

RESEARCH CENTRE

**Inria Center  
at Rennes University**

IN PARTNERSHIP WITH:

**Ecole Nationale Supérieure  
Mines-Télécom Atlantique Bretagne Pays  
de la Loire**

2022

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

|   |           |
|---|-----------|
| <b>Project-Team STACK</b>                                     | <b>1</b>  |
| <b>1 Team members, visitors, external collaborators</b>       | <b>2</b>  |
| <b>2 Overall objectives</b>                                   | <b>3</b>  |
| 2.1 STACK in a Nutshell . . . . .                             | 3         |
| 2.2 Context & Objectives . . . . .                            | 3         |
| 2.3 Scientific Foundations . . . . .                          | 4         |
| <b>3 Research program</b>                                     | <b>5</b>  |
| 3.1 Overview . . . . .  | 5         |
| <b>4 Application domains</b>                                  | <b>7</b>  |
| 4.1 Overview . . . . .  | 7         |
| 4.2 Industrial Internet . . . . .                             | 7         |
| 4.3 Internet of Skills . . . . .                              | 7         |
| 4.4 e-Health . . . . .  | 8         |
| 4.5 Network Virtualization and Mobile Edge Services . . . . . | 8         |
| <b>5 Social and environmental responsibility</b>              | <b>8</b>  |
| 5.1 Footprint of research activities . . . . .                | 8         |
| 5.2 Impact of research results . . . . .                      | 8         |
| <b>6 Highlights of the year</b>                               | <b>9</b>  |
| 6.1 Awards . . . . .  | 9         |
| <b>7 New software and platforms</b>                           | <b>9</b>  |
| 7.1 New software . . . . .                                    | 9         |
| 7.1.1 ENOS . . . . .  | 9         |
| 7.1.2 EnOSlib . . . . .                                       | 10        |
| 7.1.3 Concerto . . . . .                                      | 10        |
| 7.1.4 6TiSCH Simulator . . . . .                              | 11        |
| 7.2 New platforms . . . . .                                   | 11        |
| 7.2.1 OpenStack . . . . .                                     | 11        |
| 7.2.2 Grid'5000 . . . . .                                     | 11        |
| 7.2.3 SeDuCe . . . . .  | 12        |
| 7.2.4 PiSeDuCe . . . . .                                      | 12        |
| 7.2.5 SAMURAI . . . . .                                       | 12        |
| 7.2.6 SLICES-FR/SLICES . . . . .                              | 12        |
| <b>8 New results</b>  | <b>13</b> |
| 8.1 Resource Management . . . . .                             | 13        |
| 8.2 Programming Support . . . . .                             | 14        |
| 8.3 Energy-aware computing . . . . .                          | 16        |
| 8.4 Security and Privacy . . . . .                            | 16        |
| <b>9 Bilateral contracts and grants with industry</b>         | <b>18</b> |
| 9.1 Bilateral contracts with industry . . . . .               | 18        |
| 9.2 Bilateral grants with industry . . . . .                  | 18        |
| <b>10 Partnerships and cooperations</b>                       | <b>19</b> |
| 10.1 International research visitors . . . . .                | 19        |
| 10.1.1 Visits to international teams . . . . .                | 19        |
| 10.2 European initiatives . . . . .                           | 20        |
| 10.2.1 H2020 projects . . . . .                               | 20        |

|  |           |
|--|-----------|
| 10.3 National initiatives              | 20        |
| 10.3.1 ANR                             | 20        |
| 10.3.2 PIA 4                           | 21        |
| 10.3.3 CPER                            | 21        |
| 10.3.4 Local and regional projects     | 22        |
| 10.3.5 Inria Challenges                | 22        |
| <b>11 Dissemination</b>                | <b>22</b> |
| 11.1 Promoting scientific activities   | 23        |
| 11.1.1 Scientific events: organisation | 23        |
| 11.1.2 Scientific events: selection    | 23        |
| 11.1.3 Journal                         | 23        |
| 11.1.4 Invited talks                   | 24        |
| 11.1.5 Scientific expertise            | 24        |
| 11.1.6 Research administration         | 24        |
| 11.2 Teaching - Supervision - Juries   | 24        |
| 11.2.1 Teaching                        | 24        |
| 11.2.2 Supervision                     | 24        |
| 11.2.3 Juries                          | 25        |
| 11.3 Popularization                    | 26        |
| 11.3.1 Articles and contents           | 26        |
| 11.3.2 Education                       | 26        |
| 11.3.3 Interventions                   | 26        |
| <b>12 Scientific production</b>        | <b>26</b> |
| 12.1 Major publications                | 26        |
| 12.2 Publications of the year          | 27        |
| 12.3 Cited publications                | 29        |

