UMR IRISA

Activity Report 2024

Team MAGELLAN

Reliable and Responsible Decentralized Computing Infrastructures

Joint team with Centre Inria de l'Université de Rennes

D1 – Large Scale an Secure Systems

Université de Rennes | rennes | INSA INSTITUT NATIONAL DES SCIENCES APPLIQUÉES RENNES | cnrs | Inria | Inserm | CentraleSupélec | IMT Atlantique Bretagne-Pays de la Loire École Mines-Télécom | Université Bretagne Sud ubs

Contents

Project-Team MAGELLAN **1**

1 Team members, visitors, external collaborators **3**

2 Overall objectives **4**

3 Research program **5**

4 Application domains **7**

5 Social and environmental responsibility **7**
5.1 Footprint of research activities . . . 7
5.2 Impact of research results . . . 7

6 Highlights of the year **8**
6.1 Awards . . . 8

7 New software, platforms, open data **8**
7.1 New software . . . 8
7.1.1 Tansiv . . . 8
7.1.2 SimGrid . . . 8
7.1.3 EnOSlib . . . 9

8 New results **10**
8.1 Reliable decentralized computing infrastructures . . . 10
8.1.1 Federated fog/cloud infrastructure design . . . 10
8.1.2 Fog/cloud storage services . . . 11
8.1.3 Reliable fog platforms in adverse natural environments . . . 11
8.2 Reliable decentralized application runtimes . . . 12
8.2.1 Reliable serverless runtimes . . . 12
8.2.2 Dynamic platform adaptation for Federated Learning . . . 12
8.2.3 Reliable data stream processing runtimes . . . 12
8.2.4 Decentralized resource provisioning for serverless applications . . . 13
8.2.5 Data processing at scale . . . 13
8.2.6 Mitigation of tail processing latency . . . 14
8.3 Distributed infrastructure frugality . . . 14
8.3.1 Energy accounting of digital services . . . 14
8.3.2 Water consumption of Cloud data centers . . . 15
8.3.3 Steering efficiency for distributed infrastructures . . . 15
8.3.4 Frugal mobile computing . . . 16
8.4 Evaluation methodologies and tools . . . 17
8.4.1 Understanding software infrastructures . . . 17
8.4.2 Predictive models and simulation frameworks . . . 17
8.4.3 Sustainable research ecosystem . . . 17

9 Bilateral contracts and grants with industry **18**
9.1 Bilateral contracts with industry . . . 18

10 Partnerships and cooperations **19**
10.1 International initiatives . . . 19
10.1.1 Associate Teams in the framework of an Inria International Lab or in the framework of an Inria International Program . . . 19
10.2 International research visitors . . . 19
10.2.1 Visits of international scientists . . . 19
10.3 European initiatives . . . 20

10.3.1 Digital Europe . . . 20
10.4 National initiatives . . . 21
10.5 Regional initiatives . . . 27

11 Dissemination **28**
11.1 Promoting scientific activities . . . 28
11.1.1 Scientific events: organisation . . . 28
11.1.2 Scientific events: selection . . . 28
11.1.3 Journal . . . 29
11.1.4 Invited talks . . . 29
11.1.5 Leadership within the scientific community . . . 30
11.1.6 Scientific expertise . . . 30
11.1.7 Research administration . . . 30
11.2 Teaching - Supervision - Juries . . . 30
11.2.1 Teaching . . . 30
11.2.2 Supervision . . . 31
11.2.3 Juries . . . 32
11.2.4 Productions (articles, videos, podcasts, serious games, ...) . . . 33
11.2.5 Participation in Live events . . . 33
11.2.6 Others science outreach relevant activities . . . 33

12 Scientific production **33**
12.1 Major publications . . . 33
12.2 Publications of the year . . . 34
12.3 Cited publications . . . 36

Project-Team MAGELLAN

Creation of the Project-Team: 2024 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, Researcher]

Faculty Members

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

Post-Doctoral Fellows

- Meysam Mayahi [CNRS, Post-Doctoral Fellow, from Nov 2024]
- Ophélie Renaud [ENS RENNES, Post-Doctoral Fellow, from Oct 2024]

PhD Students

- Quentin Acher [INRIA, hosted with the COAST team]
- Maxime Agusti [OVH, hosted with ENS Lyon]
- Khaled Arsalane [UNIV RENNES]
- Matteo Chancerel [CNRS, from Sep 2024]
- Leo Cosseron [ENS RENNES]
- Wedan Emmanuel Gnibga [CNRS]
- Chih-Kai Huang [UNIV RENNES]
- Ammar Kazem [UNIV RENNES]
- Govind Kovilkkatt Panickerveetil [INRIA]
- Mohamed Cherif Zouaoui Latreche [INSA RENNES]
- Mathieu Laurent [ENS RENNES]
- Pablo Leboulanger [CNRS, from Sep 2024]
- Vladimir Ostapenco [INRIA, hosted at ENS Lyon]
- Volodia Parol-Guarino [INRIA]
- Mohammad Rizk [INRIA]
- Matthieu Silard [CNRS]
- Thomas Stavis [INRIA, from Oct 2024, hosted at ENS Lyon]
- Marc Tranzer [INRIA]
- Mathis Valli [INRIA, hosted with the KerData team]

Technical Staff

- Victorien Elvinger [INRIA, Engineer, from Dec 2024]
- Matthieu Nicolas [INRIA, Engineer]
- Loris Penven [UNIV RENNES, Engineer, in collaboration with SED]
- Matthieu Simonin [INRIA, Engineer, in collaboration with SED]

Interns and Apprentices

- Corentin Durand [INRIA, Intern, from Jun 2024 until Jul 2024]

Administrative Assistant

- Amandine Seigneur [INRIA, from Sep 2024]

2 Overall objectives

Large-scale computing platforms constitute an essential infrastructure for the development of modern societies. Similar to transportation infrastructure and various other utilities such as water and electricity supply, computation has become indispensable for providing citizens, companies and public services with countless services in a wide variety of sectors including information, entertainment, education, commerce, industry, health, government, defense, science and many others [33, 35].

Over the past decades these large-scale computing infrastructures have undergone many evolutions which made them alternate between centralized and decentralized architectures. Starting with (centralized) mainframe computers, we can cite (decentralized) personal computers, followed by (centralized) client-server and cluster-based systems, (decentralized) grid computing platforms, and (centralized) cloud data centers. Each of these evolutions provided users with additional capabilities and required addressing a fresh set of scientific and technical challenges.

The latest of these evolutions is the current very strong trend toward decentralized geo-distributed virtualized infrastructures. Instead of solely hosting their compute, storage and networking resources in a handful of very large data centers located far from their end users, decentralized infrastructures extend the data centers with a myriad of additional resources broadly distributed across a wide geographical area. The main motivations for these evolutions are:

- **Proximity:** the development of highly interactive services such as virtual and augmented reality, and of Internet-of-Things applications, require the presence of computing resources located in the vicinity of the end users and their devices [30]. It has been predicted that in 2025, 75% of enterprise-generated data will be created and processed outside a centralized cloud, or a data center; in 2018 it was only 10% [38]. To process these data streams close to their origin, and thereby avoid wasting precious resources such as long-distance network bandwidth, fog and edge computing platforms extend traditional cloud data centers with additional resources located as close as possible from the end users.
- **Resource aggregation:** as many enterprises are dispersed in many geographical locations while providing global online services, it is becoming increasingly necessary to federate the companies' resources in their different premices to form a single global infrastructure. Most enterprises further extend their private infrastructures for short or long periods of time with additional resources from one or more public cloud(s), thereby creating decentralized hybrid cloud infrastructures [34].
- **Disaster tolerance:** centralized infrastructures have proven to be vulnerable to major outages in one or more co-located datacenters due to accidents (e.g., the OVHCloud data center fire in March 2021 [39]) or natural disasters (e.g., Hurricane Sandy in October 2012 [31]). Decentralizing the computing infrastructures is an efficient way to mitigate such events, as usually a significant fraction of the system remains operational despite the presence of local outages and can ensure business continuity.

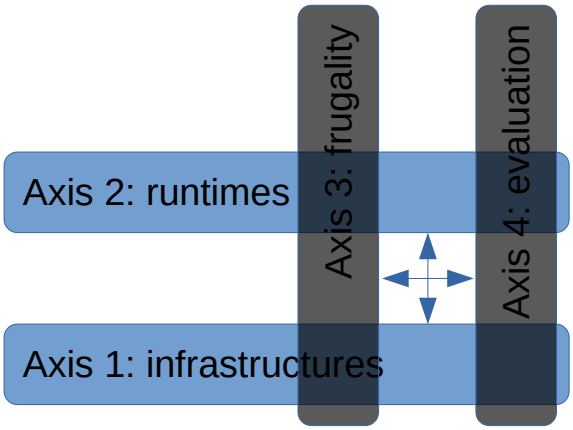


Figure 1: Magellan research axes organization.

- **Regulatory constraints:** with the increasing public sensitivity about security and privacy concerns, a large number of national or regional regulations are being enacted to restrict the location where sensitive data in domains such as military, health, finance and employment may be stored and processed [36]. Decentralized computing infrastructures are a natural solution to respect such regulations while maintaining a global management of such data across the geographical boundaries [37].

The development of reliable large-scale decentralized virtualized computing infrastructures however remains in its infancy [32]. Depending on the use cases and the motivations for decentralization, very different and often ad-hoc solutions are being used. However, dedicating a specific set of devices to a single service negates the economies of scale delivered by the multi-tenancy and statistical multiplexing principles of cloud computing, as it essentially brings these services back to a pre-cloud era where every application had to be provisioned individually with its own dedicated hardware.

Developing complex applications capable of exploiting the possibilities created by decentralized computing infrastructures also remains extremely difficult. In a geo-distributed infrastructure the application resources are often located close to their users but necessarily far from each other (in geographical as well as networking distance terms). Failures and large performance variability are the norm, whether they apply to a single node, a group of nodes, or the networking infrastructure between them. Large-scale infrastructures often host many unrelated applications which may compete for the usage of finite resources. This creates a pressure to design and develop new data management and middleware solutions which provide simple and powerful abstractions to their users while automating the difficult resource management issues and competition for resources created by geo-distribution.

