

2025 Activity Report

RESEARCH CENTRE: Inria Centre at Rennes University


IN PARTNERSHIP WITH: Ecole Nationale Supérieure Mines-Télécom Atlantique
Bretagne Pays de la Loire, Nantes Université, Orange SA

Project-Team

STACK

Software Stack for Massively Geo-Distributed
Infrastructures

In collaboration with Laboratoire des Sciences du numérique de Nantes



Project-Team STACK

Creation of the Project-Team: 2019 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

- A1.1.8. – Security of architectures
- A1.1.10. – Reconfigurable architectures
- A1.1.13. – Virtualization
- A1.2.1. – Dynamic reconfiguration
- A1.2.2. – Supervision
- A1.2.4. – QoS, performance evaluation
- A1.2.8. – Network security
- A1.3.4. – Peer to peer
- A1.3.5. – Cloud
- A1.3.6. – Fog, Edge
- A1.5.1. – Systems of systems
- A1.6. – Green Computing
- A2.1.7. – Distributed programming
- A2.1.10. – Domain-specific languages
- A2.5.2. – Component-based Design
- A2.6. – Infrastructure software
- A2.6.1. – Operating systems
- A2.6.2. – Middleware
- A2.6.3. – Virtual machines
- A2.6.4. – Ressource management
- A3.1.2. – Data management, quering and storage
- A3.1.3. – Distributed data
- A3.1.8. – Big data (production, storage, transfer)
- A4.1. – Threat analysis
- A4.4. – Security of equipment and software
- A4.9. – Security supervision

Other research topics and application domains

- B2. – Digital health
- B4. – Energy
- B4.5.1. – Green computing
- B5.1. – Factory of the future
- B6.3. – Network functions
- B6.4. – Internet of things
- B6.5. – Information systems
- B7. – Transport and logistics
- B8. – Smart Cities and Territories

Contents

Project-Team STACK	1
1 Team members, visitors, external collaborators	5
2 Overall objectives	7
2.1 STACK in a Nutshell	7
2.2 Context & Objectives	7
2.3 Scientific Foundations	8
3 Research program	9
3.1 Overview	9
4 Application domains	11
4.1 Overview	11
4.2 Industrial Internet	12
4.3 Internet of Skills	12
4.4 e-Health	12
4.5 Network Virtualization and Mobile Edge Services	12
4.6 Urgent Computing	12
5 Social and environmental responsibility	13
5.1 Footprint of research activities	13
5.2 Impact of research results	13
6 Highlights of the year	13
6.1 Awards	14
7 Latest software developments, platforms, open data	14
7.1 Latest software developments	14
7.1.1 ENOS	14
7.1.2 EnOSlib	14
7.1.3 Concerto	15
7.1.4 StreamBed	15
7.1.5 Cheops	16
7.1.6 Edge-to-cloud video processing	16
7.2 New platforms	17
7.2.1 Grid'5000	17
7.2.2 PiSeDuCe	17
7.2.3 SAMURAI	17
7.2.4 SLICES-FR/SLICES	17
8 New results	18
8.1 Resource Management	18
8.2 Programming Support	20
8.3 Energy-aware computing	21
8.4 Security and Privacy	21
9 Bilateral contracts and grants with industry	22
9.1 Bilateral contracts with industry	22
9.2 Bilateral grants with industry	23

10 Partnerships and cooperations	24
10.1 International initiatives	24
10.1.1 Inria associate team not involved in an IIL or an international program	24
10.2 International research visitors	24
10.2.1 Visits of international scientists	24
10.3 European initiatives	25
10.3.1 Horizon Europe	25
10.3.2 Other european programs/initiatives	25
10.4 National initiatives	26
10.4.1 ANR	26
10.4.2 PIA 4	29
10.4.3 CPER	30
10.4.4 Local and regional projects	30
10.4.5 Inria Challenges	30
10.4.6 Inria transfer projects (ADT)	31
10.5 Regional initiatives	31
10.5.1 A Cloud-Edge-IoT continuum site for SLICES-FR	31
11 Dissemination	31
11.1 Promoting scientific activities	31
11.1.1 Scientific events: organisation	31
11.1.2 Scientific events: selection	32
11.1.3 Journal	32
11.1.4 Invited talks	33
11.1.5 Scientific expertise	33
11.1.6 Research administration	33
11.2 Teaching - Supervision - Juries - Educational and pedagogical outreach	34
11.2.1 Supervision	34
11.2.2 Juries	36
11.3 Popularization	36
11.3.1 Participation in Live events	36
12 Scientific production	37
12.1 Major publications	37
12.2 Publications of the year	37
12.3 Cited publications	41

1 Team members, visitors, external collaborators

Research Scientists

- Adrien Lebre [Team leader, INRIA, Professor Detachement, until Apr 2025, HDR]
- Thierry Coupaye [Orange SA, Industrial member, HDR]
- Thomas Hassan [Orange SA, Industrial member]
- Daniel Balouek [INRIA, ISFP]
- Sébastien Bolle [Orange SA, Industrial member]
- Abdelhadi Chari [Orange SA, Industrial member]
- Philippe Raipin Parved [Orange SA, Industrial member]

Faculty Members

- Hélène Coullon [IMT ATLANTIQUE, Associate Professor, HDR]
- Carlos Gonzalez [IMT ATLANTIQUE, Associate Professor]
- Remous Koutsiamanis [IMT ATLANTIQUE, Associate Professor]
- Thomas Ledoux [IMT ATLANTIQUE, Professor, HDR]
- Jean-Marc Menaud [IMT ATLANTIQUE, Professor, HDR]
- Jacques Noyé [IMT ATLANTIQUE, Associate Professor]
- Kandaraj Piamrat Lerebours [UNIV NANTES, Associate Professor, HDR]
- Guillaume Rosinosky [IMT ATLANTIQUE, Associate Professor]
- Mario Südholt [IMT ATLANTIQUE, Professor, HDR]

Post-Doctoral Fellows

- Sidna Jeddou [IMT ATLANTIQUE, Post-Doctoral Fellow, from Sep 2025]
- Eloi Perdereau [IMT ATLANTIQUE, Post-Doctoral Fellow]

PhD Students

- Tengfei An [INRIA, until Sep 2025]
- Severin Bradley Anzie [IMT ATLANTIQUE, from May 2025]
- Simon Artus [ORANGE, CIFRE, from Mar 2025]
- Lucien Astie [IMT ATLANTIQUE, from Oct 2025]
- Hiba Awad [SMILE, CIFRE, until Feb 2025]
- Aymene Boucha [IMT ATLANTIQUE, from Sep 2025]
- Samia Boutalbi [Ericsson, CIFRE]
- Christophe Dion [ORANGE LABS]
- Celeste Precil Guimapi Guefack [IMT ATLANTIQUE]

- Mouheb Jemai [CEA, from Apr 2025]
- Housseem Jmal [UNIV NANTES]
- Mohammed Abdrrahim Lahmar [IMT ATLANTIQUE]
- Wedwang Romial Menra [INRIA]
- Martin Molli [IMT ATLANTIQUE]
- Duc Thinh Ngo [ORANGE, CIFRE]
- Lomig Piette [INRIA, from Oct 2025]
- Gaëtan Plisson [IMT ATLANTIQUE, from Oct 2025]
- Olivia Proust [IMT Atlantique, from Sep 2025]
- Nathan Rabier [INRIA, from Nov 2025]
- Irina Samus [IMT ATLANTIQUE, from Oct 2025]
- Abdou Seck [IMT ATLANTIQUE, until Jun 2025]
- Cherif Si Mohammed [ADEME]
- Lylian Siffre [Kapela, CIFRE]

Technical Staff

- Alexis Bitailou [IMT ATLANTIQUE, Engineer]
- Lucas Gazeau [IMT ATLANTIQUE, Engineer, from Mar 2025]
- Baptiste Jonglez [INRIA, Engineer]
- Sidi Mohammed Kaddour [INRIA, Engineer]
- Anas Mokhtari [IMT ATLANTIQUE, Engineer, until Jan 2025]
- Matthieu Rakotojaona Rainimangavelo [IMT ATLANTIQUE, Engineer, from Feb 2025 until Jun 2025]

Interns and Apprentices

- Hervé Mbailassem [IMT ATLANTIQUE, Intern, from Mar 2025 until Aug 2025]
- Sambo Mohamadou Bachirou [IMT ATLANTIQUE, Intern, from Mar 2025 until Aug 2025]
- Martin Plaud [INRIA, Apprentice, from Sep 2025]
- Gaëtan Plisson [IMT ATLANTIQUE, Intern, from Feb 2025 until Aug 2025]
- Nathan Rabier [IMT ATLANTIQUE, Intern, from Apr 2025 until Sep 2025]
- Laid Rahmoune [INRIA, Intern, from May 2025 until Jul 2025]
- Franc Abel Zogning Tedongmouo [IMT ATLANTIQUE, Intern, from Mar 2025 until Aug 2025]

Administrative Assistant

- Anne-Claire Binetruy [INRIA]

Visiting Scientists

- Lorenzo Carnevale [University of Messina]
- Maurice Djibril Faye [The Cheikh Hamidou Kane Digital University]

2 Overall objectives

2.1 STACK in a Nutshell

The STACK team addresses challenges related to the management and advanced usages of the Cloud to IoT continuum (infrastructures on the Cloud, Fog, Edge, and IoT). More specifically, the team is interested in delivering appropriate system abstractions to operate and use massively geo-distributed infrastructures, from the lowest to the highest levels of abstraction (i.e. system to application development, respectively), and addressing crosscutting dimensions such as energy or security. These infrastructures are critical for the emergence of new kinds of applications related to the digitalization of the industry and the public sector (a.k.a. the Industrial Internet, smart cities, e-medicine, etc.).

2.2 Context & Objectives

Initially proposed to interconnect computers worldwide, the Internet has significantly evolved to become in two decades a key element in almost all our activities. This (r)evolution mainly relies on the progress that has been achieved in the computation and communication fields which in turn has led to the well-known and widely spread Cloud Computing paradigm. Nowadays, most Internet exchanges occur between endpoints ranging from small-scale devices, such as smart-phones, to large-scale facilities, *i.e.*, cloud computing platforms, in charge of hosting modern information systems.

With the emergence of the Internet of Things (IoT), stakeholders expect a new revolution that will push, once again, the limits of the Internet, in particular by favouring the convergence between physical and virtual worlds into an *augmented world* or *cyber-physical world*. This convergence is about to be made possible thanks to the development of minimalist sensors as well as complex industrial physical machines that can be connected to the Internet through edge computing infrastructures. Edge computing is an extension of the cloud computing model that consists in deploying a federation (or cooperation) of smaller data centers at the edge of the network, thus closer to sensors, devices, machines, and end-users that produce and consume data [96, 94]. This new kind of digital infrastructure, which covers resources from the “center” to the extreme edge of the network, is expected to improve almost all aspects of daily life and the decision processes in various domains such as industry, transportation, health, training and education. The corresponding applications target the control and optimization of the business processes of most companies thanks to the intensive use of ICT systems and real-time data collected by geo-distributed physical devices (video, sensors, ...).

Among the obstacles to this new generation of Internet services is the development of a convenient and powerful software stack, *i.e.*, a set of system mechanisms and software abstractions capable of operating and exposing a significant number of diverse computational resources in a unified, efficient and sustainable way.

In other words, this framework should allow operators, and devops, to manage the life-cycle of both the digital infrastructure and the applications deployed on top of this infrastructure, throughout the **cloud to IoT continuum**. These include operations such as the initial configuration but also all the reconfigurations that can be required in response to particular events (maintenance operations, equipment failures, application load variation, user mobility, energy shortage, etc.).

The existing software stacks that have been proposed to manage Cloud Computing platforms are not appropriate for handling the specifics of the next generation of digital infrastructure (in terms of scale, heterogeneity, dynamicity, security threats, and energy opportunities). For example, this infrastructure have to be operated remotely and automatically as much as possible as it will impossible to have human presence on all locations. Due to their number, it will be necessary to allow operations not on a single site but on sets defined on the fly as needed. Moreover, the management mechanisms must have been designed to cope with intermittent network access to the sites. That is to say, offering on the one hand safety properties and on the

other hand autonomy in order to allow each site to remain as operational as possible in the event of network partitioning. Finally, currently existing interfaces (APIs) should be extended to turn location into a first-class citizen. In particular, the locality aspects should be redefined from the core system building blocks to the high-level application programming interfaces.

The STACK activities cover the full Cloud to IoT continuum, including recent challenges related to the network dimension and urgent computing. An enlargement of STACK core activities has started with the arrivals of Ass. Prof Koutsiamanis, Ass. Prof Piamrat and Dr. Balouek who respectively joined the team in 2021, 2022 and 2023. In 2024, the expertise of STACK has been further strengthened with the arrival of Orange members, and Ass. Prof Rosinosky and Ass. Prof Gonzalez.

2.3 Scientific Foundations

Through the ongoing integration of Orange members, STACK consolidates its expertise in distributed systems, networks, cyber-physical systems, IoT, device management, and software programming as well as combining significant skills in the design, practical development and evaluation of large-scale systems. More precisely, our research activities mainly rely on a set of scientific foundations detailed below.

- **(Distributed) Systems.** The first scientific foundation of the team is related to our strong expertise in resource management and capacity planning of large-scale infrastructure [76, 88, 80, 98, 92, 93]. This includes the design and evaluation of system mechanisms and algorithms to operate and use computation, network, storage, and IoT resources in an efficient and sustainable manner. Our knowledge is based on traditional as well as distributed system fundamentals, covering virtualization technologies, storage, security, energy, and distributed/parallel algorithms.
- **Networks.** Another set of expertise in the team concerns network related topics. This includes intelligent analysis and management in wireless and mobile networks using artificial intelligence and machine learning techniques, with particular focus on graph neural networks [5, 31], federated learning [14, 15] for traffic prediction and network security. It also includes the optimisation of wireless low-power and lossy networks (LLN), typically wireless Industrial IoT networks, through energy-aware network resource and communications scheduling and routing [79, 77, 54, 75, 57]. Additionally, using SDN enhances network security through fine-grained security policies and efficient control plane/data plane management of complex routing decisions. This enables high-performance networks and scalable management of large-scale environments, ensuring optimized resource usage and robust communication [70, 72].
- **Digital Twins, Network and Device Management.** Based on initial expertise in IoT platforms and Cyber-Physical-Systems, the management of connected devices and sensor data, and especially on distributed and autonomic architectures of such platforms [89, 46, 45], the team has developed a broader vision of network management and operation with a strong expertise in *digital twins* as a pivotal technology. This includes graph-based modelling of digital twins [87], semantic modelling and ontology mapping [47], graph storage distribution, federation and historisation [73, 74] — and the application of these concepts and technologies in different use cases in the domain of smart building (e.g., localisation [57], dynamic wireless IoT network resource allocation [75]), smart industry (e.g., support for reliable and low-latency wireless Industrial IoT networks [79, 77]), logistics [84] around the Thing' in digital twin platform [64] (*Thing' in the future*).
- **Autonomic and Self-Adaptive Systems.** Considering the high (and ever increasing complexity) of ICT systems, autonomic and self-adaptive policies have become the *de facto* standard for designing and building large-scale systems. This second family includes, for example, research approaches that have been harnessed to tackle system modularity, configuration and reconfiguration of dynamic and distributed systems, as well as retroaction and autonomic loops. All these concepts enable administrators and developers to deal with various objectives such as performance, high availability, low energy consumption, etc. STACK members have provided several relevant contributions in the last couple of years [67, 68, 83, 69, 82, 53, 55, 56].
- **Software Engineering and Programming.** Similarly, software engineering and advances in programming are highly valuable to correctly design complex systems such as the software stack we target.

Leveraging the expertise of software programming of the team, STACK contributions leverage various techniques including component-based programming models [65, 50, 66, 81, 62, 63], event-driven [71, 97], data-driven and workflow models, as well as models for Utility Computing (Service Level Agreement, aka. SLA) [95], and more generally, distributed and parallel programming models.


- **Experiment-Driven Research.** Finally, the last important domain of expertise of the future team consists in the evaluation of complex software stacks at large scale through simulations and in-vivo experiments. This includes knowledge on experimental methodology, measuring/monitoring/tracing tools [48] and more recently aspects related to software-defined experiments and reproducible research [59, 58]. Team members are also in charge of the animation of the LASCARE working group (LArge SCAle ARchitecture Experimentation and Simulation) of the IOLab.

We aim at strengthening the knowledge in these different areas through two kinds of contributions: First through scientific articles as a regular project team, and second, through concrete pieces of software that can be transferred to major opensource communities.

3 Research program

3.1 Overview

STACK activities have been focused on the management and programming of geo-distributed data centers with a work program defined around four research topics as depicted in Figure 1a. The first two ones are related to the resource management mechanisms and the programming support that are mandatory to operate and use ICT geo-distributed resources (compute, storage, network). They are transversal to the three software layers that generally compose a software stack (System/Middleware/Application in Figure 1b) and nurture each other (*i.e.*, the resource management mechanisms will leverage abstractions/concepts proposed by the programming support axis and reciprocally). The third and fourth research topics are related to the Energy and Security dimensions (both also crosscutting the three software layers). Although they could have been merged with the first two axes, we identified them as independent research directions due to their critical aspects with respect to the societal challenges they represent.



