

2025 Activity Report

RESEARCH CENTRE: Inria Centre at Rennes University

IN PARTNERSHIP WITH: CNRS, École normale supérieure de Rennes, Institut national des sciences appliquées de Rennes, Université de Rennes


Project-Team

MAGELLAN

Reliable and Responsible Decentralized Computing
Infrastructures



In collaboration with Institut de recherche en informatique et systèmes aléatoires (IRISA)



Project-Team MAGELLAN

Creation of the Project-Team: 2024 January 01

Each year, Inria research teams publish an Activity Report presenting their work and results over the reporting period. These reports follow a common structure, with some optional sections depending on the specific team. They typically begin by outlining the overall objectives and research programme, including the main research themes, goals, and methodological approaches. They also describe the application domains targeted by the team, highlighting the scientific or societal contexts in which their work is situated. The reports then present the highlights of the year, covering major scientific achievements, software developments, or teaching contributions. When relevant, they include sections on software, platforms, and open data, detailing the tools developed and how they are shared. A substantial part is dedicated to new results, where scientific contributions are described in detail, often with subsections specifying participants and associated keywords. Finally, the Activity Report addresses funding, contracts, partnerships, and collaborations at various levels, from industrial agreements to international cooperations. It also covers dissemination and teaching activities, such as participation in scientific events, outreach, and supervision. The document concludes with a presentation of scientific production, including major publications and those produced during the year.

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
- A1.3. – 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
- A2.6. – 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
- A4.5.2. – Model-checking
- 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

Contents

Project-Team MAGELLAN	1
1 Team members, visitors, external collaborators	6
2 Overall objectives	7
3 Research program	8
4 Application domains	10
5 Social and environmental responsibility	10
5.1 Footprint of research activities	10
5.2 Impact of research results	10
6 Highlights of the year	11
6.1 Awards	11
7 Latest software developments, platforms, open data	11
7.1 Latest software developments	11
7.1.1 Tansiv	11
7.1.2 SimGrid	12
7.1.3 EnOSlib	13
8 New results	13
8.1 Reliable decentralized computing infrastructures	13
8.1.1 Fog/cloud storage services	13
8.1.2 Fog/cloud infrastructure design	14
8.1.3 Fog/cloud networking services	14
8.1.4 Reliable fog platforms in resource-constrained natural environments	15
8.2 Reliable decentralized application runtimes	15
8.2.1 Reliable serverless runtimes	15
8.2.2 Decentralized resource provisioning for serverless applications	16
8.2.3 Simulation of FaaS applications at the edge	17
8.2.4 Dynamic platform adaptation for Federated Learning	17
8.2.5 Reliable data stream processing runtimes	17
8.3 Distributed infrastructure frugality	18
8.3.1 Energy monitoring of Cloud servers	18
8.3.2 Estimating carbon footprint of computing infrastructures	19
8.3.3 Energy-efficiency in distributed systems	19
8.3.4 Setting limits to reduce the energy consumption of digital infrastructures	20
8.3.5 Frugal mobile computing	21
8.4 Evaluation methodologies and tools	21
8.4.1 Simulating distributed IT systems	21
8.4.2 Understanding software infrastructures	22
8.4.3 Verifying distributed system’s correctness	22
8.4.4 Designing problem-specific computing infrastructures	22
8.4.5 Supporting experiments on distributed systems’ energy consumption	23
8.4.6 Predictive models and simulation frameworks for Wi-Fi networks	23
8.4.7 Sustainable research ecosystem	23
9 Bilateral contracts and grants with industry	24
9.1 Bilateral contracts with industry	24

10 Partnerships and cooperations	25
10.1 International initiatives	25
10.1.1 Associate Teams in the framework of an Inria International Lab or in the framework of an Inria International Program	25
10.2 European initiatives	26
10.2.1 Digital Europe	26
10.3 National initiatives	27
10.4 Regional initiatives	34
11 Dissemination	34
11.1 Promoting scientific activities	34
11.1.1 Scientific events: organisation	34
11.1.2 Scientific events: selection	35
11.1.3 Journal	36
11.1.4 Invited talks	36
11.1.5 Leadership within the scientific community	36
11.1.6 Scientific expertise	36
11.1.7 Research administration	37
11.2 Teaching - Supervision - Juries - Educational and pedagogical outreach	37
11.2.1 Teaching	37
11.2.2 Supervision	38
11.2.3 Juries	39
11.3 Popularization	40
11.3.1 Productions (articles, videos, podcasts, serious games, ...)	40
11.3.2 Participation in Live events	40
11.3.3 Others science outreach relevant activities	40
12 Scientific production	40
12.1 Major publications	40
12.2 Publications of the year	40
12.3 Cited publications	45

1 Team members, visitors, external collaborators

Research Scientists

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

Faculty Members

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

Post-Doctoral Fellows

- Tiago Da Silva Barros [INRIA, Post-Doctoral Fellow, from Dec 2025]
- Armel Jeatsa Toulepi [INRIA, Post-Doctoral Fellow, from Apr 2025]
- Anudipa Mondal [CNRS, Post-Doctoral Fellow, from Apr 2025]
- Ophelie Renaud [ENS RENNES, Post-Doctoral Fellow]
- Aya Shehata [CNRS, Post-Doctoral Fellow, from Sep 2025]
- Hayfa Tayeb [CNRS, Post-Doctoral Fellow, from Jun 2025]

PhD Students

- Khaled Arsalane [UNIV RENNES, until Nov 2025]
- Matteo Chancerel [CNRS]
- Leo Cosserson [ENS RENNES]
- Haraesh Jayasethu Ramachandran [UNIV RENNES, from Oct 2025]
- Ammar Kazem [UNIV RENNES]
- Govind Kovilkkatt Panickerveetil [INRIA]
- Mohamed Cherif Zouaoui Latreche [INSA RENNES]
- Mathieu Laurent [ENS RENNES]
- Minh Thong Le Viet [UNIV RENNES, from Oct 2025]
- Pablo Leboulanger [CNRS]
- Volodia Parol-Guarino [INRIA, until Sep 2025]
- Mohammad Rizk [INRIA]

- Ahmed Rjiba [INRIA, from Nov 2025]
- Matthieu Silard [CNRS]
- Marc Tranzer [INRIA]

Technical Staff

- Victorien Elvinger [INRIA, Engineer]
- Matthieu Nicolas [INRIA, Engineer]
- Loris Penven [INRIA, Engineer, until Sep 2025]
- Matthieu Simonin [INRIA, Engineer]

Interns and Apprentices

- Haraesh Jayasethu Ramachandran [UNIV RENNES, Intern, from Feb 2025 until Jul 2025]
- Yann Laine Odic [CNRS, Intern, from Mar 2025 until Aug 2025]
- Elsa Mühl [UNIV RENNES, Intern, from Jun 2025 until Aug 2025]
- Mahfoud Sayah [CNRS, Intern, from May 2025 until Aug 2025]

Administrative Assistant

- Amandine Seigneur [INRIA]

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 [51, 53].

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 [47]. 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% [56]. 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 premises 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 [52].
- **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 [57]) or natural disasters (e.g., Hurricane Sandy in October 2012 [48]). 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.
- **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 [54]. Decentralized computing infrastructures are a natural solution to respect such regulations while maintaining a global management of such data across the geographical boundaries [55].

The development of reliable large-scale decentralized virtualized computing infrastructures however remains in its infancy [50]. 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 [58]. 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

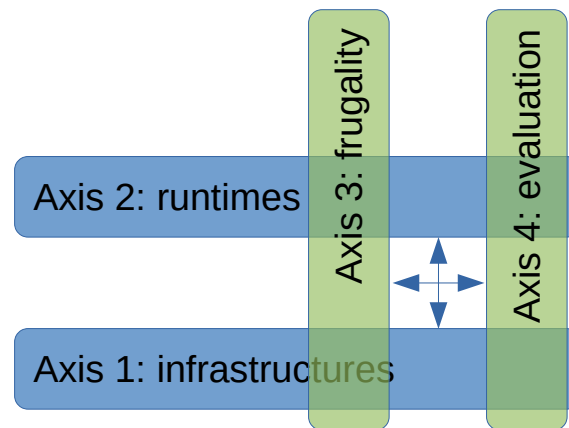


Figure 1: Magellan research axes organization.

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 GDRS 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. For the GreenDays, multi-GDR thematic days, she wrote the code of conduct of the event [42].

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. In the local context Matthieu Simonin has been appointed by the head of Inria to assess the carbon footprint of the local activities. Work is in progress for producing relevant indicators in the field of computing infrastructure usage and travels[39].

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 OVHcloud 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

- Shadi Ibrahim defended his “Habilitation à diriger des recherches” on May 14th 2025.
- Shadi Ibrahim co-organized a Dagstuhl Seminar on [Storage Systems and I/O for Emerging Workloads on HPC Systems](#) from Jan 11 to Jan 16, 2026.
- Cédric Tedeschi was elected director of ISTIC, the University of Rennes' department of computer science and electronics.
- The [13th IEEE International Conference on Cloud Engineering \(IC2E\)](#) took place in Rennes on September 23rd-26th 2025. Several team members were part of the organization committee: Guillaume Pierre as general co-chair, Anne-Cécile Orgerie as program co-chair, Shadi Ibrahim as workshop co-chair, Ammar Kazem as Web chair, Amandine Seigneur as local arrangement co-chair.

6.1 Awards

- Wedan Emmanuel Gnibga received a [runner-up PhD thesis award](#) from GDR RSD.

7 Latest software developments, platforms, open data

7.1 Latest software developments

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 guarantee 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, Fogs, Clusters and HPC, allowing multi-domain studies. It can be used either to simulate algorithms and prototypes of applications, or to emulate real MPI applications through the virtualization of their communication, or to formally assess algorithms and applications that can run in the framework.

The formal verification module explores all possible message interleavings in the application, searching for states violating the provided properties. This tool can be used to assess safety 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 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: Breaking the seal: v4.0 was not the final release.