Nowadays, one cannot design large-scale computing infrastructures without considering their social and environmental impact. Decentralized computing infrastructures inherit many energy-related properties from their centralized counterparts, but they also feature a number of unique challenges and opportunities. On the one hand, decentralization promotes the usage of smaller (and potentially less energy-efficient) data centers; but on the other hand it creates opportunities for strategically selecting resources, for example based on the (un)availability of renewable energy in different locations.

Finally, our field of science designs systems that are so complex that the only ways to fully understand them are to treat them as natural phenomena and to make use of advanced experimental methodologies to evaluate their strengths and weaknesses [40]. It is therefore part of our scientific responsibilities to participate in the development of robust, open and reproducible evaluation methodologies for our systems, being based on experimentation or simulations.

3 Research program

The Magellan team is organized along four main research axes. As shown in Figure 1, the first two axes respectively address lower-level and higher-level systems-related topics whereas the other two are

transversal and cross-cutting.

Axis 1: Reliable decentralized computing infrastructure. Decentralized computing infrastructures constitute the foundation which supports virtualized computing, storage and networking services upon which a wide range of sophisticated geo-distributed applications may be developed and executed. They should be scalable according to the number of nodes and their broad geographical distribution, reliable in the face of fluctuating node and network performance and availability, and easy to maintain and to operate.

Within this axis we plan to focus on the design of scalable and reliable virtualized execution platforms, application lifecycle management, federated infrastructure design, multi-tenancy, networking services and data management. A long-term goal is to enable the development of generic fog computing platforms which may eventually become public services operated by local authorities (e.g., cities, regions) to serve a large number of requirements stemming from their own usage as well as those from local companies and citizens. Beyond the performance- and efficiency-related motivations for such platforms (and their associated research challenges regarding scalability, stability, usability, resource management, data management etc.), we expect that sovereignty requirements regarding application's data and the infrastructures which manage them are going to become increasingly important from a social as well as scientific point of view.

Axis 2: Reliable decentralized application runtimes. Decentralized application runtimes should support developing, deploying and managing complex applications in a simple and natural way, while hiding the complexities of operating the underlying infrastructure. The runtimes should support automatically managing the quality requirements of applications and responding to changes in operating conditions, while cost-effectively provisioning resources from the decentralized infrastructure.

Within this axis we plan to pursue the ongoing work which aims to deliver a set of middleware environments to facilitate the development and operation of decentralized applications which exploit the underlying decentralized infrastructures to the fullest extent. The long-term objective is to make the development and operation of these applications no more complex than the development of cloud-based applications is today. We expect to invest our efforts on middlewares which support programming models such as data stream processing and function-as-a-service computing, where the main research challenges deal with automated application management, resource provisioning, auto-scaling, resource and data sharing within and across data centers, handling data bursts and mitigating stragglers.

Axis 3: Distributed infrastructure frugality. Within this axis, the Magellan team aims to reduce as much as possible the environmental impact of large-scale fog/cloud platforms. Understanding and reducing the environmental impact of the full life-cycle of our hardware and software resources is a difficult challenge. For instance, the manufacturing phase is dominating in the environmental impact of many ICT devices. Reducing their energy consumption during their use phase therefore does not necessarily guarantee lower environmental impacts, and increasing the device lifetime is crucial. As many such devices are deeply integrated within complex Cloud infrastructures, reducing the energy consumption of one part may in fact increase it for another part. In this context, we plan to further characterize the energy consumption of digital infrastructures, and to leverage the collected information to improve the efficiency of these infrastructures.

Axis 4: Evaluation methodologies and tools. Cloud and Fog 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.

Within this axis, the Magellan team aims to help professionalize the scientific evaluations in our research domain. We plan to push further our expertise in experimentation and simulation by combining these approaches. Building a mathematical surrogate of a given system (i.e., a simulation model of

the real system) constitutes a highly interesting challenge: the assessment and improvement of the model provides a deeper understanding of the system, while the resulting surrogate enables accurate performance predictions without any execution on real platforms. Co-developing the real system and its surrogate model in the long term is a very promising direction to efficiently design, build and operate the complex systems that constitute the modern large-scale distributed infrastructures.

4 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,
- Natural environment monitoring,
- Smart grids,
- Energy and sustainable development,
- Home IoT applications,
- Bio-informatics applications,
- Data science applications.

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.

Matthieu Simonin is involved in the CNRS GDR Labos 1point5. The GDR launches a scientific study on the environmental footprint of French public research and is making several tools available on a dedicated platform. 1000 labs in France are assessing their carbon footprint using the labos1point5's tools.

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.

In the context of the ANR EDEN4SG project, 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 ANR 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 Magellan 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.

Finally, the Magellan team involved in the Inria Alvearium challenge, in collaboration with the Hive company. The goal is to participate in the design of a scalable and reliable peer-to-peer storage system that is deployed on users' machines, thereby extending cloud storage and reducing the energy overhead and the environmental impact of the stored data.

6 Highlights of the year

- The Magellan team was created on January 1st 2024 following the termination of the Myriads team.
- Shadi Ibrahim served as General Chair of the 36th International Conference on Scientific and Statistical Database Management (SSDBM) which took place in Rennes on July 10-12th 2024 [20].

6.1 Awards

- Anne-Cécile Orgerie received the Lovelace-Babbage award from the French Academy of Sciences, in collaboration with the Société Informatique de France [1].

7 New software, platforms, open data

7.1 New software

7.1.1 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 garantee accurate arrival and transfer of network packets while ensuring realistic time flow within the VMs. The project can leverage several frameworks for simulating the data (SimGrid or ns-3) and several virtualization solutions to encapsulate the application, intercept the network traffic and enforce the interruption decision (Qemu in emulation mode, KVM and Xen with hardware-accelerated virtualization). Tansiv can be used in various situations: malware analysis (e.g. to defeat malware evasion technique based on network timing measures) or analysis of an application on a geo-distributed context.

Contact: Louis Rilling

Partner: DGA-MI

7.1.2 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.

[1] Anne-Cécile Orgerie : le numérique au défi des économies d'énergie

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.

Release Contributions: The Anne of Brittany release (she became a Duchess 536 years ago).

* Various improvements and unification of the simulation APIs

* MC: Enable the verification of Python programs, and of condvars and iprobe calls.

* MC: Exhibit the critical transition when a failure is found.

* (+ internal refactoring and bug fixes)

News of the Year: There was one major release in 2024. On modeling aspects, we further improved and unified the relevant APIs. The MPI and S4U interfaces are now more consistent and integrated, and a new plugin was introduced to model disk arrays (JBODs). We also pursued our efforts to improve the overall framework, through bug fixes, code refactoring and other software quality improvements. In particular, interfaces that were deprecated since almost a decade were removed to ease the maintenance burden on our community.

On the model checking front, we are still working hard to find our first "wild bug", a.k.a. existing bug in an arbitrary piece of code while all bugs found so far with Mc SimGrid are bugs that we purposely added to a given code base. To that extend, our efforts span upon three axes. We first extended the programming model, to allow for the verification of more programs and hopefully of programs containing bugs. The most notable extensions are that the verified programs can now be written in Python, involve MPI_iprobe, or use condition variables. Mc SimGrid can now exhibit the critical transition when a failure is found. Before that transition, at least one exploration is correct, After it, all explorations are faulty. We added this feature to help us characterize whether the found issue is an application bug (and help us understand its root cause), or whether it is yet another bug in Mc SimGrid itself. Finally, we fixed dozens of bugs and vastly optimized the verification code to improve our chances to find a wild bug. Still, we did not find any such bug yet, so the chase continues.

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

7.1.3 EnOSlib

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 on bare-metal testbeds. 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, IoT-LAB and more) - Reusable software provisioning: In order to configure your nodes, EnOSlib exposes different APIs with different level of expressivity - Reusable services: Install common services such as Docker,

monitoring stacks, network emulation... - Reusable experiment facilities: Tasks help you to iterate faster on your experimentation workflow

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

Release Contributions: To reduce dependencies, the default pip package no longer includes Jupyter support.

Add support for Ansible 8, 9 and 10

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

Publications: hal-01664515, hal-01689726

Contact: Mathieu Simonin

Participants: Alexis Bitaillou, Volodia Parol-Guarino, Thomas Badts, Matthieu Rakotojaona Rainimangavelo, Mathieu Simonin, Baptiste Jonglez, Marie Delavergne

8 New results

8.1 Reliable decentralized computing infrastructures

8.1.1 Federated fog/cloud infrastructure design

Participants: Guillaume Pierre, Chih-Kai Huang.

The decentralization of computing infrastructures creates new opportunities for designing extremely large-scale infrastructures that are not under the control of a single operator (as all the current cloud providers propose), but rather organized as a federation of large numbers of smaller, local infrastructure operators [64]. Besides the obvious issue of establishing trust between independent operators (which may be addressed by technological or legal means), realizing this vision will require a major redesign of current infrastructure orchestrators. Instead of a monolithic platform composed of a small number of "control plane"' nodes which manage large numbers of "worker"' nodes, it becomes necessary to rethink our computing infrastructures as a horizontal collection of independent entities which may collaborate with each other while retaining local control over their resource management policies.

We proposed UnBound, a framework for meta-federations which addresses the difficult multi-tenancy issue in this context to isolate workloads belonging to different users and fog providers. Extensive evaluations with federations of up to 500 geo-distributed Kubernetes clusters demonstrate that UnBound maintains comparable application deployment times to the state of the art in a single member cluster scenario, avoids increasing cross-cluster network traffic, keeps resource consumption within acceptable limits, and exhibits stability and scalability, making it a suitable solution for large-scale fog computing deployments [13].

Monitoring large fog computing federations is also a challenge as standard solutions such as Prometheus Federations transmit detailed monitoring data about each cluster member at fixed intervals, which generates large network bandwidth usage. We proposed Acala, an aggregate monitoring framework for geo-distributed Kubernetes cluster federations which aims to provide the management cluster with aggregated information about the entire cluster instead of individual servers. Based on actual deployment under a controlled environment in the geo-distributed Grid'5000 testbed, our evaluations show that Acala reduces the cross-cluster network traffic by up to 97% and the scrape duration by up to 55% in the single member cluster experiment. Our solution also decreases cross-cluster network traffic by 95% and memory resource consumption by 83% in multiple member cluster scenarios [6].

These contributions and others from previous years are also reported in Chih-Kai Huang's PhD thesis [21].

