

RESEARCH CENTRE

**Inria Centre
at Rennes University**

IN PARTNERSHIP WITH:

Ecole Nationale Supérieure
Mines-Télécom Atlantique Bretagne Pays
de la Loire, Nantes Université

2023

ACTIVITY REPORT

Project-Team
STACK

Software Stack for Massively Geo-Distributed Infrastructures

IN COLLABORATION WITH: Laboratoire des Sciences du numérique de
Nantes

DOMAIN

**Networks, Systems and Services,
Distributed Computing**

THEME

Distributed Systems and middleware

Inria

Contents

Project-Team STACK	1
1 Team members, visitors, external collaborators	3
2 Overall objectives	4
2.1 STACK in a Nutshell	4
2.2 Context & Objectives	4
2.3 Scientific Foundations	5
3 Research program	6
3.1 Overview	6
4 Application domains	8
4.1 Overview	9
4.2 Industrial Internet	9
4.3 Internet of Skills	9
4.4 e-Health	9
4.5 Network Virtualization and Mobile Edge Services	9
4.6 Urgent Computing	10
5 Social and environmental responsibility	10
5.1 Footprint of research activities	10
5.2 Impact of research results	10
6 Highlights of the year	11
6.1 Awards	11
7 New software, platforms, open data	11
7.1 New software	11
7.1.1 ENOS	11
7.1.2 EnOSlib	11
7.1.3 Concerto	12
7.1.4 6TiSCH Simulator	12
7.2 New platforms	13
7.2.1 OpenStack	13
7.2.2 Grid'5000	13
7.2.3 SeDuCe	13
7.2.4 PiSeDuCe	14
7.2.5 SAMURAI	14
7.2.6 SLICES-FR/SLICES	14
8 New results	14
8.1 Resource Management	14
8.2 Programming Support	16
8.3 Energy-aware computing	18
8.4 Security and Privacy	18
9 Bilateral contracts and grants with industry	19
9.1 Bilateral contracts with industry	19
9.2 Bilateral grants with industry	20

10 Partnerships and cooperations	21
10.1 International research visitors	21
10.1.1 Visits of international scientists	21
10.2 European initiatives	22
10.2.1 Horizon Europe	22
10.2.2 Other european programs/initiatives	22
10.3 National initiatives	22
10.3.1 ANR	22
10.3.2 PIA 4	23
10.3.3 CPER	24
10.3.4 Local and regional projects	24
10.3.5 Inria Challenges	24
11 Dissemination	25
11.1 Promoting scientific activities	25
11.1.1 Scientific events: organisation	25
11.1.2 Scientific events: selection	25
11.1.3 Journal	26
11.1.4 Invited talks	26
11.1.5 Leadership within the scientific community	26
11.1.6 Scientific expertise	27
11.1.7 Research administration	27
11.2 Teaching - Supervision - Juries	27
11.2.1 Teaching	27
11.2.2 Supervision	28
11.2.3 Juries	29
11.3 Popularization	29
11.3.1 Articles and contents	30
11.3.2 Education	30
11.3.3 Interventions	30
12 Scientific production	30
12.1 Major publications	30
12.2 Publications of the year	30
12.3 Other	33
12.4 Cited publications	33

Project-Team STACK

Creation of the Project-Team: 2019 January 01

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
- A1.2.5. – Internet of things
- 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. – 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

1 Team members, visitors, external collaborators

Research Scientists

- Adrien Lebre [Team leader, INRIA, Senior Researcher, from Dec 2023, HDR]
- Daniel Balouek-Thomert [INRIA, ISFP, from Feb 2023]

Faculty Members

- H el ene Coullon [IMT ATLANTIQUE, Associate Professor]
- Mohamed Graiet [IMT ATLANTIQUE, ATER]
- Remous-Aris Koutsiamanis [IMT ATLANTIQUE, Associate Professor]
- Thomas Ledoux [IMT ATLANTIQUE, Professor, HDR]
- Adrien L ebre [IMT ATLANTIQUE, Professor, until Nov 2023, HDR]
- Jean-Marc Menaud [IMT ATLANTIQUE, Professor, HDR]
- Jacques Noy e [IMT ATLANTIQUE, Associate Professor]
- Kandaraj Piamrat [NANTES UNIVERSITE, Associate Professor]
- Mario S udholt [IMT ATLANTIQUE, Professor, HDR]

Post-Doctoral Fellows

- Ons Aouedi [IMT ATLANTIQUE, until Nov 2023]
- Farid Arfi [IMT ATLANTIQUE, from Feb 2023]
- Alexis Bitailou [IMT ATLANTIQUE, until Jan 2023]
- Yasmina Bouizem [INRIA, Post-Doctoral Fellow]
- Marie Delavergne [IMT ATLANTIQUE, from Apr 2023]
- Tayeb Diab [IMT ATLANTIQUE]
- Jabran Khan [IMT ATLANTIQUE, from Feb 2023]
- Cl ement Mommessin [IMT ATLANTIQUE]
- Eloi Perdereau [IMT ATLANTIQUE, Post-Doctoral Fellow]
- Jolan Philippe [IMT ATLANTIQUE, Post-Doctoral Fellow, from Feb 2023]

PhD Students

- Tengfei An [INRIA, from Nov 2023]
- Geo Johns Antony [INRIA]
- Hiba Awad [SMILE, CIFRE]
- Samia Boutalbi [ERICSSON, CIFRE]
- Divi De Lacour [ORANGE, CIFRE]
- Dan Freeman Mahoro [ORANGE, CIFRE, until Oct 2023]

- Duc Thinh Ngo [ORANGE, CIFRE]
- Antoine Omond [IMT ATLANTIQUE]
- Abdou Seck [IMT ATLANTIQUE]
- Cherif Si Mohammed [ADEME, from Dec 2023]

Technical Staff

- Baptiste Jonglez [INRIA, Engineer]
- Sidi Mohammed Kaddour [IMT ATLANTIQUE, Engineer, from Dec 2023]
- Anas Mokhtari [IMT ATLANTIQUE, Engineer, from Sep 2023]
- Matthieu Rakotojaona Rainimangavelo [IMT ATLANTIQUE, Engineer, from May 2023]

Interns and Apprentices

- Mohamed Khaled Ben Mna [LS2N, from Mar 2023 until Aug 2023, Master Intern]

Administrative Assistant

- Anne-Claire Binetruy [INRIA]

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 [74, 72]. 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 reified 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. Starting with Ass. Prof Koutsiamanis, Ass. Prof Piamrat and Dr. Balouek (who respectively joined the team in 2021, 2022 and 2023), this enlargement of STACK core activities will be further strengthened with the arrival of Orange members expected in 2024.

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 [58, 68, 62, 78, 70, 71]. 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 such as approaches based on supervised/semi-supervised/non-supervised and federated learning [7, 31, 32, 11]. 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 [61, 59, 40, 57, 43].
- **Cyber-Physical Systems, Digital Twins, IoT, Device Management** Based on initial expertise in Device Management (DM) and IoT platforms for the management of connected devices and sensor data, and especially on distributed and autonomic architectures of such platforms [69, 34, 33], the team has developed a broader vision of cyber-physical systems with a strong expertise in *digital twins* as a pivotal technology. This includes graph-based modelling of cyber-physical systems [67], semantic modelling and ontology mapping [35], graph storage distribution and historisation [55, 56] — and the application of these concepts and technologies in different use cases in the domain of smart building (e.g., localisation [43], dynamic wireless IoT network resource allocation[57]), smart industry (e.g.,

support for reliable and low-latency wireless Industrial IoT networks [61, 59]), logistics [66] around the Thing'in digital twin platform [48] (Thing'in).

- **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 [51, 52, 65, 53, 64, 39, 41, 42].
- **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 [49, 38, 50, 63, 47, 9], event-driven [54, 77], data-driven and workflow models, as well as models for Utility Computing (Service Level Agreement, aka. SLA) [73], 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 [36] and more recently aspects related to software-defined experiments and reproducible research [46, 44]. 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

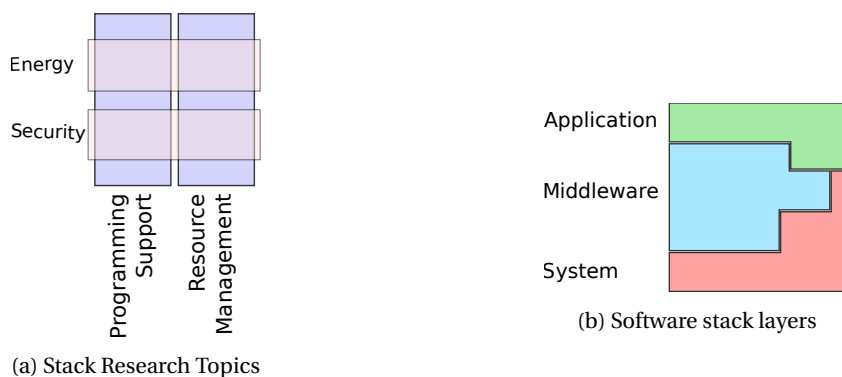
STACK activities have been focused on the management and programming of geo-distributed data centers with a work program defined around four research topics as depicted in Figure 1a. The first two ones are related to the resource management mechanisms and the programming support that are mandatory to 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.