(a) Stack Research Topics

(b) Software stack layers

This scientific roadmap to address challenges related to the management and programming of geo-distributed infrastructures applies to the Cloud to IoT continuum and continues to have significant scientific and socio-economic impact. Hence, STACK organizes its activities around these four crosscutting lines that form a unique approach.

Additionally, our activities extend to the management of IoT devices with the ultimate goal of covering the entire Cloud-to-IoT continuum through a common software stack.

Our vision is to base this computing continuum software stack on control loops following the MAPE-K model¹, which can be seen as an infinite loop that monitors the infrastructure as well as the state of the applications in order to maintain in an autonomous manner the expected objectives (in terms of performance, robustness, etc.).

¹MAPE-K stands for Monitoring, Analyzing, Planning, Execution and Knowledge [78].

Although it is largely adopted in Cloud orchestrators such as Kubernetes, delivering a MAPE-K software stack for the computing continuum faces multiple challenges. The first one is related to the diversity of resources to consider. KubeEdge [100], for instance, proposes extensions to integrate servers and IoT devices under the same framework. However, the supported operations are rather limited as they only cover communication between software components running on servers at the edge and the connected devices. In other words, the IoT devices are not considered in the control loops. From our viewpoint, this weak integration is linked to an incomplete understanding of the needs that such a system must take into account (in particular on the IoT device management side). To favor the integration of both dimensions into a common system, we aim to identify major structural properties as well as management operations necessary at the operator and DevOps levels and to implement them when needed. A second important challenge is related to the geo-distribution property (and so the intermittent network connectivity) of this type of infrastructure. This increased complexity implies revising the way control loops are designed in order to handle frequent disconnections that can occur at any time (for instance due to the low energy level of IoT devices). Here, our approach is to combine autonomous loops with well-adapted formal methods to guarantee their verification or to synthesize correct-by-design decisions. Additionally, we study their performance models to be able to automatically, safely, efficiently and in a timely manner adapt Cloud-to-IoT infrastructures and their hosted applications according to different objectives (performance, energy, security, etc.). Finally, we are working towards partitioning a Cloud-to-IoT infrastructure into several areas and delivering the illusion of a single system through a federated approach: each area is managed by an independent controller, and collaborations between areas are done through dedicated middleware. The innovative aspect relies in the way of developing this middleware so that it is reliable in spite of increasing scale and faults (collaborations will be triggered only on-demand without maintaining a global knowledge base of the entire infrastructure).

Some of the research questions we address in the medium term are:

- *How to specify and model the dynamics of Cloud-to-IoT infrastructures and the dynamics of associated systems and applications in a generic way and how to leverage this dynamics for reconfiguration purposes?* In particular, this is done through studying how existing languages such as SysML, ThingML or TOSCA may be revised for this purpose. In addition, we focus on the exploration of the properties of safety, separation of concerns and efficiency when reconfiguring distributed systems with Concerto [56] extended to IoT devices and to network resources. Finally, we address the dynamicity of systems described by such Architecture Description Languages (ADL) considering a convergence between the *Model@Runtime* and Digital Twins approaches, i.e. implementing models@runtime as digital twins.
- *How to design and deploy decentralized autonomic loops, from monitoring to the execution of reconfiguration plans?* An important challenge is related to the development of mechanisms capable of rebuilding, on-demand, a knowledge base according to the functional and non-functional properties to be satisfied. From our viewpoint, it is crucial to propose alternative approaches to avoid maintaining such a global knowledge base through time and at the scale of a Cloud-to-IoT infrastructure. Regarding the monitoring/supervision of the infrastructure, we study the latest results on complex event processing as well as machine learning techniques. The former enables the triggering of actions based on predefined events while the latter allow the management to evolve from reactive to predictive strategies. Regarding the analysis and planning phases we study how to combine and leverage satisfiability solvers (SMT and SAT solvers), constraint programming, and distributed algorithms. We are also interested in the formal verification of reconfiguration procedures.
- *How to enhance legacy distributed software in an easy and non-intrusive manner.* In particular, we focus on handling the geo-distribution of applications, optimize their processing or include security features using service meshes and code injection techniques.
- *How to increase the responsiveness of data analysis algorithms and accelerate responses of AI-enabled scenarios across the Cloud-to-IoT continuum?* Another important challenge is related to the steering of computation considering data and events measured by the IoT infrastructure, coupled with historical information and the Quality of Service (QoS) needed. We investigate application formulations that allow developers to balance requirements and costs, along with programming abstractions to define policies that can react to unforeseen events and constraints.

- *On the energy dimension, the main questions are related to the generalization of the usage of renewable energies in the Cloud-to-IoT continuum while guaranteeing availability and reliability properties.* We investigate, in particular, whether energy harvesting devices could be used at the extreme edge and how they complicate the placement challenge that we largely studied in a multi-cloud context. Besides, we are working on extending our work to include green energy awareness for users (e.g., DevOps engineers, web application end-users, etc.).
- *How can telemetry metrics collected by P4 be integrated into a distributed architecture for both local and global decision-making?* Integrating telemetry metrics collected by P4 into a distributed architecture enables real-time monitoring and adaptive decision making. P4's programmable data plane allows the extraction of fine-grained metrics such as latency, packet loss, and throughput. These metrics can be fed into a centralized or hierarchical control system, enabling local decisions to optimize immediate performance and global decisions for long-term resource allocation and policy enforcement. By leveraging SDN or ICN principles, these telemetry insights can improve traffic management, enhance fault tolerance, and support scalable and efficient operation across distributed environments.
- *Finally, on the security side, we investigate the new threats resulting from an externalized management of geo-distribution.* This includes, in particular, the identification of new possible attack channels as well as counter measures to guarantee a satisfactory level of security through the whole continuum. Furthermore, we are making efforts to extend our work on kernel security policies [49] in order to also take into account the network dimension and ensure strong isolation from Cloud/Edge servers to IoT devices.

All the aforementioned research questions are addressed through several application fields: telecommunications operators and smart buildings in the first place through this privileged partnership with the Orange colleagues who are going to join this new team but also in health, in particular, biomedical research in order to allow the execution of analyses, currently emerging, in large-scale geo-distributed environments.

4 Application domains

Industrial/Tactile Internet/Cyber-Physical applications highlight the importance of the computing continuum model. Hence, the use-cases of STACK activities are driven and nurtured by these application domains. Besides, it is noteworthy to mention that Telecom operators such as Orange have been among the first ones to advocate the deployment of Fog/Edge infrastructure. The initial reason is that a geo-distributed infrastructure enables them to virtualize a large part of their resources and thus reduce capital and operational costs. As an example, several researchers have been investigating through the IOLab, the joint lab between Orange and Inria, how 5G networks can be managed. We highlight that while our expertise does partially include the network side, the main focus is rather on how we can deploy, locate and reconfigure the software components that are mandatory to operate next generation of network/computing infrastructure. The main challenges are related to the high dynamicity of the infrastructure, the way of defining Quality of Service of applications and how it can be guaranteed. We expect our contributions will deliver advances in location based services, optimized local content distribution (data-caching) and Mobile Edge Computing ². In addition to bringing resources close to end-users, massively geo-distributed infrastructures should favor the development of more advanced networks as well as mobile services.

4.1 Overview

Supporting industrial actors and open-source communities in building an advanced software management stack is a key element to favor the advent of new kinds of information systems as well as web applications. Augmented reality, telemedicine and e-health services, smart-city, smart-factory, smart-transportation and remote security applications are under investigations. Although, STACK does not intend to address directly the development of such applications, understanding their requirements is critical to identify how the next generation of ICT infrastructure should evolve and what are the appropriate software abstractions for operators, developers and end-users. STACK team members have been exchanging since 2015 with a number of

²ETSI MEC specifications.

industrial groups (notably Orange Labs and Airbus), a few medical institutes (public and private ones) and several telecommunication operators in order to identify both opportunities and challenges in each of these domains, described hereafter.

4.2 Industrial Internet

The Industrial Internet domain gathers applications related to the convergence between the physical and the virtual world. This convergence has been made possible by the development of small, lightweight and cheap sensors as well as complex industrial physical machines that can be connected to the Internet. It is expected to improve most processes of daily life and decision processes in all societal domains, affecting all corresponding actors, be they individuals and user groups, large companies, SMEs or public institutions. The corresponding applications cover: the improvement of business processes of companies and the management of institutions (*e.g.*, accounting, marketing, cloud manufacturing, etc.); the development of large “smart” applications handling large amounts of geo-distributed data and a large set of resources (video analytics, augmented reality, etc.); the advent of future medical prevention and treatment techniques thanks to the intensive use of ICT systems, etc. We expect our contributions favor the rise of efficient, correct and sustainable massively geo-distributed infrastructure that are mandatory to design and develop such applications.

4.3 Internet of Skills

The Internet of Skills is an extension of the Industrial Internet to human activities. It can be seen as the ability to deliver physical experiences remotely (*i.e.*, via the Tactile Internet). Its main supporters advocate that it will revolutionize the way we teach, learn, and interact with pervasive resources. As most applications of the Internet of Skills are related to real time experiences, latency may be even more critical than for the Industrial Internet and raise the locality of computations and resources as a priority. In addition to identifying how an Utility Computing infrastructure can cope with this requirement, it is important to determine how the quality of service of such applications should be defined and how latency and bandwidth constraints can be guaranteed at the infrastructure level.

4.4 e-Health

The e-Health domain constitutes an important societal application domain of the two previous areas. The STACK teams is investigating distribution, security and privacy issues in the fields of systems and personalized (aka. precision) medicine. The overall goal in these fields is the development of medication and treatment methods that are tailored towards small groups or even individual patients.

We have been working on different projects since the beginning of STACK (*e.g.*, PrivGen CominLabs). In general, we are applying and developing corresponding techniques for the medical domains of genomics, immunobiology and transplantology (see Section 10).

The STACK team continue to contribute to the e-Health domain by harnessing advanced architectures, applications and infrastructure for the Fog/Edge, Cloud/Edge, and Cloud/Edge/IoT continuum.

4.5 Network Virtualization and Mobile Edge Services

Telecom operators have been among the first to advocate the deployment of massively geo-distributed infrastructure, in particular through working groups such as the [Mobile Edge Computing at the European Telecommunication Standards Institute](#). The initial reason is that a geo-distributed infrastructure enables Telecom operators to virtualize a large part of their resources and thus reduces capital and operational costs. Through the Sylva project, we aim to expand our collaborative efforts with Orange and other key stakeholders in the network community. We focus on exploring how Cloud-Native Functions (CNFs) and telco cloud infrastructure can support innovative use cases over the SLICES-FR experimental infrastructure.

4.6 Urgent Computing

Urgent Computing refers to a class of time-critical scientific applications that leverage distributed data sources to facilitate important decision-making in a timely manner. The overall goal of Urgent Computing is

to predict the outcome of scenarios early enough to prevent critical situations or to mitigate their negative consequences. Motivating use cases refers to rapid response scenarios across the Cloud-to-IoT Continuum, such as in natural disaster management, which implies to gather the local state of each device, transform it into a global knowledge of the network, characterize the observed phenomenon according to an applied model, and finally, trigger appropriate actions. The STACK team investigates Urgent Computing through the realization of a fluid ecosystem where distributed computing resources and services are aggregated on-demand to support delay-sensitive and data-driven workflows.

5 Social and environmental responsibility

5.1 Footprint of research activities

In addition to the international travels, the environmental footprint of our research activities is linked to our intensive use of large-scale testbeds such as Grid'5000 (STACK members are often in the top 10 list of the largest consumers). Although the access to such facilities is critical to move forward in our research roadmap, it is important to recognize that they have a strong environmental impact as described in the next paragraph.

5.2 Impact of research results

The environmental impact of digital technology is a major scientific and societal challenge. Even though the software looks to be virtual objects, it is executed on very real hardware contributing to the carbon footprint. This impact materializes during the manufacturing and destruction of hardware infrastructure (estimated at 45% of digital consumption in 2018 by the Shift Project) and during the software use phase via terminals, networks and data centers (estimated at 55%). Stack members have been studying various approaches for several years to reduce the energy footprint of digital infrastructures during the use phase. The work carried out revolves around two main axes: (i) reducing the energy footprint of infrastructures and (ii) adapting the software applications hosted by these infrastructures according to the energy available. More precisely, this second axis investigates possible improvements that could be made by the end-users of the software themselves. At scale, involving end-users in decision-making processes concerning energy consumption would lead to more frugal Cloud computing.

In 2025, the team has taken part to the Inria Pulse challenge (*i.e.*, “Défis”). Started in Oct 2022, this 4 years program between Inria and Qarnot Computing, with the support of ADEME aims to develop and promote best practices in terms of reducing and recycling emissions of intensive computing infrastructures.

Moreover, the team pursues its activities related to the Samurai platform, a project to design an innovative hardware infrastructure for the scientific study of the cross-cutting issues of computing infrastructures supporting artificial intelligence and their energy autonomy (see the Samurai project Section in 10).

Finally, the team has initiated and taken part to several debates, panels and workshops related to the environmental impact and the illusion of unlimited resources of cloud computing platforms (e.g. [Infrastructure workshop in Dec 2025](#))

6 Highlights of the year

Regarding scientific contributions, the team has produced major results on the management of large-scale infrastructures, in particular the team continued to study the relevance of AI technics [2, 3, 1]

On the software side, the team has pursued its efforts on the development of the EnosLib library and the resulting artifacts to help researchers perform experiment campaigns with direct contributions from several research engineers of the team.

On the platform side, we continued our effort and took part in the different actions around the SLICES and SLICES-FR, see Section 7.2.

Finally, we would like to highlight our continued efforts to maintain the strong visibility of the team on the convergence of telco/cloud and edge-related topics. In particular, the team has taken part to two european submissions (both under review, HORIZON-JU-SNS-2025-01-STREAM-C-01 and HORIZON-CL4-2025-03-DATA-08). The second submission is a major initiative with a €70 million budget, bringing together over

60 partners across Europe. Being solicited to take part to such actions are prime examples of the recognition of our work, alongside our significant involvement in the French Cloud PEPR initiative.

6.1 Awards

At the 25th IFIP WG 6.1 International Conference on Distributed Applications and Interoperable Systems (DAIS 2025), held as part of DisCoTec 2025 in Lille, France (June 16-20, 2025), the paper "Justin: Hybrid CPU/Memory Elastic Scaling for Distributed Stream Processing" [20] received the best-paper award recognition for both the DAIS and the DisCoTec conference. Authored by researchers Donatien Schmitz, Guillaume Rosinsky, and Etienne Rivière the paper introduces Justin, a novel auto-scaling system for distributed stream processing engines like Apache Flink.

The ANR **SeMaFoR project** lead by Thomas Ledoux was awarded the Innovation Trophy in the Digital Sector by the Pôle Universitaire d'Innovation (PUI) of Nantes Université (award ceremony on December 18, 2025). It was selected by the PUI's steering committee.

7 Latest software developments, platforms, open data

7.1 Latest software developments

7.1.1 ENOS

Name: Experimental eNvironment for OpenStack

Keywords: OpenStack, Experimentation, Reproducibility

Functional Description: A typical experiment workflow using Enos is the sequence of several phases:

- enos up : Enos will read the configuration file, get machines from the resource provider and will prepare the next phase
- enos os : Enos will deploy OpenStack on the machines. This phase rely highly on Kolla deployment.
- enos init-os : Enos will bootstrap the OpenStack installation (default quotas, security rules, ...)
- enos bench : Enos will run a list of benchmarks. Enos support Rally and Shaker benchmarks.
- enos backup : Enos will backup metrics gathered, logs and configuration files from the experiment.

Release Contributions: - Install a fixed version of Docker on nodes - Add support for kolla-ansible 12 (Openstack Wallaby), which is now the default - Add support for Debian 11 base image - Add support for Python 3.10 and 3.11 - Drop support for Python 3.7 - Update to Enoslib 8 - When creating a configuration template, use a fixed version of kolla-ansible (so that templates are not affected when future versions of Enos update the default version of kolla-ansible)

URL: <https://github.com/BeyondTheClouds/enos>

Publication: [hal-01664515](#)

Contact: Baptiste Jonglez

Participant: 4 anonymous participants

Partner: Orange Labs

7.1.2 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

7.1.3 Concerto

Name: Concerto

Keywords: Reconfiguration, Distributed Software, Component models, Dynamic software architecture

Functional Description: Concerto is an implementation of the formal model Concerto written in Python. Concerto allows:

1. the description of the life cycle and the dependencies of software components,
2. the description of a component assembly that forms the overall life cycle of a distributed software,
3. the automatic reconfiguration of a Concerto assembly of components by using a set of reconfiguration instructions as well as a formal operational semantics.

URL: <https://gitlab.inria.fr/VerDi-project/concerto>

Publications: [hal-03103714](#), [hal-02535077](#), [hal-01897803](#)

Contact: H el ene Coullon

Participant: 4 anonymous participants

Partners: IMT Atlantique, LS2N, LIP

7.1.4 StreamBed

Name: StreamBed capacity planning for steam processing

Keywords: Big data, Data stream, Performance measure

Functional Description: StreamBed is a capacity planning system for stream processing. It predicts, ahead of any production deployment, the resources that a query will require to process an incoming data rate sustainably, and the appropriate configuration of these resources. StreamBed builds a capacity planning model by piloting a series of runs of the target query in a small-scale, controlled testbed. We implement StreamBed for the popular Flink DSP engine. Our evaluation with large-scale queries of the Nexmark benchmark demonstrates that StreamBed can effectively and accurately predict capacity requirements for jobs spanning more than 1,000 cores using a testbed of only 48 cores.