8.1.2 Fog/cloud storage services

Participants: Shadi Ibrahim, Marc Tranzer, Mohammad Rizk.

Replication has been successfully deployed and practiced to ensure high data availability in large-scale distributed storage systems. However, with the relentless growth of generated and collected data, replication has become expensive not only in terms of storage cost, but also in terms of network cost, hardware cost and energy cost. While already popular for archived data in peer-to-peer (P2P) systems and cold data, in recent years, erasure coding (EC) has been progressively performed on the critical path of data access and has been has been increasingly used in practical storage systems such as Ceph and HDFS. The main reason for such an adoption is that EC can provide the same fault tolerance guarantee as replication, but at a much lower storage cost.

Storage devices account for nearly 11% of the total energy consumed by a data center, so using EC could result in significant energy savings. However, EC requires large data exchanges when reading data or recovering from failures, leading to higher energy consumption. Accordingly, and in an effort to understand the tradeoffs between data access performance, repair bandwidth and energy overhead, we conducted large-scale experiments on a CephFS cluster to evaluate the energy and performance of read, write and data repair under EC and replication [27]. A more detailed article on the same topic is in preparation.

We also started working on the design and implementation of a scalable and reliable P2P storage system using EC as a fault tolerance technique. We have We implemented and evaluated the performance of EC in IPFS – an open source and is widely used P2P storage system – and show that our implementation can provide a good performance compared to replication when reading and writing data to IPFS [26]. A more detailed article on the same topic is in preparation.

8.1.3 Reliable fog platforms in adverse natural environments

Participants: Guillaume Pierre, Ammar Kazem, Matthieu Nicolas.

Natural environment observatories allow a wide range of scientists such as biologists, botanists and hydrologists, to observe a zone of particular interest from the points of view of their different scientific disciplines. They follow a data-driven approach based on a variety of sensors and/or actuators deployed in the natural environment, coupled with different techniques to report the produced data to a public or private cloud for further analysis. The constraints which stem from the specificities of such observatories however deviate from traditional IoT use-cases deployed in urban environments. Observatories are often created in remote locations, where energy supply, cellular networking coverage and human maintenance are more challenging than usual. To address these challenges we are designing new fog computing platforms that can process observation data in-situ while respecting the specific constraints of these environments.

In this context we conducted a survey of 35 observatories in France and abroad to better understand their existing systems and practice as well as avenues for improvement using fog computing technologies [23]. A more detailed article on the same topic is in preparation.

We also started addressing the difficult question of designing a fog computing cluster from a hardware and software point of view that has sufficient processing capacity to handle a pre-defined set of data processing applications while reducing its energy consumption as much as possible. For this we exploit heterogeneous clusters of single-board computers (Raspberry Pi or equivalent) that can offer interesting tradeoffs between processing power and necessary energy.

Finally we started to redesign and improve the already-existing LivingFog platform [2]. The objective is to specialize it to become a viable platform to host a set of data processing applications in a natural environment observatory. The short-time goals are to help detect sensor failures in order to issue timely alarms. In a longe run we will extend the set of applications to include more specialized processing functions.

8.2 Reliable decentralized application runtimes

8.2.1 Reliable serverless runtimes

Participants: Nikos Parlavantzas, Mohamed Cherif Zouaoui Latreche.

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 at any location on the Edge-Cloud continuum, can be rapidly triggered, and consume resources only when needed. However, supporting fog applications places stringent requirements on FaaS runtimes. First, the runtimes should support latency objectives for functions, critical for latency-sensitive fog applications. Second, the runtimes should support reducing the energy consumption of the underlying infrastructure, particularly important for resource-constrained fog nodes. A major limitation of current FaaS runtimes is their lack of support for meeting latency and energy consumption requirements.

To address this limitation, we proposed an approach for FaaS scheduling, called FoRLess [14]. The approach uses Deep Reinforcement Learning, specifically Deep Q-Networks (DQN), to learn how to schedule functions in order to jointly optimize latency and energy consumption. We implemented the approach by extending Fission, a popular Kubernetes-based open-source FaaS runtime, and evaluated its effectiveness through experiments on the Grid'5000 testbed. The experimental results demonstrated that FoRLess significantly outperforms baseline scheduling methods in both latency and energy efficiency. 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.

8.2.2 Dynamic platform adaptation for Federated Learning

Participants: Cédric Tedeschi, Mathis Valli.

Federated Learning (FL) has emerged as a paradigm shift enabling heterogeneous clients and devices to collaborate on training a shared global model while preserving the privacy of their local data. However, a common yet impractical assumption in existing FL approaches is that the deployment environment is static, which is rarely true in heterogeneous and highly-volatile environments like the Edge-Cloud Continuum, where FL is typically executed. While most of the current FL approaches process data in an online fashion, and are therefore adaptive by nature, they only support adaptation at the ML/DL level (e.g., through continual learning to tackle data and concept drift), putting aside the effects of system variance. Moreover, the study and validation of FL approaches strongly rely on simulations, which, although informative, tends to overlook the real-world complexities and dynamics of actual deployments, in particular with respect to changing network conditions, varying client resources, and security threats.

In a paper published at the FlexScience workshop [17], we made a first step to address these challenges. We investigated the shortcomings of traditional, static FL models and identified areas of adaptation to tackle real-life deployment challenges. We devised a set of design principles for FL systems that can smartly adjust their strategies for aggregation, communication, privacy, and security in response to changing system conditions. To illustrate the benefits envisioned by these strategies, we presented the results of a set of initial experiments on a 25-node testbed. The experiments, which vary both the number of participating clients and the network conditions, show how existing FL systems are strongly affected by changes in their operational environment. Based on these insights, we proposed a set of take-aways for the FL community, towards further research into FL systems that are not only accurate and scalable but also able to dynamically adapt to the real-world deployment unpredictability.

8.2.3 Reliable data stream processing runtimes

Participants: Guillaume Pierre, Shadi Ibrahim, Khaled Arsalane, Govind Kovilkkatt Panickerveetil.

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 this context we started exploring the limitations of stream processing autoscaling systems when dealing with stateful data processing operators in geo-distributed environments [8]. We demonstrate that Flink's backpressure mechanism should not be used as the only trigger for rescaling operations in heterogeneous network conditions. Raw performance, as well as performance predictability, also degrade quickly in the presence of stateful data processing operators and/or high network latency between the processing nodes. This work is still ongoing and a second publication on this topic is in preparation.

Another issue related to data stream processing focuses on the message brokers such as Apache Kafka that are used to ingest incoming data and reliably distribute them to a data stream processing engine. It is very difficult to right-size such systems because their performance and resource efficiency varies a lot based on many hardware- and configuration-related aspects. Because of their stateful nature they are also difficult to dynamically rescale after their initial deployment, which leads many users to drastically over-provision them "just in case." We are currently working on a methodology to better assist Kafka users in exploring the space of workable configurations and in selecting the optimal one for a given expected workload. A publication on this topic is in preparation.

8.2.4 Decentralized resource provisioning for serverless applications

Participants: Nikos Parlavantzas, Volodia Parol–Guarino
.

Application resource provisioning in decentralised environments poses significant challenges. First, resources are typically owned and managed by multiple, independent providers, ranging from individuals and small-scale edge cloud operators to large, traditional cloud providers. All these providers should be incentivized to make their resources available to applications. Second, resources are highly heterogeneous and volatile, which, combined with the dynamism of application workloads, makes it difficult to meet application quality of service (QoS) requirements.

We aim to address these challenges in the context of FaaS applications. To this end, we have proposed a market-based approach, called GIRAFF, for placing FaaS applications in multi-provider fog infrastructures, along with an open-source implementation. In our approach, clients submit application placement requests associated with SLAs that specify guarantees over network latency and allocated resources. A marketplace then organizes an auction where fog nodes bid on the SLA to determine the nodes that will host the functions and the revenue for fog node owners.

In 2024, we extended GIRAFF to add support for function chains, proposed a fog-specific cost model, and performed reproducible experiments on the Grid'5000 testbed at large scale (up to 663 fog nodes on 60 servers) in order to compare GIRAFF with baseline placement methods. The evaluation demonstrates GIRAFF's effectiveness in optimizing FaaS application placement. An article about these results is under submission. This research is part of the thesis project of Volodia Parol-Guarino, started in October 2022.

8.2.5 Data processing at scale

Participants: Shadi Ibrahim.

In computing continuum workflow, the data is processed in stages, with the completion of the later stages being highly dependent on the completion of the previous stage. Within each stage, there

are tasks that run in parallel and perform complex computations. The completion time of a stage is highly dependent on the completion of the last task within that stage. In an attempt to optimize the data-intensive workflows in the computing continuum, this work investigates the factors that can affect performance within and between stages for future optimization strategies [24]. In addition, an HPC performance analysis methodology is adapted based on our findings to validate that our results can be applied across platforms. A more detailed article on the same topic is in preparation.

Deep Neural Network (DNN) inference has been widely adopted in today's intelligent applications, such as autonomous driving, virtual reality, image recognition. Usually, inference services are implemented as a workflow consisting of multiple DNN models. As the demand for parallel computing in DNN models is continuously growing, inference computations are gradually moving to GPUs. Serverless inference (i.e., deploying inference workflows on serverless computing) is gaining increasing popularity for proven high elasticity, cost efficiency (i.e., pay-as-you-go), and transparent deployment. In serverless inference, each DNN model is encapsulated as a "function", which can automatically scale up or down according to the change in number of requests. Existing serverless inference systems isolate functions in separate monolithic GPU runtimes (e.g., CUDA context), which is too heavy for short-lived and fine-grained functions, leading to a high startup latency, a large memory footprint, and expensive inter-function communication. Accordingly, we propose a new lightweight GPU sandbox for serverless inference workflow, called StreamBox [18]. StreamBox unleashes the potential of streams and efficiently realizes them for serverless inference by implementing fine-grain and auto-scaling memory management, allowing transparent and efficient intra-GPU communication across functions, and enabling PCIe bandwidth sharing among concurrent streams. Evaluations over real-world workloads show that StreamBox reduces the GPU memory footprint by up to 82% and improves throughput by 6.7X compared to state-of-the-art serverless inference systems.

8.2.6 Mitigation of tail processing latency

Participants: Shadi Ibrahim.