## 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
- A1.3.4. – 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

## Faculty Members

- Adrien Lebre [Team leader, IMT ATLANTIQUE, Professor, HDR]
- H el ene Coullon [IMT ATLANTIQUE, Associate Professor]
- Mohamed Graiet [IMT ATLANTIQUE, Associate Professor, from Sep 2022]
- Remous-Aris Koutsiamanis [IMT ATLANTIQUE, Associate Professor]
- Thomas Ledoux [IMT ATLANTIQUE, Professor, HDR]
- Jean-Marc Menaud [IMT ATLANTIQUE, Professor, HDR]
- Jacques Noy e [IMT ATLANTIQUE, Associate Professor]
- Kandaraj Piamrat [UNIV NANTES, Associate Professor]
- Mario S udholt [IMT ATLANTIQUE, Professor, HDR]

## Post-Doctoral Fellows

- Abdelghani Alidra [IMT Atlantique, until Oct 2022]
- Yasmina Bouizem [INRIA, from Sep 2022]
- Eloi Perdereau [IMT ATLANTIQUE, from Sep 2022]

## PhD Students

- Geo Johns Antony [INRIA]
- Ons Aouedi [UNIV Nantes]
- Hiba Awad [Alter Way, CIFRE]
- Samia Boutalbi [ERICSSON, CIFRE]
- Divi De Lacour [ORANGE, CIFRE, from Feb 2022]
- Marie Delavergne [INRIA]
- Pierre Jacquet [INRIA]
- Dan Freeman Mahoro [ORANGE LABS, CIFRE, from Apr 2022]
- Duc-Thinh Ngo [ORANGE LABS, from Dec 2022]
- Antoine Omond [IMT ATLANTIQUE]
- Abdou Seck [IMT ATLANTIQUE, from Nov 2022]

## Technical Staff

- Alexis Bitailou [IMT ATLANTIQUE, Engineer, from Sep 2022]
- Baptiste Jonglez [INRIA, Engineer, from Mar 2022]

## Administrative Assistant

- Anne-Claire Binetruy [INRIA]

## 2 Overall objectives

### 2.1 STACK in a Nutshell

The STACK team addresses challenges related to the management of 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 (system) levels to the highest (application development) ones, 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 and Tactile Internet).

### 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 on a range 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 [78, 76]. 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 will be largely without human presence and will therefore have to be operated remotely and automatically as much as possible. 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 a few challenges related to the network dimension). Starting with Ass. Prof Koutsiamanis and Ass. Prof Piamrat (who respectively joined

the team in 2021 and 2022), this enlargement of STACK core activities will be further strengthened with the arrival of Orange members expected in 2023.

### 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.

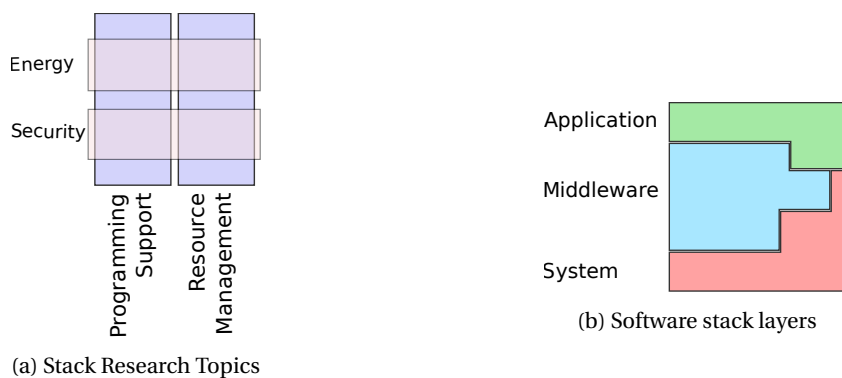
- **(Distributed) Systems.** The first scientific foundation of the team is related to our strong expertise in resource management and capacity planning of large-scale infrastructure [62, 72, 66, 82, 74, 75]. This includes the design and evaluation of system mechanisms and algorithms to operate and use 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 [37, 35, 34, 36]. 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 [65, 63, 45, 61, 48].
- **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 [73, 39, 38], 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 [71], semantic modelling and ontology mapping [40], graph storage distribution and historisation [59, 60] — and the application of these concepts and technologies in different use cases in the domain of smart building (e.g., localisation [48], dynamic wireless IoT network resource allocation [61]), smart industry (e.g., support for reliable and low-latency wireless Industrial IoT networks [65, 63]), logistics [70] around the Thing'in digital twin platform [52] (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 [55, 56, 69, 57, 68, 44, 46].
- **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 [53, 43, 54, 67, 51], event-driven [58, 81], data-driven and workflow models, as well as models for Utility Computing (SLA) [77], 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 [41] and more recently aspects related to software-defined experiments and reproducible research [50, 18]. 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 piece of software that can be transfer 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 transverse 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 model<sup>1</sup>, 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.).**