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 [60].



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 [79], 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 [42] extended to IoT devices and to network resources. Finally, we address the dynamicity of systems described by such Architecture Description-like 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.

- *How to handle in an easy and non-intrusive manner the geo-distribution of complex legacy software stacks to avoid re-implementing from scratch large open source projects such as OpenStack or Kubernetes?* In particular, we are pushing our recent proposal [45] further to handle geo-distribution concerns through a dedicated service mesh.
- *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.
- *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.).
- *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 [37] 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 dynamics 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 Computing². 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.

²ETSI MEC specifications.

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 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, light-weight 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 Comin-Labs). 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.

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. As an example, we are investigating through the I/O Lab, the joint lab between Orange and Inria, how can a Centralized Radio Access Networks (*a.k.a.* C-RAN or Cloud-RAN) be supported for 5G networks. We highlight that our expertise is not on the network side but rather on where and how we can deploy, allocate and reconfigure software components, which are mandatory to operate a C-RAN infrastructure, in order to guarantee the quality of service expected by the end-users. Finally, working with actors from the network community is a valuable advantage for a distributed system research group such as STACK. Indeed, achievements made within one of the two communities serve the other.

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 2023, the team has taken part in two Inria Challenges:

- The first one is built around a partnership between Inria and OVHCloud. It aims to study end-to-end eco-design of a cloud in order to reduce its environmental impact.
- The second one involved Inria and Qarnot Computing, with the support of ADEME. Entitled Pulse, this challenge aims to develop and promote best practices in terms of reducing and recycling emissions of intensive computing infrastructures.

Last but not least, STACK started in 2022 a new platform 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).

6 Highlights of the year

Regarding scientific contributions, the team has produced major results on the management of large-scale infrastructures, in particular one publication in the prestigious journal ACM Computing Surveys [9]. The team also continued to consolidate its presence in the network community (e.g. in Elsevier Computer Networks [11])

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.

Finally, on the platform side, we continued our effort and took part in the different actions around the SLICES and SLICES-FR, see Section 10.

6.1 Awards

The team has received the 3rd Best paper award at the 16th IEEE/ACM International Conference on Utility and Cloud Computing (UCC 2023) for its paper [24].

Finally, it is noteworthy to mention the new position of Adrien Lebre as the leader of the Cloud, Network and system program at Inria since December 2023 as well as the co-leader of the French Cloud PEPR action.

7 New software, platforms, open data

7.1 New software

7.1.1 ENOS

Name: Experimental eNvironment for OpenStack

Keywords: OpenStack, Experimentation, Reproducibility

Functional Description: Enos workflow :

A typical experiment 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.

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

Publication: hal-01664515

Contact: Adrien Lebre

Participants: Mathieu Simonin, Marie Delavergne, Adrien Lebre, Baptiste Jonglez

Partner: Orange Labs

7.1.2 EnOSlib

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

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

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

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

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

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

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

Contact: Mathieu Simonin

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

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 to : 1. describe the life-cycle and the dependencies of software components, 2. describe a components assembly that forms the overall life-cycle of a distributed software, 3. automatically reconfigure 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

Participants: Christian Perez, H el ene Coullon, Maverick Chardet, Simon Robillard

Partners: IMT Atlantique, LS2N, LIP

7.1.4 6TiSCH Simulator

Name: High-level simulator of a 6TiSCH network

Keywords: Network simulator, 6TiSCH

Functional Description: The simulator is written in Python. While it doesn't provide a cycle-accurate emulation, it does implement the functional behavior of a node running the full 6TiSCH protocol stack. This includes RPL, 6LoWPAN, CoAP and 6P. The implementation work tracks the progress of the standardization process at the IETF.

Publication: [hal-03919553](#)

Contact: Malisa Vucinic

Participant: Remous Koutsiamanis

7.2 New platforms

7.2.1 OpenStack

Participants: Geo Johns Antony, Marie Delavergne, Baptiste Jonglez, Matthieu Rako-tojaona Rainimangavelo, Adrien Lebre.

OpenStack is the de facto open-source management system to operate and use Cloud Computing infrastructure. Started in 2012, the OpenStack foundation gathers 500 organizations including groups such as Intel, AT&T, RedHat, etc. The software platform relies on tens of services with a 6-month development cycle. It is composed of more than 2 millions of lines of code, mainly in Python, just for the core services. While these aspects make the whole ecosystem quite swift, they are also good signs of maturity for this community.

We created and animated between 2016 and 2018 the **Fog/Edge/Massively Distributed (FEMDC) Special Interest Group** and have been contributing to the Performance working group since 2015. The former investigates how OpenStack can address Fog/Edge Computing use cases whereas the latter addresses scalability, reactivity and high-availability challenges. In addition to releasing white papers and guidelines [75], the major result from the academic view point is the aforementioned EnOS solution, a holistic framework to conduct performance evaluations of OpenStack (control and data plane). In May 2018, the FEMDC SiG turned into a larger group under the control of the OpenStack foundation. This group gathers large companies such as Verizon, ATT, etc. Although our involvement has been less important since 2020, our participation is still significant. For instance, we co-signed the second white paper delivered by the edge working group in 2020 [76] and are still taking part to the working group, investigating new use-cases as well as resulting challenges.

7.2.2 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.

7.2.3 SeDuCe

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

The SeDuCe Project aims to deliver a research testbed dedicated to holistic research studies on energetical aspects of datacenters. Part of the Grid'5000 Nantes site, this infrastructure is composed of probes that measure the power consumption of each server, each switch and each cooling system, and also measure the temperature at the front and the back of each servers. These sensors enable reasearch to cover the full spectrum of the energetical aspect of datacenters, such as cooling and power consumption depending of experimental conditions.

The testbed is connected to renewable energy sources (solar panels). This “green” datacenter enables researchers to perform real experiment-driven studies on fields such as temperature based scheduling or “green” aware software (*i.e.*, software that take into account renewable energies and weather conditions).

7.2.4 PiSeDuCe

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

In 2023, we consolidated the development of PiSeDuCe, a deployment and reservation system for Edge Computing infrastructures composed of multiple Raspberry Pi Cluster started in 2020. Typically, a cluster of 8 Raspberry Pi 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 PiSeducu 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.

7.2.5 SAMURAI

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

In 2022 the SAMURAI (Sustainable And autoNoMous gReen computing for AI) project was accepted, 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 will extend SeDuCe into energy autonomy by adding a smart and clean energy storage system. Additionally, it will extend the capabilities of the platform by adding AI computing nodes (GPUs) for the scientific study of AI tools. Finally, it will also add new sensor nodes within the Nantes connected object platform (Nantes nodes of the national FIT IoT-Lab platform) to support future work on embedded AI. In 2023, We built the IoT platform and network integration with the SeDuCe and G5K platform.

7.2.6 SLICES-FR/SLICES

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

STACK Members are 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, clouds, HPC, big data, etc.). In 2022, STACK contributions are mainly related to SLICES, the European initiative relative. In particular, members have taken part to the SLICES-DS project (Design Study) as well as the SLICES-PP project (Preparatory Phase) of the [SLICES-RI action](#).

8 New results

8.1 Resource Management

Participants: Ons Aouedi, Daniel Balouek, H el ene Coullon, Remous-Aris Koutsiamanis, Duc Thinh Ngo, Kandaraj Piamrat, Thomas Ledoux, Pierre Jacquet.

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), may not be 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 2023, we have continued our effort to answer this question and delivered multiple contributions. We also start new activities around Urgent Computing (see Section 4).