The large scale and heterogeneity of distributed processing infrastructures often result in partial failures and unavoidable performance variability. This, in turn, causes unexpected long tails in job execution and the appearance of stragglers (i.e., tasks that take significantly longer time to finish than the normal execution time). Big data analytic systems employ speculative execution to mitigate stragglers at large scale. Specifically, new copies of detected stragglers are launched on available machines. Straggler detection plays an important role in the effectiveness of speculative execution. The methods used to detect stragglers use the information extracted from the last received heartbeats which may be outdated when the detection is triggered. In this ongoing work [22], we investigate how heartbeat arrivals and the use of outdated information in straggler detection affect the accuracy of straggler detection. Our preliminary results show that considering heartbeat timestamps can improve the accuracy of straggler detection, especially when combined with rate based detection methods which typically mark as stragglers tasks whose estimated elapsed times are 20% longer than the mean or median.

8.3 Distributed infrastructure frugality

8.3.1 Energy accounting of digital services

Participants: Anne-Cécile Orgerie, Maxime Agusti, Vladimir Ostapenco.

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. Existing energy models are generally based on system usage data, which is incompatible with the general privacy policies of bare-metal server contracts, as proposed by OVHCloud. To deal with these problems, we proposed an original non-intrusive method for estimating the energy consumption of a server cooled by direct-chip liquid-cooling, based on the coolant temperature and the processor temperature obtained via IPMI [7]. Our approach, developed in collaboration with OVHCloud, is evaluated on an experiment carried out on 19 bare metal servers of a production infrastructure equipped with physical wattmeters.

Data centers are very energy-intensive facilities whose power provision is challenging and constrained by power bounds. In modern data centers, servers account for a significant portion of the total power consumption. In this context, the ability to limit the instant power consumption of an individual computing node is an important requirement. There are several energy and power capping techniques that can be used to limit compute node power consumption, such as Intel RAPL. Although it is nowadays mainly utilized for energy measurement, Intel RAPL (Running Average Power Limit) was originally designed for power limitation purposes. It is unclear exactly how Intel RAPL technology operates and what effects it has on application performance and power consumption. We conducted a thorough analysis of Intel RAPL technology as a power capping leverage on a variety of heterogeneous nodes for a selection of CPU and memory intensive workloads [15]. For this purpose, we first validated Intel RAPL power capping mechanism using a high-precision external power meter and investigated properties such as accuracy, power limit granularity, and settling time. Then, we attempted to determine which mechanisms are employed by RAPL to adjust power consumption.

8.3.2 Water consumption of Cloud data centers

Participants: Anne-Cécile Orgerie, Emmanuel Wedan Gnibga.

In an era of growing water demand and amplifying climatic disruptions, water is an increasingly limited and precious resource. As data centers continue to proliferate due in particular to growing artificial intelligence demand, datacenter reliance on water for cooling is drawing increasing criticism. In this work, we investigated the water and energy consumption, carbon emissions, and cost associated with the energy demand (including its cooling) from a datacenter, in four diverse locations (California, Texas, Germany, and France) [12]. We considered water use both in evaporative cooling and power generation, employing a hybrid datacenter cooling model that is operable with and without evaporative cooling. Subsequently, we examined solutions that expand the equipment's operational range and actively exploit the cooling design redundancy. Moreover, we investigated the benefits of making data centers dynamic, that is changing operating balance hourly.

Going further, we looked at the environmental impact of digital infrastructures through the particular prism of water-related issues in the data centres (DCs) that make up the clouds [4]. We attempted to draw up an overview of the use of water in DCs, whether this use is direct, for cooling and humidifying electronic equipment, or indirect, for electricity generation, production and the end-of-life of equipment. This overview is illustrated with figures, mainly from the cloud service providers themselves, to give some orders of magnitude on water consumption and trends. The aim is also to discuss these figures, the relevance of the usual indicators and the associated potential impact transfers, and to recall the major spatiotemporal issues and conflicts of use linked to water resources.

8.3.3 Steering efficiency for distributed infrastructures

Participants: Anne-Cécile Orgerie, Martin Quinson, Emmanuel Wedan Gnibga, Matthieu Silard.

Integrating larger shares of renewables in data centers' electrical mix is mandatory to reduce their carbon footprint. However, as they are intermittent and fluctuating, renewable energies alone cannot

provide a 24/7 supply and should be combined with a secondary source. Finding the optimal infrastructure configuration for both renewable production and financial costs remains difficult. In this work, we examined three scenarios with on-site renewable energy sources combined respectively with the electrical grid, batteries alone and batteries with hydrogen storage systems. The objectives were first, to size optimally the electric infrastructure using combinations of standard microgrids approaches, secondly to quantify the level of grid utilization when data centers consume/export electricity from/to the grid, to determine the level of effort required from the grid operator, and finally to analyze the cost of 100% autonomy provided by the battery-based configurations and to discuss their economic viability [5].

Self-consumption is developing worldwide to increase renewable electricity consumption and reduce electricity bills. It can be carried out individually or collectively (grouping several entities), but is generally restricted geographically. In datacenters, one may use load shifting to benefit from the self-consumption tariffs, at the cost of increasing energy consumption. An alternative would be to extend the self-consumption rules to wider perimeters. In this work, we proposed a comparative study on several aspects influencing the collective self-consumption (CSC) of an Edge infrastructure, including spatial load shifting, temporal load shifting and extending the current rules to encompass wider geographical boundaries [11].

Distributed infrastructures also include other systems than ICT. In particular, we are interested in smart grids that 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. Driven by governmental policies, particularly in Europe, Electric Vehicles (EVs) are gradually replacing the current fleet. Since they are only used during a small portion of the time for transportation, their batteries can be used for other purposes during the rest of the time. The power grid constantly requires to maintain the balance between electricity production and consumption. Controlling the charge of EVs can avoid the need to invest in grid reinforcement. This research work focuses on flattening the consumption in residential areas, also known as peak shaving, which reduces the need for reactive power plants on the grid. Existing solutions often rely on coordination among vehicles, meaning that they require a communication network between the various consumers in the area (e.g., vehicles, houses). However, such a dedicated communication network is not yet deployed in practice. It would be costly to deploy and operate, and could potentially constitute a privacy breach for consumers. In this work, we proposed a local smart charging solution only relying on information exchanges between each house's smartmeter and its EV's charging station [16]. We compared our local peak shaving to solutions without communication infrastructure, and to an optimal theoretical solution provided by a linear programming solver. Simulation campaigns show that our local peak shaving algorithm offers a solution close to the optimum, with advantages in terms of voltage stability, fairness among users and vehicle battery lifetime preservation.

8.3.4 Frugal mobile computing

Participants: Martin Quinson, Victorien Elvinger.

Mobile computing is very representative of the ever increasing computational offering. In the recent decades, the hardware capacity has increased exponentially in both computational power, memory and storage space, and display size. This increase is in pace with the software offering that went from scarce text messages to augmented reality and advanced features leveraging embedded neural networks. The only decreasing metric is maybe the battery which was reduced from one week or more to one day or less, despite the steady increase of the battery capacity.

In this context, the SmolPhone project is research action exploring other potential designs and their respective advantages for mobile computing. Its practical aspects consist in designing a sort of low-tech smartphone offering some services of a classical smartphone with a one-week battery lifetime. The goal is not to optimize a typical smartphone but rather to reconsider the design space of mobile computing and beyond. We envision the designed solution not as a marketable product but as an effective workbench

for future research in the domain of frugal computing.

The AEx SmolPhone started this year. On the hardware side, we evaluated the practical feasibility of the SmolPhone vision, which involves heterogeneous multiprocessors and low-power technologies for the screen and network. We identified several usage modes according to the provided features, and quantified the energy consumption of the processing units in each mode, as well as the time required to switch between modes [25]. A similar evaluation of the broadband network energy consumption is still to be done in this context.

On the software side, our work on the evaluation of alternate solutions to browse the web on a device too limited to display complex HTML5 pages is still ongoing. We are also currently exploring the design of an applicative framework enabling one to easily develop graphical applications targeting such hardware.

8.4 Evaluation methodologies and tools

8.4.1 Understanding software infrastructures

Participants: Léo Cosseron, Martin Quinson, Matthieu Simonin.

One of the most challenging infrastructure to study is the one deployed by malware operators. The difficulties classically posed by distributed infrastructures add to the fact that the infrastructure conceptors explicitly try to evade the analysis and/or attack the analysis system. As an answer, Virtual Machine Introspection (VMI) is used by sandbox-based analysis frameworks to observe malware samples while staying isolated and stealthy. Sandbox detection and evasion techniques based on hypervisor introspection are becoming less of an issue since running server and workstation environments on hypervisors is becoming standard. However, the fake network environment around a sandbox VM offers opportunities similar to hypervisor introspection for malware to evade. Malware can evaluate the discrepancy between observed performances and a real, presumed network environment of infected targets. VMI pauses also cause visible network performance glitches. To solve these issues, we proposed to extend virtual clock manipulation to synchronize hardware-accelerated virtual machines with a discrete-event network simulator [9].

8.4.2 Predictive models and simulation frameworks

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

The deployment of applications closer to end-users through fog computing has shown promise in improving network communication times and reducing contention. However, the use of fog applications such as microservices necessitates intricate network interactions among heterogeneous devices. Consequently, understanding the impact of different application and infrastructure parameters on performance becomes crucial. Current literature either offers end-to-end models that lack granularity and validation or fine-grained models that only consider a portion of the infrastructure. We first compared experimentally the accuracy of the existing simulation frameworks. Then, we used SimGrid to obtain comprehensive metrics regarding microservice applications operating in the fog, both for performance evaluation (e.g. application latency) and environmental evaluation (e.g. energy consumption and greenhouse gas emissions) [10].

8.4.3 Sustainable research ecosystem

Participants: Anne-Cécile Orgerie, Martin Quinson, Matthieu Simonin.

Scientific computing platforms offer computing and storage resources that are shared between multiple users. They generally do not have access to information that would enable them to assess the environmental footprint of their use of these resources. However, calculating their footprint could raise awareness of the environmental impacts generated by the use of these platforms and, ideally, lead to the organisation of a reduction in this footprint. The environmental footprint is often reduced to measuring its carbon footprint, or even the carbon emissions linked to its use alone. However, the reality of environmental impacts is much more complex, being multi-factorial and exacerbated by technical acceleration and intensification of use. As a complete analysis of environmental impacts is hampered by the lack of data that can be systematically obtained throughout the life cycle of digital devices, our study nevertheless focuses on a practical case: the calculation of carbon costs in a computing centre dedicated to research (the Rennes site of Grid'5000) [19]. We first explored the various methods of calculating the environmental footprint that can be applied to the management of a computing infrastructure, including embodied energy; and secondly we presented and discussed a calculation method that can be used to attribute a carbon cost to the use of a computing core for one hour.