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 [83], 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

<sup>1</sup>MAPE-K stands for Monitoring, Analyzing, Planning, Execution and Knowledge [64].



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 infrastructure 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 of Concerto [47] extended to IoT devices and to network resources. Finally, we address the dynamicity of systems described by such ADL-like languages 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.
- *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 [49] further to handle geo-distribution concerns through a dedicated service mesh.
- *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 [42] 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 enable them to virtualize a large part of their resources and thus reduce capital and operational costs. As an example, several researchers have been investigating through the IOLab, the joint lab between Orange and Inria, how 5G networks can be managed. We highlight that while our expertise does partially include the network side, the main focus is rather on how we can deploy, locate and reconfigure the software components that are mandatory to operate next generation of network/computing infrastructure. The main challenges are related to the high dynamicity of the infrastructure, the way of defining Quality of Service of applications and how it can be guaranteed. We expect our contributions enable advances in location based services, optimized local content distribution (data-caching) and Mobile Edge Computing<sup>2</sup> In addition to bringing resources close to end-users, massively geo-distributed infrastructure should favor the development of more advanced network as well as mobile services.

### 4.1 Overview

Supporting industrial actors and open-source communities in building an advanced software management stack is a key element to favor the advent of new kinds of information systems as well as web applications. Augmented reality, telemedicine and e-health services, smart-city, smart-factory, smart-transportation and remote security applications are under investigations. Although, STACK does not intend to address directly the development of such applications, understanding their requirements is critical to identify how the next generation of ICT infrastructure should evolve and what are the appropriate software abstractions for operators, developers and end-users. STACK team members have been exchanging since 2015 with a number of industrial groups (notably Orange Labs and Airbus), a few medical institutes (public and private ones) and several telecommunication operators in order to identify both opportunities and challenges in each of these domains, described hereafter.

### 4.2 Industrial Internet

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

### 4.3 Internet of Skills

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

---

<sup>2</sup>ETSI MEC specifications.

is important to determine how the quality of service of such applications should be defined and how latency and bandwidth constraints can be guaranteed at the infrastructure level.

#### 4.4 e-Health

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

We have been working on different projects since the beginning of STACK (e.g., PrivGen CominLabs, ANR PICNIC). In general, we are applying and developing corresponding techniques for the medical domains of genomics, immunobiology and transplantology (see Section 10): there, we investigate how to secure and preserve privacy if potentially sensitive personal data is moved and processed by distributed biomedical analyses.

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 enable Telecom operators to virtualize a large part of their resources and thus reduce 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.

## 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 top10 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 remains virtual objects, it is executed on very real hardware contributing to the carbon footprint. This impact materializes during the manufacture / destruction of hardware infrastructure (estimated at 45% of digital consumption in 2018 by the 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 axe 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 2022, 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 impacts.
- The second one is involved Inria and Qarnot Computing, with the support of ADEME. Entitled Pulse, it 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 autonomies (see the Samurai project Section in 10).

## 6 Highlights of the year

Regarding scientific results, the team has produced major results on the management of large-scale infrastructures, in particular on the network side. We underline this point because it was an axis where the team was absent in previous years. The team also demonstrated once again its expertise on reproducible experiment driven research: the paper entitled "SMT-Based Planning Synthesis for Distributed System Reconfigurations" from H el ene Coullon and Simon Robillard has received the Springer badges "artifacts available" and "artifacts reusable", which is not yet usual in the software engineering community.