Network resource management: The team has continued activities covering the Cloud–IoT continuum. First, an architecture that cover the whole continuum has been proposed in [14]. It deploys hierarchical federated learning (HFL) to cope with limitations of traditional FL, which suffers from several issues such as (i) robustness, due to a single point of failure, (ii) latency, as it requires a significant amount of communication resources, and (iii) convergence, due to system and statistical heterogeneity. HFL adds the edge servers as an intermediate layer for sub-model aggregation, several iterations will be performed before the global aggregation at the cloud server takes place, thus making the overall process more efficient, especially with non-independent and identically distributed (non-IID) data. Moreover, using traditional Artificial Neural Networks (ANNs) with HFL consumes a significant amount of energy, further hindering the application of decentralized FL on energy-constrained mobile devices. Therefore, this paper presents HFedSNN: an energy-efficient and fast-convergence model by incorporating Spike Neural Networks (SNNs) within HFL. SNN is a generation of neural networks, which promises tremendous energy and computation efficiency improvements. Taking advantage of HFL and SNN, numerical results demonstrate that HFedSNN outperforms FL with SNN (FedSNN) in terms of accuracy and communication overhead by 4.48% and 26×, respectively. Furthermore, HFedSNN significantly reduces energy consumption by 4.3× compared to FL with ANN (FedANN).

Moreover, as the complexity and scale of modern networks continue to grow, the need for efficient, secure management, and optimization becomes increasingly vital. Digital twin (DT) technology has emerged as a promising approach to address these challenges by providing a virtual representation of the physical network, enabling analysis, diagnosis, emulation, and control. The emergence of Software-defined network (SDN) has facilitated a holistic view of the network topology, enabling the use of Graph neural network (GNN) as a data-driven technique to solve diverse problems in future networks. The survey [10] explores the intersection of GNNs and Network digital twins (NDTs), providing an overview of their applications, enabling technologies, challenges, and opportunities. We discuss how GNNs and NDTs can be leveraged to improve network performance, optimize routing, enable network slicing, and enhance security in future networks. Additionally, we highlight certain advantages of incorporating GNNs into NDTs and present two case studies. Finally, we address the key challenges and promising directions in the field, aiming to inspire further advancements and foster innovation in GNN-based NDTs for future networks.

With an orientation specific to 5G networks and cloud computing, in [18] we introduce the Mobile Edge Slice Broker (MESB), a novel solution for orchestrating the deployment and operation of mobile edge slices in multi-cloud environments. We focus on reconciling the challenges faced by mobile edge slice tenants in selecting Cloud Infrastructure Providers (CIPs), ensuring seamless integration of network functions, services, and applications while maintaining quality of service (QoS). Our MESB architectural design enables dynamic monitoring, reconfiguration, and optimization of mobile edge slices to meet evolving requirements. This work represents a first step in enhancing the efficiency and effectiveness of resource management in the convergence of telecom and cloud services in 5G networks.

In line with the previous works in network resource management, in [11] we continue to focus on efficient and secure network management by applying federated learning specifically to 5G cellular traffic prediction. In this work we emphasize the viability of various neural network models, including RNNs and CNNs, in a federated setting for time-series forecasting, while also exploring the effectiveness of different federated aggregation functions in handling non-identically distributed data. Additionally, we demonstrate that preprocessing techniques at local base stations significantly improve prediction accuracy, and that federated learning offers substantial benefits in reducing carbon emissions, computational demands, and communication costs, making it a viable and efficient approach for large-scale and intelligent resource management in network systems.

CloudFactory: [19] introduces CloudFactory, an IaaS workload generator, as a solution to address the representativeness problem. This tool includes a library for sharing IaaS statistics and a generator for producing realistic VM workloads based on these statistics. CloudFactory is offered as open-source software for use by Cloud providers and researchers to facilitate the evaluation of new contributions. To demonstrate the relevance of CloudFactory, we discuss an analysis of the scheduling evolution for different IaaS workload intensities of two different cloud providers (Microsoft Azure and Chameleon). OVHcloud statistics, computed from CloudFactory, are also reported, and a comparison is made with other Cloud providers.

Urgent Computing: The Continuum Computing encompasses the integration of diverse infrastructures, including cloud, edge, and fog, to facilitate seamless migration of applications based on their specific needs, ensuring optimal satisfaction of their requirements. In particular, Urgent applications aim at processing distributed data sources while considering the uncertainty of infrastructure and timeliness constraints of data-driven workflows. The primary obstacles in this particular context mostly pertain to the incapacity to promptly respond to changes in the environment or the quality of service (QoS) constraints of the application, as well as the incapability to maintain an application in a stateless manner, hence impeding its relocation without the risk of data loss. This line of work is highly motivated by natural disaster case studies. In [17], we leverage three different platforms to combine smoke detection at the edge using camera imagery, wildfire simulations in the cloud, and finally heavy HPC models for generating pollution concentration maps. The resulting workflow is a prime example of a Cloud-to-IoT continuum use case that monitors and manages the air quality impacts of remote wildfires. In this context, we frame challenges and discuss how the R-Pulsar programming system can support urgent data-processing pipelines that tradeoff the content of data, cost of computation, and urgency of the results.

In [24], we tackle the aforementioned issues through the introduction of a framework based on a Function-as-a-Service (FaaS) and event-driven architecture. This framework enables the decomposition, localization, and relocation of applications inside a Continuum infrastructure, facilitated by a rule engine that is both system and data-aware. Results highlight the use of the execution history as a basis for determining the optimal location for executing a function, hence guaranteeing the desired Quality of Service (QoS). The evaluation was carried out by simulating a smoke detection use case as a representative scenario for Urgent Computing. We extended this approach in the context of the TEMA Horizon Europe project that aims at addressing Natural Disaster Management [25]. We introduced the Urgent Function Enabler (UFE) platform, a fully distributed architecture able to define, spread, and manage FaaS functions, using local IoT data and a computing infrastructure composed of mobile and stable nodes.

This approach presents significant potential in numerous fields. We are currently exploring a climate risk management framework based on real data from the Danish Meteorological and Hydrological Institutes (DMI, and DHI), coupled with road condition (pavement bearing capacity) data from authorities [23]. The aim of this project is to enhance a risk assessment model for evaluating pavement structure functionality during extreme weather events.

8.2 Programming Support

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

Fog Modeling: In [6], we propose a detailed overview of the current state-of-the-art in terms of Fog modeling languages. We relied on our long-term experience in Cloud Computing and Cloud Modeling to contribute a feature model describing what we believe to be the most important characteristics of Fog modeling languages. We also performed a systematic scientific literature search and selection process to obtain a list of already existing Fog modeling languages. Then, we evaluated and compared these Fog modeling languages according to the characteristics expressed in our feature model. As a result, we

discuss in this paper the main capabilities of these Fog modeling languages and propose a corresponding set of open research challenges in this area.

The size, complexity, and heterogeneity of Fog systems throughout their life cycle poses challenges and potential costs. To address this, [15] proposes a generic model-based approach called VeriFog, utilizing a customizable Fog Modeling Language (FML) for verifying Fog systems during the design phase. The approach was tested with three use cases from different application domains, considering various non-functional properties for verification (energy, performance, security). Developed in collaboration with the industrial partner Smile, this approach represents a crucial step towards a comprehensive model-based support for the entire life cycle of Fog systems.

Configuration languages: In the context of the OTPaaS project, we have started to study the anatomy of configuration languages with the objective of unveiling the building blocks of such languages. In particular, we have focused on a very expressive, and declarative, new configuration language, called CUE (CUE stands for *Configure Unify Execute*). CUE is open source, inspired by previous work at Google. It represents configurations in a way inspired by *typed feature structures*, a concept inherited from computational linguistics. This is not the first time that a language is built on top of this concept: LIFE, built in the 80s, revolves around *psi-terms*, a variant of types feature structures. Whereas CUE is a domain-specific language, LIFE is a generic language, which combines logic programming, functional programming and object-oriented programming. As such, it potentially provides a general framework to relate configuration as data and computations. We have presented our initial findings on the relationship between CUE and LIFE at CONFLANG '23, a SPLASH workshop [29].

Software deployment/reconfiguration: For a few years, the team has been working on deployment and dynamic reconfiguration of distributed software systems through the tool Concerto [42]. By offering a programmable lifecycle of resources (services, devices etc.), the team has shown multiple times that the level of parallelism and asynchronism when executing reconfigurations can be increased, thus reducing the execution time of reconfigurations. The accumulated knowledge on component-based distributed software reconfiguration has made possible the publication in 2023 of a survey in ACM Computing Surveys [9], a great success. In particular, this survey analyzes the state of the art on reconfiguration through the spectrum of verification. Indeed, given the complexity of distributed software and the adverse effects of incorrect reconfigurations, a suitable methodology is needed to ensure the correctness of reconfigurations in component-based systems. This survey gives the reader a global perspective over existing formal techniques that pursue this goal. It distinguishes different ways in which formal methods can improve the reliability of reconfigurations, and lists techniques that contribute to solving each of these particular scientific challenges.

