program co-chair of the Workshop on Challenges and Opportunities of Efficient and Performant Storage Systems (CHEOPS@EuroSys2023), Rome, Italy.

Member of the conference program committees

- Anne-Cécile Orgerie was a member of the program committees of IEEE Cluster 2023, IEEE HiPC 2023, and HotCarbon 2023.
- Guillaume Pierre was a member of the program committees of EuroSys 2023, MASCOTS 2023, IEEE/ACM SEC 2023 and WSCC 2023.
- Shadi Ibrahim was a member of the program committees of SC 2023 (Research papers and research posters), IEEE IPDPS 2023, IEEE Cluster 2023, IEEE/ACM CCGrid 2023 (SCALE Challenge), Euro-Par 2023, IEEE/ACM UCC 2023, IEEE/ACM PDSW@SC 2023, and QuickPar 2023.
- Nikos Parlavantzas was a member of the program committees of IEEE/ACM UCC 2023, IEEE CloudCom 2023, IEEE ISPD 2023, JSSPP'23, and VHPC'23
- François Lemerrier was a member of the program committees of Algotel/Cores 2023

11.1.3 Journal

Member of the editorial boards

- Anne-Cécile Orgerie is a member of the editorial board of IEEE Transactions on Parallel and Distributed Systems.
- Shadi Ibrahim is an associate editor of IEEE Internet Computing Magazine.
- Shadi Ibrahim is an associate editor for High Performance Big Data Systems of Frontiers in High Performance Computing Journal.
- Shadi Ibrahim was a guest editor of IEEE Network – The Magazine of Global Internetworking: Special Issue on Interplay Between Machine Learning and Networking Systems.

Reviewer - reviewing activities

- Shadi Ibrahim was a reviewer for ACM Transactions on Storage and Springer Future Generation Computer Systems.
- Nikos Parlavantzas was a reviewer for the Journal of Grid Computing.

11.1.4 Invited talks

- Anne-Cécile Orgerie: “Measuring and modeling the energy consumption of servers”, invited presentation at the DIPOpt (Deep learning, image analysis, inverse problems, and optimization) workshop, Lyon, France, November 30, 2023.
- Anne-Cécile Orgerie: “Network-aware energy-efficient virtual machine management in distributed Cloud infrastructures with on-site photovoltaic production”, invited presentation at the Workshop on Scheduling Variable Capacity Resources for Sustainability, Paris, France, March 29, 2023.
- Anne-Cécile Orgerie: “Consommation énergétique et impacts environnementaux des systèmes distribués”, seminar at IMAG (Institut Montpellierain Alexander Grothendieck), virtually, December 4, 2023.
- Anne-Cécile Orgerie: “Consommation énergétique et impacts environnementaux des systèmes distribués”, seminar at ETIS (Equipes Traitement de l’Information et Systèmes), Cergy, March 9, 2023.
- Anne-Cécile Orgerie: “Consommation énergétique et impacts environnementaux des systèmes distribués”, seminar at Colloquium of LORIA (Laboratoire lorrain de recherche en informatique et ses applications), Nancy, March 2, 2023.
- Anne-Cécile Orgerie: “Consommation énergétique et impacts environnementaux des systèmes distribués”, seminar at Federation Normastic day, Caen, February 7, 2023.
- Anne-Cécile Orgerie: “Sciences informatiques écoresponsables”, invited presentation at the INS2I CNRS conference “Vers une informatique plus durable”, Paris, France, November 27, 2023.
- Anne-Cécile Orgerie: “Impacts environnementaux du numérique”, invited presentation with Olivier Ridoux at the cati SICPA meeting of INRAE, Saint-Gilles, France, October 4, 2023.
- Anne-Cécile Orgerie: “Sobriété et enjeux technologiques : vers quelle durabilité?”, workshop chair at the seminar on planetary boundaries and sustainability stakes of CNRS, Paris, France, July 4, 2023.
- Anne-Cécile Orgerie: “Impacts environnementaux du numérique et structures de recherche associées”, invited presentation at the seminar on energy transition and society of CNRS, Paris, France, April 28, 2023.