* Allow one to unseal netzones to modify the platform even after the simulation start. * The model-checker can now report memory race conditions (see tutorial). * Pip builds should now work out of the box. * (+ the usual bug fixes overall, and improvements to the Java/Python bindings).

News of the Year: There were 2 major releases in 2025. We released v4.0 in March, embodying 10 years of development. This turns SimGrid into a mature and stable research instrument. The users can easily extend this tool to adapt it to their specific research, while trusting its software implementation. Release v4.1 was published in November, showing that the development did not stall even if the framework is mostly in maintenance mode. The performance simulation mode was extended to allow modifications of the platform topology during the simulation.

Beyond these technical development, most of the scientific work occurred in the embedded model checker. We first worked on improving the performance of the verification process, and fixed some bugs. Our preferred algorithm (ODPOR reduction + Best-first exploration to enable random walk) is now much faster and consumes much less memory, being 5x faster on small scenarios. But since performance is now linear with the number of states, while it used to be polynomial, 5x faster is the lowest performance boost you can expect from the new version. The number of states to explore for a given scenario is still the same (ODPOR was not improved), but we now can explore these states much faster.

In addition, we added the ability to check for race conditions in the user code. This feature relies on a specific pass to the clang LLVM compiler, to instrument the memory accesses. We use this feature for teaching purposes at our institutions.

We are working on a parallel explorer leveraging all cores to accelerate the exploration, but unfortunately, we did not manage to find all the bugs in our parallel explorer yet. We are working on verifying SimGrid with itself to hammer this bug out.

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

Publication: hal-04909441

Contact: Martin Quinson

Participants: Mathieu Laurent, Anne-Cécile Orgerie, Arnaud Legrand, Augustin Degomme, Arnaud Giersch, Frédéric Suter, 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: Mathieu Simonin, 6 anonymous participants

8 New results

8.1 Reliable decentralized computing infrastructures

8.1.1 Fog/cloud storage services

Participants: Shadi Ibrahim, Marc Tranzer, Mohammad Rizk, Armel Jeatsa Toulepi.

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 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.

In this context, to effectively and efficiently adopt EC in distributed storage systems, we study the effect of data chunk distribution on the performance of reads and data-intensive applications. We also present the design and evaluation of an EC-aware block placement that balances the distribution of data chunks across nodes [18]. In addition, we continued working on (i) understanding the tradeoffs between data access

performance, repair bandwidth, and energy overhead in a CephFS cluster; and (ii) addressing the challenge of data retrieval in geo-distributed erasure-coded storage systems. Two more detailed articles on these topics are in preparation. We also designed and implemented an erasure-coded storage system that efficiently and effectively integrates Reed-Solomon codes on top of the InterPlanetary File System (IPFS). A thorough paper on this topic is under review.

Another issue related to storage services is in-memory caching. We study how to use compression to increase the effective memory capacity, enabling the node to handle a larger load while ensuring a fast response time. A detailed article on this topic is in preparation.

8.1.2 Fog/cloud infrastructure design

Participants: Shadi Ibrahim.

Containers, renowned for their lightweight nature and flexibility, have seen growing adoption for deploying edge services such as web applications. However, existing cloud-oriented container deployment frameworks fail to address the unique challenges of edge environments, including geographical distribution, device heterogeneity, and resource constraints. This oversight leads to suboptimal performance for latency-sensitive edge services like HPC/AI-powered autonomous driving and edge gaming, which demand rapid startup and immediate responsiveness. Our investigation demonstrates that current on-demand image solutions require excessive client-registry communication, resulting in prolonged Round-Trip Time (RTT) - a particularly severe limitation in geographically distributed edge platforms. Furthermore, we observe that the user-space file system (e.g., FUSE), typically employed to handle device heterogeneity, introduces substantial overhead to the native I/O stack. More critically, our findings reveal that on-demand image solutions exacerbate storage pressure on resource-constrained edge devices. To overcome these challenges, we introduce EDDE, an edge-optimized container deployment framework that redesigns the on-demand image pipeline [16].

Container orchestration systems, such as Kubernetes, streamline containerized application deployment. As more and more applications are being deployed in Kubernetes, there is an increasing need for rescheduling - relocating a running pod to different nodes - due to system upgrades, node failures, and load-balancing optimizations. Live migration, which transfers services from source nodes to target nodes with minimal downtime, is the ideal support for rescheduling. However, implementing live migration for pods that run stateful services is challenging, because Kubernetes manages pods as stateless. First, the current pod's network namespace initialization process causes a mismatch in the network state between the migrated pod and internal containers. Second, migrating the memory state results in extended downtime. Third, Kubernetes operations on pods do not consider preserving the state of the pods. Therefore, we propose KubeSPT to achieve live migration of stateful pods in rescheduling scenarios [8]. First, we synchronize the network state of pods and internal containers by controlling packet flow and implement fast service redirection. Second, we introduce a Hot Data and Lazy-Restore method for memory restoration to reduce migration downtime.

8.1.3 Fog/cloud networking services

Participants: François Lemercier, Guillaume Pierre, Haraesh Jayasetu Ramachandran, Yann Lainé-Odic.

In the context of Industrial IoT applications (i.e. critical response time environments such as manufacturing or monitoring of critical equipment), latency is at the center of a tremendous number of studies aiming to optimize the placement of resources in distributed architectures. To ensure that the quality of service is guaranteed, several solutions exist to reconfigure component placement (migration) and can reduce the overall latency by changing components and routes. However, knowing precisely which component is the source of problematic latency remains scarcely addressed. When taking a decision for a reconfiguration or a migration, which can be triggered due to latency issues, it can be beneficial to check whether the source of the latency can be solved before instantiating a migration or a full reconfiguration. Proper measurement protocols exist but generally refer to specific case studies and would not allow straightforward integration

into edge systems. In this context, we have started work to characterize and evaluate the different sources of latency in such environments, before proposing optimization solutions, notably routing-based approaches, to reduce latency.

“Clusterize” was an Artistic / Scientific project driven by François Lemerrier and Loïc Nguyen. It aimed to create a fluid and immersive soundscape, without a beginning or an end, embodying the dynamics of a geo-distributed cluster. This musical installation served as a living mirror of digital events, where each piece of data within the cluster played a role, translated into sounds that evolved with the data flows as well as the gestures and interactions of the audience. A dedicated platform was developed to highlight the artifacts inherent to the execution of a geo-distributed Kubernetes cluster, in order to feed and control the sonic material of the artist participating in the Clusterize project.

8.1.4 Reliable fog platforms in resource-constrained natural environments

Participants: Guillaume Pierre, Ammar Kazem, Matthieu Nicolas, Elsa Mühl, Har-aesh Jayasethu Ramachandran.

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 [43].

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 [4]. 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. We presented a first demo of our solutions [21], and a more thorough paper on this topic was submitted.

We continued the work started last year to redesign and improve the already-existing LivingFog platform [1]. 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. A first version of this platform was deployed in the Kaligandaki observatory in Lete, Nepal (see Figure 2).

Finally, we developed the LivingBench benchmarking tool which is designed to exercise an IoT/Edge computing platform based on a use-case derived from our work on the LivingFog platform [20].

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



Figure 2: LivingFog installation in Lete, Nepal.

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 are developing FaaS scheduling approaches based on Reinforcement Learning (RL). Our initial work proposed an RL-based approach for optimizing the placement of FaaS functions within a single cluster. Building on this work, in 2025, we have developed a FaaS control plane that spans across multiple geo-distributed clusters, jointly optimizing latency and energy efficiency under dynamically evolving FaaS workloads. Experiments conducted on Grid’ 5000 allowed us to compare RL-based scheduling methods with baseline scheduling methods and explore the trade-offs between centralised and decentralised RL-based control. An article on this topic is currently under submission.

Another limitation of FaaS runtimes in edge environments with resource-constrained devices is memory waste caused by keeping function sandboxes idle to avoid cold starts. To address this, we applied a mechanism that unloads libraries from idle sandboxes and restores them on demand, achieving significant memory savings and improved warm-start performance [25].

8.2.2 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. To evaluate our approach, we performed reproducible experiments on the Grid’5000 testbed at large scale (up to 663 fog nodes on 60 servers). The evaluation demonstrated GIRAFF’s effectiveness in optimizing FaaS application placement [26]

In 2025, we extended GIRAFF with a client library for building fog-native FaaS applications. The library enables application owners to manage scaling, control costs, and handle resource shortages.

8.2.3 Simulation of FaaS applications at the edge

Participants: Nikos Parlavantzas, François Lemercier, Sayah Mahfoud Abd El Ali.

Serverless computing is an increasingly popular model for delivering cloud services, isolating developers from infrastructure management and allowing them to focus on application logic. At its core, the Function-as-a-Service (FaaS) model runs short-lived, event-triggered functions that consume resources only during execution.

Applying FaaS to edge computing has gained attention, as it reduces resource usage on constrained edge nodes and supports IoT applications and virtual network functions. However, few simulation tools exist for executing FaaS applications in edge environments. In this context, we benchmarked existing solutions and modified a simulator to include an energy model and a use case tailored to our research project.

8.2.4 Dynamic platform adaptation for Federated Learning

Participants: Cédric Tedeschi, Mathis Valli, Minh Thong Le Viet.

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.

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.

We also started to investigate alternate ways where Edge devices learn locally and opportunistically exchange their model so as to build a more precise model. Finding the right trade-off between accuracy of the models obtained within an area (composed of one or several interrelated nodes) and the network traffic generated will constitute the main driver for the solution.

8.2.5 Reliable data stream processing runtimes

Participants: Guillaume Pierre, Khaled Arsalane, Govind Kovilkkatt Panickerveetil, Shadi Ibrahim, Tiago Da Silva Barros.

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 explored the limitations of stream processing autoscaling systems when dealing with stateful data processing operators in geo-distributed environments where the processing capacity of the compute nodes are heterogeneous. We identified the challenges and limitations of current approaches and proposed directions for addressing them [12]. Khaled Arsalane defended his PhD thesis on this topic on December 15th 2025 [36].

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 developed a methodology to assist Kafka users in exploring the space of workable configurations and in selecting the optimal one for a given expected workload. A paper on this topic was recently accepted for publication.