In addition to this survey, last year the team has worked on the decentralized version of Concerto, namely Concerto-D. Such decentralized reconfiguration tool is very useful when building global knowledge on the state of the system is difficult, impossible, or very costly because of faults, disconnection or scale. If a publication on Concerto-D itself is not yet published, the usage of Concerto-D has been explored and published on three use cases describe below.

First, Cyber-Physical Systems (CPS) deployed in scarce resource environments like the Arctic Tundra (AT) face extreme conditions. Nodes deployed in such environments have to carefully manage a limited energy budget, forcing them to alternate long sleeping and brief uptime periods. During uptimes, nodes can collaborate for data exchanges or computations by providing services to other nodes. Deploying or updating such services on nodes requires coordination to prevent failures (e.g., sending new/updated API, waiting for service activation/deactivation, etc.). In [20], we study analytically and experimentally, to what extent using relay nodes for communications reduces the deployment and update durations (reconfiguration), and which factors influence this reduction. Intuitively, as dealing with sleeping nodes with low chances of uptime overlaps (no uptime synchronization), a coordinated deployment or update takes very long to finish if nodes have to communicate directly. In [21] we evaluate and study nodes' energy consumption during reconfiguration tasks coordination according to different CPS configurations (i.e., number of nodes, uptime duration, radio technology, or relay node availability). Results show high energy consumption in scenarios where nodes wake up specifically to deploy/update. It is shown that it is beneficial to do adaptation tasks while overlapping with existing uptimes (i.e., reserved for observation

activities). This paper also evaluates and studies how nodes' uptime duration and relay node availability influence energy consumption. Increasing uptime duration can reduce energy consumption, up to 12

Second, Concerto-D has been studied and used in the context of the SeMaFoR ANR project. The context of this project is Fog Computing, a paradigm aiming to decentralize the Cloud by geographically distributing away computation, storage and network resources as well as related services. This notably reduces bottlenecks and data movement. However, managing Fog resources is a major challenge because the targeted systems are large, geographically distributed, unreliable and very dynamic. Cloud systems are generally managed via centralized autonomic controllers automatically optimizing both application QoS and resource usage. To leverage the self-management of Fog resources, we propose to orchestrate a fleet of autonomic controllers in a decentralized manner, each with a local view of its own resources. In [13], we present our SeMaFoR (Self-Management of Fog Resources) vision that aims at collaboratively operating Fog resources. SeMaFoR is a generic approach made of three cornerstones: an Architecture Description Language for the Fog, a collaborative and consensual decision-making process, and an automatic coordination mechanism for reconfiguration (based on Concerto-D).

Finally, the usage of Concerto-D has been studied on the particular case of Urgent Computing. This field considers sudden situations of danger such as climate disasters, fires and air pollution, which happen without control in a possibly short window frame, and puts people's lives in danger. In such situation, some applications to help solve the situation have to be dynamically deployed, or existing applications have to be updated in a short period of time and according to new unanticipated constraints. The study in [16] provides a justification for the necessity of dynamic adaptation in applications that are time-sensitive. Furthermore, we present our viewpoint on time-sensitive applications and undertake a thorough analysis of the underlying principles and challenges that need to be resolved in order to accomplish this goal.

Decentralized federated learning: Federated Learning (FL) collaboratively trains machine learning models on the data of local devices without having to move the data itself: a central server aggregates models, with privacy and performance benefits but also scalability and resilience challenges. We have presented FDFL [citedelacour:hal-03946638](https://hal.archives-ouvertes.fr/hal-03946638), a new fully decentralized FL model and architecture that improves standard FL scalability and resilience with no loss of convergence speed. FDFL provides an aggregator-based model that enables scalability benefits and features an election process to tolerate node failures. Simulation results show that FDFL scales well with network size in terms of computing, memory, and communication compared to related FL approaches

8.3 Energy-aware computing

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

The activities on this axis are mainly related to the design, development and deployment of the SAMURAI project (7.2.5), a testbed that will allow researchers to investigate energy related challenges over the computing continuum (from the Cloud to IoT devices/cyber physical systems).

Additionally, in [22], we aim to address the sustainability of machine learning (ML) models in the context of cellular traffic forecasting, specifically in the emerging 5G networks. We introduce a novel sustainability indicator to assess the trade-off between energy consumption and predictive accuracy of deep learning (DL) models in a federated learning (FL) framework. We then use it to demonstrate that while larger, more complex models offer slightly better accuracy, their significantly higher energy consumption and environmental impact make simpler models more viable for real-world applications, highlighting the importance of energy-aware computing in the development of distributed AI technologies for network management.

8.4 Security and Privacy

Participants: Ons Aouedi, Wilmer Edicson Garzon Alfonso, Mohamed Graiet, Kandraj 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 in both domains.

Network security and privacy: The rapid development of network communication along with the drastic increase in the number of smart devices has triggered a surge in network traffic, which can contain private data and in turn affect user privacy. Recently, Federated Learning (FL) has been proposed in Intrusion Detection Systems (IDS) to ensure attack detection, privacy preservation, and cost reduction, which are crucial issues in traditional centralized machine-learning-based IDS. However, FL-based approaches still exhibit vulnerabilities that can be exploited by adversaries to compromise user data. At the same time, meta-models (including the blending models) have been recognized as one of the solutions to improve generalization for attack detection and classification since they enhance generalization and predictive performances by combining multiple base models. Therefore, we have proposed in [7] a Federated Blending model-driven IDS framework for the Internet of Things (IoT) and Industrial IoT (IIoT), called F-BIDS, in order to further protect the privacy of existing ML-based IDS. It consists of a Decision Tree (DT) and Random Forest (RF) as base classifiers to first produce the meta-data. Then, the meta-classifier, which is a Neural Networks (NN) model, uses the meta-data during the federated training step, and finally, it makes the final classification on the test set. Specifically, in contrast to the classical FL approaches, the federated meta-classifier is trained on the meta-data (composite data) instead of user-sensitive data to further enhance privacy. To evaluate the performance of F-BIDS, we used the most recent and open cyber-security datasets, called Edge-IIoTset (published in 2022) and InSDN (in 2020). We chose these datasets because they are recent datasets and contain a large amount of network traffic including both malicious and benign traffic.

Secure biomedical analyzes: In Wilmer Garzon's PhD thesis [28], he has shown that researchers require new tools and techniques to address the restrictions and needs of global scientific collaborations over geodistributed biomedical data. He has identified three kinds of constraints related to global collaborations, namely, technical, legal, and socioeconomic constraints.

He has proposed Fully Distributed Collaborations (FDC), which are research endeavors that provide means to analyze massive biomedical information collaboratively while respecting legal and socioeconomic restrictions.

A new random-forest based AI-algorithm MuSiForest improves over other existing AI approaches by ameliorating computation time and reducing the amount of shared data while having a training precision close to that of centralized random forest techniques.

Finally, he has proposed FeDeRa, a language to specify, deploy, and execute FDC-compliant multi-site scientific workflows.

9 Bilateral contracts and grants with industry

9.1 Bilateral contracts with industry

Kelio (formerly 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 2023, several activities were conducted, such as a hackathon with 100 participants, an event promoting the integration of international students at IMT Atlantique and preliminary discussions on the topic of a joint PhD 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 "A model-based approach for dynamic, multi-scale distributed systems" started in Nov. 2021.

The first publication to emerge from this collaboration will be presented at next year's SAC conference [15].

OVHcloud

Participants: Yasmina Bouziem, Helene Coullon, Thomas Ledoux, Pierre Jacquet.

In 2021, INRIA and OVHcloud have signed an agreement to jointly study the problem of a more frugal Cloud. They have identified 3 axes : (i) Software eco-design of Cloud services and applications; (ii) Energy efficiency leverages; (iii) Impact reduction and support for Cloud users. The Stack team obtained a PhD grant and a 24-month post-doc grant.

Pierre Jacquet started his PhD in October 2021, under a co-supervision with the Spirals team with the subject "Fostering the Frugal Design of Cloud Native Applications". Pierre presented a poster, "Sampling Resource Requirements to Optimize Virtual Machines Sizing", at the EuroSys conference in April 2022.

Yasmina Bouziem started her postdoc in September 2022 for two years, under a co-supervision with the Avalon team in Lyon with the subject "pulling the energy cost of a reconfiguration execution within reconfiguration decisions".