News of the Year: Publication and artefact available on GitHub.

URL: <https://github.com/CloudLargeScale-UCLouvain/StreamBed>

Publication: [hal-04708354](#)

Contact: Guillaume Rosinosky

Participant: an anonymous participant

Partner: Université Catholique de Louvain (UCL), Louvain-la-Neuve, Belgium

7.1.5 Cheops

Name: Cheops for the edge

Keywords: Edge Computing, Geo-distribution, Infrastructure software

Functional Description: Cheops handles the task of synchronizing and replicating your application resources (in the sense of REST).

The user will interact with Cheops to perform operations on resources, and Cheops will then interact with replicas of your application to make sure that all resource copies eventually converge to the same state.

It is assumed that operations are always associated to a specific resource. Cheops gives the possibility to specify the exact distribution of each resource manually so that operators can define how they want them to be spread.

URL: <https://cheops-for-the-edge-discovery-e233bd9558e1939580edccc727171f286.gitlabpages.inria.fr/>

Publications: [hal-04886168](#), [hal-04522961](#), [tel-04081084](#)

Contact: Adrien Lebre

7.1.6 Edge-to-cloud video processing

Keywords: Edge Computing, Video analysis

Functional Description: The software performs distributed video processing to identify animals in video feeds, in order to warn populations if a dangerous animal is detected.

The primary goal is to provide a use-case to perform research on edge-to-cloud infrastructure. To achieve this, the software is instrumented with many metrics.

The software is composed of three components:

- video capture: either from a camera, or by replaying a pre-recorded video file that can be parametrized
- motion detection: detects motion in video feeds. If a motion is detected, the feed is forwarded to the recognizer component
- object recognizer: uses a YOLO model to determine the kind of animal visible in the video feed

These three components are designed to execute in different places in the edge-to-cloud infrastructure, so that researchers can explore the trade-off between performance, latency, transferred data, and quality.

URL: <https://gitlab.inria.fr/STACK-RESEARCH-GROUP/software/edge-to-cloud-video-processing>

Contact: Baptiste Jonglez

Partner: IMT Atlantique

7.2 New platforms

7.2.1 Grid'5000

Participants: Remous-Aris Koutsiamanis, Baptiste Jonglez, Adrien Lebre, Jean Marc Menaud.

Grid'5000 is a large-scale and versatile testbed for experiment-driven research in all areas of computer science, with a focus on parallel and distributed computing including Cloud, HPC and Big Data. It provides access to a large amount of resources: 12000 cores, 800 compute-nodes grouped in homogeneous clusters, and featuring various technologies (GPU, SSD, NVMe, 10G and 25G Ethernet, Infiniband, Omni-Path) and advanced monitoring and measurement features for traces collection of networking and power consumption, providing a deep understanding of experiments. It is highly reconfigurable and controllable. STACK members are strongly involved into the management and the supervision of the testbed, notably through the steering committee or the SeDuCe testbed described hereafter.

7.2.2 PiSeDuCe

Participants: Remous-Aris Koutsiamanis, Baptiste Jonglez, Jean Marc Menaud.

We continue to manage and extend the PiSeDuCe platform, a deployment and reservation system for Edge Computing infrastructures composed of multiple Raspberry Pi Cluster started in 2020. The platform is typically composed of a cluster of 8 Raspberry Pi, which costs less than 900 euros and only needs an electrical outlet and a wifi connection for its installation and configuration. Funded by the CNRS through the Kabuto project, and in connection with the SLICES-FR initiative, we have extended PiSeduce to propose a device to cloud deployment system (from devices on Fit IoTLab to servers in Grid'5000). PiSeDuCe and SeDuCe led us to submit the Samurai CPER proposal. Recent developments have made the platform more performant, being able to manage 45 worker Raspberry Pis using just one controller Raspberry Pi. The platform has also been demonstrated and used in a hands-on tutorial at the "Green IT - Numérique responsable" 2024 summer school in Nantes.

7.2.3 SAMURAI

Participants: Remous-Aris Koutsiamanis, Baptiste Jonglez, Jean Marc Menaud.

The SAMURAI (Sustainable And autoNoMoUs gReen computing for AI) project is currently being financed as a part of the energy and digital transition theme. The project aims at reinforcing an innovative hardware infrastructure for the scientific study of the intersecting problems of computing infrastructure that supports artificial intelligence and its energy autonomy. SAMURAI is focused on extending SeDuCe into energy autonomy by adding a smart and clean energy storage system. Additionally, it has extended the capabilities of the platform by adding AI computing nodes (servers with GPUs) for the scientific study of AI tools. Finally, it will also add new sensor nodes within the Nantes connected object platform (Nantes site of the national SLICES-FR platform) to support future work on embedded AI as well as moregenerally on the Cloud-Edge-IoT continuum. As the majority of the hardware for the IoT, Edge and Cloud with GPUs has been procured, the project is focusing on building the platform and network integration with the SeDuCe and SLICES-FR platform.

7.2.4 SLICES-FR/SLICES

Participants: Remous-Aris Koutsiamanis, Baptiste Jonglez, Adrien Lebre, Jean Marc Menaud.

In 2025, STACK participated in the SLICES-PP project (Preparatory Phase) as part of the SLICES-RI European infrastructure initiative. STACK was particularly active within SLICES-FR, the French component of the infrastructure. STACK Members have been involved in the definition and bootstrapping of the SLICES-FR infrastructure. This infrastructure can be seen as a merge of the Grid'5000 and FIT testbeds with the goal of providing a common platform for experimental Computer Science (Next Generation Internet, Internet of things, network functions, clouds, HPC, big data, etc.). Adrien Lèbre and Remous-Aris Koutsiamanis are part of the SLICES-FR Board (currently provisional, pending the official creation of the GIS legal structure). Additionally, Remous-Aris Koutsiamanis and Baptiste Jonglez are members of the Architects Committee of SLICES-FR, where they contributed key design documents focused on networking and interoperability, aligning with the overall SLICES-RI design. Finally, Adrien Lèbre serves as the French scientific representative on the SLICES Interim Supervisory Board, a provisional body awaiting the formal establishment of the ERIC legal structure.

8 New results

8.1 Resource Management

Participants: Daniel Balouek, H el ene Coullon, Remous-Aris Koutsiamanis, Adrien Lebre, Duc Thinh Ngo, Kandaraj Piamrat, Thomas Ledoux, Guillaume Rosinosky, Cherif Si Mohammed, Mario S udholt.

The evolution of the cloud computing paradigm in the last decade has amplified the access to on-demand services (economically attractive, easy-to-use manner, etc.). However, the current model, built upon a few large datacenters (DCs), is not suited to guarantee the needs of new use cases, notably the boom of the Internet of Things (IoT). To better respond to the new requirements (in terms of delay, traffic, etc.), compute and storage resources should be deployed closer to the end-user, forming with the national and regional data centers a new computing continuum. The question is then how to manage such a continuum to provide end-users the same services that made the current cloud computing model so successful. In 2025, we have continued our effort to answer this question and delivered multiple contributions, including additional results related to Urgent Computing (see Section 4).

Network resource management: In telecommunications, growing users and devices in next-generation networks (beyond 5G) intensify traffic demands, stressing limited resources. O-RAN architectures enable flexible and cost-efficient networks but introduce complex, large-scale resource management challenges. Our proposed Graph-Augmented PPO (GPPO)[31], which combines Graph Neural Networks for topology-aware representations with action masking to jointly optimize functional splits and virtualized unit placement. Experiments on small- and large-scale O-RAN scenarios show that GPPO outperforms state-of-the-art methods, achieving up to 18% lower deployment cost and 25% higher reward with perfect reliability. Additionally, cellular traffic prediction faces challenges like dynamic base station deployment. Network Digital Twins (NDT) rely on accurate traffic forecasting, but existing spatiotemporal models often depend on fixed graphs, limiting adaptability in dynamic networks. Our proposed Flex+ [5], an inductive graph-based model that predicts eNodeB traffic using local k-hop spatial correlations and temporal features, enabling operation on unseen nodes and in data-scarce settings. Experiments on large-scale cellular data show 5.9% accuracy improvement in inductive scenarios, 22% error reduction with only 3 days of training data, and up to 10× faster inference via knowledge distillation without loss of accuracy. These efforts are part of the work that have been done since a decade and presented during the HDR of Kandaraj Piamrat [38]. Finally, on the resource-constrained side of networks, reliable communication in Low-power and Lossy Networks (LLNs) is a key requirement for Industrial IoT, yet single-path RPL often fails to provide deterministic

reliability and latency guarantees under lossy or congested conditions. While multipath RPL combined with PAREO functions (ARQ, packet replication and elimination, overhearing) can improve delivery by leveraging controlled redundancy, its efficiency crucially depends on how Preferred and Alternative Parents are selected, as overly permissive choices may trigger excessive flooding and energy waste. In this context, we proposed ODeSe (On-Demand Selection) [16], a two-hop parent selection algorithm that aligns relay choices across nodes at the same DODAG layer by encouraging shared upstream relays, while dynamically adjusting parent assignments when necessary (and falling back to Soft Common Ancestor selection when no suitable alignment is possible). Experiments on a 32-node, five-hop topology with 50% link quality show that ODeSe reaches 99.14% packet delivery while reducing redundant transmissions, achieving comparable reliability to the most permissive baseline but with lower power consumption (0.09 vs. 0.10 mW in the evaluated setting).

Data resource management: Big data stream processing approaches often do not consider the fact that state has to be persisted in jobs: they mainly consider CPU-bound jobs. In this context, following our previous work on capacity planning estimation [91], where we pinpointed the non-linear behaviour of join-based jobs, and proved the behaviour of jobs is predictable at very high rates with simple regression methods, we have done an analysis of the kind of jobs that are used in a typical industrial data lake provider [19]. Our findings are that queries are highly heterogeneous in size and complexity, and that there is an high quantity of low complexity jobs involving high state needs. Benchmarks currently used in the community do not consider this fact. We have also proposed a new autoscaler [20] for Apache Flink permitting to take into account CPU and memory needs for the aforementioned scenarios. It identifies memory pressures and chooses whether scaling CPU or memory should be done.

The second activity on data management is related to the Inria/Qarnot computing Pulse Challenge. More precisely, we explored data replication, a key strategy for improving system performance in geo-distributed computing environments. While replication can improve efficiency, naive strategies often result in excessive and unnecessary data transfers, leading to inefficient resource utilization, particularly in infrastructures interconnected via heterogeneous network links. To address these limitations, we have proposed a replication strategy that estimates the utility of each potential replica before deployment. The approach has been first evaluated in homogeneous environments and then extended to heterogeneous settings with varying network characteristics. Simulation results show that our method significantly reduces data transfers while maintaining high execution efficiency, achieving a balanced trade-off between performance and resource consumption [22, 32, 44].

(Bio)medical analyses: In the course of our long-running cooperation with researchers from Nantes University Hospital (CHU Nantes), we have worked on clinical data for lung transplantations in order to develop medical analyses for the survival rate after such transplantations. Although not directly aligned with STACK's objectives, the group's expertise, in particular on data management has proven valuable in finding the right techniques to analyse such a large data repository.

The main limitation to long-term lung transplant (LT) survival is chronic lung allograft dysfunction (CLAD), which leads to irreversible lung damage and significant mortality. Individual factors can impact CLAD, but no large genetic investigation has been conducted to date. In 2024, we have established the multicentric Genetic COhort in Lung Transplantation (GenCOLT) biobank from a rich and homogeneous sub-part of COLT cohort [52]. We continued this work in 2025. Precisely, we applied statistical data analyses and learning techniques to the data bank on lung transplantations. This work that has been presented as part of Simon Brocard's PhD thesis, has resulted, in particular, in the first definition of links between two gene markers and graft rejection complications [43].

Urgent Computing: Urgent computing scenarios describe challenges in promptly responding to changes in the Edge-Cloud Continuum or adapting the quality of service (QoS) constraints of the application. This line of work is highly motivated by decision-making systems or natural disaster case studies such as earthquake early warning and wildfires.

In 2025, we proposed an approach focusing on leveraging Artificial Intelligence for Data-Driven Natural Disaster Management [21], to facilitate dynamic resource management and adaptive system modeling,

thereby addressing the decision-making challenges posed by disaster scenarios. We then focused on anticipating system states in such architectures by predicting CPU utilization in the computing continuum [9]. We extended this work by implementing transfer learning from Virtual Machines to Containers using Transformers [18]. This work enables the system to anticipate CPU states with very little prior information, and thus accelerate decision-making for time-critical applications.

In parallel tasks, we contributed to a systematic mapping study to understand the different uses of AI in microservices lifecycle [4]. Outcomes of this study highlight the impact of AI during the design and deployment phases of microservices-based implementation, along with on-going practises. Finally, we targeted adaptation mechanisms along two dimensions: (1) software variability for the potential of capturing an extensive number of configuration with the ambition to manage advanced tradeoffs between cost and quality [17, 30], and (2) feedback-oriented architecture for Application-level observability [23].

8.2 Programming Support

Participants: Daniel Balouek, H el ene Coullon, Thomas Ledoux, Jacques Noy e, Antoine Omond, Eloi Perdereau, Jolan Philippe, Hiba Awad, Mario S udholt, Divi De Lacour.

Fog Modeling: Fog Computing moves some Cloud functions closer to where data is generated. This approach cuts bandwidth use, reduces delays, and minimizes data transfers. However, designing and building Fog systems is complex and often leads to mistakes.

To address this, we can apply software engineering best practices, such as verifying system properties before deployment. Previous research has focused on checking non-functional aspects of Fog systems at early stages. In [7], we introduce VeriFogOps, an approach that automatically selects the right deployment tools based on Quality of Service (QoS) needs and then creates CI/CD pipelines to support Fog system deployment. We tested and validated VeriFogOps with two real-world use cases, using different QoS solutions and deployment tools. Developed with our industry partner Smile, this work helps support the entire lifecycle of Fog systems.

In 2025, Hiba Awad defended her PhD thesis [33] on the subject of this important question dealing with Quality of service assurance before deployment of Fog systems with model-based engineering and DevOps.

Configuration languages: In 2025, in the context of the PEPR Taranis, we have pursued our work on the anatomy on configuration languages, in collaboration with Philippe Merle (Spirals), as well as on the semantics of the configuration language CUE.

Software deployment/reconfiguration: For a few years, the team has been working on deployment and dynamic reconfiguration of distributed software systems through the Concerto tool suite [56, 90, 85, 60, 86]. Compared to the literature, Concerto is a component model closer to Infrastructure-as-Code (IaC) approaches of the DevOps community (as Aeolus [61]). In Concerto, the lifecycle of pieces of software (application or infrastructure) are modeled in a programmable manner with fine grain dependencies, enhancing the flexibility and speed of deployments and management procedures (i.e., reconfiguration) compared to existing approaches. Notably, the PhD of Antoine Omond [37] and the HDR of H el ene Coullon [34] has been defended in 2025.

In 2024, while not being published yet, we have started to work more directly on Infrastructure-as-Code tools and languages. After working on Concerto, Concerto-D and Ballet (i.e., well constructed research models and prototypes), we think it is important to start from existing production and more complex languages, with the aim of bringing together both sides. First, in collaboration with Daniel Sokolowski (who visited the team in December 2023) and Guido Salvaneschi, and in the context of the PEPR Cloud (taranis project), Eloi Perdereau and H el ene Coullon are working on a submission around the formal semantics and verification of Terraform and Pulumi (provisioning tools). Second, the Inria transfer action (ADT) project CoAnsible has started to develop an Ansible extension that uses Concerto as a coordination backend. Finally, the ANR project For-CoaLa coordinated by H el ene Coullon in collaboration with Fr ed eric Loulergue (University of

Orléans) has started. In this context the PhD of Olivia Proust focuses on the formal semantics of Ansible and CoAnsible (including Concerto) to verify general theorems on associated languages. For now, none of these initiative got accepted papers but we are very active. Jolan Philippe, who is a former PhD student and postdoc of the team is now associate professor at the University of Orléans and has join the project.

In addition to this, in the context of the PEPR Cloud (Taranis project), in collaboration with Christian Perez, we co-supervise the postdoc of Quentin Guilloteau. From this collaboration we also have submitted a first attempt to add the component model Legato on top of Concerto. The goal of Legato is to handle heterogeneous deployment environments and encapsulate them in a hierarchical way. This model also leads to a new Nested Doll Placement problem introduced in the submission and solved with the Gurobi interger linear programming solver. The CIFRE PhD of Simon Artus also targets to extend Concerto or Legato to be able to handle heterogeneous hardware, in particular deployment on devices with specific protocols.

Declarative learning: We have motivated the need for a generic definitional framework and implementation support for transfer learning. We have then introduced Generic Transfer Learning (GTL). GTL supports the declarative definition of transfers through neural network transformations and dataset manipulations and includes corresponding Python implementation support. We have also presented a case study demonstrating how to define and implement a transfer using GTL in the health domain. [11]

In the domain of distributed learning, we have provided a first version of a model for monitoring techniques dedicated to support distributed learning algorithms. that .ite...boucha:hal-05460554

8.3 Energy-aware computing

Participants: Remous-Aris Koutsiamanis, Thomas Ledoux, Jean-Marc Menaud, Guillaume Rosinosky, Thierry Coupaye, Lylian Siffre.