Finally, we are investigating how to optimize the allocation of resources and tasks when deploying data streaming processing applications in the Fog, with a focus on latency, throughput, energy consumption, and maximum sustainable throughput.

8.3 Distributed infrastructure frugality

8.3.1 Energy monitoring of Cloud servers

Participants: Anne-Cécile Orgerie, Maxime Agusti, Anudipa Mondal, Hayfa Tayeb.

In an effort to raise awareness on the increasing carbon emissions of Cloud computing, the European Corporate Sustainability Reporting Directive effectively requires providers to supply their customers with an assessment of the carbon impact associated with their use. This represents a challenge for bare metal servers, where the deployment of dedicated power meters is often unfeasible at scale. To address this, we presented PPEM-BM [11], a novel sensor-driven modeling approach to estimate the power consumption of bare metal servers using CPU temperature data acquired via IPMI. PPEM-BM enhances and generalizes the existing POWERHEAT method, which correlates CPU temperature with power. Our methodology involves training individual power models, performing cross-evaluation to determine their portability, and then using a Learning to Rank (LTR) model to select the most appropriate pre-trained model for a target server based on its hardware configuration and CPU temperature statistics. An experiment conducted on 1,076 production servers at OVHcloud shows that PPEM-BM demonstrates a significant improvement compared to models based solely on hardware profiles. The approach offers a practical, scalable, and cost-effective solution for hosting providers to monitor energy consumption without widespread sensor deployment.

As the demand for AI-driven workloads increases, the energy consumption of Graphics Processing Units (GPUs) devices has come under intense scrutiny, particularly in hyperscale data centers where large numbers of accelerators are centralized and leased to diverse clients. In the context of cloud hyperscalers, GPU power monitoring presents several challenges that vary depending on the product offered. The monitoring capabilities of physical devices may be limited or even absent for some products. However, given the substantial energy demands of GPUs, power monitoring is essential for both cloud providers and clients. Operators require tools to manage power distribution effectively, such as balancing workloads across Power Distribution Units (PDUs), while clients need visibility into power usage to optimize their workloads for energy efficiency. To address these challenges, we proposed methods for estimating the energy consumption

of jobs running on GPU devices in cloud environments, spanning from shared and managed offerings like ML-as-a-Service (MLaaS) to less managed products (e.g., Infrastructure-as-a-Service (IaaS)) [19]. Our models demonstrate the benefits of sharing GPUs for small AI workloads, as well as the current sub-optimal utilization of GPUs in cloud hyperscalers, based on insights from an IaaS GPU cluster.

8.3.2 Estimating carbon footprint of computing infrastructures

Participants: Anne-Cécile Orgerie, Matthieu Simonin, Matteo Chancerel.

Attributing the carbon costs of shared ICT infrastructures to its end-users is frequently promoted as a way to encourage awareness of environmental impacts and advocate for more sustainable practices. We explored the intricacies of this approach by focusing on shared ICT infrastructures specifically dedicated to academic research, several of which having recently introduced carbon intensity values for their users [17]. This scenario serves as a practical case study for examining the methodologies and challenges associated with evaluating the carbon intensity of shared ICT infrastructures. We explore the choices with their limitations, discuss the objectives behind their implementation of this type of environmental indicator and offer actionable insights. This analysis aims to contribute to the broader discussion on sustainable computing practices and the role of environmental indicators in driving meaningful change.

To reduce their carbon footprint, some Fog operators power their nodes with renewable energy combined with batteries to prevent intermittency. Yet, often they do not consider the embodied carbon footprint of the power infrastructure. Moreover, having 24h/7d uptime for Fog services comes at a high embodied carbon cost and do not favor the right to disconnect. We modeled a Fog infrastructure, geographically spread, powered by solar panels, wind turbines and batteries for its whole life time [13]. Its nodes are turned off during the night and can be intermittent during the day. We provide a theoretical framework to analyze the balance between a satisfying uptime of the service, and a small carbon footprint. Our simulation-based evaluation using real traces shows that reaching 95% of uptime during work hours rather than 99% divides by two the size of the powering infrastructure, thus decreasing the carbon footprint of the Fog and providing a usecase for combining both environmental and social concerns of Fog users.

Finally, we studied the uniqueness of large scale real-life fine-grained electrical consumption time-series and show its link to privacy threats [7].

8.3.3 Energy-efficiency in distributed systems

Participants: Anne-Cécile Orgerie, Matthieu Silard.

Diverse computing paradigms are continuously emerging, transforming how a device interacts with data and completes tasks. These advancements offer functional benefits, like allowing even basic devices to run powerful applications efficiently by offloading demanding tasks to computing centers with available processing power and storage. This ease of task delegation has led to a boom in the number of users, more powerful and critical applications, and a constant drive to expand computing infrastructures to keep up with intensifying demands. However, this growing capacity and geographical distribution comes with a rising environmental cost. Some solutions are being developed to mitigate the environmental impact, including infrastructure optimization (hardware and design choices), finding dynamic ways for lower energy use (both in the hardware and software), using eco-friendly energy sources, etc. We proposed a state of the art that delves into the evolution of the computing landscape, the environmental challenges it presents, and the solutions being developed to address them [34].

Due to some overprovisioning policies and variable usage, data centers in production can face low average resource utilization. This can result in a waste of underused servers and energy. In this context, virtual machine (VM) consolidation combined with shutdown policies can be a pertinent approach for improving resource utilization and reducing energy consumption of the entire cloud infrastructure. However, VM consolidation requires expensive migration techniques, which can potentially affect performance.

Consolidation of workload has been proposed and studied as a core capability since the invention of the Cloud. But after two decades of deployment of Cloud infrastructures, VM consolidation is still rarely used in production for small and large-scale environments. We explored and revisit the potential of savings that can be achieved through a versatile and efficient Virtual Machine consolidation in small and large-scale production infrastructures through usage analysis of two Cloud providers infrastructures [41, 5]. We show that potential benefits in terms of saved cloud resources and energy usage reduction can occur for systems in production.

Energy consumption in 5G/6G wireless networks presents a fundamental challenge, encompassing various aspects related to energy, availability, cost, and environmental concerns. According to the literature, the use of sleep modes (SMs) has emerged as an effective strategy to improve the energy efficiency of base stations (BS) within recently proposed 5G frameworks, while also considering future 6G developments. We proposed a comprehensive review to elucidate what SMs are, the decision process to switch to and from SM BSs, and the main challenges associated with them [6]. Additionally, power consumption models are analyzed to understand how to evaluate the performance of SMs. Following an overview of the primary SM strategies developed in recent years, we introduce a matrix model classification of these strategies. This classification is based on the policies and tools implemented for SMs in wireless networks. Finally, based on this analysis, we highlight the most promising strategies and present ideas for future research aimed at achieving an energy-efficient and sustainable use of SMs in 5G/6G wireless networks.

The massive integration of non-dispatchable Distributed Energy Resource (DER) and the growing electrification of uses, notably via Electric Vehicles (EVs), poses major challenges for Distribution System Operator (DSO). These evolutions are causing congestion on Medium Voltage (MV) networks, and increase uncertainty of production periods. To address these challenges without costly infrastructure investment, we proposed an innovative demand management solution based on direct, intermediary-free communication between the DSO and users [28]. Leveraging existing Advanced Metering Infrastructures (AMIs), this approach broadcasts dynamic incentive signals to Electric Vehicle Supply Equipments (EVSEs), which enables optimized management of energy flows, and thus provides a rapid response to network needs. Our solution has the advantage of using existing power and network infrastructures and does not require to gather private information. In addition, it allows users to choose whether or not to provide their devices as flexible loads. The results show significant potential for improving grid stability and reducing congestion, offering a cost-effective, scalable alternative for smart grids.

Smart grids are transforming electricity network management through the use of Advanced Metering Infrastructure (AMI). They facilitate the balance between production and consumption of electricity by providing real-time data and consumption control for the Distribution System Operator (DSO). However, the implementation details, which are often proprietary, are not readily available. We proposed a demo that presents an innovative visualization platform to simulate a Low Voltage (LV) residential network of 48 smart meters [29]. The platform provides a visual understanding of the communication dynamics in an AMI by reproducing message exchanges and simulating link disruptions. This demo presents the state-of-the-art G3-PLC protocol stack currently employed in the French AMI deployment and demonstrates the transmission of routing and application messages generated by LOADng and DLMS/COSEM, respectively.

8.3.4 Setting limits to reduce the energy consumption of digital infrastructures

Participants: Anne-Cécile Orgerie, Pablo Leboulanger, Robin Richard.

Ten years after the drafting of the Paris Agreement, the objective to reduce greenhouse gas emissions by half by 2030 seems difficult to reach. The impact of ICT technologies is growing year after year. Data centers have a large impact in this domain. As a lot of research is focused on improving energy efficiency of data centers, one could ask if this improvement would suffice to reduce greenhouse gas emissions [45]. To study the impact of the energy efficiency on the long term, we proposed to model a data center that has a limited power capacity and is regularly refreshed with new servers, more energy-efficient than the previous ones, to cope with its workload [23]. Our results explore various growth rates for the load, and the energy efficiency and show that without reduction of usage of data center resources, the improvements in energy efficiency will not be not enough to reach the Paris Agreement's objective.

To reduce the ecological impact of a technology, scientists often focus on energy efficiency issues, ignoring the complex rebound effect generated. We focused on the video transmission technology, and discuss the urgent need to be able to set limits [30]. We show that these limits can provoke opposition or circumvention, illustrating the difficulty of the task. Finally, we conclude that the question of limits must be considered as a research problem in its own right, and that it is intrinsically multidisciplinary.

8.3.5 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 entered its second year in 2025. On the hardware side, Pierre Nozet (hosted in the TARAN team) designed the first prototypes of the PCB for a future devboard involving heterogeneous multiprocessors and low-power technologies for the screen and network. More work is needed to design a fully functional prototype. On the software side, Victorien Elvinger implemented an adapted runtime that can be used to implement graphical applications targetting limited hardware, but still making it possible to write the applicative logic in a scripting language compiled on the board.