9 Bilateral contracts and grants with industry

9.1 Bilateral contracts with industry

Défi Inria OVHCloud

Participants: Anne-Cécile Orgerie, Shadi Ibrahim, Govind Kovilkkatt Panickerveetil, Guillaume Pierre, Maxime Agusti, Vladimir Ostapenco, Marc Tranzer.

Title: FrugalCloud: Eco-conception de bout en bout d'un cloud pour en réduire les impacts environnementaux

Partner Institution(s):

- Inria, France
- OVHCloud, France

Date/Duration: 2021-2025

Additionnal info/keywords: 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, Magellan, Spirals, Stack. Members of the Magellan team contribute to four sub-challenges including (1) Software ecodesign 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

Participants: Shadi Ibrahim, Mohammad Rizk, Quentin Acher.

Title: Alvearium: Large Scale Secure and Reliable Peer-to-Peer Cloud Storage.

Partner Institution(s):

- Inria, France

- Hive, France

Date/Duration: 2023-2026

Additionnal info/keywords: 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, Magellan, WIDE, COATI. Members of the Myriads team will contribute to two axes. Specifically, the Magellan team coordinates the axis on reliable and cost-efficient data placement and repair in P2P storage over immutable data; and contributes to the axis on the management of mutable data over P2P storage.

10 Partnerships and cooperations

10.1 International initiatives

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

Hermes@Scale

Participants: Shadi Ibrahim.

Title: Hermes@Scale: Scalable and Energy Efficient Data Management for Scientific Workloads in Computing Continuum.

Partner Institutions:

- Inria, France
- Lawrence Berkeley National Lab (LBNL), USA

Date/Duration: 2024-2026

Additionnal info/keywords: Hermes@Scale is a follow-up of the Hermes associate team (2019-2022) which was mainly focused on accelerating the performance of multi-site scientific applications while considering coordinated data and metadata management between a few homogeneous supercomputing facilities. The new Hermes@Scale project is still interested in introducing advanced data management techniques to optimize the performance of scientific workloads, but considers new problems related to the heterogeneity and the scale (beyond supercomputing facilities) of the emerging Computing Continuum platforms and ever growing complexity and variety of compute and data intensive workloads. As a result, in the Hermes@Scale project, we plan to investigate new interference-aware scheduling strategies that consider the whole I/O path of various applications when sharing different platforms from the Edge to the Clouds and HPC systems; explore solutions to enable easy-to program/tune, and elastic compute- and data-intensive workloads in distributed Computing Continuum; leverage erasure codes and data compression to reduce the storage and network footprint of data storage and transfer; and perform seamless data movement across storage tiers and platforms. Finally, we plan to optimize the energy consumption of scientific workloads by considering the energy efficiency of data management solutions, thus reducing the environmental overhead and high carbon emissions.

10.2 International research visitors

10.2.1 Visits of international scientists

Other international visits to the team

Frédéric Sutter :

Status Senior Research Scientist.

Institution of origin: Oak Ridge National Laboratory.

Country: USA

Dates: December 6-10th 2024.

Context of the visit: Discuss the future orientations of the SimGrid framework.

Mobility program/type of mobility: Research stay

10.3 European initiatives

10.3.1 Digital Europe

ACHIEVE

Participants: Guillaume Pierre, Cédric Tedeschi.

Title: ACHIEVE: Advanced Cloud and High-performance computing Education for a Valiant Europe

Partner Institution(s):

- EIT Digital
- EIT Digital Spain
- Universita Degli Studi di Trento
- Middle East Technical University
- Université de Rennes
- Kungliga Tekniska Hoegskolan
- Universitate Babes Bolyai
- Aalto University
- Politecnico di Milano
- University of Novi Sad
- Evolutionary Archetypes Consulting
- Techvalley Management
- Infineon
- Odtu Teknokent
- Scientific and Technological Research center council of Turkiye
- Research Institute of Sweden
- Amazon Web Services

Date/Duration: 2024-2028

Additionnal info/keywords: ACHIEVE aims to implement a double-degree master's program focusing on Cloud and Networking Infrastructure with specializations that align with the strategic development of advanced, green, and efficient HPC systems. Additionally, it includes a minor in Innovation and Entrepreneurship. This program will be collaboratively designed and delivered by eight higher education institutions from six different countries, in partnership with major industry associations and companies specializing in cloud computing and HPC, along with an innovative SME for educational program delivery, a non-profit association, and EIT Digital, a leader in advanced digital skills education in Europe.

10.4 National initiatives

AEx Smolphone

Participants: Martin Quinson, Victorien Elvinger.

Title: SmolPhone: un smartphone conscient des limites énergétiques.

Partner Institution(s):

- Inria Magellan team
- Inria Taran team

Date/Duration: 2024-2026

Additionnal info/keywords: The SmolPhone project is research action exploring potential frugal and low-tech solutions in the domain of mobile computing. Its practical aspects consist in designing a sort of low-tech smartphone offering some services of a classical smartphone with a one-week battery lifetime. The goal of the AEx project is to unlock several technical difficulties by conceiving a software and hardware development workbench enabling future research on the topic of frugal computing and low-tech IT services.

ADT SmartObs

Participants: Guillaume Pierre, Matthieu Nicolas.

Title: Une plateforme originale pour le monitoring environemental

Partner Institution(s):

- Inria Magellan team
- Inria Taran team

Date/Duration: 2024-2025

Additionnal info/keywords: The goal of the SmartObs project is to help develop the SmartSense and LivingFog platforms dedicated to natural environment monitoring, and to enable interdisciplinary collaborations with partners from the environmental sciences domain. The objective is to generate, analyse and transmit smart sensor data (SmartSense) and further process these data in situ (LivingFog) before sending them to a remote cloud.

ANR Dark-Era

Participants: Martin Quinson, .

Title: Dark-Era: Dataflow Algorithm aRchitecture co-design of SKA pipeline for Exascale RadioAstronomy

Partner Institution(s):

- CentraleSupélec, Laboratoire des Signaux et Systèmes (LSN)

- INSA Rennes, Institut d'Electronique et de Télécommunication de Rennes (IETR)
- Observatoire Paris, Galaxies, Etoiles, Physique, Instrumentation
- Observatoire de la Côte d'Azur Nice, Laboratoire J-L. Lagrange
- ENS Rennes, Institut de Recherche en Informatique et Systèmes Aléatoires (IRISA)

Date/Duration: 2021-2025

Additionnal info/keywords: 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.

ANR FACTO

Participants: Anne-Cécile Orgerie, François Lemercier.

Title: FACTO: A Multi-Purpose Wi-Fi Network for a Low-Consumption Smart Home

Partner Institution(s):

- CNRS, IRISA lab
- University of Lyon 1, LIP lab
- Orange
- Fondation Blaise Pascal

Date/Duration: 2021-2024

Additionnal info/keywords: 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

Participants: Anne-Cécile Orgerie, Matthieu Silard.

Title: EDEN4SG: Efficient and Dynamic ENergy Management for Large-Scale Smart Grids

Partner Institution(s):

- CNRS, IRISA lab
- CNRS, SATIE lab
- University of Toulouse, IRIT lab
- EDF
- SRD Energies
- Orange

Date/Duration: 2023-2027

Additionnal info/keywords: 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.

IPCEI-CIS E2CC (BPI)

Participants: Anne-Cécile Orgerie.

Title: Eco Edge to Cloud Continuum

Partner Institution(s):

- CNRS, IRISA lab
- Inria, LIG lab
- University of Toulouse, IRIT lab
- CNRS, LAAS lab
- ATOS (Bull SAS)
- Armadillo
- CGI France
- NBS
- Ningaloo
- Provenrun
- Ryax

Date/Duration: 2024-2028

Additionnal info/keywords: The aim of the IPCEI-CIS E2CC project is to provide a standardized end-to-end platform that can be connected natively to any cloud provider, offering the right level of interoperability and service portability and creating a continuum from Edge to Cloud, including cybersecurity, decarbonisation, orchestration and platform functions. This platform will provide hardware and software solutions, exhibiting vertical services (computer vision, decarbonisation services, etc.), cybersecurity services, AI/MLOPs, and Baremetal & Edge infrastructures.

CARECloud (PEPR Cloud)

Participants: Anne-Cécile Orgerie, Matteo Chancerel, Pablo Leboulanger, Thomas Stavis.

Title: CARECLOUD: Understanding, improving, reducing the environmental impacts of Cloud Computing.

Partner Institution(s):

- CNRS, IRISA lab
- Inria, CRIStAL lab
- CNRS, I3S lab
- University of Toulouse, IRIT lab
- Inria, LIP lab
- IMT, LS2N lab
- IMT, SAMOVAR lab

Date/Duration: 2023-2030

Additionnal info/keywords: 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.

CNRS DISTANT

Participants: Guillaume Pierre, Ammar Kazem, Matthieu Nicolas.

Title: Automated **D**ata qual**I**ty as**S**urance for a cri**T**ical zone observ**A**tory i**N** **T**he Himalayas – Kaligandaki River Nepal

Partner Institution(s):

- Université de Rennes – Géosciences Rennes, France
- Université de Rennes – IRISA, France
- IRD, France
- German Research Centre for Geosciences GFZ, Germany

- Tribhuvan University, Nepal
- Smartphones4Water Citizen Science NGO, Nepal

Date/Duration: 2024-2025

Additionnal info/keywords: Mountain landscapes respond disproportionately quickly to changes in external forcing by tectonics and climate. Erosion and associated extreme events are common and present a continuous threat to local populations. At the same time, mountains act as resources and supply 1/3 of the world population with fresh water. Yet mountains are notoriously under studied and are widely underrepresented in global monitoring networks. One particular problem of monitoring such landscapes is the data quality, especially in remote areas where manned support is difficult to maintain. The dedicated monitoring area is the Kaligandaki Catchment in Nepal, spanning an ideal transect across the Himalayan Mountain Range from the edge of the Tibetan Plateau, through the high mountains, to the low-lying Gangetic foreland. We will implement a LivingFog platform computing and data logging system that can detect bogus data and issue alerts to the operators. This system is designed to serve as a near real-time data processing system for early warning of catastrophic events in the future.