The activities on this axis are mainly related to the design, development and deployment of the SAMURAI project(7.2.3), a testbed that will allow researchers to investigate energy related challenges over the computing continuum (from the Cloud to IoT devices/cyber physical systems).

In [99], we propose to address the inherent complexities of effective application scheduling in cloud environments by focusing on reducing operational costs more specifically in terms of energy consumption and application migration expenses. In this intent, we propose a two-stage decision-making process supplemented by heuristics and metaheuristics. Extensive experiments executed on a public dataset and industry-standard contests demonstrate the significant advancements provided by our approach.

Local-first Software: The growing energy footprint of ICT services has become a critical environmental concern. While current approaches mainly focus on optimizing existing architectures, we advovate in exploring a more fundamental transformation: moving from cloud-centric architectures to local-first software architectures, where data and computing reside primarily on end-user devices. Through a preliminary study [25], we examine the implications of such an architectural shift on energy consumption. We analyze these impacts through three concrete examples, demonstrating how local software approaches could reduce energy consumption. Our analysis reveals both opportunities and challenges associated with this architectural transformation.

In [24], we go further. Through a controlled empirical study using SeekTune, an open-source music recognition application, we compare the client-side and server-side energy and data consumption of a local processing architecture versus a traditional SaaS model for a computationally intensive task. Our results demonstrate that, while local processing consumes more energy on the client endpoint, the combined energy cost of data transmission and server-side computation in the SaaS model often results in a larger overall energy footprint. Furthermore, the local version transmits significantly less data over the network than the SaaS model.

8.4 Security and Privacy

Participants: Wilmer Edicson Garzon Alfonso, Housseem Jmal, Remous-Aris Koutsiamanis, Kandaraj Piamrat, Mario Südholt.

This year the STACK team has provided new results on security and privacy issues in the networking and biomedical domains. We have developed new AI-based methods that support new means of property analysis and dynamic adaptation.

Systems security: As cloud and edge infrastructures increasingly follow the confidential computing paradigm, tenants can no longer assume that privileged software such as the host OS or the hypervisor is trustworthy. In this context, vTPMs are attractive trust anchors for integrity measurement and secure key storage, but many existing designs still depend on the host environment or on a shared physical TPM, which both enlarges the trusted computing base and weakens isolation. To address these limitations, we propose a new vTPM architecture that executes entirely inside an Intel SGX enclave hosted within the tenant VM, providing strong isolation from the underlying infrastructure [27]. Our design includes a trust establishment mechanism based on SGX remote attestation via a Privacy CA, which binds the vTPM identity to a specific enclave and prevents fake vTPM and cuckoo attacks. We also protect the vTPM persistent state by sealing the NVRAM within the enclave and by leveraging SGX monotonic counters to detect rollback attempts, thus mitigating NVRAM tampering and state rollback. Experimental results show that this strong isolation comes with minimal overhead, and in several cases improves performance compared to a baseline vTPM without SGX protection.

Network security and privacy: AI has enabled intelligent way of management network security and privacy. Particularly, Federated Learning (FL) enables distributed intelligence at the network edge, while semi-decentralized FL (SDFL) improves robustness by coordinating multiple servers instead of a single central one. To benefit of both, we propose TUNE-FL [14], an adaptive semi-synchronous SDFL framework that ensures consensus under arbitrary topologies and efficiently handles client heterogeneity in computation and data distribution. Experiments on intrusion detection datasets show that TUNE-FL outperforms representative baselines in accuracy while reducing training time by up to 97x. FL also enables privacy-preserving collaborative intrusion detection system (IDS) for IoT networks but suffers from efficiency issues due to client heterogeneity and non-IID data. To tackle this issue, we propose FLAIR [15], an adaptive semi-synchronous FL framework that uses a Decision Transformer to intelligently select clients based on current and historical information, reducing communication and computation overhead. Experiments on IDS datasets show that FLAIR reduces computation and communication time by up to 94% and 93%, respectively, while maintaining comparable detection accuracy to baseline methods. In addition, we begin investigating hierarchical supervision and monitoring in the continuum [26] in order to tackle heterogeneity issues while preserving privacy and optimizing overhead.

Secure and cooperative autonomous systems: We have proposed a unified reference architecture for cooperative autonomous systems (CAS), called RACAS, to support the design of efficient and modular CASes. We have also shown how such systems can be secured on-demand. In addition, we have illustrated how RACAS can be applied to an intelligent transportation system including networks of autonomous and connected vehicles. [10]

9 Bilateral contracts and grants with industry

9.1 Bilateral contracts with industry

Kelio (formely Bodet Software)

Participants: Thomas Ledoux.

The ArchOps 2 Chair (for Architecture, Deployment and Administration of Agile IT Infrastructures) is an industrial chair of IMT Atlantique, in partnership with **Kelio**, an SME specialized in solutions for time and attendance management. It is dedicated to all IMT Atlantique students in the field of IT. It is also a channel for the transfer of high-level skills: researchers, experts and industrials.

In 2025, several activities were conducted, such as a provocatively named conference "2022-2026: Comment l'IA "a mangé" le Développement Logiciel" with approximately 120 attendees, and preliminary discussions on the topic of a joint doctoral thesis with the Stack team.

9.2 Bilateral grants with industry

Alterway/Smile

Participants: Thomas Ledoux, Hiba Awad.

In 2020, during the preparation of the ANR SeMaFoR project, we started a cooperation with **Alterway/Smile**, an SME specialized in Cloud and DevOps technologies. This cooperation resulted in a joint PhD thesis (called Cifre) entitled "Quality of Service Assurance Before Deployment of Fog Systems with Model-Based Engineering and DevOps" started in Nov. 2021.

In 2025, Hiba Awad defended her PhD thesis [33] and published one last article [7].

Kapela

Participants: Thomas Ledoux, Lylian Siffre.

Lylian Siffre started his PhD in Nov. 2024 on the subject "Impacts and Uses of Local-First Software for Energy Optimization of IT Services", under a co-supervision with **Kapela** (Constellation group), an SME specialized in IT eco-design.

In 2025, Lylian Siffre published two articles in international conferences ([25, 24]) where he demonstrated the interest of local-first software over classic SaaS software to reduce the energy footprint of distributed architectures in certain use cases.

Orange

Participants: Simon Artus, Paul Bori, H  l  ne Coullon, Divi de Lacour, Adrien Lebre, Thomas Ledoux, Jean-Marc Menaud, Duc-Thinh Ngo, Kandaraj Piamrat, Mario S  dholt.

Since 2022, Orange Labs and the Stack team have launched several PhD grants.

Paul Bori started his PhD in January 2023, with the subject "Container application security: a programmable OS-level approach to monitoring network flows and process executions".

Duc-Thinh Ngo started his PhD in December 2022, on the subject "Dynamic graph learning algorithms for the digital twin in edge-cloud continuum", under a co-supervision with Orange team in Rennes. He defended his thesis on 27 Nov. 2025 [36].

Divi de Lacour has defended his PhD thesis [35] on 16 June 2025 on the subject "Architecture et services pour la protection des donn  es pour syst  mes coop  ratifs autonomes", under a co-supervision with the Orange team in Chatillon (Paris south region).

Since April 2024, the STACK research group has officially been established as a joint team between Inria, IMT Atlantique, Nantes University, and Orange Labs. This collaboration includes several PhD grants funded by the three institutions. For clarity, we no longer distinguish between PhD candidates based on their funding source (Inria, IMT Atlantique or Orange) in this context.

Ericsson

Participants: Samia Boutalbi, Mario Südholt, Remous-Aris Koutsiamanis.

Samia Boutalbi started her PhD in January 2022, on the subject "Secure deployment of micro-services in a shared Cloud RAN/MEC environment", under a co-supervision with the Ericsson team in Paris. She has finalized her PhD manuscript [51] this year and her defense is planned for March/April 2026.

10 Partnerships and cooperations

10.1 International initiatives

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

RANMA - Resource Adaptation and Network Management for continuum computing Applications with the National Institute of Informatics (Tokyo, Japan) from 2025 to 2027. The objective of the RANMA associate team will focus on contributing to the programming support and the resource management of infrastructure and services across the network and the application layers of the Computing Continuum. Led by D. Balouek and K. Piamrat.

10.2 International research visitors

10.2.1 Visits of international scientists

Other international visits to the team

Maurice Djibril Faye

Status (Assistant professor)

Institution of origin: Université numérique Cheikh Hamidou KANE item[Country:Senegal]

Dates: Dec 1st to Dec 21st

Context of the visit: Natural disasters, conflicts, and public emergencies in Africa often disrupt communication networks, delaying responses, and putting lives at risk. Emergency data is also vulnerable to tampering and interception. The goal of this visit is to explore research themes, and possibilities of collaboration for rapid, resilient, and trustworthy emergency communications.

Mobility program/type of mobility: Research stay

Theodoros TSIOLAKIS

Status: (PhD candidate)

Institution of origin: Democritus University of Thrace item[Country:] Greece

Dates: 12/05/2025 - 21/05/2025

Context of the visit: Energy efficiency of machine learning implementations within geo-distributed systems.

Mobility program/type of mobility: ERASMUS+ Short Term Mobility

10.3 European initiatives

10.3.1 Horizon Europe

SLICES-PP / SLICES-RI (Scientific Large Scale Infrastructure for Computing/Communication Experimental Studies - Research Instrument)

Participants: Adrien Lebre, Remous-Aris Koutsiamanis.

The STACK Research team is actively involved in SLICES-RI, the European initiative aimed at developing a large-scale research infrastructure for digital sciences. This project brings together partners from 16 countries, including France, Greece, Poland, Norway, the Netherlands, Italy, and Switzerland, all committed to contributing resources and expertise.

Members of our team are currently actively participating in the Preparatory Phase (SLICES-PP) of the project, helping to shape its direction and implementation. Our goal is to continue this involvement at both the European level and the national level via SLICES-FR, contributing to the creation of a flexible platform that supports large-scale experimental research in areas like networking protocols, radio technologies, and cloud and edge computing architectures.

Once established, SLICES-RI aims provide researchers across Europe with long-term access to advanced computing, storage, and network resources. This infrastructure will facilitate experiments in various domains, including information theory, networking, distributed systems, and software engineering, thereby enabling experimentation and innovation in multiple domains such as smart cities, e-health, industrial internet, transport, and energy solutions.

SEED: Training the next generation of technological scientists to achieve Societal, Energy, Environmental, industrial and Digital transitions

Participants: Mario Südholt (*Coordinator*).

SEED, which stands for Societal, Energy, Environmental, industrial and Digital transitions, is a 60-month interdisciplinary, international and intersectoral doctoral training programme offered by IMT Atlantique and co-funded by the European Union. Its overall budget, managed entirely by IMT Atlantique, amounts to 8 M€. The programme itself is designed to nurture four key dimensions: thesis interdisciplinarity, internationality, cross-sector experience, and promotion of innovation. It offers 40 fully funded early-stage researcher (ESR) positions within three different tracks. Each track builds on the same fundamental excellence trainings implementing a 4i approach (Interdisciplinarity, Internationality, Intersectorial, Innovation), while providing a different degree of mobility and focus. Further information at: [SEED website](#).

10.3.2 Other european programs/initiatives

DI4SPDS (Distributed Intelligence for Enhancing Security and Privacy of Decentralised and Distributed Systems)

Participants: Housseem Jmal, Kandaraj Piamrat.

Decentralised systems face challenges from sophisticated cyber-attacks that evolve and propagate to disrupt different parts. Additionally, communication overhead makes implementing authentication and access control complex. Existing approaches unlikely provide effective access control and multi-stage attack detection due to limited event capture and information correlation. This project offers a framework to improve security and privacy of decentralised systems through cross-domain access control, collaborative intrusion detection, and dynamic risk management considering resource consumption. It facilitates subsystem collaboration to

prevent widespread disruption from attacks and share threat awareness. The project will develop methods and prototypes utilizing blockchain, federated learning, and multi-agent architecture to enhance access control, detection, risk management, and response capabilities. The consortium is composed of four partners: Nantes University, LUT University (Finland), Universidad de Castilla - La Mancha (Spain), and Firat University (Turkey). LUT is the coordinator of the project DI4SPDS has been accepted in July 2023, started from March 2024 for 36 months, with an allocated budget of 874k€ (230K€ for Stack).

DISCOVER-US: Collaboration with NSF on fundamental research on new concepts for distributed computing and swarm intelligence

Participants: Daniel Balouek.

DISCOVER-US represents a collaboration between EU and US research institutions, focusing on advancing distributed computing and swarm intelligence. At the core of DISCOVER-US is the commitment to develop a practical infrastructure that supports a dynamic research ecosystem. The stack team has significantly contributed to a vision white paper highlighting requirements and potential approaches for next-generation distributed computing. This work is aligned with key objectives from Horizon Europe and the US National Science Foundation, driving progress in cloud-to-edge processing technologies, AI, and cybersecurity.

10.4 National initiatives

10.4.1 ANR

NET4AI (Network Acceleration for Generative Artificial Intelligence)

Participants: Remous-Aris Koutsiamanis (*coordinator*), Mario Südholt.

Generative AI now relies on very large foundation models whose training and inference stretch compute clusters and—crucially—the interconnect fabric. At scale, collective communications, congestion, and failures/stragglers limit Model FLOPs Utilization (MFU) far more than raw FLOPs. NET4AI addresses this by making the AI fabric (network + compute) explicitly observable and controllable for GenAI workloads, so that orchestration and communications adapt to load, faults, and topology rather than treating the network as a black box. The project pursues three objectives: (1) *High-accuracy AI fabric awareness* via advanced monitoring and workload–infrastructure co-observation to feed scheduling and resilience decisions; (2) *Optimized traffic scheduling algorithms*: design of efficient collectives tightly integrated with adaptive routing, traffic control, fast failover, packet scheduling, and network/compute-aware orchestration to improve completion time, stability, reliability and fairness; (3) *Large-scale evaluation and experimentation* with new metrics and tools that assess end-to-end network and power efficiency, combining scalable simulation and cutting-edge testbeds (SLICES-RI). The consortium is composed of five partners: Laboratoire CEDRIC (CNAM), Laboratoire Informatique d'Avignon Université (AU), LS2N - IMT Atlantique (STACK team), Scalnyx, and Huawei. NET4AI started in April 2025 for a period of 48 months, is supported by ANR under AAPG 2024 (PRCE), and has an overall budget of 934K€ (203K€ for STACK). See the [NET4AI](#) web site for more information.

For-CoaLa (Formalization of Configuration Languages)

Participants: Hélène Coullon (*coordinator*), Olivia Proust.

Large distributed software systems (applications or infrastructures) are now ubiquitous, with component-based systems (e.g., service-oriented architectures or microservices) offering a convenient way to structure large systems, in particular distributed systems deployed in the Cloud, in the core, or at the edge of the network.

DevOps operations, that include system configurations and reconfigurations, are required to handle various kinds of scenarios such as fault tolerance, scalability, software updates, or various optimizations, etc. Such changes may lead to faults. A study of 597 unplanned outages that affected popular cloud services between 2009 and 2015 found that 16% of them were caused by a system upgrade.

On the one hand, many configuration tools and languages exist in the DevOps community, some of them being specific to the provisioning of resources in Cloud providers, packaging problems, containerized deployments, configuration of applications or infrastructures, etc. The main advantage of these tools is their full integration and adoption in the DevOps community. Their disadvantage is they lack both formal and textual specifications. Moreover, their contours are blurred. On the other hand, many initiatives have been studied in academia to contribute to the deployment, configuration, and reconfiguration of distributed software, bringing improvements such as expressivity, speed, safety, etc. Many come with precise and sometimes formal definitions. However, they lack the breadth of the mainstream DevOps tools.

The goal of For-CoaLa is twofold: (1) understand and bridge the gap between a popular tool from the DevOps community (Ansible) and a tool from academia (Concerto); (2) improve the understanding of these languages based on mechanized formal semantics and develop verified semantic-preserving cross-language transformations. See the [For-CoaLa](#) web site for more information.

SeMaFoR (Self-Management of Fog Resources)

Participants: Thomas Ledoux (*coordinator*), H el ene Coullon, Matthieu Rakotojaona Rainimangavelo.

Fog Computing is a paradigm designed to decentralize cloud infrastructure by extending computing and storage resources to the network’s edge, enabling their geographic distribution along with associated services. The SeMaFoR project aims to model, design, and implement a generic, decentralized solution for the self-management of Fog resources. The consortium comprises three partners: LS2N-IMT Atlantique (Stack, NaoMod, TASC), LIP6-Sorbonne University (Delys), and Alterway/Smile (SME). The Stack team oversees the project with an allocated budget of €506,000, of which €230,000 is for Stack.

SeMaFoR, which began in March 2021, concluded in August 2025. The highlight of the year was the SeMaFoR project winning the 2025 Innovation Trophy for the digital sector at the PUI University of Nantes.

For more information, visit the SeMaFoR website: [Semafor](#).

PicNic (Transfert de grands volumes de donn ees entre datacenters)

Participants: Jean-Marc Menaud (*STACK representative*), Remous-Aris Koutsiamanis, Adrien Lebre, Abdou Seck, Guillaume Rosinosky.