8.4 Evaluation methodologies and tools

8.4.1 Simulating distributed IT systems

Participants: Martin Quinson.

Our team plays a major role in the advance of the SimGrid simulator of IT systems. This framework has a major impact on the community. Cited by over 600 scientific papers (including 60 PhD thesis and 150 journal articles), it was used as a scientific instrument by thousands of researchers over the years. This year, we coordinated the release of SimGrid v4.0 [46], concluding ten years of development. We also published a journal article describing resulting framework and its contributions to the field of distributed system simulation [3].

This year, we used this tool to propose an approach to the problem of measuring the performance of one-sided and collective operations on a set of distributed resources. This problem is not trivial because of the complexity of distributed time measurement on non-synchronized hardware. Simulation is attractive because distributed applications can be executed on a controlled set of resources and in a reproducible manner. Writing a fully calibrated simulator for a new software library represents however a time- and labor-intensive effort that is usually only undertaken for well established software solutions such as the MPI standard. Our approach, detailed in [15] makes it possible to study an existing standard for high-performance computing parallel applications, featuring one-sided communications and atomic and collective operations. We present the several mechanisms that are mandatory to integrate the behavior of the real distributed applications into the simulated platform. Leveraging this simulation framework greatly simplifies this endeavor, opening the way to the study of more distributed middleware and libraries through realistic simulations.

8.4.2 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 dynamic malware detection and 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 and high-end sandboxes manipulate virtual clocks to mask VM execution pauses caused by VMI. However masking VM execution pauses on multicore VMs assumes that pauses involve all virtual cores, which breaks introspection practices observed with popular tools like libVMI. Indeed some introspection procedures require to pause a virtual core while waiting for an event on another virtual core. Moreover, despite virtualization not being a reliable hint for a malware to be under analysis, single-core VMs are less and less common in production workloads and having only one core is a reliable hint of being analyzed. To lower the VMI detection ability of malware on multi-core VMs, we first proposed several metrics to characterize the side effects caused by a VMI monitor in multi-core VMs. Then we introduced a new strategy to conceal VMI pauses in multi-core guests, and asserted its performance compared to the regular approaches used by sandboxes. Our results show that this new strategy improves the stealthiness of VMI pauses on a multi-core system for some metrics. This work is published in [14].

This work on hiding VMI pauses was also presented in details in the more general framework of Léo Cosseron's PhD thesis defended in December 2025 [2], where controlling the flow of time in VMs allows to interconnect VMs with a network simulator. This combination of hardware virtualization and network simulation provides both a tool to analyze real distributed applications in reproducible setups and the infrastructure for a performance-wide stealth network environment in a malware analysis sandboxes.

8.4.3 Verifying distributed system's correctness

Participants: Mathieu Laurent, Martin Quinson.

Assessing the correctness of distributed and parallel applications is notoriously difficult due to the complexity of the concurrent behaviors and the difficulty to reproduce bugs. Model checking and systematic debugging is an appealing solution, but it remains limited to small programs because of the states space explosion problem. In this context, Dynamic Partial Order Reduction (DPOR) techniques have proved successful in exploiting concurrency to verify applications without exploring all their behaviors. However, they may lack of efficiency when tracking non-systematic bugs of real size applications. In this work [22], we suggest two adaptations of the Optimal Dynamic Partial Order Reduction (ODPOR) algorithm with a particular focus on bug finding and explanation. The first adaptation is an out-of-order version called RFS ODPOR which avoids being stuck in uninteresting large parts of the state space. Once a bug is found, the second adaptation takes advantage of ODPOR principles to efficiently find the origins of the bug.

8.4.4 Designing problem-specific computing infrastructures

Participants: Ophelie Renaud, Martin Quinson.

The Square Kilometre Array Observatory (SKAO) project aims to build the two largest telescope arrays in the world, located in South Africa and Australia. These instruments will generate unprecedented volumes of data, requiring its computing infrastructure (called Science Data Processor – SDP) to operate at terabyte rates under strict energy and performance constraints. Anticipating the computational cost of imaging algorithms

is therefore critical for both astronomers and system designers. Yet, testing new algorithmic strategies directly on large HPC platforms is often impractical, due to cost, limited availability, and long development cycles. To address this, we proposed SimSDP [44, 27], a rapid prototyping tool designed to support the development and optimization of radio-astronomy imaging pipelines for the future SKAO SDP. The tool enables early exploration of algorithmic choices and their impact on computational performance and energy consumption without requiring access to large-scale systems.

This year, we illustrated the use of SimSDP through three case studies: (1) modeling different strategies for spectral parallelism (joint vs. distributed deconvolution), (2) simulating pipeline execution on large-scale multi-core, multi-node systems with realistic observation sizes, and (3) prototyping pipelines with real-world data from the NenuFAR instrument. These experiments highlight how SimSDP can help astronomers and instrument designers anticipate the computational cost of imaging algorithms, compare different processing strategies, and better align algorithmic development with the capabilities of future HPC infrastructures [33, 32].

8.4.5 Supporting experiments on distributed systems' energy consumption

Participants: Matthieu Simonin.

EnOSlib [49] software has been recently extended to ease the collection of energy metrics. This allows experimenters to gather energy metrics from different sources being software probes or environmental metrics offered by SLICES-FR platform. This new feature is crucial for supporting experiment driven research in the field of distributed system experiments on energy consumption. This has been exemplified in various workshops (e.g., [10]).

8.4.6 Predictive models and simulation frameworks for Wi-Fi networks

Participants: François Lemercier, Anne-Cécile Orgerie, Aya Shehata.

We proposed a study of the current consumption of various Medium Access Control (MAC) states for Wi-Fi devices, focusing on the latest IEEE 802.11ax (Wi-Fi 6) and 802.11ac standards [24]. Through detailed experimental measurements, we provide new, precise values for the current drawn in MAC states such as idle, transmission, reception, and sleep. We capture real-time current consumption in different scenarios with various traffic loads, offering an understanding of the power usage across different operational modes. The results reveal significant differences in power consumption patterns compared to earlier Wi-Fi standards (e.g., 802.11n), with notable improvements in energy efficiency in sleep and idle states, but increased current draw in transmission and reception states. These findings provide valuable insights into the energy behavior of recent Wi-Fi devices, offering a solid basis for the simulation of power management optimization strategies in next-generation wireless systems.

8.4.7 Sustainable research ecosystem

Participants: Anne-Cécile Orgerie, Matthieu Simonin.

Research on the energy transition raises a series of major challenges: it is a problem-oriented field of research (that of decarbonising the energy mix), which must be resolutely ethical and include principles of social and environmental justice from the very beginning... and this despite the problematic nature of the notion of transition which gives a misleading notion of the process that will lead to desired changes - if there are any. Taking this observation as its starting point, we set out a number of avenues for research into key issues [9]. It begins by examining the way in which the energy transition is transforming the interplay between

actors in the energy sector and the scales at which decisions are taken and implemented. While inviting to take account of the wide variety of national contexts and levels of development, we acknowledge the certainties and uncertainties surrounding the composition of the energy mix. It also stresses the importance of taking account of the externalities of technological choices, particularly in environmental terms. This enables us to draw up a research agenda focusing not only on the spatial aspects of transitions, but also on the ways in which economic models need to incorporate new parameters and new constraints. We also invite to consider how the relationship between energy production and consumption is changing, and how consumption needs to be reduced. In so doing, we defend the idea that research, which is necessarily interdisciplinary and reflexive, has a strong role to play in basing transition action on scientifically sound findings.

We produced a Manifesto from the Perspectives Workshop 25122 entitled “Climate Change: What is Computing’s Responsibility?” held March 16-19, 2025 at Schloss Dagstuhl, Germany [40]. The workshop provided a forum for world-leading computer scientists and expert consultants on environmental policy and sustainable transition to engage in a critical and urgent conversation about computing’s responsibilities in addressing climate change - or more aptly, climate crisis.

We also produced a quiz on the environmental impacts of ICT that was used during several conferences and thematic days [37].

We worked with the Paris Observatory to establish their carbon footprint due to ICT infrastructures and devices [38].

We are involved in the Labos1point5 community in which we actively contribute to the implementation of tools and also the animation of the network of laboratories in transition (started in early 2024). The latter offers several opportunities of studies made in conjunction with sociologists.

We have helped several conferences to assess their carbon footprint (IEEEVR - 1500 attendees, SFRBM - 200 attendees, GreenDays and IEEE IC2E - 70 attendee). Initial results show that attending abroad conferences is very carbon intensive: emissions per day and per capita might be 2 orders of magnitude more than the objective set by the Paris agreements. This advocates for finding ways of disseminating scientific results in a sustainable manner in the future.

9 Bilateral contracts and grants with industry

9.1 Bilateral contracts with industry

Défi Inria OVHCloud - FrugalCloud

Participants: Anne-Cécile Orgerie, Shadi Ibrahim, Govind Kovilkkatt Panickerveetil, Guillaume Pierre, Maxime Agusti, 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

Additional 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 ecodeign 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 - Alvearium

Participants: Shadi Ibrahim, Mohammad Rizk, Quentin Acher, Armel Jeatsa Toulepi.

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

Partner Institution(s):

- Inria, France
- Hive, France

Date/Duration: 2023-2026

Additional 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.

Défi Inria Hive - Cupseli

Participants: Shadi Ibrahim.

Title: **Cupseli:** Collaborative Unified Platform for a Scalable and Efficient Learning Infrastructure.

Partner Institution(s):

- Inria, France
- Hive, France

Date/Duration: 2025-2029

Additional info/keywords: The goal of this collaborative framework between Hive and Inria is to demonstrate that even the most demanding AI and Big Data applications can run efficiently on heterogeneous, distributed, and volatile resources — while maintaining accuracy, ensuring privacy, and reducing environmental impact. Eleven Inria project-teams are involved in this challenge including ARGO, MIMOVE, COAST, MAGELLAN, STACK, WIDE, OCKHAM, COATI, NEO, TADAAM and TOPAL. Members of the MAGELLAN team will contribute to the axis on large-scale computing on intermittent resources.