Orange

Participants: Paul Bori, Divi de Lacour, Adrien Lebre, Thomas Ledoux, Dan Freeman Mahoro, Jean-Marc Menaud, Duc-Thanh Ngo, Kandaraj Piamrat, Mario Südholt.

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

Dan Freeman Mahoro has initiated a new collaboration on the theme of "digital twins". This co-operation has led to a joint doctoral thesis entitled "Digital twins of complex cyber-physical systems", which began in April 2022, co-supervised with the Orange team in Rennes, and has a first publication [26]. However, Dan Freeman Mahoro stopped his PhD for several reasons in October, preferring to join the industry as an engineer.

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-Thanh 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.

Divi de Lacour started his PhD in January 2022, on the subject "Architecture et services pour la protection des données pour systèmes coopératifs autonomes", under a co-supervision with the Orange team in Chatillon (Paris south region).

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.

10 Partnerships and cooperations

10.1 International research visitors

10.1.1 Visits of international scientists

Other international visits to the team

Daniel Sokolowski

Status PhD

Institution of origin: School of Computer Science, St. Gallen

Country: Switzerland

Dates: Dec 12 - 15, 2023

Context of the visit: Invited talk to the VELVET days, and collaboration

Mobility program/type of mobility: lecture

Shashikant Ilager

Status Postdoctoral Researcher

Institution of origin: Vienna University Of Technology

Country: Austria

Dates: Oct 15 - Nov 31, 2023

Context of the visit: Exploring collaboration on urgent computing and energy efficiency; Writing of collaborative project between the two groups

Mobility program/type of mobility: N/A

Manish Parashar

Status Professor

Institution of origin: University of Utah

Country: USA

Dates: Nov 23 - Nov 24, 2023

Context of the visit: Keynote; Exploring associate team frameworks with Inria teams

Mobility program/type of mobility: N/A

10.2 European initiatives

10.2.1 Horizon Europe

SLICES-PP

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

SLICES-RI (Research Infrastructure), which was recently included in the 2021 ESFRI roadmap, aims at building a large infrastructure needed for the experimental research on various aspects of distributed computing, networking, IoT and 5/6G networks. It will provide the resources needed to continuously design, experiment, operate and automate the full lifecycle management of digital infrastructures, data, applications, and services. Based on the two preceding projects within SLICES-RI, SLICES-DS (Design Study) and SLICES-SC (Starting Community), the SLICES-PP (Preparatory Phase) project will validate the requirements to engage into the implementation phase of the RI lifecycle. It will set the policies and decision processes for the governance of SLICES-RI: i.e., the legal and financial frameworks, the business model, the required human resource capacities and training programme. It will also settle the final technical architecture design for implementation. It will engage member states and stakeholders to secure commitment and funding needed for the platform to operate. It will position SLICES as an impactful instrument to support European advanced research, industrial competitiveness and societal impact in the digital era.

The involvement of the group is rather low if we consider the allocated budget (4K€) but the group is strongly involved at the national level taking part to different discussions related to SLICES-FR, the French node of SLICES.

10.2.2 Other european programs/initiatives

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

Participants: 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, will be running from March 2024 for 36 months, with an allocated budget of 874k€ (230K€ for Stack).

10.3 National initiatives

10.3.1 ANR

SeMaFoR (Self-Management of Fog Resources)

Participants: Thomas Ledoux (*coordinator*), H el ene Coullon, Jolan Philippe.

Fog Computing is a paradigm that aims to decentralize the Cloud at the edge of the network to geographically distribute computing/storage resources and their associated services. It reduces bottlenecks and data movement. But managing a Fog is a major challenge: the system is larger, unreliable, highly dynamic and does not offer a global view for decision making. The objective of the SeMaFoR project is to model, design and develop a generic and decentralized solution for the self-management of Fog resources.

The consortium is composed of three partners: LS2N-IMT Atlantique (Stack, NaoMod, TASC), LIP6-Sorbonne Universit  (Delys), Alter way/Smile (SME). The Stack team supervises the project. SeMaFoR is running for 42 months (starting in March 2021 with an allocated budget of 506k , 230K  for Stack). See the [Semafor](#) web site for more information.

The main results of the year 2023 were the survey "A feature-based survey of Fog modeling languages" [6] in the FGCS journal and a position paper [1] in the SEAMS conference ranked 'A'.

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

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

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 lisme, Institut de Canc rologie de l'Ouest / Informatique, Institut de Recherche en Informatique de Toulouse, Laboratoire des Sciences du Num rique 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).

10.3.2 PIA 4

OTPaas

Participants: Adrien Lebre (*STACK representative*), Farid Arfi, Alexis Bitailleu, H el ene Coullon, Marie Delavergne, Mohamed Graiet, Jabran Khan, Remous-Aris Koutsiamanis, Thomas Ledoux, Jean-Marc Menaud, Anas Mokhtari, Cl ment Mommessin, Jacques Noy , Kandaraj Piamrat, Eloi Perdereau, Matthieu Rakotojaona Rainimangavelo, Mario S udholt, Ons Aouedi, Tayeb Diab.

The OTPaaS project targets the design and development of a complete software stack to administrate and use edge infrastructures for the industry sector. The consortium brings together national and user technology suppliers from major groups (Atos / Bull, Schneider Electric, Valeo) and SMEs / ETIs (Agile 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).

The OTPaaS platform objectives are:

- 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.

10.3.3 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).

10.3.4 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.

10.3.5 Inria Challenges

FrugalCloud (Inria-OVHCloud)

Participants: Yasmina Bouziem, Hélène Coullon, Thomas Ledoux, Pierre Jacquet.

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.

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.

11 Dissemination

11.1 Promoting scientific activities

Participants: Daniel Balouek, H el ene Coullon, Remous-Aris Koutsiamanis, Adrien Lebre, Jean-Marc Menaud, Jacques Noy e, Kandaraj Piamrat, S udholt Mario, Thomas Ledoux.

11.1.1 Scientific events: organisation

General chair, scientific chair

- H el ene Coullon: VELVET days at IMT Atlantique in Nantes the 13th and 14th of December: helene-coullon.fr/pages/velvet
- Thomas Ledoux: Retirement Ceremony - Pierre Cointe (14th of September) webia.lip6.fr/briot/hommage-cointe-140923
- Daniel Balouek: International Workshop on Urgent Analytics for Distributed Computing (collocated with Euro-Par 2023). (Aug 28th) quickpar.github.io

Member of the organizing committees

- Daniel Balouek: Publicity chair, ICPP23
- Adrien Lebre: co-organizer of Middlewedge 2023 (workshop co-located with Middleware)

11.1.2 Scientific events: selection

Chair of conference program committees

- Daniel Balouek: Middleware for the Edge Workshop - Dec 11 (MiddleWedge 2023) at Middleware23

Member of the conference program committees

- Kandaraj Piamrat: IEEE ICC (CSM symposium) 2023, IEEE CCNC (IoT: From Sensors to Vertical Applications Track) 2023, IEEE IWQoS 2023, IEEE WiMob 2023, IEEE ICC Workshop on Data Driven Intelligence for Networks and Systems (DDINS) 2023, ACM CoNEXT 2023 workshop SAFE, Algotel/Cores 2023

- Hélène Coullon: SBAC-PAD, ICE workshop, QuickPar (workshop), algotel item Adrien Lebre: IEEE CloudCom, IEEE IC2E, IEEE ICFEC, MiddleWedge workshop (co-located with ACM Middleware)
- Thomas Ledoux: IEEE ASMECC 2023 (Workshop on Autonomic and Self-* Management for the Edge-Cloud Continuum)
- Daniel Balouek: Euro-Par 2023, SuperComputing 2023 Reproducibility challenge, ICCS 2023, UCC 2023
- Remous-Aris Koutsiamanis: IEEE ICC 2023, IEEE ISCC 2023, ICCS (International Conference on Computational Science) 2023, MiddleWedge workshop (co-located with ACM Middleware), International Workshop on Urgent Analytics for Distributed Computing (collocated with Euro-Par 2023), FCN'2023 MWN (Mobile and Wireless Networking Symposium), Algotel/Cores 2023
- Mario Südholt: CloudCom'23, ICSSP'23

Reviewer

- Helene Coullon: JLAMP
- Daniel Balouek: QuickPar Workshop 2023

11.1.3 Journal

Member of the editorial boards

- Adrien Lebre : Associate Editor IEEE Transactions on Cloud Computing
- Mario Südholt: Advisory Board member of the Programming Journal