- Anne-Cécile Orgerie: “Sciences informatiques éco-responsables”, invited presentation at the new recruits day of INS2I, Paris, France, March 20, 2023.
- Shadi Ibrahim: "Scalable and Efficient Big Data Management in Distributed Systems: Addressing performance variability for Data processing in the Cloud ", invited presentation at the Workshop on Scheduling Variable Capacity Resources for Sustainability, Paris, France, March 29, 2023.

11.1.5 Leadership within the scientific community

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

11.1.6 Scientific expertise

- Anne-Cécile Orgerie was a member of the selection committees for an assistant professor position at Université Grenoble Alpes and Inria research scientist positions (DR) in 2023.
- Shadi Ibrahim was a member of the [ACM Heidelberg Laureate Forum \(HLF\) 2023 Young Researcher Selection Committee](#).
- Shadi Ibrahim was a project reviewer for the French National Research Agency (ANR) AAPG 2023.

11.1.7 Research administration

- Anne-Cécile Orgerie is an officer (chargée de mission) for the IRISA cross-cutting axis on Green IT.
- Anne-Cécile Orgerie was member of the Inria Evaluation Committee.
- Anne-Cécile Orgerie is member of the steering committee of the CNRS GDR RSD.

11.2 Teaching - Supervision - Juries

Participants: Marin Bertier, Shadi Ibrahim, François Lemercier, Anne-Cécile Orgerie, Nikos Parlavantzas, Jean-Louis Pazat, Guillaume Pierre, Martin Quinson, Cédric Tedeschi.

11.2.1 Teaching

- Bachelor: Marin Bertier, Networks, Département Informatique L3, Insa Rennes.
- Bachelor: Marin Bertier, C Language Département Informatique L3, Insa Rennes.
- Bachelor: Marin Bertier, C Language, Département Mathématique L3, Insa Rennes.
- Bachelor: Nikos Parlavantzas, Theoretical and practical study, Département Informatique L3, Insa Rennes.
- Bachelor: Nikos Parlavantzas, Networks, Département Informatique L3, Insa Rennes.
- Bachelor: Nikos Parlavantzas, Multi-core architectures, Département Informatique L3, Insa Rennes.
- Bachelor: Jean-Louis Pazat, Introduction to programming L1, Département STPI, INSA de Rennes.
- Bachelor: Jean-Louis Pazat, High Performance Computing, Département Informatique L3, Insa Rennes.
- Bachelor: Jean-Louis Pazat, High Performance Computing, Département Mathématiques L3, Insa Rennes.

- Bachelor: Guillaume Pierre, Systèmes Informatiques, L3 MIAGE, Univ. Rennes 1.
- Bachelor: Guillaume Pierre, Systèmes d'exploitation, L3 Informatique, Univ. Rennes 1.
- Bachelor: Martin Quinson, Architecture et Systèmes, 60 hETD, L3 Informatique, ENS Rennes.
- Bachelor: Martin Quinson, Pedagogy, 15 hETD, L3 Informatique, ENS Rennes.
- Bachelor: Cédric Tedeschi, Cloud and networks, L3, Univ. Rennes 1.
- Bachelor: François Lemerrier, Networking, Services and Protocols, L3, Univ. Rennes 1.
- Master: Marin Bertier, Operating Systems, Département Informatique M1, INSA de Rennes
- Master: Marin Bertier, Distributed systems, Département Informatique M2, INSA de Rennes
- Master: Anne-Cécile Orgerie, Green ICT, 4.5 hETD, M2, Telecom SudParis Evry.
- 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, Master for Smart Data Science, ENSAI, Bruz.
- Master: Nikos Parlavantzas, SQL, M2, Master for Smart Data Science, 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, C++ system programming (20h ETD), ENS Rennes.
- Master: Martin Quinson, Préparation à l'Agrégation d'Informatique (Networking, 20h ETD), ENS Rennes.
- Master: Martin Quinson, Scientific Outreach, M2, 30 hEDT, ENS Rennes.
- Master: Cédric Tedeschi, Concurrency in Systems and Networks, M1, Univ. Rennes 1.
- Master: Cédric Tedeschi, Service Technology, M1, Univ. Rennes 1.
- Master: Cédric Tedeschi, Parallel Programming, M1, Univ. Rennes 1.
- Master: Cédric Tedeschi, Advanced Cloud Infrastructures, M2, Univ. Rennes 1.
- Master: Shadi Ibrahim, Cloud Computing and Hadoop Technologies, 36hETD, M2 : Statistics for Smart Data, ENSAI, Bruz.
- Master: Shadi Ibrahim, Cloud1 (Map-Reduce), 17.5hETD, M2 , IMT-Atlantique, Nantes.
- Master: Shadi Ibrahim, Distributed Big Data, 45hETD, M2, ENSAI, Bruz.
- Master: François Lemerrier, Networking, Services and Protocols, M1, Univ. Rennes 1.
- Master: François Lemerrier, Software Engineering Project, M1, Univ. Rennes 1.