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

Title: Scalable and Energy Efficient Data Management for Scientific Workloads in Computing Continuum

Duration: 2024 - 2026

Coordinator: Kesheng WU (kwu@lbl.gov)

Partners:

- Lawrence Berkeley National Laboratory Berkeley (États-Unis)

Inria contact: Shadi Ibrahim

Summary: The previous Hermes project 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 European initiatives

10.2.1 Digital Europe

Participants: Cédric Tedeschi, Guillaume Pierre.

- **Title:** ACHIEVE: Advanced Cloud and High-performance computing Education for a Valiant Europe
- **Partner Institution(s):**
 - EIT Digital
 - EIT Digital Spain
 - Università Degli Studi di Trento
 - Middle East Technical University
 - Université de Rennes
 - Kungliga Tekniska Högskolan
 - Universitate Babeş 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 Türkiye

- Research Institute of Sweden
- Amazon Web Services
- **Date/Duration:** 2024-2028
- **Additional 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.

10.3 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

Additional 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 environnemental

Partner Institution(s):

- Inria Magellan team
- Inria Taran team

Date/Duration: 2024-2025

Additional 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, Ophélie Renaud.

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

Additional 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 is developed for the prospective profiling of data-oriented applications during their development.

ANR FACTO

Participants: Anne-Cécile Orgerie, François Lemercier, Aya Shehata.

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

Additional 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

Additional 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, Anudipa Mondal, Hayfa Tayeb.

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

Additional 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, CRISAL lab
- CNRS, I3S lab
- University of Toulouse, IRIT lab
- Inria, LIP lab
- IMT, LS2N lab
- IMT, SAMOVAR lab

Date/Duration: 2023-2030

Additional 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 Data quality assurance for a critical zone observatory in The 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

Additional 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, Ibtissem Oueslati.

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

Additional 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, Minh Thong Le Viet.

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

Additional 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, Ahmed Rjiba, Tiago Da Silva Barros.

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

Additional 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.

IPCEI-CIS DXP

Participants: Shadi Ibrahim, Cédric Tedeschi.

Title: Data Exchange**Partner Institution(s):**

- Inria
- Amadeus

Date/Duration: 2024-2029

Additional info/keywords: The aim of the IPCEI-CIS DXP project is to design and develop an open source management solution for a federated and distributed data exchange platform operating in an open, scalable, and massively distributed Cloud-Edge Continuum.

10.4 Regional initiatives

VideoImpact (Labex Cominlabs)

Participants: Anne-Cécile Orgerie, Robin Richard, Natacha Lapeyroux.

Title: **VideoImpact:** Model, Deploy, Orchestrate, and Optimize Cloud Applications and Infrastructure

Partner Institution(s):

- Inria
- CNRS, IRISA
- Université de Rennes, ARENES
- INSA de Rennes, IETR
- Université Catholique de l'Ouest de Nantes, ARENES
- IMT Atlantique, LEGO

Date/Duration: 2 years

Additional info/keywords: Recent studies forecast a global warming of 3.1°C in 2100 if the GHG emissions do not decrease. Hence, every part of our society must urgently aim sobriety, including the digital world, that is not intangible, contrary to popular belief. Video consumption takes a significant part among the emissions of the digital world and constitutes a representative example of unbounded and energy-consuming digital system. In that context, a crucial question to tackle is how to set limits to the deployment of a digital system, and for example to video delivery systems? This question is, by nature, lying at the crossroad of many fields (including human and social sciences). Interestingly, many initiatives have recently emerged at the regional level, e.g., the rapprochement between the GIS Marousin and video processing scientists of INSA and IRISA, and set interesting perspectives of wide collaborative user experiments. In that context, the VideoImpact project proposes to answer the following questions: In order to set a sobriety policy, what should we limit in priority? the number of hours spent by a user watching videos? The TV screen size? The video resolutions? The deployment of more efficient digital infrastructure? The VideoImpact project aims at developing i) an environmental footprint model for the video delivery chain to identify the clear levers to sobriety, ii) a solid network of industrial and academic partners of the Rennes' neighborhood around the goal of reducing the environmental impact of video consumption and iii) to launch a concrete experimentation in collaboration with Human and Social scientists. The conclusions will be used in the context of further collaborations with Human and Social Scientists to set real user experiments to assess the feasibility and acceptance of such levers.

11 Dissemination

11.1 Promoting scientific activities

11.1.1 Scientific events: organisation

General chair, scientific chair

- Guillaume Pierre was General co-Chair of the 13th IEEE International Conference on Cloud Engineering which took place in Rennes on September 23rd-26th 2025.
- Shadi Ibrahim was a member of the steering committee for the 9th edition of the **Workshop Performance and Scalability of Storage Systems(Per3S)**, which took place in Paris on May 23 2025 [31].
- Shadi Ibrahim is a member of the steering committee of the **IEEE International Conference on Cluster Computing**.

- 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 \(CHEOPS\)](#).
- Shadi Ibrahim is a member of the steering committee of the of the [International Conference on Scalable Scientific Data Managementt \(SSDBM\)](#).

Member of the organizing committees

- Amandine Seigneur was local arrangements co-chair of the 13th IEEE International Conference on Cloud Engineering which took place in Rennes on September 23rd-26th 2025.
- Anne-Cécile Orgerie, Amandine Seigneur and Matthieu Simonin were member of the organizing committee of the [GreenDays](#).
- Shadi Ibrahim was the Workshop and Tutorial Co-Chair of the 13th IEEE International Conference on Cloud Engineering which took place in Rennes on September 23rd-26th 2025.
- François Lemerrier was co-organizer of the Shadow Program committee of the [27èmes Rencontres Francophones sur les aspects Algorithmiques des Télécommunications](#).

11.1.2 Scientific events: selection

Chair of conference program committees

- Guillaume Pierre is program co-chair of the track on Scheduling, Resource Management, Cloud, Edge Computing, and Workflows at Euro-Par 2026.
- Anne-Cécile Orgerie was Program co-Chair of the 13th IEEE International Conference on Cloud Engineering which took place in Rennes on September 23rd-26th 2025.
- Anne-Cécile Orgerie was co-chair of the track on Sustainable IT Systems: abstraction, architecture, design and implementation at the IEEE International Symposium on Cluster, Cloud, and Internet Computing (CCGrid) 2025.

Member of the conference program committees

- Guillaume Pierre was a program committee member of the ACM/SPEC ICPE 2025, IEEE ICFEC 2025, IEEE ICDCS 2025 (poster track), ACM DEBS 2025, ACM HPDC 2025, MASCOTS 2025, DAIS 2025, IEEE ICFEC 2026, and ACM/SPEC ICPE 2026 conferences.
- Anne-Cécile Orgerie was a program committee member of the AlgoTel 2025 conference.
- Shadi Ibrahim was a program committee member of the ACM/IEEE SC 2025, IEEE Cluster 2025, ACM DEBS 2025, ACM/IEEE CCGrid 2025, ACM/IEEE UCC 2025, IEEE ISPDC 2025, Euro-Par 2025 conferences.
- Nikos Parlavantzas was a program committee member of the IEEE CLOUD 2025, JSSPP 2025, and VHPC 2025 conferences.
- François Lemerrier was a program committee member of the Cores 2025 conference.

11.1.3 Journal

Member of the editorial boards

- Anne-Cécile Orgerie was a member of the editorial board of the Special Issue on Sustainability and Computing of the Communications of the ACM published in 2025.
- Shadi Ibrahim is a member of the editorial board of IEEE Transactions on Cloud Computing.
- Shadi Ibrahim is an associate editor of IEEE Internet Computing Magazine.
- Shadi Ibrahim is an associate editor for High Performance Big Data Systems of Frontiers in High Performance Computing Journal.
- Shadi Ibrahim was a guest editor of IEEE Internet Computing Magazine: Special Issue on Serverless Computing (Nov/Dec 2024). The special issue was published in January 2025 [35].

11.1.4 Invited talks

- Guillaume Pierre gave a keynote on “Fog computing, from smart city services to natural environment monitoring” at the 43rd Brazilian Symposium on Computer Networks and Distributed Systems, May 22nd 2025, Natal, Brazil [43].
- Anne-Cécile Orgerie gave a keynote on “Carbon and water footprints of computing infrastructures: evaluations and tendencies” at the Infrastructure Workshop, Nantes, France, December 5, 2025.
- Anne-Cécile Orgerie gave a keynote on “Carbon Footprint Allocation Models in Distributed Systems” at UCC (IEEE/ACM International Conference on Utility and Cloud Computing), Nantes, France, December 4, 2025.
- Anne-Cécile Orgerie gave a keynote on “Greening ICT: challenges and research directions in the field of distributed systems towards sustainability” at the Green ICT and ICT for Green workshop co-located with Informatics Europe Summit (ECSS), Rennes, France, October 29, 2025.
- Anne-Cécile Orgerie gave a keynote on “Energy Consumption and Environmental Impact of Distributed Systems” at ISPDC (IEEE International Symposium on Parallel and Distributed Computing), Rennes, France, July 10, 2025.
- Shadi Ibrahim gave an invited talk on “Scalable and Efficient Big Data Processing in Clouds: Addressing Performance Variability” at the Fifth Workshop on Challenges and Opportunities of Efficient and Performant Storage Systems (CHEOPS’25), Rotterdam, Netherlands, March 31 2025

11.1.5 Leadership within the scientific community

- Anne-Cécile Orgerie is the director of the CNRS research and service group on ICT environmental impact (GDRS EcoInfo).
- Anne-Cécile Orgerie is the chief scientist for the Rennes site of Grid’5000 and co-responsible for the Rennes site of SLICES-FR.

11.1.6 Scientific expertise

- Anne-Cécile Orgerie was an HCERES expert for the evaluation of a laboratory in 2025.
- Anne-Cécile Orgerie was a reviewer for a research proposal for the FWF Austrian Science Fund in 2025.
- Anne-Cécile Orgerie was a reviewer for a research proposal for the Norwegian Velux Foundation in 2025.
- Shadi Ibrahim was a member of the [ACM Heidelberg Laureate Forum \(HLF\) 2025 Young Researcher Selection Committee](#).