Large dataset transfer from one datacenter to another is still an open issue. Currently, the most efficient solution is the exchange of a hard drive with an express carrier, as proposed by Amazon with its SnowBall offer. Recent evolutions regarding datacenter interconnects announce bandwidths from 100 to 400 Gb/s. The contention point is not the network anymore, but the applications which centralize data transfers and do not exploit parallelism capacities from datacenters which include many servers (and especially many network interfaces – NIC). The PicNic project addresses this issue by allowing applications to exploit network cards available in a datacenter, remotely, in order to optimize transfers (hence the acronym PicNic). The objective is to design a set of system services for massive data transfer between datacenters, exploiting distribution and parallelisation of networks flows.

The consortium is composed of several partners: Laboratoire d’Informatique du Parall elisme, Institut de Canc erologie de l’Ouest / Informatique, Institut de Recherche en Informatique de Toulouse, Laboratoire des Sciences du Num erique de Nantes, Laboratoire d’Informatique de Grenoble, and Nutanix France.

PicNiC will be running for 42 months (starting in Sept 2021 with an allocated budget of 495k , 170k  for STACK).

Taranis

Participants: Daniel Balouek, Helene Coullon (*STACK Representative*), Adrien Lebre, Thomas Ledoux, Jacques Noyé, Eloi Perdereau, Nathan Rabier, Guillaume Rosinosky, Gaëtan Plisson.

New infrastructures, such as Edge Computing or the Cloud-Edge-IoT computing continuum, make cloud issues even more complex, as they add new challenges linked to the diversity and heterogeneity of resources (from small sensors to data centers/HPCs, from low-power networks to core networks), geographical distribution, as well as increased requirements for dynamicity and security, all under constraints such as energy consumption.

To exploit these new infrastructures efficiently, the Taranis project is based on a strategy aimed at abstracting the description of the structure of applications and resources in order to automate their management even further. In this way, it will be possible to globally optimize the resources used with regard to multi-criteria objectives (price, deadline, performance, energy, etc.) on both the user side (applications) and the resource provider side (infrastructures). Taranis also addresses the challenges of abstracting application reconfiguration and dynamically adapting resource usage.

The consortium is composed of 6 partners (Inria, CNRS, IMT, University of Grenoble Alpes, CEA and University of Rennes) for a budget of 7.2M€ (470K€ for STACK) overall.

Spirec

Participants: Mario Südholt (*coordinator*), Remous-Aris Koutsiamanis (*STACK Representative*), Kandaraj Piamrat, Carlos Gonzalez, Jean-Marc Menaud.

The SPIREC project will meet the challenges of supervising services of the Cloud-Edge-IoT continuum, detecting their execution anomalies and predicting their resource usage. The project aims to define methods and techniques, notably using distributed machine learning, to enable its efficient management, provide means to secure them and, more generally, ensure a variety of quality of service properties. The partners will also develop software components and tools in order to integrate these functionalities in existing infrastructures and applications, in particular SLICES, industrial systems and future software ecosystems.

The consortium is composed of 6 partners (Inria, CNRS, CEA, IMT Atlantique, Télécom SudParis and Université de Versailles, Saint Quentin) for a budget of 2,7M€ (580K€ for STACK) overall.

Steel

Participants: Daniel Balouek, Carlos Gonzalez (*STACK Representative*), Jeddou Sidna, Guillaume Rosinosky, Adrien Lebre.

The STEEL project aims to provide solutions for efficient and secure data storage and processing on cloud-based infrastructures. The consortium is composed of 8 partners (Inria, CNRS, University of Grenoble, University of Bordeaux, University of Rennes, IMT Atlantique, IMT TeraLab and IN2P3) for a budget of 3 millions € (300K€ for STACK) overall. activities are organized around three technical work packages. A fourth work package is dedicated to management, communication and dissemination of results. STACK members are involved in the second wp, addressing the challenges related to the management of data sets in presence of node failures and network partitions.

CareCloud

Participants: Jean-Marc Menaud (*STACK representative*), Remous-Aris Koutsiamanis, Thomas Ledoux.

At a time when climate change is a growing concern, with serious consequences for people and the planet worldwide, all sectors (transport, construction, agriculture, industry, etc.) must contribute to the effort to reduce GHG emissions. Clouds, despite their ability to optimize processes in other sectors, are no exception to this observation: the increasing slope of their GHG emissions must be reversed, or their potential benefits in other sectors will be wiped out. The CARECloud project aims to drastically reduce the environmental impact of cloud infrastructures. The consortium is composed of 4 partners (CNRS, IMT Atlantique, Inria and Univeristy of Paul Sabatier - Toulouse), for a budget of 5.5M€ (600k€ for STACK) overall.

Jérémy Woirhaye began his PhD in October 2025, under a joint supervision between the Spirals team (Romain Rouvoy) and Thomas Ledoux, on the topic "Towards an intelligent slicing of microservices applications in the cloud".

SILECS

Participants: Baptiste Jonglez, Remous-Aris Koutsiamanis (*STACK Representative*), Adrien Lebre, Jean-Marc Menaud.

Digital transformation relies on a sophisticated infrastructure of networks, computing and services. The availability, reliability, performance, interoperability and energy efficiency of these systems are major challenges that the digital sciences must meet to foster innovation, sovereignty and industrial competitiveness.

SILECS, the Cloud/Fog/Edge/IoT part of the SLICES-FR platform, enables prototyping and reproducible experiments at all levels of the Cloud IoT continuum. It meets the experimental needs of researchers in networks, systems, telecoms, IoT and other fields. The main objective of SILECS and SLICES-FR is to build a tool for experimentation that fosters the design of new services and applications in distributed computing, edge computing, reprogrammable wired or wireless networks and IoT, using a diversity of technologies on all aspects of the data chain, software or hardware, to meet the needs of the community.

The consortium is composed of 3 partners (Inria, CNRS, and IMT) for a budget of 12 millions € (272k€ for STACK) overall .

10.4.2 PIA 4

OTPaaS

Participants: Farid Arfi, Daniel Balouek, Hélène Coullon, Marie Delavergne, Tayeb Diab, Mohamed Graiet, Housseem Jmal, Sidi Mohammed Kadour, Remous-Aris Koutsiamanis, Adrien Lebre (*STACK representative, until Feb 2024*), Thomas Ledoux, Jean-Marc Menaud, Anas Mokhtari, Jacques Noyé (*STACK representative, from Feb 2024*), Kandaraj Piamrat, Eloi Perdereau, Matthieu Rakotojaona Rainimangavelo, Mario Südholt.

The OTPaaS project targeted the design and development of a complete software stack to administrate and use edge infrastructures for the industry sector. The consortium brought together national and user technology suppliers from major groups (Atos / Bull, Schneider Electric, Valeo) and SMEs / ETIs (Agileo Automation, Mydatamodels, Duplprint, Solem, Tridimeo, Prosyst, Soben), with a strong support from major French research institutes (CEA, Inria, IMT, CAPTRONIC). The project started in October 2021 for a period of 36 months with an overall budget of 56M€ (1.2M€ for STACK) and was extended until April 2025 (with very limited involvement of STACK in 2025).

The OTPaaS platform objectives were:

- To be built on national and sovereign technologies for the edge cloud.
- To be validated by industrial demonstrators of multisectoral use cases.
- To be followed and supported by ambitious industrialization programs.
- To be accompanied by a massive campaign to promote its use by SMEs / midcaps.

- To integrate solutions for controlling energy consumption.
- To be compliant with the Gaia-X ecosystem.

10.4.3 CPER

SAMURAI

Participants: Jean-Marc Menaud (*coordinator*), Remous-Aris Koutsiamanis.

The SAMURAI (Sustainable And autoMoUs gReen computing for AI) infrastructure aims to design an innovative hardware infrastructure for the scientific study of the cross-cutting issues of computing infrastructures supporting artificial intelligence and their energy autonomies.

This project paves the way toward a larger infrastructure at the national level in the context of the SLICES-FR initiative.

The project started in 2022 for a period of 5 years with an overall budget of 730K€ (500K€ for STACK).

10.4.4 Local and regional projects

SysMics network

Participants: Mario Südholt.

SysMics is an integrated cluster of research that is part of the Nantes Excellence Initiative in Medicine and Engineering. Its main objective is the development of new methods for precision medicine, in particular, based on genomic analyses. In this context, we have worked on new large-scale distributed biomedical analyses and provided several results on how to distributed popular statistical analyses, such as FAMD-based and EM-based analyses.

10.4.5 Inria Challenges

FrugalCloud (Inria-OVHCloud)

Participants: H el ene Coullon, Thomas Ledoux.

A joint collaboration between Inria and OVHcloud company on the challenge of frugal cloud has been launched in October 2021 with a budget of 2 M€. It addresses several scientific challenges on the eco-design of cloud frameworks and services for large scale energy and environmental impact reduction, across three axes: i) Software eco-design of services and applications; ii) Efficiency leverages; iii) Reducing the impact and supporting users of the Cloud.

The main activities of the Stack team were at the start of the project between 2021 and 2024 (e.g., PhD thesis of Pierre Jacquet).

CUPSELI (Inria-Hivenet)

Participants: Guillaume Rosinosky, Mario Südholt, Lomig Piette.

The d efi CUPSELI is a joint collaboration between INRIA and the Hivenet company. Its aim is to push the limits of distributed AI computing. Its goal 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. The project is structured amongst three axes: frugality, security and confidentiality, and volatility.

For now, the main activities of the STACK team on this challenge is on the security and confidentiality axis (PhD thesis of Lomig Piette).

Pulse (Inria-Qarnot Computing)

Participants: Adrien Lebre.

The joint challenge between Inria and Qarnot computing is called PULSE, for "PUshing Low-carbon Services towards the Edge". It aims to develop and promote best practices in geo-repaired hardware and software infrastructures for more environmentally friendly intensive computing.

The challenge is structured around two complementary research axes to address this technological and environmental issue:

Axis 1: "Holistic analysis of the environmental impact of intensive computing".

Axis 2: "Implementing more virtuous edge services"

The STACK group is mainly involved in the second axis, addressing the challenges related to data management.

10.4.6 Inria transfer projects (ADT)

CoAnsible

Participants: H el ene Coullon, Baptiste Jonglez, Sidi-Mohammed Kaddour.

The transfer actions CoAnsible is a two year project to develop an Ansible extension able to use the Concerto language [56] as an efficient backend tasks coordinator to handle the lifecycle of infrastructure resources. A 24-months software engineer is hired starting in December 2024.

10.5 Regional initiatives

10.5.1 A Cloud-Edge-IoT continuum site for SLICES-FR

Participants: Remous-Aris Koutsiamanis, Jean-Marc Menaud, Baptiste Jonglez.

The SLICES-FR site in Nantes, comprising the STACK Research team and the LS2N and IETR UMRs, is advancing the development of a Cloud-Edge-IoT continuum-supporting site. This infrastructure envisions to enable seamless, low-latency experimentation across the continuum, contributing a regional site to the national SLICES-FR and European SLICES-RI platforms.

The current efforts are focused on integrating new hardware and collaborating within the architecture and technical teams to align with the project's vision, funded by multiple national projects for funding hardware acquisition and providing engineering resources for management, integration, and operations, ensuring the sites's robustness and readiness for experimentation.

11 Dissemination

11.1 Promoting scientific activities

11.1.1 Scientific events: organisation

- H el ene Coullon: Program co-chair of UCC 2025, workshops co-chair of CCGrid 2025.

- Guillaume Rosinosky: Workshop co-chair of UCC/BDCAT 2025, INSPIRE 2025 workshop co-located with UCC/BDCAT 2025.
- Kandaraj Piamrat: Program co-chair of BDCAT 2025.
- Remous-Aris Koutsiamanis: Tutorial co-chair of UCC 2025.

General chair, scientific chair, steering committees

- Daniel Balouek: General co-chair of UCC/BDCAT 2025, QUICK 2025, workshop co-located with CCGrid, publicity chair of IC2E 2025.

Member of the steering committees Adrien Lebre IEEE ICFEC (International Conference on Fod and Edge Computing)

11.1.2 Scientific events: selection

Chair of conference program committees

- Carlos Gonzalez: MedComNet 2025, Session 5: Special Topics in Networked Applications

Member of the conference program committees

- Daniel Balouek: DAIS 2025, IC2E 2025, UCC 2025, HPDC 2025, PAISE 2025, CCGRID 2025, SAC 2025.
- Thomas Ledoux: IEEE Cloud 2025, UCC 2025.
- Adrien Lebre: ICFEC 2025, UCC 2025.
- Kandaraj Piamrat: IEEE CCNC 2025, IEEE ICCE 2025, IEEE ICC 2025, IEEE WiMob 2025, IEEE VCC 2025, IEEE Globecom 2025 (SAC-CECN, SAC-BD), IEEE BDCAT 2025.
- Guillaume Rosinosky: Dais 2025, Middleware 2025.
- Remous-Aris Koutsiamanis: UCC 2025, IEEE CSCN 2025, IEEE Globecom 2025 CISS, CoRes 2025.
- Carlos Gonzalez : UCC/BDCAT2025.
- Mario Südholt : ICISSP 2022-25

11.1.3 Journal

Member of the editorial boards

- Mario Südholt is a member of the advisory board of "The Programming Journal"

Reviewer - reviewing activities

- Kandaraj Piamrat: IEEE Transactions on Network and Service Management (TNSM), Future Generation Computer Systems (FGCS), Elsevier Computer Communication (COMCOM)
- Carlos Gonzalez: IEEE Access, IEEE Sensors Letters, Future Internet, Sensors, Electronics, International Journal of Communication Systems
- Daniel Balouek: IEEE Internet Computing, Future Generation Computer Systems
- Remous-Aris Koutsiamanis: MDPI Machine Learning and Knowledge Extraction, Elsevier Expert Systems With Applications, Elsevier Ad Hoc Networks, IEEE Transactions on Network and Service Management
- Guillaume Rosinosky: International Journal of Information Security, Future Generation Computer Systems

11.1.4 Invited talks

- Hélène Coullon: Invited "wide-"speaker at DisCoTeC 2025 (20th International Federated Conference on Distributed Computing Techniques). Challenges in Infrastructure-as-Code: efficiency, decentralization, and formalization.
- Baptiste Jonglez: Invited tutorials on EnOSlib at DAIS 2025 and the SLICES-FR summer school 2025.
- Adrien Lebre: Invited talk on “Le programme Systeme, Réseau et Cloud (SRC) et la continuité numérique”, academies des technologies, inter agences du numérique”Academie des technologies, avril 2025.
- Adrien Lebre: Invited talk on “Digital Computing Continuum, the shared vision of two French institutes: Inria and CEA”, DG Connect (Future Network), DG Connect, Brussels, June 2025.
- Adrien Lebre: Invited talk on “La continuité de la connectivité : le projet inter agences du numérique” CEA workshop, Dec 2025.
- Mario Südholt: Invited talk on the declarative definition and effective implementation of transfer learning algorithms at the Symposium on Systems Science at University of Monastir, Tunisia, 2 Jul. 2025

11.1.5 Scientific expertise

- Remous-Aris Koutsiamanis: SLICES-FR Summer School 2025: Member of the Program committee, Session chair of session "L'expérimentation en question entre complexité et multiculturalisme", and Panel member of session "Construisons ensemble Slices-FR"
- Adrien Lebre: French representative for the discussion on the possibility of a new European IPCEI related to the computing continuum, April 2025.
- Mario Südholt: Expert participating in the selection process of the Tiburtius-Preis of the best MSc-thesis of all universities of the Bundesland Berlin, Germany.

11.1.6 Research administration

- Kandaraj Piamrat is the international coordinator of the LS2N (Laboratoire des sciences du numérique de Nantes).
- Hélène Coullon is the vice-president of the ACM SigOps France
- Hélène Coullon is co-chair of the working group YODA (trustworthyY and Optimal Dynamic Adaptation) in the national GDR GPL (software engineering and languages)
- Remous-Aris Koutsiamanis is co-chair of the IETF ROLL (Routing Over Low power and Lossy networks) Working Group
- Remous-Aris Koutsiamanis : Member of the executive committee of the Grid'5000 GIS (Groupement d'intérêt scientifique), Member of the SLICES-FR temporary executive board, Member of the SLICES-FR Architects committee, Member of the SLICES-FR Site managers, Scientific representative of IMT to SLICES.
- Adrien Lebre: Member of the executive committee of the Grid'5000 GIS (Groupement d'intérêt scientifique), Member of the SLICES-FR temporary executive board, Scientific French representative of the SLICES ISB, Co-director of the <I/O> Lab, a joint lab between Inria and Orange Labs, co-leader the French CLOUD PEPR, and since December 2023, head of the Cloud, Network and System program at the French agency "Algorithmes, Logiciels et Usages" operated by Inria.

11.2 Teaching - Supervision - Juries - Educational and pedagogical outreach

As a team mainly composed of Associate Prof. and Prof., the amount of teaching activities is significant (around 150hours per person). We present here only the management activities.