11.2.2 Supervision

- PhD defended: Clément Courageux-Sudan, “Reducing the energy consumption of Internet of Things”, defended in December 2023, supervised by Anne-Cécile Orgerie and Martin Quinson.
- PhD defended: Adrien Gougeon, “Designing an energy-efficient communication network for the dynamic and distributed control of the electrical grid”, defended in January 2023, supervised by Anne-Cécile Orgerie and Martin Quinson.
- PhD in progress: Emmanuel Gnibga, “Modeling and optimizing edge computing infrastructures and their electrical system”, started in November 2021, supervised by Anne-Cécile Orgerie and Anne Blavette.
- PhD in progress: Vladimir Ostapenco, "Modeling and design of a framework and its Gantt Chart to manage physical and logical environmental levers", started in December 2021, supervised by Laurent Lefèvre and Anne-Cécile Orgerie.
- PhD in progress: Maxime Agusti, "Observation of baremetal co-location platforms, models and catalog proposal to reduce energy consumption", started in December 2021, supervised by Eddy Caron, Laurent Lefèvre and Anne-Cécile Orgerie.
- PhD in progress: Matthieu Silard, "Co-optimization of electrical and communication networks", started in February 2023, supervised by Anne-Cécile Orgerie, Nicolas Montavont and Georgios Papadopoulos.
- PhD in progress: Chih-Kai Huang, "Scalable decentralized fog computing platforms", started in 2021, supervised by Guillaume Pierre.
- PhD in progress: Khaled Arsalane, "Performance modeling of agregation operators in geo-distributed data stream processing systems", started in October 2022, supervised by Guillaume Pierre.
- PhD in progress: Govind Kovilkatt Panickerveetil, "Energy-efficient data stream processing", started in April 2023, co-supervised by Guillaume Pierre and Romain Rouvoy.
- PhD in progress: Ammar Kazem, "In-natura data processing systems for environmental observation under energy constraints", started in October 2023, co-supervised by Guillaume Pierre and Laurent Longuevergne.
- PhD in progress: Davaadorj Battulga, "Stream Processing Pipelines in Fog Environment", started in September 2018, supervised by Cédric Tedeschi and Daniele Miorandi.
- PhD in progress: Volodia Parol-Guarino, "Flexible resource allocation for FaaS applications in the fog", started in October 2022, supervised by Nikos Parlavantzas.
- PhD in progress: Mohamed Cherif Zouaoui Latreche, "Balancing Performance and Sustainability for FaaS in the Fog", started in October 2023, supervised by Nikos Parlavantzas and Hector Duran-Limon.
- PhD in progress: Mohammad Rizk, “Reliable and cost-efficient data placement and repair in P2P storage over immutable data”, started in November 2023, supervised by Shadi Ibrahim, Thomas Lambert, and Guillaume Pierre.
- PhD in progress: Quentin Acher, “Management of mutable data over P2P storage”, started in September 2023, supervised by Shadi Ibrahim and Claudia-Lavinia Ignat.
- PhD in progress: Mathieu Laurent, “Efficient verification of asynchronous distributed programs“, started in October 2023, supervised by Martin Quinson and Thierry Jéron.
- PhD in progress: Leo Cosseron, “Time-Accurate Network Simulation Interconnecting VMs with Hardware Virtualization Towards Stealth Analysis“, started in October 2022, supervised by Martin Quinson and Louis Rilling.