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.
- Shadi Ibrahim co-organizes the SCI-Rennes seminar: a monthly series of scientific seminars for all staff at Inria research centre at Rennes University.
- Guillaume Pierre is a member of Inria Rennes center's "bureau du comité des projets"

11.2 Teaching - Supervision - Juries - Educational and pedagogical outreach

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: François Lemerrier, Système Cloud et Réseaux, L3 Informatique, Univ. Rennes.
- Bachelor: François Lemerrier, Mineure Cloud et Réseaux, L3 Informatique, Univ. Rennes.
- Bachelor: Anne-Cécile Orgerie, Impacts environnementaux du numérique, classe CPES, Lycée Chateaubriand, 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.
- Bachelor: Cédric Tedeschi, Systèmes, Cloud et Réseaux, L3 Informatique, Univ. Rennes.
- 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: François Lemerrier, Cloud, M1 RSSI, Univ. Rennes.
- Master: François Lemerrier, Network Administration and new technologies, M1 CNI, Univ. Rennes.
- Master: François Lemerrier, Advanced Network Infrastructure, M2 Cloud et Réseaux, Univ. Rennes.
- Master: Anne-Cécile Orgerie, Energy efficiency in distributed systems, TelecomSudParis, Palaiseau.
- Master: Anne-Cécile Orgerie, Consommation énergétique des ordinateurs, formation for computer science preparatory class teachers, CIRM, Marseille.
- 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: Guillaume Pierre, Distributed Systems, M1 CR&CNI, Univ. Rennes.
- Master: Guillaume Pierre, Systèmes répartis, ENSAI.
- Master: Guillaume Pierre, Services for Cloud technology, M1 CR&CNI, Univ. Rennes.
- Master: Guillaume Pierre, Advanced Cloud Infrastructures, M2 CR&CNI&SIF, Univ. Rennes.
- 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 Informatique, Univ. Rennes.
- Master: Cédric Tedeschi, Parallel Programming with Python, ENSAI

11.2.2 Supervision

- PhD defended: Khaled Aarsalane, “Scalable Data Stream Processing in Heterogeneous Environments” defended on December 15th 2025 at the Université de Rennes [36].
- PhD defended: Léo Cosseron, “Time-accurate network simulation interconnecting virtual machines with hardware virtualization towards stealth analysis” defended on December 8th 2025 at ENS Rennes [2].
- PhD in progress: Maxime Agusti, “Observation of baremetal co-location platforms, models and catalog proposal to reduce energy consumption”, started in 2021, supervised by Eddy Caron, Laurent Lefèvre and Anne-Cécile Orgerie.
- PhD in progress: Matteo Chancerel, “Optimization of a fully distributed Fog powered by renewable energy sources”, started in 2024, supervised by Anne-Cécile Orgerie.
- 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: Govind Kovilkkatt Panickerveetil, “Traitement de flux de données économe en énergie”, started in 2023, supervised by Guillaume Pierre and Romain Rouvov.
- PhD in progress: Haraesh Jayasethu Ramachandran, “Latency-driven container network optimization in edge industrial IoT,” started in 2025, supervised by François Lemercier and Guillaume Pierre.
- PhD in progress: Pablo Leboulanger, “Minimalist cloud, low on energy, hardware and software resources”, started in 2024, supervised by Anne-Cécile Orgerie.
- PhD in progress: Robin Richard, “Limits for the global video consumption”, started in 2025, supervised by Thomas Maugey and Anne-Cécile Orgerie.

- PhD in progress: Matthieu Silard, “Co-optimization of electrical and communication networks”, started in 2023, supervised by Anne-Cécile Orgerie, Nicolas Montavont and Georgios Papadopoulos.
- PhD in progress: Thomas Stavis, “Replay of environmental leverages in Cloud infrastructures”, started in 2024, supervised by Laurent Lefèvre and Anne-Cécile Orgerie.
- PhD in progress: Quentin Acher, “Management of mutable data over P2P storage”, started in 2023, supervised by Shadi Ibrahim and Claudia-Lavinia Ignat.
- 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: 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: Volodia Parol-Guarino, “Flexible resource allocation for FaaS applications in the fog”, started in October 2022, supervised by Nikolaos Parlavantzas.
- PhD in progress: Mohamed Cherif Zouaoui Latreche, “Balancing Performance and Sustainability for FaaS in the Fog”, started in October 2023, supervised by Nikolaos Parlavantzas and Hector Duran-Limon.
- PhD in progress: Ahmed Rjiba, “Decentralised, market-based application orchestration in Fog and IoT environments”, started in November 2025, supervised by Nikos Parlavantzas and Remous-Aris Koutsiamanis.
- PhD in progress: Stella Lafortune Tchoutcha Sandjong, “Energy-Efficient FaaS for Low-Power Edge and IoT Deployments”, started in December 2025, supervised by Barbe Thystere Mvondo Djob and Nikolaos Parlavantzas.
- PhD in progress: Mathieu Laurent on Efficient verification of asynchronous distributed systems, started in October 2023, supervised by Thierry Jéron (Devine) and Martin Quinson.
- PhD in progress: Minh Thong Le Viet on “Towards a programmable autonomic platform for decentralized learning”, started in October 2025, supervised by Cédric Tedeschi.

11.2.3 Juries

- Anne-Cécile Orgerie was chairperson at the PhD defense of Robin Boezennec (Université de Rennes), December 10th 2025.
- Anne-Cécile Orgerie was examiner at the PhD defense of Roblex Nana Tchakouté (Mines de Paris), December 5th 2025.
- Anne-Cécile Orgerie was examiner at the PhD defense of Killian Castillon du Perron (Université de Côte d’Azur), November 21st 2025.
- Anne-Cécile Orgerie was examiner at the PhD defense of Tiago da Silva Barros (Université de Côte d’Azur), November 3rd 2025.
- Anne-Cécile Orgerie was chairperson at the PhD defense of Meriem Ghali (ENS de Lyon), October 16th 2025.
- Anne-Cécile Orgerie was examiner at the PhD defense of Meven Mognol (Université de Rennes), July 2nd 2025.
- Anne-Cécile Orgerie was opponent and examiner at the PhD defense of Antoine Omond (Arctic University of Norway, Tromsø and Université de Nantes), May 23rd 2025.

- Shadi Ibrahim was examinator at the PhD defense of Mazen Ezzeddine (Université Côte d’Azur), September 16th 2025.
- Martin Quinson was chairperson at the PhD defense of Khaled Arsalane (Université de Rennes), December 15th 2025.
- Guillaume Pierre was reviewer at the PhD defense of Aymeric Agon (Sorbonne Université), September 15th 2025.

11.3 Popularization

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

- Guillaume Pierre gave an interview to Le Nouvel Observateur which was mentioned in article “[Que dit la panne d’AWS, le cloud d’Amazon, de notre dépendance aux géants du numérique ?](#)” on October 21st 2025.
- Guillaume Pierre gave an interview to Brief.Science which was mentioned in article “[Le cloud, moteur invisible de nos vies connectées](#)” on October 31st 2025.
- Anne-Cécile Orgerie gave an interview to Les Echos which was mentioned in the article “[Environnement : faut-il moraliser l’usage du numérique](#)” on February 10th 2025.
- Anne-Cécile Orgerie gave an interview to Mediapart which was mentioned in the article “[Les déchets du numérique, un fléau pour la ville canadienne de Rouyn-Noranda](#)” on February 10th 2025.
- Anne-Cécile Orgerie gave an interview to Le Parisien which was mentioned in the article “[Dans la course à l’IA, des machines françaises parmi les plus puissantes au monde](#)” on February 23rd 2025.

11.3.2 Participation in Live events

- Anne-Cécile Orgerie proposed a session at the [popularization day](#) on “les décodeuses” with 550 students from 15 high schools.

11.3.3 Others science outreach relevant activities

- François Lemerrier proposed “Clusterize” in collaboration with artist Loïg Nguyen in the context of the IRISA program “[1 Artist 1 Scientist](#).”

12 Scientific production

12.1 Major publications

- [1] 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. 15).
- [2] L. Cosseron. ‘Time-accurate network simulation interconnecting virtual machines with hardware virtualization towards stealth analysis’. Université de Rennes, 8th Dec. 2025. URL: <https://theses.hal.science/tel-05446019> (cit. on pp. 22, 38).

12.2 Publications of the year

International journals

- [3] H. Casanova, A. Giersch, A. Legrand, M. Quinson and F. Suter. ‘Lowering entry barriers to developing custom simulators of distributed applications and platforms with SimGrid’. In: *Parallel Computing* 123 (Mar. 2025), p. 103125. DOI: [10.1016/j.parco.2025.103125](https://doi.org/10.1016/j.parco.2025.103125). URL: <https://hal.science/hal-04909441> (cit. on p. 21).

- [4] A. Kazem, G. Pierre and L. Longuevergne. ‘Enhancing Environmental Observatories with Fog Computing’. In: *Frontiers in Environmental Science* 13 (4th Apr. 2025), pp. 1–15. DOI: [10.3389/fenvs.2025.1568016](https://doi.org/10.3389/fenvs.2025.1568016). URL: <https://inria.hal.science/hal-05007858> (cit. on p. 15).
- [5] S. Lambert, V. Ostapenco, L. Lefèvre, E. Caron, A.-C. Orgerie, B. Fichel and R. Grivel. ‘Revisiting virtual machine consolidation to save resources and energy in heterogeneous production cloud infrastructures’. In: *International Journal of High Performance Computing Applications* (18th Dec. 2025). DOI: [10.1177/10943420251399694](https://doi.org/10.1177/10943420251399694). URL: <https://hal.science/hal-05457159> (cit. on p. 20).
- [6] J. C. Sánchez Huertas, G. Vallero, L. Nuaymi and A.-C. Orgerie. ‘Review and classification of the use of SMs for energy saving in next-generation radio access networks’. In: *Computer Networks* 273 (Dec. 2025), p. 111745. DOI: [10.1016/j.comnet.2025.111745](https://doi.org/10.1016/j.comnet.2025.111745). URL: <https://hal.science/hal-05310147> (cit. on p. 20).
- [7] A. Voyez, T. Allard, G. Avoine, P. Cauchois, E. Fromont and M. Simonin. ‘The Privacy Cost of Fine-Grained Electrical Consumption Data’. In: *Scientific Reports* 15.17391 (19th May 2025), pp. 1–8. DOI: [10.1038/s41598-024-78285-7](https://doi.org/10.1038/s41598-024-78285-7). URL: <https://hal.science/hal-03833605> (cit. on p. 19).
- [8] H. Zhang, S. Wu, H. Fan, Z. Huang, W. Xue, C. Yu, S. Ibrahim and H. Jin. ‘KubeSPT: Stateful Pod Teleportation for Service Resilience With Live Migration’. In: *IEEE Transactions on Services Computing* 18.3 (May 2025), pp. 1500–1514. DOI: [10.1109/TSC.2025.3564888](https://doi.org/10.1109/TSC.2025.3564888). URL: <https://hal.science/hal-05464256> (cit. on p. 14).