NF-JEN (PEPR 5G and future networks)

Participants: Anne-Cécile Orgerie.

Title: NF-JEN: Just Enough Networks

Partner Institution(s):

- CNRS, IRISA lab
- CEA, LETI
- INSA de Lyon, CITI lab
- Inria, CRIStAL lab
- IMT, IEMN lab
- CNRS, IETR lab
- INP Bordeaux, IMS lab
- IMT, Lab-STICC lab
- ESIEE - Université Gustave Eiffel, LIGM lab
- Inria, LIP lab
- IMT, LTCI lab
- CNRS, XLIM lab

Date/Duration: 2023-2027

Additionnal info/keywords: 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.

STEEL (PEPR Cloud)

Participants: Cédric Tedeschi, Mathis Valli.

Title: STEEL: Secure and Efficient Data Storage and Processing on Cloud-based Infrastructures

Partner Institution(s):

- DRIM (CNRS, INSA Lyon)
- ERODS (Université Grenoble Alpes)
- Kerdata (Inria, INSA Rennes)
- IN2P3 (CNRS)
- MAGELLAN (Université de Rennes, Inria)
- PROGRESS (Université de Bordeaux, INP)
- STACK (IMT Atlantique, Inria, LS2N)
- TeraLab (IMT)

Date/Duration: 7 years

Additionnal info/keywords: The strong development of cloud computing and its massive adoption for the storage of unprecedented volumes of data in a growing number of domains has brought to light major technological challenges. In this project we will address several of these challenges, organized in three research directions:

- Exploitation of emerging technologies for efficient storage on cloud infrastructures. We will address this challenge through NVRAM-based distributed performance storage solutions, as close as possible to data production and consumption locations (disaggregation principle) and develop strategies to optimize the trade-off between data consistency and access performance.
- Efficient storage and processing of data on hybrid, heterogeneous infrastructures within the digital edge-cloud-supercomputer continuum. The execution of hybrid workflows combining simulations, analysis of sensor data flows and machine learning requires storage resources ranging from the edge to cloud infrastructures, and even to supercomputers, which poses challenges for unified data storage and processing.
- Confidential storage, in connection with the need to store and analyze large volumes of data of strategic interest or of a personal nature.

Taranis (PEPR Cloud)

Participants: Shadi Ibrahim, Nikos Parlavantzas.

Title: Taranis: Model, Deploy, Orchestrate, and Optimize Cloud Applications and Infrastructure

Partner Institution(s):

- IMT, France
- Université Grenoble Alpes, France
- Université de Rennes, France

- Inria, France
- CNRS, France
- CEA, France
- ENS Lyon, France
- Université Claude Bernard Lyon 1, France
- Université de Lille, France
- INSA Rennes, France

Date/Duration: 7 years

Additionnal info/keywords: 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 Magellan team 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.

10.5 Regional initiatives

ARED TDFDE

Participants: Khaled Arsalane, Guillaume Pierre.

Title: Traitement de données en flux distribué et à état

Date/Duration: 2022-2025

Additionnal info/keywords: 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

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

Title: Time-Accurate Network Simulation Interconnecting VMs with Hardware Virtualization Towards Stealth Analysis

Date/Duration: 2022-2025

Additionnal info/keywords: 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

11.1.1 Scientific events: organisation

General chair, scientific chair

- Anne-Cécile Orgerie was co-organizer of the eco-ICT research school on eco-responsible ICT and digital sufficiency in networks and distributed systems.
- Shadi Ibrahim was the General Chair of the 36th International Conference on Scientific and Statistical Database Management (SSDBM) which took place in Rennes on July 10-12th 2024.
- Shadi Ibrahim was the General Chair of the Workshop on Challenges and Opportunities of Efficient and Performant Storage Systems (CHEOPS), held in conjunction with EuroSys 2024.
- Shadi Ibrahim was the Tutorials Chair for ISC High Performance 2024, which took place in Hamburg, Germany.
- Shadi Ibrahim was a member of the steering committee for the 8th edition of the Workshop Performance and Scalability of Storage Systems(Per3S), which took place in Paris on May 28 2024.
- Shadi Ibrahim is a member of the steering committee of the IEEE International Conference on Cluster Computing.
- Shadi Ibrahim was a member of the steering committee for ISC High Performance 2024.
- Shadi Ibrahim is a member of the steering committee of the International Parallel Data Systems Workshop (PDSW).
- Shadi Ibrahim is a member of the steering committee of the workshop on Challenges and Opportunities of Efficient and Performant Storage Systems (CHEOPS).

Member of the organizing committees

- Anne-Cécile Orgerie and Matthieu Simonin were members of the local arrangements team of ACM REP 2024 conference.
- François Lemercier was co-organizer of the 26èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications.

11.1.2 Scientific events: selection

Chair of conference program committees

- Anne-Cécile Orgerie was chair of the track on System Software and Cloud Computing at the International Conference for High Performance Computing, Networking, Storage, and Analysis (SuperComputing 2024).

Member of the conference program committees

- Guillaume Pierre was a program committee member of the IEEE ICFEC 2024 and DAIS 2024 conferences.
- Anne-Cécile Orgerie was a program committee member of the IEEE IPDPS 2024 and ICT4S 2024 conferences.
- Nikos Parlavantzas was a program committee member of IEEE CloudCom 2024, IEEE ISPDC 2024, JSSPP 2024, and VHPC 2024.
- François Lemercier was a program committee member of the Cores 2024 Conference.
- Shadi Ibrahim was a program committee member of the IEEE DAC 2024, SC 2024 (Research papers and Birds of a Feather), IEEE IPDPS 2024, IEEE Cluster, CCGrid 2024 (scale challenge), and Euro-Par 2024 conferences.

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.
- Anne-Cécile Orgerie is a member of the editorial board of the Special Issue on Sustainability and Computing of the Communications of the ACM in 2024.
- Shadi Ibrahim is an associate editor of IEEE Internet Computing Magazine.
- Shadi Ibrahim is a guest editor of IEEE Internet Computing Magazine: Special Issue on Serverless Computing in 2024.
- Shadi Ibrahim an associate editor for High Performance Big Data Systems of Frontiers in High Performance Computing Journal.

Reviewer - reviewing activities

- Nikos Parlavantzas was a reviewer for IEEE Internet Computing and ACM Transactions on Autonomous and Adaptive Systems journals.
- Shadi Ibrahim was a reviewer for IEEE Internet Computing Magazine, IEEE Transactions on Parallel and Distributed Systems journal, and Springer Future Generation Computer Systems journal.

11.1.4 Invited talks

- Anne-Cécile Orgerie gave a keynote at the EGC (Extraction et la Gestion des Connaissances) 2024 conference in Dijon, France.
- Guillaume Pierre gave an invited talk “Fog Computing Challenges in Natural Environment Observatories” at the Cloud Control Workshop on June 20th 2024 in Skåvsjöholm, Sweden.
- Martin Quinson gave an invited talk “SmolPhone : a smartphone with energy limits.” at the Eco-ICT winter school on October 11th 2024 in Beg Meil, France.

11.1.5 Leadership within the scientific community

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

11.1.6 Scientific expertise

- Guillaume Pierre was a reviewer for a research proposal for the Austrian Science Fund (FWF).
- Anne-Cécile Orgerie was a reviewer for a research proposal for the Université Catholique de Louvain in Belgium.
- Anne-Cécile Orgerie was a reviewer for a research proposal for the Mitacs program in Canada.
- Anne-Cécile Orgerie was a member of the selection committees for an assistant professor position at University of Rennes in 2024.
- Cédric Tedeschi was a reviewer for two research proposals for the Natural Sciences and Engineering Research Council of Canada.
- Shadi Ibrahim was a member of the ACM Heidelberg Laureate Forum (HLF) 2024 Young Researcher Selection Committee.
- Shadi Ibrahim was a reviewer for a reserach proposal for the CEFIPRA Programme 2024: Indo-French Centre for the Promotion of Advanced Research.
- Shadi Ibrahim was a member of the jury (selection committee) for an assistant professor position at ENSAI Bruz in 2024.

11.1.7 Research administration

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

11.2 Teaching - Supervision - Juries

11.2.1 Teaching

- Bachelor: Marin Bertier, Programmation en Java, L1, INSA Rennes.
- Bachelor: Marin Bertier, Langage C, L3, INSA Rennes.
- Bachelor: Marin Bertier, Réseaux, L3, INSA Rennes.
- Bachelor: Shadi Ibrahim, Introduction to Big Data, ENSAI, Bruz.
- Bachelor: Francois Lemercier, Système Cloud et Réseaux, L3 Informatique, Univ. Rennes.
- Bachelor: Nikos Parlavantzas, Databases, L2, INSA Rennes.
- Bachelor: Guillaume Pierre, Systèmes Informatiques, L3 MIAGE, Univ. Rennes.
- Bachelor: Guillaume Pierre, Systèmes d'exploitation, L3 Informatique, Univ. Rennes.
- Bachelor: Martin Quinson, Architecture et Systèmes, L3 Informatique, ENS Rennes.
- Bachelor: Martin Quinson, Pédagogie, L3 Informatique, ENS Rennes.
- Master: Guillaume Pierre, Distributed Systems, M1 CR&CNI, Univ. Rennes.