Reviewer - reviewing activities

- Kandaraj Piamrat: IEEE Vehicular Technology Magazine, IEEE Network Magazine, International journal of intelligent systems, Elsevier-Journal of systems architecture, Elsevier-Computer communications
- Remous-Aris Koutsiamanis: IEEE Access, Elsevier Internet of Things, IEEE Transactions on Mobile Computing, Elsevier Computer Networks, Elsevier Ad Hoc Networks, Elsevier Results in Engineering, Elsevier Expert Systems With Applications, Elsevier Wireless Networks, Springer Journal of Network and Systems Management,

11.1.4 Invited talks

- Thomas Ledoux: SIF (Société Informatique de France) annual conference : "Frugalité numérique : vers une approche holistique et centrée sur l'humain pour les systèmes distribués"
- Daniel Balouek: Distributed Computing Continuum (DCC) at IEEE Cloud 2023 : "Leveraging the Computing Continuum for Urgent Science"
- Adrien Lebre: Cloud Computing lecture at the RSD Winter school (France)

11.1.5 Leadership within the scientific community

- Helene Coullon is co-chair of the working group YODA (trustworthyY and Optimal Dynamic Adaptation) in the national GDR GPL (software engineering and languages)
- Helene Coullon is co-chair of the working group CyclOps (DevOps, deployment and reconfiguration, lifecycle etc.) in the joint laboratory between Inria and Orange
- Helene Coullon is Vice-president of the French ACM SIGOPS group

11.1.6 Scientific expertise

- Jacques Noyé has acted as a reviewer for the Swiss National Science Foundation.
- Remous-Aris Koutsiamanis has acted as a reviewer for the National Science Centre Poland.
- Daniel Balouek has acted as a reviewer for the USA National Science Foundation.

11.1.7 Research administration

- Adrien Lebre : Member of the executive committee of the Grid'5000 GIS (Groupement d'intérêt scientifique), Member of the SLICES-FR temporary executive board, 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 Inria.
- Kandaraj Piamrat is international coordinator of LS2N (Laboratoire des sciences du numérique de Nantes).
- Jean-Marc Menaud is deputy head of LS2N (Laboratoire des sciences du numérique de Nantes)
- 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 Architects committee, Member of the SLICES-FR Site managers, Scientific representative of IMT to SLICES.

11.2 Teaching - Supervision - Juries

Participants: Daniel Balouek, H  l  ne Coullon, Remous-Aris Koutsiamanis, Adrien Lebre, Jean-Marc Menaud, Jacques Noy  , Kandaraj Piamrat, S  d-holt Mario, Thomas Ledoux.

11.2.1 Teaching

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.

- H  l  ne Coullon: responsible for the LOGIN training in computer science at IMT Atlantique (last year)
- Adrien Lebre: responsible for the FISE research program at IMT Atlantique (Nantes campus).
- Thomas Ledoux: 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 redesign of the common core of the 1st year of the IT curriculum at IMT Atlantique
- Thomas Ledoux: head of the Fili  re informatique nantaise since Sept. 2020. This entity, created by the University of Nantes, 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 num  rique**.
- Jacques Noy  : 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.
- Kandaraj Piamrat : elected member of scientific council at Faculty of Sciences and Techniques, Nantes University
- Kandaraj Piamrat : responsible for Licence MIAGE, CS Department, Nantes University

11.2.2 Supervision

- PhD: Marie Delavergne, Cheops, a service-mesh to geo-distribute micro-service applications at the Edge [27], Sept 2019 - March 2023, Director: A. Lebre.
- PhD: Wilmer Garzon, Secure distributed workflows for biomedical data analytics, Sept 2019 - Jan 2023, Director: M. Südholt, Advisor: D. Benavides.
- PhD in progress: Geo Johns Antony, Resource management in geo-distributed infrastructures: the Kubernetes case, April 2021 - Sept 2024, Director: A. Lebre.
- PhD in progress: Samia Boutalbi, Secure deployment of microservices in a shared RAN/MEC Cloud environment, Jan 2022 - Jan 2025, Director: M. Südholt, Advisor: R.-A. Koutsiamanis.
- PhD in progress: Abdou Seck, Parallel transfer service for the exchange of large volumes of data between Datacenters, June 2022 - June 2025, Director: J.-M. Menaud, Noel De Palma, Advisor: R.-A. Koutsiamanis.
- PhD in progress: Pierre Jacquet, Fostering the Frugal Design of Cloud Native Applications, Oct 2021 - Oct 2024, Director R. Rouvoy (SPIRALS Lille), Advisor: T. Ledoux.
- PhD in progress: Hiba Awad, A Model-based Approach for Multi-Scale and Dynamic Distributed Systems, Nov. 2021 - March 2025. Director: T. Ledoux.
- PhD in progress: Antoine Omond, Safe, efficient and low-energy self-adaptation for Cyber Physical Systems - Application to a scientific observatory in the Arctic tundra, Dec 2021 - Dec 2024. Director: T. Ledoux, Advisor: H. Coullon.
- PhD in progress: Dan Freeman Mahoro, Digital twins of complex cyber-physical systems, Apr. 2022 - March 2025, Director: T. Ledoux.
- PhD in progress: 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: Tengfei An, Modeling and Studying Self-Stabilization within Kubernetes, Nov 2023 - Sept 2026, Director: A. Lebre, Advisors: H. Coullon, J. Noyé
- PhD in progress: Divi De Lacour, "Architecture et services pour la protection des données pourses stèmes coopératifs autonomes", Jan 2022 - Dec 2025, Director: M. Südholt.
- Post-doc: Yasmina Bouizem, Measure and model the energy consumption of the procedures of deployment and reconfiguration at OVH, Sept 2022 - Sept 2023, Advisor: H. Coullon.
- Post-doc: Eloi Perdereau, Configuration languages, OTPaaS, Sept 2022 - March 2024, Advisor: J. Noyé.
- Post-doc: Ons Aouedi, Advanced AI for infrastructures and applications in the IoT-Cloud continuum, Jan.- Nov. 2023, Advisor: K. Piamrat, M. Südholt
- Post-doc: Farid Arfi, Safety of decentralized reconfigurations, Feb. 2023 - Aug. 2024, Advisor: H. Coullon

- Post-doc: Tayeb Diab, Decentralized AI for the continuum Cloud-Edge-IoT, Nov. 2022-Apr. 2024, Advisor: M. Südholt
- Post-doc: Jabran Khan, Placement in Cloud, Edge and Far Edge devices in the IoT to Cloud continuum, Feb. 2023 - Jul. 2024, Advisor: R.-A. Koutsiamanis, J.-M. Menaud
- Post-doc: Anas Mokhtari, A Holistic Approach for Designing Carbon-Aware and Energy-Aware Cloud applications, Sept. 2023 - Aug. 2024, Advisor: T. Ledoux.
- Post-doc: Clément Mommessin, Model of Energy consumption in Cloud, Edge and Far Edge devices in the IoT to Cloud continuum, Jan. 2023 - Jun. 2024, Advisor: J.-M. Menaud, R.-A. Koutsiamanis
- Post-doc: Jolan Philippe, Decentralization of distributed systems reconfiguration by using constraint programming, Feb. 2023 - Aug. 2024, Advisor: H. Coullon

11.2.3 Juries

- Kandaraj Piamrat and Helene Coullon were members of the selection committee for an associate professor positions in Rennes.
- Kandaraj Piamrat was an examiner of Sabra Ben Saad Thesis - Security architectures for network slice management for 5G and beyond. EURECOM, 02/23.
- Kandaraj Piamrat was an examiner of Mustafa Al Samara Thesis - Network Anomaly Detection in IoT: from Data Efficiency to Enhancing QoS, 06/23.
- Helene Coullon was member of the selection committee for research positions at Centre Inria de l'Université de Rennes
- Helene Coullon was an examiner of the PhD defense of Edouard Guégain at Université de Lille - Configuration-driven Software Optimization, 2023-09-29
- Helene Coullon was an examiner of the PhD defense of Davaadorj Battulga at Université de Rennes 1 - Contributions to the Management of Stream Processing Pipelines in Fog Computing Environments, 2023-03-23
- Adrien Lebre was a reviewer for the PhD thesis of Nihel Kaboubi, "Système de partitionnement hybride pour une inférence distribuée et embarquée des CNNs sur les équipements en bordure de réseau", Université de Grenoble Alpes, March 2023
- Adrien Lebre was reviewer for the Phd Thesis of Thomas Feltin, "Distributed Computing at the Smart Device Edge", Ecole Doctorale de l'Institut Polytechnique de Paris, Dec 2023
- Thomas Ledoux was member of the selection committee for an associate professor position in IMT Atlantique
- Thomas Ledoux was chairman of the PhD thesis jury of Marie Delavergne ("Cheops, a service-mesh to geo-distribute micro-service applications at the Edge") [27], 2023-03-16
- Jean-Marc Menaud was member of the selection committee for two associate professor positions in Toulouse and Lyon.