- Hélène Coullon: responsible for the LOGIN training in computer science at IMT Atlantique.
- Guillaume Rosinsky: head of the [apprenticeship program in Software Engineering FIL](#). This 3-year program leads to the award of a Master degree in Software Engineering from the IMT Atlantique.
- Thomas Ledoux: member of the teaching committee at IMT Atlantique
- Thomas Ledoux: co-pilot of the core curriculum of the 1st year of the IT program at IMT Atlantique
- Thomas Ledoux: head of the Filière informatique nantaise since Sept. 2020. This entity, created by Nantes University, Centrale Nantes and IMT Atlantique, aims to bring together the main players in Computer Science training in Nantes to ensure a coherent and ambitious training offer that meets the present and future challenges of Computer Science. It is organized around a Council made up of representatives from the academic and socio-economic worlds.
- Thomas Ledoux : member of the board of directors of [Talents du numérique](#).
- Jacques Noyé: deputy head of the Automation, Production and Computer Sciences Department of IMT Atlantique.
- Mario Südholt: representative for MSc-level and PhD-level studies of the API department of IMT Atlantique.
- Mario Südholt: coordinator of the international PhD program SEED (see [10](#)).
- Kandaraj Piamrat : elected member of scientific council at Faculty of Sciences and Techniques, Nantes University

11.2.1 Supervision

- PhD in progress: Olivia Proust, Towards formally verified configuration management languages, Sept. 2025 - Aug. 2028, Director: Hélène Coullon
- PhD in progress: Simon Artus, Intelligent and dynamic adaptation of a Cloud2IoT service infrastructure, May 2025 - April 2028, Director: Hélène Coullon
- PhD in progress: Nathan Rabier, Handling dynamic constraints and deadlines in distributed software reconfiguration - Application to power transmission, Nov. 2025 - Oct. 2028, Director: Hélène Coullon
- PhD in progress: Lucien Astié, Resilience of jointly managed digital service infrastructures, Oct. 2025 - Sept. 2028, Director: Hélène Coullon, Advisor: Baptiste Jonglez
- PhD defended: Duc Thinh Ngo, Dynamic graph learning algorithms for the digital twin in edge-cloud continuum, Dec. 2022 - Nov. 2025, Director: A. Lebre, Advisor: K. Piamrat
- PhD in progress: Housseem Jmal, Federated Learning for Enhancing Security and Privacy of Decentralized and Distributed Systems, Apr. 2024 - Mar. 2027, Director: JM. Menaud, Advisor: K. Piamrat
- PhD in progress: Martin Molli, Decision models for the Edge-Cloud Computing Continuum, Nov. 2024 - Sept. 2027, Director: T. Ledoux, Advisor: D. Balouek
- PhD in progress: Lylian Siffre, Impacts and Uses of Local-First Software for Energy Optimization of IT Services, Nov. 2024 - Sept. 2027, Director: T. Ledoux

- PhD defended: Antoine Omond, Safe, efficient and low-energy self-adaptation for Cyber Physical Systems - Application to a scientific observatory in the Arctic tundra, Dec. 2021 - May 2025, Director: H el ene Coullon
- PhD change of team/direction: Tengfei An, Modeling and Studying Self-Stabilization within Kubernetes, Oct. 2023 - Sept. 2026, Previous Director: A. Lebre, Previous Advisor: H. Coullon, J. Noy e, New Laboratory: LIP6
- PhD in progress: Abdou Seck, Parallel transfer service for the exchange of large volumes of data between Datacenters, Jun. 2022 - Jun. 2025, Director: J.-M. Menaud, Noel De Palma, Advisor: R.-A. Koutsiamanis.
- PhD in progress: Samia Boutalbi, Secure deployment of microservices in a shared RAN/MEC Cloud environment, Jan 2022 -, Director: M. S udholt, M. Dammak, Advisor: R.-A. Koutsiamanis.
- PhD in progress: Mohammed Abdrrahim Lahmar, Contracts for Distributed ML-Intensive Systems, Oct. 2024 - Oct. 2027, Directors: M. S udholt, Coen De Rover (VU Brussel), Advisor: R.-A. Koutsiamanis.
- PhD in progress: Irina Samus, Energy-aware actor-based distributed programming, Oct. 25 -, Directors: M. S udholt, Jens Nicolay (VU Brussel).
- PhD in progress: Mouheb Jemai, “: Orchestration intelligente et scalable pour garantir les performances et la fiabilit e des syst emes Cloud Native”, April 2025 - March 2028, Director: A. Lebre, Advisor: F. Baligand (CEA)
- PhD defended: Divi Delacour, Architecture and security of cooperative autonomous systems, Jan. 2022 - Jan. 2025, Director: M. S udholt, Marc Lacoste (Orange).
- PhD defended: Simon Brocard, Premi eres analyses g enomiques de la transplantation pulmonaire, Jan. 22 - Dec. 25, Director: M. S udholt, Sophie Limou (Centrale Nantes), Adrian Tissot (CHU Nantes).
- PhD in progress: Cherif Si Mohammed, Eco-responsible management of data storage, Dec. 2023, Dec. 2026, Director: A. Lebre, Advisor: A. Van Kempen (Qarnot Computing)
- PhD defended: Hiba Awad, A Model-based Approach for Multi-Scale and Dynamic Distributed Systems, Nov. 2021 - Mar. 2025. Director: T. Ledoux.
- PhD in progress: Wedwang Romial Menra, Deployment and updating electronic equipment, Nov. 2024 - Nov. 2027. Director: J.-M. Menaud
- PhD in progress: Severin Bradley Anzie, Decentralised and sustainable optimisation of inter-centre data transfers via the QUIC protocol, May 2025, Director: J.-M. Menaud, Advisor: R.-A. Koutsiamanis.
- PhD in progress: Christophe Dion, (industrial “Cifre” thesis with Orange, since Dec. 2024), Director: A. Lebre
- Celeste Precil Guimapi Guefack, Efficacit e  nerg etique des applications dans le cloud, Jan. 2025, Director: J.-M. Menaud, Advisor: R.-A. Koutsiamanis.
- PhD in progress: Aymene Boucha, Distributed Machine Learning for the Cloud-Edge-IoT (CEI) continuum, Sep. 2025 - Aug. 2028. Director: M. S udholt, Advisor: K. Piamrat, C. Gonzalez.
- PhD in progress: Ga etan Plisson, Transparent service continuity for distributed applications on the edge-cloud continuum, Dec. 2025 - Nov. 2028 Director: A. Lebre, Advisor: G. Rosinosky, D. Balouek
- PhD in progress: Lomig Piette, Privacy-Enabled AI Job Execution on Heterogeneous Consumer Hardware Architectures, Oct. 2025 - Sept. 2028 Director: M. S udholt, Advisor: G. Rosinosky
- PhD in progress: Ahmed RJIBA, Flexible and Decentralised Application Orchestration in Fog Environments, Nov. 2025, Director: N. Parlavantzias (INRIA MAGELLAN), Advisor: R.-A. Koutsiamanis.

- Post-doc: Anas Mokhtari, A Holistic Approach for Designing Carbon-Aware and Energy-Aware Cloud applications, Sept. 2023 - Jan. 2025, Advisor: T. Ledoux, B. Jonglez.
- Post-doc: Sidna Jeddou, in Data Management and Network/Computing Continuum, Sep. 2025 - Aug. 2026, Advisor: A. Lebre, C. Gonzalez.

11.2.2 Juries

- Carlos Gonzalez was a jury member of the PhD of Abdulkadir Dauda, Université de Reims Champagne Ardenne - A Secure Edge Gateway for IoT: An Adaptive Approach for Multi-Application and Multi-Protocol Integration, oct. 2025.
- Remous-Aris Koutsiamanis was a jury member of the PhD of Nikolaos Pavlidis, Democritus University of Thrace, Greece - Optimizing Responsible Decentralized Machine Learning Methods for AI Applications, 3 Nov. 2025.
- Adrien Lebre was a reviewer of the PhD defense of Lise Jolicoeur at University of Bordeaux - Towards secure cluster architectures for HPC workflowse, Library, Toolbox, and Evaluation, Dec. 2025.
- Thomas Ledoux was a reviewer of the PhD defense of Imane Taleb at Rochelle Université - "Analyse de l'efficacité énergétique d'applications basées sur des microservices", june 2025.
- Thomas Ledoux was a reviewer of the PhD defense of Hugo Montfleur at University of Lille - "Concern-Oriented MicroService Architecture: Language, Library, Toolbox, and Evaluation", nov. 2025.
- Kandaraj Piamrat was a reviewer of the PhD defense of Hugo De Oliveira at Université Paris-Saclay - "Reinforcement Learning and Federated Learning-based Multi-Band Assignment for IoT Short Packet Communications", Sept. 2025.
- Kandaraj Piamrat was a reviewer of the PhD defense of Johann HUGON at ENS Lyon - "Pipelines d'extraction de métriques pour la supervision du trafic réseau sous contraintes systèmes", Dec. 2025.
- Kandaraj Piamrat was a reviewer of the PhD defense of Javier ERREA MORENO at EUROCOM - "Reinforcement Learning and Federated Learning-based Multi-Band Assignment for IoT Short Packet Communications", Dec. 2025.
- Kandaraj Piamrat was an examiner of the PhD defense of Bouchra FAKHER at Université de Haute-Alsace - "Enhanced Federated Learning for Intelligent Energy Management in IoT- Enabled Smart Buildings", July 2025.
- Kandaraj Piamrat was an examiner of the PhD defense of Ahcene BOUMHAND at University of Rennes - "Network traffic classification for identifying multi-activity situations in home environments.", Nov. 2025.
- Kandaraj Piamrat was an examiner of the PhD defense of Abdelmounaim BOUROUDI at University of Rennes - "Apprentissage par renforcement profond multi-agents pour l'allocation et la planification des ressources en 6G", Nov. 2025.
- Mario Südholt was a reviewer of the PhD defense of Mathis Manthe at INSA-Lyon - "Federated learning in neuroimage segmentation", Jun. 2025.
- Mario Südholt was an examiner of the PhD defense of Hamza KCHOCK at University of Versailles, Saint Quentin - "Edge AI and 5G Network as Enablers for Immersive Virtual Try-On Applications", Dec. 2025.

11.3 Popularization

11.3.1 Participation in Live events

Thomas Ledoux was a speaker at the Techno conference "Numérique Responsable : Le Déclat pour un Futur Durable" organized by the Pole Images & Réseaux (Nantes, 13/01/2025)

12 Scientific production

12.1 Major publications

- [1] D. De Lacour, M. A. Lacoste, M. Südholt and J. Traoré. ‘Ti-skol: A Modular Federated Learning Framework Supporting Security Countermeasure Composition’. In: *Big Data 2024: IEEE International Conference on Big Data*. Washington DC, United States, 2024, pp. 7860–7869. DOI: [10.1109/BigData62323.2024.10825127](https://doi.org/10.1109/BigData62323.2024.10825127). URL: <https://hal.science/hal-04760901> (cit. on p. 13).
- [2] S. Moreschini, S. Pour, I. Lanese, D. Balouek, J. Bogner, X. Li, F. Pecorelli, J. Soldani, E. Truyen and D. Taibi. ‘AI Techniques in the Microservices Life-Cycle: a Systematic Mapping Study’. In: *Computing* 107.4 (29th Mar. 2025), p. 100. DOI: [10.1007/s00607-025-01432-z](https://doi.org/10.1007/s00607-025-01432-z). URL: <https://hal.science/hal-05196358> (cit. on p. 13).
- [3] D.-T. Ngo, O. Aouedi, K. Piamrat, T. Hassan and P. Raipin. ‘Rethinking Traffic Prediction in Mobile Network Digital Twins: A flexible inductive graph-based learning model for data-scarce scenarios’. In: *Computer Networks* 271 (2025), pp. 1–13. DOI: [10.1016/j.comnet.2025.111569](https://doi.org/10.1016/j.comnet.2025.111569). URL: <https://hal.science/hal-05174800>. In press (cit. on p. 13).

12.2 Publications of the year

International journals

- [4] S. Moreschini, S. Pour, I. Lanese, D. Balouek, J. Bogner, X. Li, F. Pecorelli, J. Soldani, E. Truyen and D. Taibi. ‘AI Techniques in the Microservices Life-Cycle: a Systematic Mapping Study’. In: *Computing* 107.4 (29th Mar. 2025), p. 100. DOI: [10.1007/s00607-025-01432-z](https://doi.org/10.1007/s00607-025-01432-z). URL: <https://hal.science/hal-05196358> (cit. on p. 20).
- [5] D.-T. Ngo, O. Aouedi, K. Piamrat, T. Hassan and P. Raipin. ‘Rethinking Traffic Prediction in Mobile Network Digital Twins: A flexible inductive graph-based learning model for data-scarce scenarios’. In: *Computer Networks* 271 (2025), pp. 1–13. DOI: [10.1016/j.comnet.2025.111569](https://doi.org/10.1016/j.comnet.2025.111569). URL: <https://hal.science/hal-05174800>. In press (cit. on pp. 8, 18).

Invited conferences

- [6] 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>.

International peer-reviewed conferences

- [7] H. Awad, T. Ledoux, H. Bruneliere and J. Rivalan. ‘VeriFogOps: Automated Deployment Tool Selection and CI/CD Pipeline Generation for Verifying Fog Systems at Deployment Time’. In: *SAC ’25: Proceedings of the 40th ACM/SIGAPP Symposium on Applied Computing*. SAC 2025 - 40th ACM/SIGAPP Symposium on Applied Computing. Catania, Italy: ACM, 2025, pp. 1674–1691. DOI: [10.1145/3672608.3707854](https://doi.org/10.1145/3672608.3707854). URL: <https://hal.science/hal-04833623> (cit. on pp. 20, 23).
- [8] J. Caposiena, O. Carrillo, F. Le Mouël, B. Jonglez, P. Neyron and T. Arrabal. ‘Towards a flexible Network Operating System Testbed for the Computing Continuum’. In: *CCGridW 2025 - 25th IEEE International Symposium on Cluster, Cloud and Internet Computing Workshops*. 2025 IEEE 25th International Symposium on Cluster, Cloud and Internet Computing Workshops (CCGridW). Tromsø Norway, Norway, 2025, pp. 148–155. DOI: [10.1109/CCGridW65158.2025.00029](https://doi.org/10.1109/CCGridW65158.2025.00029). URL: <https://hal.science/hal-05154217>.

- [9] L. Carnevale, D. Balouek, S. Sebbio, M. Parashar and M. Villari. ‘Private Distributed Resource Management Data: Predicting CPU Utilization with Bi-LSTM and Federated Learning’. In: *CCGrid 2025 IEEE 25th International Symposium on Cluster, Cloud and Internet Computing*. Tromsø, Norway: IEEE, 2025, pp. 266–275. doi: [10.1109/ccgrid64434.2025.00048](https://doi.org/10.1109/ccgrid64434.2025.00048). URL: <https://hal.science/hal-05196357> (cit. on p. 20).
- [10] D. De Lacour, M. A. Lacoste, M. Südholt and J. Traoré. ‘Towards a Reference Architecture for Secure Cooperative Autonomous Systems for IIoT’. In: *ICIN 2025 - 28th Conference on Innovation in Clouds, Internet and Networks*. Paris, France: IEEE, 2025, pp. 181–185. doi: [10.1109/ICIN64016.2025.10943103](https://doi.org/10.1109/ICIN64016.2025.10943103). URL: <https://hal.science/hal-05023265> (cit. on p. 22).
- [11] T. Diab, M. Graiet and M. Südholt. ‘Declarative, generic definition and effective implementation of transfer learning algorithms’. In: *SPRINGER IFIP AICT Series. AIAI 2025 - 21st International Conference on Artificial Intelligence Applications and Innovations*. Limassol, Cyprus, 26th June 2025, pp. 1–14. URL: <https://hal.science/hal-05005665> (cit. on p. 21).
- [12] C. Gonzalez and S. M. Charfadine. ‘CEI-Net: An open-source framework for P4-driven network management on the CEI Continuum’. In: *MedComNet 2025 - 23rd Mediterranean Communication and Computer Networking Conference*. Cagliari (Sardinia), Italy: IEEE, 2025, pp. 1–5. doi: [10.1109/MedComNet65822.2025.11103524](https://doi.org/10.1109/MedComNet65822.2025.11103524). URL: <https://inria.hal.science/hal-05176967>.
- [13] 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>.
- [14] H. Jmal, K. Piamrat and O. Aouedi. ‘TUNE-FL: adapTive semi-synchronoUs semi-deceNtralizEd Federated Learning’. In: *CCNC 2025 - IEEE Consumer Communications & Networking Conference*. Las Vegas (Nevada), United States: IEEE, 2025, pp. 1–6. URL: <https://hal.science/hal-04752941> (cit. on pp. 8, 22).
- [15] H. Jmal, K. Piamrat, O. Aouedi and Y. Ji. ‘FLAIR: Federated Learning with Adaptive and Intelligent Reasoning for Client Selection’. In: *MSWiM 2025 - 27th International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*. Barcelone (Espagne), Spain, 2025, pp. 1–5. URL: <https://hal.science/hal-05288098> (cit. on pp. 8, 22).
- [16] T. Lagos Jenschke, G. Z. Papadopoulos, R.-A. Koutsiamanis and N. Montavont. ‘ODESe: Amélioration de l’efficacité du réseau grâce aux multi-trajets dans RPL’. In: *CORES 2025 - 10èmes Rencontres Francophones sur la Conception de Protocoles, l’Evaluation de Performances et l’Expérimentation des Réseaux de Communication*. CORES 2025 - 10èmes Rencontres Francophones sur la Conception de Protocoles, l’Evaluation de Performances et l’Expérimentation des Réseaux de Communication. Saint Valéry-sur-Somme, France, 2025, pp. 1–4. URL: <https://hal.science/hal-05003847> (cit. on p. 19).
- [17] M. Molli, D. Balouek, P. Temple and T. Ledoux. ‘Event-Driven Adaptation in the Computing Continuum Using Software Variability’. In: *2025 IEEE/ACM 18th International Conference on Utility and Cloud Computing (UCC '25)*. UCC 2025 - IEEE/ACM 18th International Conference on Utility and Cloud Computing. Nantes, France: IEEE, 1st Dec. 2025, pp. 1–2. doi: [10.1145/3773274.3774672](https://doi.org/10.1145/3773274.3774672). URL: <https://hal.science/hal-05379151> (cit. on p. 20).
- [18] D. de Novi, L. Carnevale, D. Balouek-Thomert, M. Parashar and M. Villari. ‘Predictive Resource Management in the Computing Continuum: Transfer Learning from Virtual Machines to Containers using Transformers’. In: *UCC 2025 - IEEE/ACM 18th International Conference on Utility and Cloud Computing*. Nantes, France: ACM, 2025, pp. 1–8. doi: [10.1145/3773274.3774856](https://doi.org/10.1145/3773274.3774856). URL: <https://hal.science/hal-05468841> (cit. on p. 20).
- [19] D. Schmitz, L. Berrewaerts, G. Rosinosky, S. Skhiri and E. Rivière. ‘A Tale of Many Streams: Characterizing a Hybrid Batch-Stream Production Workload in Digazu, a Data Lake Supported by Apache Kafka and Flink’. In: *DEBS '25: Proceedings of the 19th ACM International Conference on Distributed and Event-based Systems*. DEBS 2025 - 19th ACM International Conference on Distributed and Event-based Systems. Gothenburg, Sweden, 9th June 2025, pp. 188–198. doi: [10.1145/3701717.3734462](https://doi.org/10.1145/3701717.3734462). URL: <https://hal.science/hal-05148331> (cit. on p. 19).