- Master: Guillaume Pierre, Services for Cloud technology, M1 CR&CNI, Univ. Rennes.
- Master: Guillaume Pierre, Advanced Cloud Infrastructures, M2 CR&CNI&SIF, Univ. Rennes.
- Bachelor: Cédric Tedeschi, Systèmes, Cloud et Réseaux, L3 Informatique, Univ. Rennes.
- Master: Francois Lemercier, Network Administration and new technologies, M1 CNI, Univ. Rennes.
- Master: Francois Lemercier, Cloud, M1 RSSI, Univ. Rennes.
- Master: Nikos Parlavantzas, 4th-year Project, M1, INSA Rennes.
- Master: Nikos Parlavantzas, Clouds, M1, INSA Rennes.
- Master: Nikos Parlavantzas, IoT, M2, INSA Rennes.
- Master: Nikos Parlavantzas, Smart city services, M2 CNI, Univ. Rennes.
- Master: Nikos Parlavantzas, NoSQL, M2, Master for Smart Data Science, ENSAI.
- Master: Nikos Parlavantzas, SQL, M2, Master for Smart Data Science, ENSAI.
- Master: Martin Quinson, C++ system programming, ENS Rennes.
- Master: Martin Quinson, HPC programming, ENS Rennes.
- Master: Martin Quinson, Préparation à l'Agrégation d'Informatique (Networking, 20h ETD, Operating systems, 20h ETD), ENS Rennes.
- Master: Martin Quinson, Scientific Outreach, M2, 30 hEDT, ENS Rennes.
- Master: Cédric Tedeschi, Concurrences et Système d'Exploitation, M1 Informatique, Univ. Rennes.
- Master: Cédric Tedeschi, Concurrence et coopération dans les systèmes et les réseaux, M1 MIAGE. Univ. Rennes.
- Master: Cédric Tedeschi, Advanced Cloud Infrastructures, M2 Informatiquie, Univ. Rennes.
- Master: Cédric Tedeschi, Parallel Programming with Python, ENSAI
- Master: Marin Bertier, Systèmes d'exploitation, M1, INSA Rennes.
- Master: Shadi Ibrahim, Hadoop and Cloud Computing, 36hETD, M2 : Statistics for Smart Data, ENSAI, Bruz.
- Master: Shadi Ibrahim, NoSQL and Cloud technologies, 45hETD, M2, ENSAI, Bruz.
- Master: Anne-Cécile Orgerie, Energy efficiency in distributed systems, TelecomSudParis, Evry.

11.2.2 Supervision

- PhD defended: Chih-Kai Huang, "Scalability of Public Geo-Distributed Fog Computing Federations", defended on December 9th 2024, supervised by Guillaume Pierre.
- PhD defended: Wedan Emmanuel Gnibga, "Modeling and optimizing edge computing infrastructures and their electrical system", defended on November 19th 2024, supervised by Anne-Cécile Orgerie and Anne Blavette.
- PhD defended: Vladimir Ostapenco, "Modeling, evaluating and orchestrating heterogeneous leverages for large-scale cloud data centers management", defended on December 18th 2024, supervised by Anne-Cécile Orgerie and Laurent Lefèvre.
- PhD in progress: Ammar Kazem, "Systèmes de traitement de données in-natura pour l'observation environnementale sous contrainte énergétique", started in 2023, supervised by Guillaume Pierre and Laurent Longuevergne.

- PhD in progress: Khaled Arsalane, “Traitement de données en flux distribué et à état”, started in 2022, supervised by Guillaume Pierre.
- PhD in progress: Govind Kovilkkatt Panickerveetil, “Traitement de flux de données économe en énergie”, started in 2023, supervised by Guillaume Pierre and Romain Rouvoy.
- 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: Matteo Chancerel, “Optimization of a fully distributed Fog powered by renewable energy sources”, started in September 2024, supervised by Anne-Cécile Orgerie.
- PhD in progress: Pablo Leboulanger, “Minimalist cloud, low on energy, hardware and software resources”, started in September 2024, supervised by Anne-Cécile Orgerie.
- PhD in progress: Thomas Stavis, “Replay of environmental leverages in Cloud infrastructures”, started in September 2024, supervised by Laurent Lefèvre and Anne-Cécile Orgerie.
- 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: Mathis Valli, “Dynamic Adaptation of Machine Learning and Deep Learning Workflows across the Cloud-Fog-Edge Continuum“, started in April 2023, supvervised by Alexandru Costan, Loïc Cudennec and Cédric Tedeschi.
- 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: Marc Tranzer, “Energy efficient data management: Data reduction and protection meet performance and energy “, started in January 2024, supervised by Shadi Ibrahim and Guillaume Pierre.
- 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 Analysi”, started in October 2022, supervised by Martin Quinson and Louis Rilling.

11.2.3 Juries

- Guillaume Pierre was a reviewer of the PhD defense of Pierre-Antoine Rault (Université de Lorraine), December 12th 2024.
- Guillaume Pierre was chairperson at the PhD defense of Arthur Rauch (Université de Rennes), December 18th 2024.

- Anne-Cécile Orgerie was chairperson at the PhD defense of Vincent Rébiscoul (University of Rennes), May 14th 2024.
- Anne-Cécile Orgerie was chairperson at the PhD defense of Pierre Jacquet (University of Lille), July 19th 2024.
- Anne-Cécile Orgerie was chairperson at the PhD defense of Vincent Lannurien (ENSTA Bretagne), November 20th 2024.
- Anne-Cécile Orgerie was chairperson at the PhD defense of Chih-Kai Huang (University of Rennes), December 9th 2024.
- Nikos Parlavantzas was a reviewer for the Ph.D. defense of Armel Jeatsa Toulepi (Institut National Polytechnique de Toulouse), 11th September 2024.
- Cédric Tedeschi was the chairperson of the PhD defense of Thomas Bouvier (INSA de Rennes), November 4th, 2024.
- Cédric Tedeschi was the chairperson of the PhD defense of Timothé Albouy (Univ. Rennes), December 16th, 2024.
- Cédric Tedeschi was an examinor of the PhD defense of Manèle Ait Habouche (Université de Bretagne Occidentale), December 19th, 2024.
- Shadi Ibrahim was a thesis reader (member of the PhD Committee) of the PhD defense of Twinkle Jain (Northeastern University), USA, April 3rd 2024.
- Shadi Ibrahim co-organizies the SCI-Rennes seminar: a monthly series of scientific seminars for all staff at Inria research centre at Rennes University.

11.2.4 Productions (articles, videos, podcasts, serious games, ...)

- Anne-Cécile Orgerie was one of the authors of two publications in the popularization journal *Bulletin 1024* of the Société Informatique de France in April 2024 [28] and in November 2024 [29].

11.2.5 Participation in Live events

- “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) and Anne-Cécile Orgerie (Magellan) are coordinating the local version of this program, initiated in Lille. The first session in Rennes occured in April 2019, and a new session (the 6th) occured in 2024. The program is currently supported by: Fondation Blaise Pascal, ED MathSTIC, ENS de Rennes, Université de Rennes and Fondation Rennes 1.

11.2.6 Others science outreach relevant activities

- Guillaume Pierre gave an invited talk “Le cloud c’est bien, le fog c’est mieux!” to bachelor and master students of the University of Rennes on December 17th 2024.

12 Scientific production

12.1 Major publications

[1] H. Arkian, G. Pierre, J. Tordsson and E. Elmroth. ‘Model-based Stream Processing Auto-scaling in Geo-Distributed Environments’. In: ICCCN 2021 - 30th International Conference on Computer Communications and Networks. Athens, Greece, 19th July 2021, pp. 1–11. URL: https://inria.hal.science/hal-03206689.

[2] D. Battulga, M. Farhadi, M. A. Tamiru, L. Wu and G. Pierre. 'LivingFog: Leveraging fog computing and LoRaWAN technologies for smart marina management (experience paper)'. In: ICIN 2022 - 25th Conference on Innovation in Clouds, Internet and Networks. Paris, France, 7th Mar. 2022, pp. 1–8. URL: https://inria.hal.science/hal-03531372 (cit. on p. 11).

[3] M. A. Tamiru, G. Pierre, J. Tordsson and E. Elmroth. 'mck8s: An orchestration platform for geo-distributed multi-cluster environments'. In: ICCCN 2021 - 30th International Conference on Computer Communications and Networks. Athens, Greece, 19th July 2021, pp. 1–12. URL: https://inria.hal.science/hal-03205743.

12.2 Publications of the year

International journals

[4] S. Bouveret, A. Bugeau, A.-C. Orgerie and S. Quinton. 'De l'eau dans les nuages'. In: *Annales des Mines - Enjeux Numériques* 27 (Sept. 2024), pp. 40–47. URL: https://hal.science/hal-04698568 (cit. on p. 15).

[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* 9.3 (2024), pp. 315–328. DOI: 10.1109/TSUSC.2023.3307790. URL: https://hal.science/hal-04189173 (cit. on p. 16).

[6] C.-K. Huang and G. Pierre. 'Aggregate Monitoring for Geo-Distributed Kubernetes Cluster Federations'. In: *IEEE Transactions on Cloud Computing* (2024), pp. 1–15. DOI: 10.1109/TCC.2024.3482574. URL: https://inria.hal.science/hal-04736577. In press (cit. on p. 10).

International peer-reviewed conferences

[7] M. Agusti, E. Caron, B. Fichel, L. Lefèvre, O. Nicol and A.-C. Orgerie. 'PowerHeat: A non-intrusive approach for estimating the power consumption of bare metal water-cooled servers'. In: 2024 IEEE International Conferences on Internet of Things (iThings) and IEEE Green Computing & Communications (GreenCom) and IEEE Cyber, Physical & Social Computing (CPSCom) and IEEE Smart Data (SmartData) and IEEE Congress on Cybermatics. Copenhagen, Denmark, 2024, pp. 1–7. URL: https://hal.science/hal-04662683 (cit. on p. 15).

[8] K. Arsalane, G. Pierre and S. Ibrahim. 'Toward Stream Processing Elasticity in Realistic Geo-Distributed Environments'. In: IC2E 2024 - 12th IEEE International Conference on Cloud Engineering. Paphos, Cyprus, 2024, pp. 1–9. URL: https://inria.hal.science/hal-04655408 (cit. on p. 13).

[9] L. Cosseron, L. Rilling, M. Simonin and M. Quinson. 'Simulating the Network Environment of Sandboxes to Hide Virtual Machine Introspection Pauses'. In: EuroSec 2024 - 17th European Workshop on Systems Security. Athène, Greece, 22nd Apr. 2024, pp. 1–7. DOI: 10.1145/3642974.3652280. URL: https://inria.hal.science/hal-04537165 (cit. on p. 17).

[10] C. Courageux-Sudan, A.-C. Orgerie and M. Quinson. 'Studying the end-to-end performance, energy consumption and carbon footprint of fog applications'. In: ISCC 2024 - 29th IEEE Symposium on Computers and Communications. Paris, France: IEEE, June 2024, pp. 1–7. URL: https://hal.science/hal-04581677 (cit. on p. 17).