11.3 Popularization

Participants: Adrien Lebre, Thomas Ledoux.

11.3.1 Articles and contents

- Adrien Lebre - "Quèsaco le cloud ?", www.imt-atlantique.fr/fr/actualites/quesaco-le-cloud
- Thomas Ledoux - "La chasse au gaspillage dans le cloud et les data centers.", *The Conversation*, January 2023 [30]

11.3.2 Education

Thomas Ledoux was co-leader of the Ecolog initiative. This national project is the result of an association between the Institut du Numérique Responsable (INR) and the IT sector in Nantes, which includes the University of Nantes, IMT Atlantique and Centrale Nantes. Its ambition was to create a body of training courses available in open source and dedicated to sustainable computing. A webinar presenting the results was held on March 23, 2023 (ecolog.isit-europe.org).

11.3.3 Interventions

Thomas Ledoux took part in a round table related to digital sustainability/green computing organized by ADN Ouest (Nov. 14, 2023).

12 Scientific production

12.1 Major publications

- [1] A. Alidra, H. Bruneliere, H. Coullon, T. Ledoux, C. Prud'Homme, J. Lejeune, P. Sens, J. Sopena and J. Rivalan. 'SeMaFoR - Self-Management of Fog Resources with Collaborative Decentralized Controllers'. In: *SEAMS 2023 - IEEE/ACM 18th Symposium on Software Engineering for Adaptive and Self-Managing Systems*. Melbourne, Australia: IEEE, 15th May 2023, pp. 25–31. DOI: [10.1109/SEAMS59076.2023.00014](https://doi.org/10.1109/SEAMS59076.2023.00014). URL: <https://hal.science/hal-04043471>.
- [2] O. Aouedi and K. Piamrat. 'F-BIDS: Federated-Blending based Intrusion Detection System'. In: *Pervasive and Mobile Computing* 89 (2023), p. 101750. DOI: [10.1016/j.pmcj.2023.101750](https://doi.org/10.1016/j.pmcj.2023.101750). URL: <https://hal.science/hal-04223515>.
- [3] H. Coullon, L. Henrio, F. Loulergue and S. Robillard. 'Component-Based Distributed Software Reconfiguration: a Verification-Oriented Survey'. In: *ACM Computing Surveys* 56.1 (3rd May 2023), pp. 1–37. DOI: [10.1145/3595376](https://doi.org/10.1145/3595376). URL: <https://inria.hal.science/hal-04067909>.
- [4] D. De Lacour, M. Lacoste, M. Südholt and J. Traoré. 'Towards Scalable Resilient Federated Learning: A Fully Decentralised Approach'. In: *PerConAI 2023: 2nd IEEE Workshop on Pervasive and Resource-constrained Artificial Intelligence*. Atlanta (GA), United States: IEEE, 2023. DOI: [10.1109/PerComWorkshops56833.2023.10150295](https://doi.org/10.1109/PerComWorkshops56833.2023.10150295). URL: <https://inria.hal.science/hal-03946638>.
- [5] V. Perifanis, N. Pavlidis, R.-A. Koutsiamanis and P. Efraimidis. 'Federated learning for 5G base station traffic forecasting'. In: *Computer Networks* 235.1389-1286 (Aug. 2023), p. 109950. DOI: [10.1016/j.comnet.2023.109950](https://doi.org/10.1016/j.comnet.2023.109950). URL: <https://hal.science/hal-04381143>.

12.2 Publications of the year

International journals

- [6] A. Alidra, H. Bruneliere and T. Ledoux. 'A feature-based survey of Fog modeling languages'. In: *Future Generation Computer Systems* 138 (Jan. 2023), pp. 104–119. DOI: [10.1016/j.future.2022.08.010](https://doi.org/10.1016/j.future.2022.08.010). URL: <https://hal.science/hal-03759010>.
- [7] O. Aouedi and K. Piamrat. 'F-BIDS: Federated-Blending based Intrusion Detection System'. In: *Pervasive and Mobile Computing* 89 (2023), p. 101750. DOI: [10.1016/j.pmcj.2023.101750](https://doi.org/10.1016/j.pmcj.2023.101750). URL: <https://hal.science/hal-04223515>.

- [8] Y. Bouizem, D. Dib, N. Parlavantzas and C. Morin. ‘Integrating request replication into FaaS platforms: an experimental evaluation’. In: *Journal of Cloud Computing: Advances, Systems and Applications* 12.1 (22nd June 2023), pp. 1–20. DOI: [10.1186/s13677-023-00457-z](https://doi.org/10.1186/s13677-023-00457-z). URL: <https://inria.hal.science/hal-04159510>.
- [9] H. Coullon, L. Henrio, F. Loulergue and S. Robillard. ‘Component-Based Distributed Software Reconfiguration: a Verification-Oriented Survey’. In: *ACM Computing Surveys* 56.1 (3rd May 2023), pp. 1–37. DOI: [10.1145/3595376](https://doi.org/10.1145/3595376). URL: <https://inria.hal.science/hal-04067909>.
- [10] D.-T. Ngo, O. Aouedi, K. Piamrat, T. Hassan and P. Raipin-Parvédy. ‘Empowering Digital Twin for Future Networks with Graph Neural Networks: Overview, Enabling Technologies, Challenges, and Opportunities’. In: *Future internet* 15.12 (24th Nov. 2023), p. 377. DOI: [10.3390/fi15120377](https://doi.org/10.3390/fi15120377). URL: <https://hal.science/hal-04331283>.
- [11] V. Perifanis, N. Pavlidis, R.-A. Koutsiamanis and P. Efraimidis. ‘Federated learning for 5G base station traffic forecasting’. In: *Computer Networks* 235.1389-1286 (Aug. 2023), p. 109950. DOI: [10.1016/j.comnet.2023.109950](https://doi.org/10.1016/j.comnet.2023.109950). URL: <https://hal.science/hal-04381143>.

Invited conferences

- [12] D. De Lacour, M. Lacoste, M. Südholt and J. Traoré. ‘Towards Scalable Resilient Federated Learning: A Fully Decentralised Approach’. In: *PerConAI 2023: 2nd IEEE Workshop on Pervasive and Resource-constrained Artificial Intelligence*. Atlanta (GA), United States: IEEE, 2023. DOI: [10.1109/PerComWorkshops56833.2023.10150295](https://doi.org/10.1109/PerComWorkshops56833.2023.10150295). URL: <https://inria.hal.science/hal-03946638>.

International peer-reviewed conferences

- [13] A. Alidra, H. Bruneliere, H. Coullon, T. Ledoux, C. Prud’Homme, J. Lejeune, P. Sens, J. Sopena and J. Rivalan. ‘SeMaFoR - Self-Management of Fog Resources with Collaborative Decentralized Controllers’. In: *SEAMS 2023 - IEEE/ACM 18th Symposium on Software Engineering for Adaptive and Self-Managing Systems*. Melbourne, Australia: IEEE, 15th May 2023, pp. 25–31. DOI: [10.1109/SEAMS59076.2023.00014](https://doi.org/10.1109/SEAMS59076.2023.00014). URL: <https://hal.science/hal-04043471>.
- [14] O. Aouedi, K. Piamrat and M. Südholt. ‘HFedSNN: Efficient Hierarchical Federated Learning using Spiking Neural Networks’. In: *MobiWac 2023: 21st ACM International Symposium on Mobility Management and Wireless Access*. Montréal, Canada, 5th Sept. 2023. URL: <https://hal.science/hal-04197115>.
- [15] H. Awad, A. Alidra, H. Bruneliere, T. Ledoux, E. Leclercq and J. Rivalan. ‘VeriFog: A Generic Model-based Approach for Verifying Fog Systems at Design Time’. In: *The 39th ACM/SIGAPP Symposium on Applied Computing (SAC ’24)*. The 39th ACM/SIGAPP Symposium on Applied Computing (SAC ’24). Avila, Spain, 8th Apr. 2024. DOI: [10.1145/3605098.3635973](https://doi.org/10.1145/3605098.3635973). URL: <https://hal.science/hal-04332046>.
- [16] D. Balouek and H. Coullon. ‘Dynamic Adaptation of Urgent Applications in the Edge-to-Cloud Continuum’. In: *International European Conference on Parallel and Distributed Computing (EuroPar’2023)*. Limassol, Cyprus, 28th Aug. 2023. URL: <https://inria.hal.science/hal-04387689>.
- [17] D. Balouek, I. Perez, S. D. Faulstich, H. A. Holmes, I. Altintas and M. Parashar. ‘Keynote Talk: Leveraging the Edge-Cloud Continuum to Manage the Impact of Wildfires on Air Quality’. In: *2023 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops)*. Atlanta, France: IEEE, 13th Mar. 2023, pp. 27–31. DOI: [10.1109/PerComWorkshops56833.2023.10150251](https://doi.org/10.1109/PerComWorkshops56833.2023.10150251). URL: <https://inria.hal.science/hal-04387579>.
- [18] S. Boutalbi, N. Kheir, R. A. Koutsiamanis, M. Südholt and Y. Dan Moussa. ‘Mobile Edge Slice Broker: Mobile Edge Slices Deployment in Multi-Cloud Environments’. In: *ICFEC 2023 - 7th IEEE International Conference on Fog and Edge Computing*. Bangalore, India, 2023, pp. 1–6. URL: <https://hal.science/hal-04052250>.