- [20] D. Schmitz, G. Rosinosky and E. Rivi re. ‘Justin: Hybrid CPU/Memory Elastic Scaling for Distributed Stream Processing \star ’. In: *LNCS. DAIS 2025 - 25th International Conference on Distributed Applications and Interoperable Systems*. Vol. 15730. 1. Lille, France, July 2025, pp. 1–17. URL: <https://hal.science/hal-05081993> (cit. on pp. 14, 19).
- [21] S. Sebbio, L. Carnevale, D. Balouek, M. Parashar and M. Villari. ‘Data-Driven Operational Artificial Intelligence for Computing Continuum: a Natural Disaster Management Use Case’. In: *CCGRID 2025 - 25th IEEE International Symposium on Cluster, Cloud, and Internet Computing. 2025 IEEE 25th International Symposium on Cluster, Cloud and Internet Computing Workshops (CCGridW)*. Tromso, Norway: IEEE, 2025, pp. 1–8. DOI: [10.1109/CCGRIDW65158.2025.00022](https://doi.org/10.1109/CCGRIDW65158.2025.00022). URL: <https://hal.science/hal-05196359> (cit. on p. 19).
- [22] C. Si Mohammed, A. Lebre and A. van Kempen. ‘Boosting Task-Driven Applications from Cloud to Edge: Leveraging Utility for Effective Data Replication’. In: *IC2E 2025 - 13th IEEE International Conference on Cloud Engineering*. Rennes, France: IEEE, 2025, pp. 1–9. URL: <https://hal.science/hal-05233840> (cit. on p. 19).
- [23] K. Sidi, D. Balouek and B. Jonglez. ‘Application-level observability for adaptive Edge to Cloud continuum systems’. In: *UCC 2025 - IEEE/ACM 18th International Conference on Utility and Cloud Computing*. NANTES, France, 2025. DOI: [10.1145/3773274.3774855](https://doi.org/10.1145/3773274.3774855). URL: <https://inria.hal.science/hal-05371388> (cit. on p. 20).
- [24] L. Siffre, T. Ledoux, R. Pawlak and J. Guery. ‘Local Computing vs. Cloud Computing: An Empirical Study of Energy Consumption’. In: *IEEE/ACM 18th International Conference on Utility and Cloud Computing*. Nantes, France: ACM, 2025, pp. 1–10. DOI: [10.1145/3773274.3774278](https://doi.org/10.1145/3773274.3774278). URL: <https://hal.science/hal-05324506> (cit. on pp. 21, 23).
- [25] L. Siffre, T. Ledoux, R. Pawlak and J. Guery. ‘Local-First Software for Green IT’. In: *ICT4S 2025 - International Conference on Information and Communications Technology for Sustainability*. Dublin, Ireland, 2025, pp. 1–10. DOI: [10.1109/ICT4S68164.2025.00015](https://doi.org/10.1109/ICT4S68164.2025.00015). URL: <https://hal.science/hal-05037019> (cit. on pp. 21, 23).

Conferences without proceedings

- [26] A. Boucha, C. Gonzalez, K. Piamrat and M. Sudholt. ‘A model for Distributed Machine Learning’. In: *UCC 2025 - IEEE/ACM 18th International Conference on Utility and Cloud Computing*. Nantes, France: ACM, 31st Dec. 2025, pp. 1–2. DOI: [10.1145/3773274.3774671](https://doi.org/10.1145/3773274.3774671). URL: <https://hal.science/hal-05460554> (cit. on p. 22).
- [27] S. Boutalbi, R.-A. Koutsiamanis, M. Dammak and M. S dholt. ‘A New vTPM Architecture with Strong Isolation for the Cloud’. In: *ICA3PP 2025 - 25th International Conference on Algorithms and Architectures for Parallel Processing*. Zhengzhou, China, 2025. URL: <https://hal.science/hal-05287039> (cit. on p. 22).
- [28] J. Caposiena, O. Carrillo, B. Jonglez, P. Neyron and T. Arrabal. ‘Vers un banc d’essai flexible pour les syst mes d’exploitation r seau dans le Computing Continuum’. In: *ComPAS*. Ed. by F. Le Mou l. Bordeaux, France, 24th June 2025. URL: <https://hal.science/hal-05128936>.
- [29] B. Jonglez and L. Asti . ‘From research to Deuxfleurs and back again: towards digital service infrastructure as commons’. In: *UCS 2026 - 2nd conference on Undone Science in Computer Science*. Luxembourg, Luxembourg, 2026. URL: <https://inria.hal.science/hal-05469666>.
- [30] M. Molli, D. Balouek, P. Temple and T. Ledoux. ‘Facilitating Heterogeneity Management on the Computing Continuum’. In: *COMPAS 2025 - Conf rence francophone d’informatique en Parall lisme, Architecture et Syst me*. Bordeaux, France, 2025, pp. 1–8. URL: <https://hal.science/hal-05138189> (cit. on p. 20).
- [31] D.-T. Ngo, K. Piamrat, O. Aouedi, T. Hassan and P. Raipin-Parv dy. ‘Towards Scalable O-RAN Resource Management: Graph-Augmented Proximal Policy Optimization’. In: *IEEE International Symposium on Network Computing and Applications*. Lisbon, Portugal: IEEE, 5th Nov. 2025, pp. 165–173. DOI: [10.1109/NCA67271.2025.00036](https://doi.org/10.1109/NCA67271.2025.00036). URL: <https://hal.science/hal-05228435> (cit. on pp. 8, 18).

- [32] C. Si Mohammed, A. Lebre and A. van Kempen. ‘An analysis for the design of an efficient replica management strategy’. In: COMPAS 2025 - Conférence francophone d’informatique en Parallélisme, Architecture et Système. Bordeaux, France, 2025, pp. 1–7. URL: <https://hal.science/hal-05079577> (cit. on p. 19).

Doctoral dissertations and habilitation theses

- [33] H. Awad. ‘Quality of service assurance before deployment of fog systems with model-based engineering and DevOps’. Ecole nationale supérieure Mines-Télécom Atlantique, 17th Mar. 2025. URL: <https://theses.hal.science/tel-05056618> (cit. on pp. 20, 23).
- [34] H. Coullon. ‘Efficient Reconfigurations with Programmable Life Cycles: Contributions to Safety, Declarativity, and Decentralization’. Nantes Université, 23rd Apr. 2025. URL: <https://hal.science/tel-05188709> (cit. on p. 20).
- [35] D. De Lacour. ‘Architecture and security for cooperative autonomous systems’. Ecole nationale supérieure Mines-Télécom Atlantique, 16th June 2025. URL: <https://theses.hal.science/tel-05165765> (cit. on p. 23).
- [36] D.-T. Ngo. ‘Dynamic graph learning algorithms for digital twins of network architectures’. Ecole nationale supérieure Mines-Télécom Atlantique, 27th Nov. 2025. URL: <https://theses.hal.science/tel-05419326> (cit. on p. 23).
- [37] A. Omond. ‘Study of the energy consumption and duration of a cyber-physical system reconfiguration in the Arctic tundra : from experiments on real infrastructure to extensive simulations’. Ecole nationale supérieure Mines-Télécom Atlantique; Universitetet i Tromsø, 23rd May 2025. URL: <https://theses.hal.science/tel-05219003> (cit. on p. 20).
- [38] K. Piamrat. ‘From network management towards network analytics: a decade journey of research study’. Nantes Université, 12th Feb. 2025. URL: <https://hal.science/tel-05002490> (cit. on p. 18).

Reports & preprints

- [39] M. Bacou, D. Beserra, E. Dedu, L. Desgeorges, D. Donsez, A. Guitton, B. Jonglez, A. Legrand, G. Papadopoulos, O. Richard, S. Si-Mohammed, N. Tamdrari and F. Theoleyre. *Journée thématique du GDR RSD : pratiques expérimentales de la communauté systèmes et réseaux*. 31st Jan. 2025. URL: <https://hal.science/hal-04924273>.
- [40] D. De Lacour. *D-ECS: Towards decentralising video games*. 2025. URL: <https://hal.science/hal-04886531>.
- [41] D. De Lacour and A. Gautier. *ReDOSN a customizable Decentralised Social Network*. 2025. URL: <https://hal.science/hal-04889747>.
- [42] D. De Lacour, M. Lacoste, M. Südholt and J. Traoré. *Towards Volunteer Deep Learning: Security Challenges and Solutions*. 2025. URL: <https://hal.science/hal-04879559>.

Other scientific publications

- [43] S. Brocard, A. Sylla, N. d. S. B. Silva, P.-a. Gourraud, M. Südholt, N. Vince, S. Limou and C. Consortium. ‘Immunogenomic exploration in a large lung transplantation genetic cohort reveals a chronic lung allograft dysfunction association with HLA-DQB1*03:01’. In: EFI 2025 - 2ème Conférence Internationale sur l’Economie et la Finance Internationale. Prague, Czech Republic, 2025. URL: <https://hal.science/hal-05095900> (cit. on p. 19).
- [44] S. Jeddou, C. Si Mohammed, A. Lebre and C. Gonzalez. *Named Data Networking for Scalable Distributed Storage and Indexation*. 17th Dec. 2025. URL: <https://hal.science/hal-05421548> (cit. on p. 19).

12.3 Cited publications

- [45] N. Ayeb, É. Rutten, S. Bolle, T. Coupaye and M. Douet. ‘Coordinated autonomic loops for target identification, load and error-aware Device Management for the IoT’. In: *Proceedings of the 2020 Federated Conference on Computer Science and Information Systems, FedCSIS 2020, Sofia, Bulgaria, September 6-9, 2020*. Ed. by M. Ganzha, L. A. Maciaszek and M. Paprzycki. Vol. 21. Annals of Computer Science and Information Systems. 2020, pp. 491–500. DOI: [10.15439/2020F154](https://doi.org/10.15439/2020F154). URL: <https://doi.org/10.15439/2020F154> (cit. on p. 8).
- [46] N. Ayeb, É. Rutten, S. Bolle, T. Coupaye and M. Douet. ‘Towards an Autonomic and Distributed Device Management for the Internet of Things’. In: *IEEE 4th International Workshops on Foundations and Applications of Self* Systems, FAS*W@SASO/ICCAC 2019, Umea, Sweden, June 16-20, 2019*. IEEE, 2019, pp. 246–248. DOI: [10.1109/FAS-W.2019.00065](https://doi.org/10.1109/FAS-W.2019.00065). URL: <https://doi.org/10.1109/FAS-W.2019.00065> (cit. on p. 8).
- [47] N. Bakerally, C. Bareau, F. Blache, S. Bolle, C. Ecrepont, P. Folz, N. Hernandez, T. Monteil, G. Privat and F. Ramparany. ‘Semi-automatic RDFization Using Automatically Generated Mappings’. In: *The Semantic Web: ESWC 2020 Satellite Events - ESWC 2020 Satellite Events, Heraklion, Crete, Greece, May 31 - June 4, 2020, Revised Selected Papers*. Ed. by A. Harth, V. Presutti, R. Troncy, M. Acosta, A. Polleres, J. D. Fernández, J. X. Parreira, O. Hartig, K. Hose and M. Cochez. Vol. 12124. Lecture Notes in Computer Science. Springer, 2020, pp. 25–31. DOI: [10.1007/978-3-030-62327-2_5](https://doi.org/10.1007/978-3-030-62327-2_5). URL: https://doi.org/10.1007/978-3-030-62327-2_5 (cit. on p. 8).
- [48] D. Balouek, A. C. Amarie, G. Charrier, F. Desprez, E. Jeannot, E. Jeanvoine, A. Lèbre, D. Margery, N. Niclausse, L. Nussbaum et al. ‘Adding virtualization capabilities to the Grid’5000 testbed’. In: *International Conference on Cloud Computing and Services Science*. Springer. 2012, pp. 3–20 (cit. on p. 9).
- [49] M. Belair, S. Lanjepce and J.-M. Menaud. ‘SNAPPY: Programmable Kernel-Level Policies for Containers’. In: *SAC 2021: 36th ACM/SIGAPP Symposium On Applied Computing. Gwangju / Virtual, South Korea, Mar. 2021*, pp. 1636–1645. DOI: [10.1145/3412841.3442037](https://hal.inria.fr/hal-03108231). URL: <https://hal.inria.fr/hal-03108231> (cit. on p. 11).
- [50] G. Blair, T. Coupaye and J.-B. Stefani. ‘Component-based architecture: the Fractal initiative’. In: *Annals of telecommunications* 64.1 (Feb. 2009), pp. 1–4. DOI: [10.1007/s12243-009-0086-1](https://doi.org/10.1007/s12243-009-0086-1). URL: <https://doi.org/10.1007/s12243-009-0086-1> (cit. on p. 9).
- [51] S. Boutalbi. ‘Secure deployment of micro-services in a shared Cloud RAN/MEC computing environment’. PhD defense planned for March/April 2026. Dec. 2025 (cit. on p. 24).
- [52] S. Brocard, M. Morin, N. dos Santos Brito Silva, B. Renaud-Picard, B. Coiffard, X. Demant, L. Falque, J. Le Pavec, A. Roux, T. Villeneuve, C. Knoop, J.-F. Mornex, M. Salpin, V. Boussaud, O. Rousseau, V. Mauduit, A. Durand, A. Magnan, P.-A. Gourraud, N. Vince, M. Südholt, A. Tissot and S. Limou. ‘Description and first insights on a large genomic biobank of lung transplantation’. In: *European Journal of Human Genetics* (Aug. 2024), pp. 1–8. DOI: [10.1038/s41431-024-01683-y](https://inserm.hal.science/inserm-04712234). URL: <https://inserm.hal.science/inserm-04712234> (cit. on p. 19).
- [53] H. Bruneliere, Z. Al-Shara, F. Alvares, J. Lejeune and T. Ledoux. ‘A Model-based Architecture for Autonomic and Heterogeneous Cloud Systems’. In: *CLOSER 2018 - 8th International Conference on Cloud Computing and Services Science*. Vol. 1. Best Paper Award. Funchal, Portugal, Mar. 2018, pp. 201–212. DOI: [10.5220/0006773002010212](https://hal.archives-ouvertes.fr/hal-01705248). URL: <https://hal.archives-ouvertes.fr/hal-01705248> (cit. on p. 8).
- [54] A. Bruniaux, R.-A. Koutsiamanis, G. Z. Papadopoulos and N. Montavont. ‘Defragmenting the 6LoWPAN Fragmentation Landscape: A Performance Evaluation’. In: *Sensors* 21.5 (2021), p. 1711 (cit. on p. 8).
- [55] M. Chardet, H. Coullon, D. Pertin and C. Pérez. ‘Madeus: a formal deployment model’. In: *2018 International Conference on High Performance Computing & Simulation (HPCS)*. IEEE. 2018, pp. 724–731 (cit. on p. 8).