11.2.3 Juries

- Anne-Cécile Orgerie was a reviewer of the PhD manuscript of Edouard Guégain (University of Lille), September 29, 2023.
- Anne-Cécile Orgerie was a member of the PhD defense of Miguel Felipe Silva Vasconcelos (University Grenoble Alpes), December 20, 2023.
- Guillaume Pierre was a reviewer of the PhD defense of Rafaela Brum (Sorbonne University), November 29th 2023.
- Guillaume Pierre was a member of the PhD defense of Kiranpreet Kaur (Conservatoire National des Arts et Métiers), September 19th, 2023.
- Martin Quinson was a reviewer of the PhD manuscript of Julien Emmanuel (ENS Lyon), March 8, 2023.

11.3 Popularization

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

11.3.1 Internal or external Inria responsibilities

- Shadi Ibrahim is co-organizing SCI-Rennes seminar: a monthly scientific seminars for all staff at Inria research centre at Rennes University.

11.3.2 Articles and contents

- Guillaume Pierre participated in an article in "CNRS le journal": [Quand le cloud se fait diffus](#), January 2023.

11.3.3 Education

- “L codent L créent” is an outreach program to send PhD students to teach Python to middle school students in 8 sessions of 45 minutes. Tassadit Bouadi (Lacodam), Camille Maumet (Empenn) and Anne-Cécile Orgerie (Myriads) are coordinating the local version of this program, initiated in Lille. The first session in Rennes occurred in April 2019, and a new session (the 5th) occurred in 2023. The program is currently supported by: Fondation Blaise Pascal, ED MathSTIC, ENS de Rennes, Université Rennes 1 and Fondation Rennes 1.

12 Scientific production

12.1 Major publications

- [1] H. Casanova, A. Legrand, M. Quinson and F. Suter. ‘SMPI Courseware: Teaching Distributed-Memory Computing with MPI in Simulation’. In: *EduHPC-18 - Workshop on Education for High-Performance Computing*. Dallas, United States, Nov. 2018, pp. 1–10. URL: <https://hal.inria.fr/hal-01891513>.

12.2 Publications of the year

International journals

- [2] J. L. Bez, S. Byna and S. Ibrahim. ‘I/O Access Patterns in HPC Applications: A 360-Degree Survey’. In: *ACM Computing Surveys* (30th July 2023). DOI: [10.1145/3611007](https://doi.org/10.1145/3611007). URL: <https://hal.science/hal-04173899>.

- [3] Y. Bouizem, D. Dib, N. Parlavantzas and C. Morin. ‘Integrating request replication into FaaS platforms: an experimental evaluation’. In: *Journal of Cloud Computing: Advances, Systems and Applications* 12.1 (22nd June 2023), pp. 1–20. DOI: [10.1186/s13677-023-00457-z](https://doi.org/10.1186/s13677-023-00457-z). URL: <https://inria.hal.science/hal-04159510>.
- [4] H. Fan, Y. Ye, S. Ibrahim, Z. Huang, X. Li, W. Xue, S. Wu, C. Yu, X. Shi and H. Jin. ‘QoS-pro: A QoS-enhanced Transaction Processing Framework for Shared SSDs’. In: *ACM Transactions on Architecture and Code Optimization* (14th Nov. 2023), pp. 1–25. DOI: [10.1145/3632955](https://doi.org/10.1145/3632955). URL: <https://hal.science/hal-04389048>.
- [5] W. E. Gnibga, A. Blavette and A.-C. Orgerie. ‘Renewable Energy in Data Centers: the Dilemma of Electrical Grid Dependency and Autonomy Costs’. In: *IEEE Transactions on Sustainable Computing* (2023), pp. 1–13. DOI: [10.1109/TSUSC.2023.3307790](https://doi.org/10.1109/TSUSC.2023.3307790). URL: <https://hal.science/hal-04189173>.
- [6] V. Ostapenco, L. Lefèvre, A.-C. Orgerie and B. Fichel. ‘Modeling, evaluating, and orchestrating heterogeneous environmental leverages for large-scale data center management’. In: *International Journal of High Performance Computing Applications* 37.3-4 (2023). DOI: [10.1177/10943420231172978](https://doi.org/10.1177/10943420231172978). URL: <https://hal.science/hal-04047008>.
- [7] K. Toczé, A. J. Fahs, G. Pierre and S. Nadjm-Tehrani. ‘VioLinn: Proximity-aware Edge Placement with Dynamic and Elastic Resource Provisioning’. In: *ACM Transactions on Internet of Things* 4.1 (23rd Feb. 2023), pp. 1–31. DOI: [10.1145/3573125](https://doi.org/10.1145/3573125). URL: <https://inria.hal.science/hal-03869221>.