- [19] P. Jacquet, T. Ledoux and R. Rouvoy. 'CLOUDFACTORY: An Open Toolkit to Generate Production-like Workloads for Cloud Infrastructures'. In: IC2E 2023 - 11th IEEE International Conference on Cloud Engineering. Boston, Massachusetts, USA, United States: IEEE, 25th Sept. 2023, pp. 1–11. URL: <https://inria.hal.science/hal-04168667>.
- [20] 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)*. CPSCom 2023 - 16th IEEE International Conference on Cyber, Physical and Social Computing. Hainan Island, China: IEEE, 17th Dec. 2023, pp. 1–8. URL: <https://hal.science/hal-04372320>.
- [21] A. Omond, I. Raïs and H. Coullon. 'Evaluating the energy consumption of adaptation tasks for a CPS in the Arctic Tundra'. 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)*. GreenCom 2023 - 9th IEEE International Conference on Green Computing and Communications. Hainan Island, China: IEEE, 2023, pp. 1–8. URL: <https://hal.science/hal-04372340>.
- [22] V. Perifanis, N. Pavlidis, S. Yilmaz, F. Wilhelmi, E. Guerra, M. Miozzo, P. Efraimidis, P. Dini and R.-A. Koutsiamanis. 'Towards Energy-Aware Federated Traffic Prediction for Cellular Networks'. In: *2023 Eighth International Conference on Fog and Mobile Edge Computing (FMEC)*. FMEC 2023 - Eighth International Conference on Fog and Mobile Edge Computing (FMEC). Tartu, Estonia: IEEE, 2023, pp. 93–100. DOI: [10.1109/FMEC59375.2023.10306017](https://doi.org/10.1109/FMEC59375.2023.10306017). URL: <https://hal.science/hal-04381142>.
- [23] S. Pour, D. Balouek, A. Masoumi and M. Brynskov. 'An urgent computing oriented architecture for dynamic climate risk management framework'. In: International Workshop on Intelligent and Adaptive Edge-Cloud Operations and Services (Intel4EC 2023). Taormina, Italy, 6th Dec. 2021. URL: <https://inria.hal.science/hal-04387718>.
- [24] C. Sicari, D. Balouek, M. Villari and M. Parashar. 'Event-Driven FaaS Workflows for Enabling IoT Data Processing at the Cloud Edge Continuum'. In: International Conference on Utility and Cloud Computing (UCC 2023). Taormina, Italy, 2023. URL: <https://inria.hal.science/hal-04387675>.
- [25] C. Sicari, A. Catalfamo, L. Carnevale, A. Galletta, D. Balouek, M. Parashar and M. Villari. 'TEMA: Event Driven Serverless Workflows Platform for Natural Disaster Management'. In: 2023 IEEE Symposium on Computers and Communications (ISCC). Gammarth, France: IEEE, 9th July 2023, pp. 1–6. DOI: [10.1109/iscc58397.2023.10217920](https://doi.org/10.1109/iscc58397.2023.10217920). URL: <https://inria.hal.science/hal-04387630>.

Conferences without proceedings

- [26] D. F. Mahoro, T. Ledoux, T. Hassan and T. Coupaye. 'Articulating Data and Control Planes for the Composition and Synchronization of Digital Twins'. In: Midd4DT 2023 - 1st International Workshop on Middleware for Digital Twin. Bologna (ITALY), Italy: ACM, 11th Dec. 2023, pp. 13–18. DOI: [10.1145/3631319.3632300](https://doi.org/10.1145/3631319.3632300). URL: <https://hal.science/hal-04355918>.

Doctoral dissertations and habilitation theses

- [27] M. Delavergne. 'Cheops, a service-mesh to geo-distribute micro-service applications at the Edge'. Ecole nationale supérieure Mines-Télécom Atlantique, 16th Mar. 2023. URL: <https://theses.hal.science/tel-04081084>.
- [28] W. Garzón-Alfonso. 'Secure distributed workflows for biomedical data analytics'. Ecole nationale supérieure Mines-Télécom Atlantique; Escuela Colombiana de Ingeniería Julio Garavito, 30th Aug. 2023. URL: <https://theses.hal.science/tel-04224597>.

Other scientific publications

- [29] E. Perdereau and J. Noyé. *The LIFE of CUE*. 24th Oct. 2023. URL: <https://inria.hal.science/hal-04389081>.

12.3 Other

Scientific popularization

- [30] P. Jacquet, R. Rouvoy and T. Ledoux. 'La chasse au gaspillage dans le cloud et les data centers'. In: *The Conversation* (29th Jan. 2023). URL: <https://inria.hal.science/hal-04168662>.

12.4 Cited publications

- [31] O. Aouedi, K. Piamrat and D. Bagadthey. 'Handling partially labeled network data: a semi-supervised approach using stacked sparse autoencoder'. In: *Computer Networks* (2022).
- [32] O. Aouedi, K. Piamrat, G. Muller and K. Singh. 'Federated Semi-Supervised Learning for Attack Detection in Industrial Internet of Things'. In: *IEEE Transactions on Industrial Informatics* (2022).
- [33] 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>.
- [34] 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>.
- [35] 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.
- [36] 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.
- [37] 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>.
- [38] 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>.
- [39] 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>.
- [40] 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.

- [41] 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.
- [42] 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>.
- [43] 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.
- [44] 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>.
- [45] R.-A. Cherrueau, M. Delavergne and A. Lebre. ‘Geo-Distribute Cloud Applications at the Edge’. In: *EURO-PAR 2021 - 27th International European Conference on Parallel and Distributed Computing*. Lisbon, Portugal, Aug. 2021, pp. 1–14. DOI: [10.1007/978-3-030-85665-6_19](https://doi.org/10.1007/978-3-030-85665-6_19). URL: <https://hal.inria.fr/hal-03212421>.
- [46] 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.
- [47] 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>.
- [48] 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.
- [49] 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.
- [50] 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>.
- [51] F. A. De Oliveira, T. Ledoux and R. Sharrock. ‘A framework for the coordination of multiple autonomous managers in cloud environments’. In: *Self-Adaptive and Self-Organizing Systems (SASO), 2013 IEEE 7th International Conference on*. IEEE, 2013, pp. 179–188.
- [52] 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>.
- [53] 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>.
- [54] 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>.

- [55] 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/csit.2020.101417](https://doi.org/10.5121/csit.2020.101417). URL: <https://hal.inria.fr/hal-03137952>.
- [56] 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>.
- [57] 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.
- [58] 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.
- [59] 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.
- [60] 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).
- [61] 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.
- [62] 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>.
- [63] 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.
- [64] 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).
- [65] 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>.
- [66] 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.
- [67] 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.
- [68] 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.

- [69] 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>.
- [70] 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.
- [71] 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.
- [72] 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).
- [73] 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>.
- [74] 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.
- [75] The OpenStack Foundation. *Cloud Edge Computing: Beyond the Data Center (White Paper)*. (Accessed: 2020-02-08). Jan. 2018. URL: <https://www.openstack.org/assets/edge/OpenStack-EdgeWhitepaper-v3-online.pdf>.
- [76] The OpenStack Foundation. *Edge Computing: Next Steps in Architecture, Design and Testing*. (Accessed: 2020-02-08). Jan. 2020. URL: <https://www.openstack.org/edge-computing/edge-computing-next-steps-in-architecture-design-and-testing>.
- [77] 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>.
- [78] 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.
- [79] 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).