- [56] M. Chardet, H. Coullon and S. Robillard. ‘Toward Safe and Efficient Reconfiguration with Concerto’. In: *Science of Computer Programming* 203 (Mar. 2021), pp. 1–31. doi: [10.1016/j.scico.2020.102582](https://doi.org/10.1016/j.scico.2020.102582). URL: <https://hal.inria.fr/hal-03103714> (cit. on pp. 8, 10, 20, 31).
- [57] M. Charlier, R.-A. Koutsiamanis and B. Quoitin. ‘Scheduling UWB Ranging and Backbone Communications in a Pure Wireless Indoor Positioning System’. In: *IoT* 3.1 (2022), pp. 219–258 (cit. on p. 8).
- [58] 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. 9).
- [59] R.-A. Cherrueau, D. Pertin, A. Simonet, A. Lebre and M. Simonin. ‘Toward a Holistic Framework for Conducting Scientific Evaluations of OpenStack’. In: *Proceedings of the 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*. IEEE Press, 2017, pp. 544–548 (cit. on p. 9).
- [60] M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Martí-Oliet, J. Meseguer and J. F. Quesada. ‘The Maude System’. In: *Proceedings of the 10th International Conference on Rewriting Techniques and Applications*. RtA ’99. Berlin, Heidelberg: Springer-Verlag, 1999 (cit. on p. 20).
- [61] R. D. Cosmo, J. Mauro, S. Zacchiroli and G. Zavattaro. ‘Aeolus: A component model for the cloud’. In: *Information and Computation* 239.Supplement C (2014), pp. 100–121. doi: <https://doi.org/10.1016/j.ic.2014.11.002>. URL: <http://www.sciencedirect.com/science/article/pii/S0890540114001424> (cit. on p. 20).
- [62] H. Coullon, J. Bigot and C. Pérez. ‘Extensibility and Composability of a Multi-Stencil Domain Specific Framework’. In: *International Journal of Parallel Programming* (Nov. 2017). doi: [10.1007/s10766-017-0539-5](https://doi.org/10.1007/s10766-017-0539-5). URL: <https://doi.org/10.1007/s10766-017-0539-5> (cit. on p. 9).
- [63] H. Coullon, L. Henrio, F. Loulergue and S. Robillard. ‘Component-Based Distributed Software Reconfiguration: a Verification-Oriented Survey’. In: *ACM Computing Surveys* 56.1 (Jan. 2024), pp. 1–37. doi: [10.1145/3595376](https://doi.org/10.1145/3595376). URL: <https://inria.hal.science/hal-04067909> (cit. on p. 9).
- [64] T. Coupaye, S. Bolle, S. Derrien, P. Folz, P. Meye, G. Privat and P. R. Parvedy. ‘A graph-based cross-vertical digital twin platform for complex cyber-physical systems’. In: *The Digital Twin*. Ed. by R. M. Noel Crespi Adam T. Drobot. In press, to appear. Springer, Cham, 2022 (cit. on p. 8).
- [65] T. Coupaye and J. Stefani. ‘Fractal Component-Based Software Engineering’. In: *Object-Oriented Technology, ECOOP 2006 Workshop Reader, ECOOP 2006 Workshops, Nantes, France, July 3-7, 2006, Final Reports*. Ed. by M. Südholt and C. Consel. Vol. 4379. Lecture Notes in Computer Science. Springer, 2006, pp. 117–129. doi: [10.1007/978-3-540-71774-4_13](https://doi.org/10.1007/978-3-540-71774-4_13). URL: https://doi.org/10.1007/978-3-540-71774-4_13 (cit. on p. 9).
- [66] P.-C. David, T. Ledoux, M. Léger and T. Coupaye. ‘FPath and FScript: Language support for navigation and reliable reconfiguration of Fractal architectures’. In: *annals of telecommunications - annales des télécommunications* 64.1 (Feb. 2009), pp. 45–63. doi: [10.1007/s12243-008-0073-y](https://doi.org/10.1007/s12243-008-0073-y). URL: <https://doi.org/10.1007/s12243-008-0073-y> (cit. on p. 9).
- [67] F. A. De Oliveira, T. Ledoux and R. Sharrock. ‘A framework for the coordination of multiple autonomic managers in cloud environments’. In: *Self-Adaptive and Self-Organizing Systems (SASO), 2013 IEEE 7th International Conference on*. IEEE, 2013, pp. 179–188 (cit. on p. 8).
- [68] X. Etchevers, G. Salaün, F. Boyer, T. Coupaye and N. D. Palma. ‘Reliable self-deployment of cloud applications’. In: *Symposium on Applied Computing, SAC 2014, Gyeongju, Republic of Korea - March 24 - 28, 2014*. Ed. by Y. Cho, S. Y. Shin, S. Kim, C. Hung and J. Hong. ACM, 2014, pp. 1331–1338. doi: [10.1145/2554850.2554951](https://doi.org/10.1145/2554850.2554951). URL: <https://doi.org/10.1145/2554850.2554951> (cit. on p. 8).

- [69] X. Etchevers, G. Salaün, F. Boyer, T. Coupaye and N. D. Palma. ‘Reliable self-deployment of distributed cloud applications’. In: *Softw. Pract. Exp.* 47.1 (2017), pp. 3–20. doi: [10.1002/spe.2400](https://doi.org/10.1002/spe.2400). URL: <https://doi.org/10.1002/spe.2400> (cit. on p. 8).
- [70] O. Flauzac, C. Javier Gonzalez Santamaria, F. Nolot and I. Woungang. ‘An SDN approach to route massive data flows of sensor networks’. In: *International Journal of Communication Systems* 33.7 (2020). e4309 dac.4309, e4309. doi: <https://doi.org/10.1002/dac.4309>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/dac.4309>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/dac.4309> (cit. on p. 8).
- [71] V. Gasiūnas, L. Satabin, M. Mezini, A. Núñez and J. Noyé. ‘EScala: Modular Event-Driven Object Interactions in Scala’. In: *Proceedings of the 10th International Conference on Aspect-Oriented Software Development, AOSD 2011*. Porto de Galinhas, Pernambuco, Brazil: ACM Press, Mar. 2011. URL: <http://doi.acm.org/10.1145/1960275.1960303> (cit. on p. 9).
- [72] C. Gonzalez and S. M. Charfadine. ‘SDN Controllers and ML-Based Anomaly Detection in Embedded Systems: A Comparative Analysis’. In: *2023 10th International Conference on Wireless Networks and Mobile Communications (WINCOM)*. 2023, pp. 1–6. doi: [10.1109/WINCOM59760.2023.10322912](https://doi.org/10.1109/WINCOM59760.2023.10322912) (cit. on p. 8).
- [73] C. Gottstein, P. R. Parvedy, M. Hurfin, T. Hassan and T. Coupaye. ‘Inverse Space Filling Curve Partitioning Applied to Wide Area Graphs’. In: *DMS 2020 - 11th International conference on Database Management Systems*. Zurich, Switzerland, Nov. 2020, pp. 223–241. doi: [10.5121/csitr.2020.101417](https://doi.org/10.5121/csitr.2020.101417). URL: <https://hal.inria.fr/hal-03137952> (cit. on p. 8).
- [74] C. Gottstein, P. R. Parvedy, M. Hurfin, T. Hassan and T. Coupaye. ‘Partitioning Wide Area Graphs Using a Space Filling Curve’. In: *International Journal of Data Mining & Knowledge Management Process* 11.1 (Jan. 2021), pp. 13–31. doi: [10.5121/ijdkp.2021.11102](https://doi.org/10.5121/ijdkp.2021.11102). URL: <https://hal.inria.fr/hal-03513456> (cit. on p. 8).
- [75] D. Hauweele, R.-A. Koutsiamanis, B. Quoitin and G. Z. Papadopoulos. ‘Thorough performance evaluation & analysis of the 6TiSCH minimal scheduling function (MSF)’. In: *Journal of Signal Processing Systems* 94.1 (2022), pp. 3–25 (cit. on p. 8).
- [76] F. Hermenier, X. Lorca, J.-M. Menaud, G. Muller and J. Lawall. ‘Entropy: a consolidation manager for clusters’. In: *Proceedings of the 2009 ACM SIGPLAN/SIGOPS international conference on Virtual execution environments*. ACM. 2009, pp. 41–50 (cit. on p. 8).
- [77] T. L. Jenschke, R.-A. Koutsiamanis, G. Z. Papadopoulos and N. Montavont. ‘ODeSe: On-Demand Selection for multi-path RPL networks’. In: *Ad Hoc Networks* 114 (2021), p. 102431 (cit. on p. 8).
- [78] J. Kephart and D. Chess. ‘The vision of autonomic computing’. In: *Computer* 36.1 (2003), pp. 41–50. doi: [10.1109/MC.2003.1160055](https://doi.org/10.1109/MC.2003.1160055) (cit. on p. 9).
- [79] R.-A. Koutsiamanis, G. Z. Papadopoulos, B. Quoitin and N. Montavont. ‘A Centralized Controller for Reliable and Available Wireless Schedules in Industrial Networks’. In: *2020 16th International Conference on Mobility, Sensing and Networking (MSN)*. IEEE. 2020, pp. 1–9 (cit. on p. 8).
- [80] A. Lebre, J. Pastor, A. Simonet and M. Südholt. ‘Putting the Next 500 VM Placement Algorithms to the Acid Test’. In: *IEEE Transactions on Parallel and Distributed Systems* (2018). URL: <https://hal.inria.fr/hal-01816248> (cit. on p. 8).
- [81] M. Léger, T. Ledoux and T. Coupaye. ‘Reliable Dynamic Reconfigurations in a Reflective Component Model’. In: *Proceedings of the 13th International Conference on Component-Based Software Engineering*. CBSE’10. Prague, Czech Republic: Springer-Verlag, 2010, pp. 74–92. doi: [10.1007/978-3-642-13238-4_5](https://doi.org/10.1007/978-3-642-13238-4_5). URL: http://dx.doi.org/10.1007/978-3-642-13238-4_5 (cit. on p. 9).
- [82] J. Lejeune, F. Alvares and T. Ledoux. ‘Towards a Generic Autonomic Model to Manage Cloud Services’. In: *Proceedings of the 7th International Conference on Cloud Computing and Services Science (Best Paper Award)*, INSTICC. SciTePress, 2017, pp. 175–186. doi: [10.5220/0006302801750186](https://doi.org/10.5220/0006302801750186) (cit. on p. 8).

- [83] L. Letondeur, X. Etchevers, T. Coupaye, F. Boyer and N. D. Palma. ‘Architectural Model and Planification Algorithm for the Self-Management of Elastic Cloud Applications’. In: *2014 International Conference on Cloud and Autonomic Computing, London, United Kingdom, September 8-12, 2014*. IEEE Computer Society, 2014, pp. 172–179. doi: [10.1109/ICAC.2014.29](https://doi.org/10.1109/ICAC.2014.29). URL: <https://doi.org/10.1109/ICAC.2014.29> (cit. on p. 8).
- [84] Y. Liu, P. Folz, S. Pan, F. Ramparany, S. Bolle, E. Ballot and T. Coupaye. ‘Digital Twin-Driven Approach for Smart City Logistics: The Case of Freight Parking Management’. In: *Advances in Production Management Systems. Artificial Intelligence for Sustainable and Resilient Production Systems - IFIP WG 5.7 International Conference, APMS 2021, Nantes, France, September 5-9, 2021, Proceedings, Part IV*. Ed. by A. Dolgui, A. Bernard, D. Lemoine, G. von Cieminski and D. Romero. Vol. 633. IFIP Advances in Information and Communication Technology. Springer, 2021, pp. 237–246. doi: [10.1007/978-3-030-85910-7_25](https://doi.org/10.1007/978-3-030-85910-7_25). URL: https://doi.org/10.1007/978-3-030-85910-7_25 (cit. on p. 8).
- [85] A. Omond, H. Coullon, I. Raïs and O. Anshus. ‘Leveraging Relay Nodes to Deploy and Update Services in a CPS with Sleeping Nodes’. In: *Proceeding of the 2023 IEEE International Conferences on Internet of Things (iThings) and IEEE Green Computing & Communications (GreenCom) and IEEE Cyber, Physical & Social Computing (CPSCom) and IEEE Smart Data (SmartData) and IEEE Congress on Cybermatics (Cybermatics)*. Dec. 2023. doi: [10.1109/iThings-GreenCom-CPSCom-SmartData-Cybermatics60724.2023.00102](https://hal.science/hal-04372320). URL: <https://hal.science/hal-04372320> (cit. on p. 20).
- [86] J. Philippe, A. Omond, H. Coullon, C. Prud’Homme and I. Raïs. ‘Fast Choreography of Cross-DevOps Reconfiguration with Ballet: A Multi-Site OpenStack Case Study’. In: *SANER 2024: IEEE International Conference on Software Analysis, Evolution and Reengineering*. Rovaniemi, Finland: IEEE, Mar. 2024, pp. 1–11. doi: [10.1109/SANER60148.2024.00007](https://hal.science/hal-04457484). URL: <https://hal.science/hal-04457484> (cit. on p. 20).
- [87] G. Privat, T. Coupaye, S. Bolle and P. Raipin Parvedy. ‘WoT Graph as Multiscale Digital-Twin for Cyber-Physical Systems-of-Systems’. In: *Proc., 2nd W3C Web of Things Workshop. Grenoble, France*. June 2019 (cit. on p. 8).
- [88] F. Quesnel, A. Lebre and M. Südholt. ‘Cooperative and reactive scheduling in large-scale virtualized platforms with DVMS’. In: *Concurrency and Computation: Practice and Experience* 25.12 (2013), pp. 1643–1655 (cit. on p. 8).
- [89] F. Ramparany, M. Thalgott, S. Bolle and S. Martin. ‘Exploiting Data Analytics for Home Automation Services’. In: *4th IEEE International Conference on Future Internet of Things and Cloud, FiCloud 2016, Vienna, Austria, August 22-24, 2016*. Ed. by M. Younas, I. Awan and W. Seah. IEEE Computer Society, 2016, pp. 228–234. doi: [10.1109/FiCloud.2016.40](https://doi.org/10.1109/FiCloud.2016.40). URL: <https://doi.org/10.1109/FiCloud.2016.40> (cit. on p. 8).
- [90] S. Robillard and H. Coullon. ‘SMT-Based Planning Synthesis for Distributed System Reconfigurations’. In: *Lecture Notes in Computer Science*. Lecture Notes in Computer Science. Munich, Germany, Apr. 2022. doi: [10.1007/978-3-030-99429-7_15](https://inria.hal.science/hal-03536643). URL: <https://inria.hal.science/hal-03536643> (cit. on p. 20).
- [91] G. Rosinosky, D. Schmitz and E. Rivière. ‘Streambed: capacity planning for stream processing’. In: *Proceedings of the 18th ACM International Conference on Distributed and Event-based Systems*. 2024, pp. 90–102 (cit. on p. 19).
- [92] F. A. Salaht, F. Desprez, A. Lebre, C. Prud’Homme and M. Abderrahim. ‘Service placement in fog computing using constraint programming’. In: *2019 IEEE International Conference on Services Computing (SCC)*. IEEE. 2019, pp. 19–27 (cit. on p. 8).
- [93] F. A. Salaht, F. Desprez and A. Lebre. ‘An overview of service placement problem in fog and edge computing’. In: *ACM Computing Surveys (CSUR)* 53.3 (2020), pp. 1–35 (cit. on p. 8).
- [94] M. Satyanarayanan. ‘The Emergence of Edge Computing’. In: *Computer* 50.1 (Jan. 2017), pp. 30–39. doi: [10.1109/MC.2017.9](https://doi.org/10.1109/MC.2017.9) (cit. on p. 7).

- [95] D. Serrano, S. Bouchenak, Y. Kouki, F. A. de Oliveira Jr., T. Ledoux, J. Lejeune, J. Sopena, L. Arantes and P. Sens. ‘SLA guarantees for cloud services’. In: *Future Generation Computer Systems* 54.Supplement C (2016), pp. 233–246. doi: <https://doi.org/10.1016/j.future.2015.03.018>. URL: <http://www.sciencedirect.com/science/article/pii/S0167739X15000801> (cit. on p. 9).
- [96] W. Shi, J. Cao, Q. Zhang, Y. Li and L. Xu. ‘Edge computing: Vision and challenges’. In: *IEEE Internet of Things Journal* 3.5 (2016), pp. 637–646 (cit. on p. 7).
- [97] J. M. Van Ham, G. Salvaneschi, M. Mezini and J. Noyé. ‘JEScala: Modular Coordination with Declarative Events and Joins’. In: *Proceedings of the 13th International Conference on Modularity (MODULARITY '14)*. Lugano, Switzerland: ACM Press, 2014, pp. 205–216. URL: <http://doi.acm.org/10.1145/2577080.2577082> (cit. on p. 9).
- [98] Y. Xia, X. Etchevers, L. Letondeur, A. Lebre, T. Coupaye and F. Desprez. ‘Combining heuristics to optimize and scale the placement of iot applications in the fog’. In: *2018 IEEE/ACM 11th International Conference on Utility and Cloud Computing (UCC)*. IEEE. 2018, pp. 153–163 (cit. on p. 8).
- [99] Y. Xia, X. Etchevers, L. Letondeur, T. Coupaye and F. Desprez. ‘Optimizing Cloud Application Scheduling: A Dual-Stage Heuristic Approach’. In: *ICCCBDA 2024 - 9th International Conference on Cloud Computing and Big Data Analytics*. Chengdu, China: IEEE, Apr. 2024, pp. 134–140. doi: [10.1109/ICCCBDA61447.2024.10569682](https://doi.org/10.1109/ICCCBDA61447.2024.10569682). URL: <https://hal.science/hal-04886644> (cit. on p. 21).
- [100] Y. Xiong, Y. Sun, L. Xing and Y. Huang. ‘Extend Cloud to Edge with KubeEdge’. In: *2018 IEEE/ACM Symposium on Edge Computing (SEC)*. 2018, pp. 373–377. doi: [10.1109/SEC.2018.00048](https://doi.org/10.1109/SEC.2018.00048) (cit. on p. 10).