International peer-reviewed conferences

- [8] Q. Acher, C.-L. Ignat and S. Ibrahim. ‘Quantifying the Performance of Conflict-free Replicated Data Types in InterPlanetary File System’. In: *Middleware 2023 Companion Proceedings. DICG 2023 - 4th International Workshop on Distributed Infrastructure for Common Good*. Bologna, Italy, 2023, pp. 1–6. DOI: [10.1145/3631310.3633488](https://doi.org/10.1145/3631310.3633488). URL: <https://inria.hal.science/hal-04337761>.
- [9] C. Courageux-Sudan, A.-C. Orgerie and M. Quinson. ‘A Wi-Fi Energy Model for Scalable Simulation’. In: *24th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*. WoWMoM 2023 - 24th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks. Boston (MA), United States, 2023, pp. 1–10. URL: <https://hal.science/hal-04055720>.
- [10] W. E. Gnibga, A. Blavette and A.-C. Orgerie. ‘Latency, Energy and Carbon Aware Collaborative Resource Allocation with Consolidation and QoS Degradation Strategies in Edge Computing’. In: *ICPADS 2023 - IEEE International Conference on Parallel and Distributed Systems*. ICPADS 2023 - IEEE International Conference on Parallel and Distributed Systems. Hainan, China: IEEE, Dec. 2023, pp. 1–10. URL: <https://hal.science/hal-04275783>.
- [11] C.-K. Huang and G. Pierre. ‘Acala: Aggregate Monitoring for Geo-Distributed Cluster Federations’. In: *SAC 2023 - 38th ACM/SIGAPP Symposium On Applied Computing*. Tallinn, Estonia: ACM, 27th Mar. 2023, pp. 1–9. URL: <https://inria.hal.science/hal-03899133>.
- [12] C.-K. Huang and G. Pierre. ‘AdapPF: Self-Adaptive Scrape Interval for Monitoring in Geo-Distributed Cluster Federations’. In: *ISCC 2023 - 28th IEEE Symposium on Computers and Communications*. Tunis, Tunisia: IEEE, 9th July 2023, pp. 1–7. DOI: [10.1145/3555776.3577716](https://doi.org/10.1145/3555776.3577716). URL: <https://inria.hal.science/hal-04103309>.
- [13] M. Jay, V. Ostapenco, L. Lefèvre, D. Trystram, A.-C. Orgerie and B. Fichel. ‘An experimental comparison of software-based power meters: focus on CPU and GPU’. In: *CCGrid 2023 - 23rd IEEE/ACM international symposium on cluster, cloud and internet computing*. Bangalore, India: IEEE, 2023, pp. 1–13. URL: <https://inria.hal.science/hal-04030223>.
- [14] G. KP, G. Pierre and R. Rouvoy. ‘Studying the Energy Consumption of Stream Processing Engines in the Cloud’. In: *IC2E 2023 - 11th IEEE International Conference on Cloud Engineering*. Boston (MA), United States: IEEE, 25th Sept. 2023, pp. 1–9. URL: <https://inria.hal.science/hal-04164074>.