National journals

- [9] X. Arnaud de Sartre, W. Frederic, C. Bouneau, F. Bouchet, G. Debizet, L. Fontenelle (de), L. Desvallées, M. Forget, N. Rodet-Kroichvili, S. Mathy, A.-C. Orgerie, D. Pecaud, J. Py, O. Vidal and S. Velut. ‘Pour une science de la Transition Énergétique au Service de la Société: Rapport du groupe interdisciplinaire ARPEGES CNRS pour la Transition énergétique’. In: *Natures Sciences Sociétés* 33.3 (2025), pp. 1–10. URL: <https://hal.science/hal-04977245> (cit. on p. 23).

Invited conferences

- [10] B. Jonglez, M. Simonin, J. Philippe and S. M. Kaddour. ‘Multi-provider capabilities in EnOSlib: driving distributed system experiments on the edge-to-cloud continuum’. In: *Springer LNCS-IFIP, Lecture Notes in Computer Science (LNCS)*. DAIS 2025: 25th International Conference on Distributed Applications and Interoperable Systems. Vol. 15730. Lille, France: Springer, 2025, pp. 25–42. DOI: [10.1007/978-3-031-95728-4_2](https://doi.org/10.1007/978-3-031-95728-4_2). URL: <https://inria.hal.science/hal-05052776> (cit. on p. 23).

International peer-reviewed conferences

- [11] M. Agusti, E. Caron, B. Fichel, L. Lefèvre, O. Nicol and A.-C. Orgerie. ‘PPEM-BM: Portable Power Estimation Methodology for Bare Metal Servers’. In: *ICPADS 2025 - 31st IEEE International Conference on Parallel and Distributed Systems*. ICPADS 2025 - 31st IEEE International Conference on Parallel and Distributed Systems. Hefei, China: IEEE, 2025, pp. 1–8. URL: <https://hal.science/hal-05365793> (cit. on p. 18).
- [12] K. Arsalane and G. Pierre. ‘Data Stream Processing Effectiveness in Heterogeneous Computing Environments’. In: *Proceedings of the 41st ACM/SIGAPP Symposium On Applied Computing*. SAC 2026 - 41st ACM/SIGAPP Symposium On Applied Computing. Tesseloniki, Greece: ACM, 2026. URL: <https://hal.science/hal-05390665> (cit. on p. 18).
- [13] M. Chancerel and A.-C. Orgerie. ‘Life-cycle carbon assessment of by-design intermittent Fog services fully powered by renewable energy sources’. In: *UCC 2025 - 18th IEEE/ACM International Conference on Utility and Cloud Computing*. UCC 2025 - 18th IEEE/ACM International Conference on Utility and Cloud Computing. Nantes, France: IEEE, 2025, pp. 1–10. URL: <https://hal.science/hal-05365825> (cit. on p. 19).

- [14] L. Cosseron, L. Rilling and M. Quinson. ‘Mitigation of the impact of Virtual Machine Introspection Pauses on Multi-core Virtual Machines’. In: *Lecture Notes in Computer Science (LNCS)*. HS3 2025 - 1st Workshop on Hardware-Supported Software Security. Toulouse, France, 2025, pp. 1–20. URL: <https://inria.hal.science/hal-05253492> (cit. on p. 22).
- [15] C. Coti and M. Quinson. ‘Simulating a communication middleware for one-sided and collective operations’. In: ISPCS 2025 - International IEEE Symposium on Precision Clock Synchronization for Measurement, Control, & Communication. Ottawa, Canada, 2025, pp. 1–6. URL: <https://inria.hal.science/hal-05267994> (cit. on p. 21).
- [16] H. Fan, Z. Huang, S. Ibrahim, L. Gu and S. Wu. ‘EDDE: Container Deployment Framework Beyond the Cloud’. In: SC 2025 - International Conference for High Performance Computing, Networking, Storage and Analysis. St. Louis MO USA, United States: ACM, 2025, pp. 570–585. DOI: [10.1145/3712285.3759854](https://doi.org/10.1145/3712285.3759854). URL: <https://hal.science/hal-05464258> (cit. on p. 14).
- [17] G. Guennebaud, A.-L. Ligozat, A.-C. Orgerie and M. Simonin. ‘Evaluating and Reporting the Carbon Footprint of Shared Computing Platforms: Choices and Limits’. In: ISPD 2025 - 24th IEEE International Symposium on Parallel and Distributed Computing. Rennes, France: IEEE, 2025, pp. 1–7. URL: <https://hal.science/hal-05195576> (cit. on p. 19).
- [18] S. Ibrahim and J. Darrous. ‘Erasure Coding Aware Block Placement for Data-Intensive Applications’. In: CHEOPS 2025 - 5th Workshop on Challenges and Opportunities of Efficient and Performant Storage Systems. Rotterdam Netherlands, Netherlands: ACM, 30th Mar. 2025, pp. 15–22. DOI: [10.1145/3719330.3721229](https://doi.org/10.1145/3719330.3721229). URL: <https://hal.science/hal-05016741> (cit. on p. 13).
- [19] P. Jacquet, M. Agusti, E. Caron, C. Coti, M. Dias de Assuncao, L. Lefèvre and A.-C. Orgerie. ‘Untangling GPU Power Consumption: Job-Level Inference in Cloud Shared Settings’. In: EUROSYS 2026 - European Conference on Computer Systems. Edinbourg, Ecosse, United Kingdom, 2026. DOI: [10.1145/3767295.3769333](https://doi.org/10.1145/3767295.3769333). URL: <https://hal.science/hal-05291033> (cit. on p. 19).
- [20] H. Jr, A. Kazem, M. Nicolas, F. Lemerrier, G. Pierre and C. Andermann. ‘LivingBench: an IoT/Edge Platform Benchmark Based on an Environmental Observation Use-Case’. In: ICIN 2026 - 29th Conference on Innovation in Clouds, Internet and Networks. Athens, Greece, 2026. URL: <https://inria.hal.science/hal-05454254> (cit. on p. 15).
- [21] A. Kazem, G. Pierre and L. Longuevergne. ‘Demo: Energy-aware Dynamic Dimensioning of IoT Edge Clusters for Natural Environment Observation’. In: IC2E 2025 - 13th IEEE International Conference on Cloud Engineering. Rennes, France: IEEE, 13th Sept. 2025, pp. 1–3. URL: <https://hal.science/hal-05218930> (cit. on p. 15).
- [22] M. Laurent, T. Jéron and M. Quinson. ‘Towards Efficient Verification of Parallel Applications with Mc SimGrid’. In: *DisCoTec 2025 - 20th International Federated Conference on Distributed Computing Techniques*. FORTE 2025 - 45th International Conference on Formal Techniques for Distributed Objects, Components, and Systems. Lecture Notes in Computer Science. Lille, France: Springer, 2025, pp. 1–21. URL: <https://hal.science/hal-05042973> (cit. on p. 22).
- [23] P. Leboulanger and A.-C. Orgerie. ‘Power limits in data centers: what can we expect from improving energy efficiency and refreshing servers?’ In: *IC2E 2025 - 13th IEEE International Conference on Cloud Engineering*. IC2E 2025 - 13th IEEE International Conference on Cloud Engineering. Rennes, France, 2025, pp. 1–12. URL: <https://hal.science/hal-05215097> (cit. on p. 20).
- [24] F. Lemerrier and A.-C. Orgerie. ‘Towards an energy-efficient Wi-Fi: An experimental study on recent standards power consumption’. In: WCNC 2025 - IEEE Wireless Communications and Networking Conference. Milan, Italy: IEEE, 2025, pp. 01–06. DOI: [10.1109/wcnc61545.2025.10978325](https://doi.org/10.1109/wcnc61545.2025.10978325). URL: <https://hal.science/hal-04928381> (cit. on p. 23).
- [25] D. Mvondo, B. Djomgwe Teabe and N. Parlavantzas. ‘Efficient Memory Usage For Edge FaaS Platforms’. In: PerCom 2026 - 24th IEEE International Conference on Pervasive Computing and Communications. Pise, Italy, 2026. URL: <https://inria.hal.science/hal-05465369> (cit. on p. 16).

- [26] V. Parol-Guarino and N. Parlavantzas. ‘Auction-based Placement of Function Chains in the Fog at Scale’. In: *Europar 2025 - 31st International European Conference on Parallel and Distributed Computing*. Dresden, Germany, 2025, pp. 1–14. URL: <https://hal.science/hal-05121317> (cit. on p. 17).
- [27] O. Renaud, N. Gac, F. Orieux and C. Dumez-Viou. ‘SimSDP, a Rapid Prototyping tool for Radio Astronomy: From NenuFAR Experiments to SKAO-Scale Simulation’. In: *SiPS 2025 - 38th IEEE Workshop on Signal Processing Systems*. Hong Kong, China: IEEE, 2025, pp. 1–5. URL: <https://hal.science/hal-05260888> (cit. on p. 23).
- [28] M. Silard, A.-C. Orgerie, N. Montavont and G. Papadopoulos. ‘Dynamic Incentive Signaling through Advanced Metering Infrastructure Networks’. In: *SmartGridComm 2025 - IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids*. Toronto, Canada, 2025, pp. 1–5. URL: <https://hal.science/hal-05310150> (cit. on p. 20).
- [29] M. Silard, G. Papadopoulos, A.-C. Orgerie and N. Montavont. ‘Demo: A Visualization Platform for Smart Grid Network’. In: *CSCN 2025 - IEEE Conference on Standards for Communications and Networking*. Bologna, Italy: IEEE, 2025, pp. 1–3. URL: <https://hal.science/hal-05215110> (cit. on p. 20).