[11] W. E. Gnibga, A. Blavette and A.-C. Orgerie. 'Energy-related Impact of Redefining Self-consumption for Distributed Edge Datacenters'. In: IGSC 2024 - 15th International Green and Sustainable Computing Conference. Austin, United States, 2024, pp. 1–7. URL: https://hal.science/hal-04770489 (cit. on p. 16).

[12] W. E. Gnibga, A. A. Chien, A. Blavette and A.-C. Orgerie. 'FlexCoolDC: Datacenter Cooling Flexibility for Harmonizing Water, Energy, Carbon, and Cost Trade-offs'. In: e-Energy 2024 - 15th ACM International Conference on Future and Sustainable Energy Systems. Singapore, Singapore: ACM, June 2024, pp. 108–122. DOI: 10.1145/3632775.3661936. URL: https://hal.science/hal-04581701 (cit. on p. 15).

[13] C.-K. Huang and G. Pierre. 'UnBound: Multi-Tenancy Management in Scalable Fog Meta-Federations'. In: UCC 2024 - 17th IEEE/ACM International Conference on Utility and Cloud Computing. Sharjah, United Arab Emirates, 2024, pp. 1–11. URL: https://inria.hal.science/hal-04760398 (cit. on p. 10).

[14] C. Latreche, N. Parlavantzas and H. A. Duran-Limon. 'FoRLess: A Deep Reinforcement Learning-based approach for FaaS Placement in Fog'. In: UCC 2024 - 17th IEEE/ACM International Conference on Utility and Cloud Computing. Sharjah, United Arab Emirates, 2024, pp. 1–9. URL: https://hal.science/hal-04791252 (cit. on p. 12).

[15] V. Ostapenco, L. Lefèvre, A.-C. Orgerie and B. Fichel. 'Exploring RAPL as a Power Capping Leverage for Power-Constrained Infrastructures'. In: ICA3PP 2024 - 24th International Conference on Algorithms and Architectures for Parallel Processing. Macau SAR, China, 2024, pp. 1–10. URL: https://hal.science/hal-04742418 (cit. on p. 15).

[16] M. Silard, A.-C. Orgerie, N. Montavont and G. Papadopoulos. 'Local Peak Shaving for Electric Vehicles: A ready-to-deploy Smart Charging Solution'. In: IGSC 2024 - 15th International Green and Sustainable Computing Conference. Austin, United States, 2024, pp. 1–7. URL: https://hal.science/hal-04770492 (cit. on p. 16).

[17] M. Valli, A. Costan, C. Tedeschi and L. Cudennec. 'Towards Efficient Learning on the Computing Continuum: Advancing Dynamic Adaptation of Federated Learning'. In: *FlexScience'24: Proceedings of the 14th Workshop on AI and Scientific Computing at Scale using Flexible Computing Infrastructures*. FlexScience 2024 - 14th Workshop on AI and Scientific Computing at Scale using Flexible Computing Infrastructures. Pisa, Italy: ACM, 2024, pp. 42–49. DOI: 10.1145/3659995.3660042. URL: https://hal.science/hal-04698619 (cit. on p. 12).

[18] H. Wu, Y. Yu, J. Deng, S. Ibrahim, S. Wu, H. Fan, Z. Cheng and H. Jin. 'StreamBox: A Lightweight GPU SandBox for Serverless Inference Workflow'. In: ATC 2024 - USENIX Annual Technical Conference. Santa Clara, United States, 2024. URL: https://inria.hal.science/hal-04857475 (cit. on p. 14).

National peer-reviewed Conferences

[19] M. Simonin and A.-C. Orgerie. 'Méthodologies de calculs de l'empreinte carbone sur une plateforme de calcul - L'exemple du site de Rennes de Grid'5000'. In: JRES 2024 – Journée Réseaux de l'Enseignement Supérieur. Rennes, France, 2024, pp. 1–13. URL: https://hal.science/hal-04762718 (cit. on p. 18).

Edition (books, proceedings, special issue of a journal)

[20] *SSDBM '24: Proceedings of the 36th International Conference on Scientific and Statistical Database Management*. SSDBM 2024 - 36th International Conference on Scientific and Statistical Database Management. ACM, 23rd Aug. 2024, pp. 1–165. DOI: 10.1145/3676288. URL: https://hal.science/hal-04813517 (cit. on p. 8).

Doctoral dissertations and habilitation theses

[21] C.-K. Huang. 'Scalability of Public Geo-Distributed Fog Computing Federations'. Université de Rennes, 9th Dec. 2024. URL: https://inria.hal.science/tel-04858046 (cit. on p. 10).

Reports & preprints

[22] T. Lambert, S. Ibrahim and T. Jain. *On the (In)Accuracy of Straggler Detection in Big Data Analytic Systems*. 13th Apr. 2024. URL: https://inria.hal.science/hal-04545304 (cit. on p. 14).

Other scientific publications

[23] A. Kazem, G. Pierre and L. Longuevergne. 'Enhancing Observatories With Fog Computing'. In: SSDBM 2024 - 36th International Conference on Scientific and Statistical Database Management. Rennes, France, 2024, pp. 1–1. URL: https://inria.hal.science/hal-04634470 (cit. on p. 11).

[24] R. Liem and S. Ibrahim. 'Towards Optimizing Data Movement in the Computing Continuum Workflow'. In: SSDBM 2024 - 36th International Conference on Scientific and Statistical Database Management. Rennes, France, 2024. URL: https://inria.hal.science/hal-04857454 (cit. on p. 14).

[25] A. Rautureau, J. Paturel, M. Quinson and S. Rokicki. 'Quantifying the tiny-Small design of the SmolPhone'. In: ICT4S 2024 - 10th International Conference on ICT for Sustainability. Stockhlom, Sweden, 2024, pp. 1–4. URL: https://inria.hal.science/hal-04589322 (cit. on p. 17).

[26] M. Rizk, S. Ibrahim and T. Lambert. 'Quantifying the Performance of Erasure Codes in P2P Storage Systems'. In: SSDBM 2024 - 36th International Conference on Scientific and Statistical Database Management. Rennes, France, 2024. URL: https://inria.hal.science/hal-04857465 (cit. on p. 11).

[27] M. Tranzer and S. Ibrahim. 'On the Energy Footprint of Erasure Codes in Ceph'. In: SSDBM 2024 - 36th International Conference on Scientific and Statistical Database Management. Rennes, France, 2024. URL: https://inria.hal.science/hal-04857461 (cit. on p. 11).

Scientific popularization

[28] P. Ciblat, J. Combaz, M. Coupechoux, K. Marquet and A.-C. Orgerie. 'Effets environnementaux de la 5G (Partie 1) : La technologie 5G'. In: *1024 : Bulletin de la Société Informatique de France* 23 (Apr. 2024), pp. 53–72. DOI: 10.48556/SIF.1024.23.53. URL: https://hal.science/hal-03810501 (cit. on p. 33).

[29] P. Ciblat, J. Combaz, M. Coupechoux, K. Marquet and A.-C. Orgerie. 'Effets environnementaux de la 5G (partie 2) : Applications envisagées et acteurs impliqués'. In: *1024 : Bulletin de la Société Informatique de France* 24 (Nov. 2024), pp. 99–127. DOI: 10.48556/SIF.1024.24.99. URL: https://hal.science/hal-04828952 (cit. on p. 33).

12.3 Cited publications

[30] A. Ahmed, H. Arkian, D. Battulga, A. J. Fahs, M. Farhadi, D. Giouroukis, A. Gougeon, F. O. Gutierrez, G. Pierre, P. R. Souza Jr, M. A. Tamiru and L. Wu. 'Fog Computing Applications: Taxonomy and Requirements'. In: *CoRR* abs/1907.11621 (2019). URL: http://arxiv.org/abs/1907.11621 (cit. on p. 4).

[31] J. Brodkin. *Hurricane Sandy takes data centers offline with flooding, power outages*. Ars Technica. https://arstechnica.com/information-technology/2012/10/hurricane-sandy-takes-data-centers-offline-with-flooding-power-outages/. Oct. 2012 (cit. on p. 4).

[32] Cloud Native Computing Foundation. *Kubernetes at the Edge: Organizations are using edge technologies, but there is room to grow*. https://www.cncf.io/blog/2021/05/04/kubernetes-at-the-edge-organizations-are-using-edge-technologies-but-there-is-room-to-grow/. May 2021 (cit. on p. 5).

[33] Eurostat. *Cloud computing – statistics on the use by enterprises*. https://ec.europa.eu/eurostat/statistics-explained/index.php/Cloud_computing_-_statistics_on_the_use_by_enterprises. Jan. 2021 (cit. on p. 4).

[34] Flexera. *Flexera 2020 State of the Cloud Report*. https://info.flexera.com/SLO-CM-REPORT-State-of-the-Cloud-2020. 2019 (cit. on p. 4).

[35] N. Galov. *Cloud Adoption Statistics for 2021*. Hosting Tribunal. https://hostingtribunal.com/blog/cloud-adoption-statistics/. Jan. 2021 (cit. on p. 4).

[36] V. Guseyva. *Data residency laws by country: An overview*. InCountry. https://incountry.com/blog/data-residency-laws-by-country-overview/. Sept. 2020 (cit. on p. 5).

[37] S. Kamal. *Data Residency Compliance with Distributed Cloud Approach*. CPO Magazine. https://www.cpomagazine.com/data-protection/data-residency-compliance-with-distributed-cloud-approach/. Jan. 2021 (cit. on p. 5).

[38] R. van der Meulen. *What Edge Computing Means for Infrastructure and Operations Leaders*. Smarter with Gartner. https://www.gartner.com/smarterwithgartner/what-edge-computing-means-for-infrastructure-and-operations-leaders/. Oct. 2018 (cit. on p. 4).

[39] OVHcloud. *Fire at our Strasbourg Site*. https://us.ovhcloud.com/press/press-releases/2021/fire-our-strasbourg-site. Apr. 2021 (cit. on p. 4).

[40] A. V. Papadopoulos, L. Versluis, A. Bauer, N. Herbst, J. Von Kistowski, A. Ali-eldin, C. Abad, J. N. Amaral, P. Tůma and A. Iosup. 'Methodological Principles for Reproducible Performance Evaluation in Cloud Computing'. In: *IEEE Transactions on Software Engineering* (2019) (cit. on p. 5).