- [15] A. Pagliari and G. Pierre. ‘TransScale: Combined-Approach Elasticity for Stream Processing in Fog Environments’. In: *Mobile Cloud 2023 - 11th IEEE International Conference on Mobile Cloud Computing, Services and Engineering*. Athens, Greece: IEEE, 17th July 2023, pp. 1–8. URL: <https://inria.hal.science/hal-04126361>.
- [16] V. Parol-Guarino and N. Parlavantzas. ‘GIRAFF: Reverse Auction-based Placement for Fog Functions’. In: *WoSC ’23: 9th International Workshop on Serverless Computing*. Bologna, Italy: ACM, 11th Dec. 2023, pp. 53–58. DOI: [10.1145/3631295.3631402](https://doi.org/10.1145/3631295.3631402). URL: <https://inria.hal.science/hal-04384516>.
- [17] J. Paturel, C. Quinson, M. Quinson and S. Rokicki. ‘SmolPhone: a smartphone with energy limits’. In: *IGSC 2023 - 14th International Green and Sustainable Computing*. Toronto, Canada, 28th Oct. 2023, p. 4. URL: <https://inria.hal.science/hal-04156447>.
- [18] D. Rosendo, K. Keahey, A. Costan, M. Simonin, P. Valduriez and G. Antoniu. ‘KheOps: Cost-effective Repeatability, Reproducibility, and Replicability of Edge-to-Cloud Experiments’. In: *ACM REP ’23: Proceedings of the 2023 ACM Conference on Reproducibility and Replicability*. REP 2023 - ACM Conference on Reproducibility and Replicability. Santa Cruz, CA, United States: ACM, 28th June 2023, pp. 62–73. DOI: [10.1145/3589806.3600032](https://doi.org/10.1145/3589806.3600032). URL: <https://hal.science/hal-04157720>.
- [19] H.-Y. Tao, C.-K. Huang and S.-H. Shen. ‘A Low-overhead Network Monitoring for SDN-Based Edge Computing’. In: *ISCC 2023 - 28th IEEE Symposium on Computers and Communications*. Tunis, Tunisia: IEEE, 9th July 2023, pp. 1–6. URL: <https://inria.hal.science/hal-04105471>.
- [20] H. Wu, J. Deng, H. Fan, S. Ibrahim, S. Wu and H. Jin. ‘QoS-Aware and Cost-Efficient Dynamic Resource Allocation for Serverless ML Workflows’. In: *2023 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IPDPS - 2023 IEEE International Parallel and Distributed Processing Symposium. St. Petersburg, United States: IEEE, 18th July 2023, pp. 886–896. DOI: [10.1109/IPDPS54959.2023.00093](https://doi.org/10.1109/IPDPS54959.2023.00093). URL: <https://hal.science/hal-04389016>.

Doctoral dissertations and habilitation theses

- [21] A. Gougeon. ‘Optimizing a dynamic and energy efficient network piloting the electrical grid’. Université de Rennes, 11th Jan. 2023. URL: <https://theses.hal.science/tel-04086397>.

Reports & preprints

- [22] L. Cosseron, M. Quinson, L. Rilling and M. Simonin. *Hiding Virtual Machine Introspection Pauses in Networked Sandboxes with Network Simulation*. RR-9528. Inria Rennes - Bretagne Atlantique & IRISA, 13th Nov. 2023, pp. 1–14. URL: <https://inria.hal.science/hal-04282253>.

12.3 Cited publications

- [23] R.-A. Cherrueau, M. Simonin and A. Van Kempen. ‘EnosStack: A LAMP-like stack for the experimenter’. In: *INFOCOM WKSHPs 2018 - IEEE International Conference on Computer Communications*. Honolulu, United States: IEEE, Apr. 2018, pp. 336–341. DOI: [10.1109/INFCOMW.2018.8407024](https://doi.org/10.1109/INFCOMW.2018.8407024). URL: <https://hal.inria.fr/hal-01916484>.
- [24] D. Rosendo, P. Silva, M. Simonin, A. Costan and G. Antoniu. ‘E2Clab: Exploring the Computing Continuum through Repeatable, Replicable and Reproducible Edge-to-Cloud Experiments’. In: *Cluster 2020 - IEEE International Conference on Cluster Computing*. Kobe, Japan, Sept. 2020, pp. 1–11. DOI: [10.1109/CLUSTER49012.2020.00028](https://doi.org/10.1109/CLUSTER49012.2020.00028). URL: <https://hal.archives-ouvertes.fr/hal-02916032>.