National peer-reviewed Conferences

- [30] T. Maugey, D. Menard, A.-C. Orgerie and R. Richard. ‘Diffusion vidéo : dépasser l’objectif d’efficacité et viser une sobriété choisie’. In: *GRETSI - XXXe Colloque Francophone de Traitement du Signal et des Images*. GRETSI 2025 - XXXe Colloque Francophone de Traitement du Signal et des Images. Strasbourg, France, 2025, pp. 1–4. URL: <https://hal.science/hal-05223926> (cit. on p. 21).

Conferences without proceedings

- [31] J.-T. Acquaviva, J. Boukhobza, P. Deniel, S. Ibrahim, P. Raipin-Parvédy and F. Trahay. ‘Minutes from the 9th edition of the Performance and Scalability of Storage Systems workshop (Per3S), 23rd May 2025, "Maison des Mines et des Ponts", Paris’. In: 9th edition of the workshop Performance and Scalability of Storage Systems (Per3S). Paris, France, 2025, pp. 1–2. URL: <https://hal.science/hal-05116147> (cit. on p. 34).
- [32] N. Gac and O. Renaud. ‘Dark-era : Dataflow Algorithm architecture co-design of SKA pipeline for Exascale Radio Astronomy’. In: *SaLF 2025 - 11th Annual conference on Science At Low Frequencies*. Orléans, France, 2025. URL: <https://hal.science/hal-05430702> (cit. on p. 23).
- [33] O. Renaud, S. Wang, N. Gac and F. Orieux. ‘Extension du prototypage rapide pour les pipelines d’imagerie radioastronomique : Simulation avancée multi-nœud HPC’. In: *GRETSI 2025 - 30e Colloque sur le traitement du signal et des images*. Strasbourg, France, 2025, pp. 1–4. URL: <https://hal.science/hal-05216024> (cit. on p. 23).

Scientific book chapters

- [34] W. E. Gribga, A. Blavette and A.-C. Orgerie. ‘Managing Distributed Cloud Infrastructures for Sustainability’. In: *Scheduling Variable Capacity Resources for Sustainability*. CRC Press, 2026, pp. 83–124. URL: <https://hal.science/hal-05399000> (cit. on p. 19).

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

- [35] *Serverless Computing* 28.6 (31st Jan. 2025). DOI: [10.1109/MIC.2024.3524507](https://doi.org/10.1109/MIC.2024.3524507). URL: <https://hal.science/hal-05464253> (cit. on p. 36).

Doctoral dissertations and habilitation theses

- [36] K. Arsalane. ‘Scalable Data Stream Processing in Heterogeneous Environments’. Université de rennes, 15th Dec. 2025. URL: <https://inria.hal.science/tel-05433201> (cit. on pp. 18, 38).

Reports & preprints

- [37] S. Bouveret, A. Bugeau, E. Frenoux, J. Lefevre, L. Lefèvre, A.-L. Ligozat, P. Marquet, A.-C. Orgerie and D. Trystram. *Quiz sur les impacts environnementaux du numérique*. EcoInfo, Feb. 2025, pp. 1–5. URL: <https://hal.science/hal-04960328> (cit. on p. 24).
- [38] K. Dassas, C. Bonamy, B. Bzeznik, R. David, E. Frenoux, G. Guennebaud, A.-L. Ligozat, D. Mallarino and A.-C. Orgerie. *Estimer l'impact carbone des activités numériques de l'Observatoire de Paris*. EcoInfo, Feb. 2025, pp. 1–47. URL: <https://hal.science/hal-05068666> (cit. on p. 24).
- [39] M. Simonin. *Bilan Carbone du centre de Inria/IRISA de Rennes pour l'année 2024. Les déplacements professionnels: dynamiques et disparités entre les équipes du centre Inria de Rennes.: Avec une étude sur les émissions liées aux déplacements professionnels : dynamiques et disparités entre équipes du centre Inria de Rennes*. Inria Centre at the University of Rennes, France; Irisa Institut de Recherche en Informatique et Systèmes Aléatoires, 10th Sept. 2025, pp. 1–13. URL: <https://inria.hal.science/hal-05248084> (cit. on p. 10).

Other scientific publications

- [40] B. Knowles, V. L. Hanson, C. Becker, M. Berners-Lee, A. A. Chien, B. Combemale, V. Coroamă, K. de Bosschere, Y. Ding, A. Friday, B. Gamazaychikov, L. Hardman, S. Hinterholzer, M. Höjer, L. Kaack, L. Kuijter, A.-L. Ligozat, J. T. Muehlberg, Y. Nah, T. Olsson, A.-C. Orgerie, D. Pargman, B. Penzenstadler, T. Romanoff, E. Strubell, C. Venters and J. Zhao. *Climate Change: What is Computing's Responsibility?* 2025. DOI: [10.4230/DagMan.11.1.1](https://doi.org/10.4230/DagMan.11.1.1). URL: <https://hal.science/hal-05369257> (cit. on p. 24).
- [41] S. Lambert, V. Ostapenco, L. Lefèvre, E. Caron, A.-C. Orgerie, B. Fichel and R. Grivel. 'Revisiting virtual machine consolidation to save resources and energy in heterogeneous production cloud infrastructures'. In: *ICT4S 2025 - International Conference on Information and Communications Technology for Sustainability*. Dublin, Ireland, 9th June 2025, pp. 1–1. URL: <https://inria.hal.science/hal-05100569> (cit. on p. 20).
- [42] A.-C. Orgerie. *Code de conduite - Journées scientifiques*. Rennes, France, 2025. URL: <https://inria.hal.science/hal-05041948> (cit. on p. 10).
- [43] G. Pierre. *Fog computing, from smart city services to natural environment monitoring*. Natal, Brazil, 19th May 2025. URL: <https://inria.hal.science/hal-05084864> (cit. on pp. 15, 36).
- [44] O. Renaud, N. Gac, F. Orioux and C. Dumez-Viou. 'SimSDP, a Rapid Prototyping tool for Radio Astronomy: From NenuFAR Experiments to SKAO-Scale Simulation'. In: *SaLF 2025 - 11th Annual conference on Science At Low Frequencies*. Orleans (France), France, 8th Dec. 2025, pp. 1–1. URL: <https://hal.science/hal-05406207> (cit. on p. 23).

Scientific popularization

- [45] A.-C. Orgerie. 'Consommation électrique et efficacité énergétique des infrastructures de calcul'. In: *Le calcul à découvert*. CNRS éditions, Jan. 2025, pp. 248–249. URL: <https://hal.science/hal-04928468> (cit. on p. 20).

Software

- [46] [SW] H. Casanova, A. Degomme, A. Giersch, A. Legrand, M. Quinson and F. Suter, *SimGrid* version 4.0, 18th Apr. 2025. LIC: GNU Lesser General Public License v2.1 or later. HAL: [hal-05040017](https://hal.science/hal-05040017), URL: <https://inria.hal.science/hal-05040017>, vcs: <https://framagit.org/simgrid/simgrid/>, SWHID: [sw:1:rev:3714440fcff4a244f3d51dc0495bab153a5c298c](https://sw.hicet.fr/rev/3714440fcff4a244f3d51dc0495bab153a5c298c); origin=<https://framagit.org/simgrid/simgrid.git>; visit=[sw:1:snp:f2414fd9d467641d7a695ed64e3d2fd779087095](https://sw.hicet.fr/snp/f2414fd9d467641d7a695ed64e3d2fd779087095)) (cit. on p. 21).

12.3 Cited publications

- [47] 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. 7).
- [48] J. Brodtkin. *Hurricane Sandy takes data centers offline with flooding, power outages*. Ars Technica. <https://arstechnica.com/information-technology/2012/10/hurricane-sandy-take-s-data-centers-offline-with-flooding-power-outages/>. Oct. 2012 (cit. on p. 8).
- [49] R.-A. Cherrueau, M. Delavergne, A. van Kempen, A. Lebre, D. Pertin, J. Rojas Balderrama, A. Simonet and M. Simonin. ‘EnosLib: A Library for Experiment-Driven Research in Distributed Computing’. In: *IEEE Transactions on Parallel and Distributed Systems* 33.6 (June 2022), pp. 1464–1477. DOI: [10.1109/TPDS.2021.3111159](https://doi.org/10.1109/TPDS.2021.3111159). URL: <https://inria.hal.science/hal-03324177> (cit. on p. 23).
- [50] 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. 8).
- [51] 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. 7).
- [52] Flexera. *Flexera 2020 State of the Cloud Report*. <https://info.flexera.com/SLO-CM-REPORT-State-of-the-Cloud-2020>. 2019 (cit. on p. 8).
- [53] N. Galov. *Cloud Adoption Statistics for 2021*. Hosting Tribunal. <https://hostingtribunal.com/blog/cloud-adoption-statistics/>. Jan. 2021 (cit. on p. 7).
- [54] 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. 8).
- [55] 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. 8).
- [56] 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. 7).
- [57] OVHcloud. *Fire at our Strasbourg Site*. <https://us.ovhcloud.com/press/press-releases/2021/fire-our-strasbourg-site>. Apr. 2021 (cit. on p. 8).
- [58] A. V. Papadopoulos, L. Versluis, A. Bauer, N. Herbst, J. Von Kistowski, A. Ali-eldin, C. Abad, J. N. Amaral, P. Tüma and A. Iosup. ‘Methodological Principles for Reproducible Performance Evaluation in Cloud Computing’. In: *IEEE Transactions on Software Engineering* (2019) (cit. on p. 8).