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 two research engineers of the team (Baptiste Jonglez and Alexis Bitailou).

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 Best paper award at IEEE International Conference on Communications (IEEE ICC 2022) for its paper [23] "Intrusion detection for Softwarized Networks with Semi-supervised Federated Learning" by Ons Aouedi, Kandaraj Piamrat, Guillaume Muller, Kamal Singh.

The team also participated in the [ITU-organised Machine Learning in 5G Challenge](#) and won the 1st place among over 29 research teams on the Problem Statement "Federated Traffic Prediction for 5G and Beyond" and 4th place in the grand challenge among the winners of all problem statements. The work consisted in presenting a benchmark for Federated Learning for 5G Base Station Traffic Forecasting, where various optimizations to existing methods were introduced.

Finally, we would like to highlight the scientific agreement in September 2022 between Inria and Orange to support the activities of STACK for the next 4 years (administrative validation is planned for the first quarter of 2023). This reorganization of the ongoing STACK research group is related to the integration of five researchers from Orange Labs leading to the first project-team under the umbrella of the IOLab, a joint lab between Inria and Orange.

## 7 New software and platforms

### 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, Alexis Bitailou, Marie Delavergne, Baptiste Jonglez, 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 [79], 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 [80] 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 a 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 2021, 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 PiSeducue 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.

### 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 bootstrap 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, Hélène Coullon, Remous-Aris Koutsiamanis, Adrien Lebre, Jean-Marc Menaud, Kandaraj Piamrat.

The evolution of the cloud computing paradigm in the last decade has amplified the access of 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. In the case of telecommunication operators, the network Point of Presence (PoP), which they have always operated, can be inexpensively extended to host these resources. The question is then how to manage such a massively Distributed Cloud Infrastructure (DCI) to provide end-users the same services that made the current cloud computing model so successful. In 2022, we have continued our effort to answer this question and delivered multiple contributions, with a new focus on the management of network devices.

**Replica indexation:** The first contribution we have made is related to how peers in a geo-distributed infrastructure could maintain an index of relevant replicas [26]. Although the holy grail of how to store and manipulate data in Edge infrastructures is yet to be found, state-of-the-art approaches demonstrated the relevance of replication strategies that bring content closer to consumers: The latter enjoy better response time while the volume of data passing through the network decreases overall. Unfortunately, locating the closest replica of a specific content requires indexing every live replica along with its location. Relying on remote services enters in contradiction with the properties of Edge infrastructures as locating replicas may effectively take more time than actually downloading the content. At the opposite, maintaining such an index at every node would prove overly costly in terms of memory and traffic, especially since nodes can create and destroy replicas at any time. In this paper, we abstract content indexing as distributed partitioning: every node only indexes its closest replica, and connected nodes with a similar index compose a partition. Our decentralized implementation AS-cast is (i) efficient, for it uses partitions to lock down the traffic generated by its operations to relevant nodes, yet it (ii) guarantees that every node eventually acknowledges its partition despite concurrent operations. Our complexity analysis supported by simulations shows that AS-cast scales well in terms of generated traffic and termination time. As such, AS-cast can constitute a new building block for geo-distributed services.

**Network resource management:** As aforementioned, the team has initiated new activities covering the Cloud-IoT continuum. Two recent topics in this context have been investigated resulting in two survey papers. First, in [15], we have surveyed the state of the art on intelligent management in next-generation networks, which can be applied to Cloud-IoT continuum. We started by presenting a comprehensive background beginning from conventional machine learning (ML) algorithms and deep learning (DL) and follow this with a focus on different dimensionality reduction techniques, which is a crucial topic for scalability. Afterward, we presented the study of ML/DL applications in a softwarized environment and highlighted the issues and challenges that should be considered. Second, as the Internet of Things (IoT) has remarkably evolved over the last few years, incentive techniques such as the blockchain, game theory, and Artificial Intelligence (AI) are highly desirable to build a sustainable IoT ecosystem. For these reasons, we have presented in [20] a systematic literature review of the incentive techniques for IoT, aiming to provide an overview of incentive-enabled IoT from background, motivations, and enabling techniques. Moreover, in the context of radio resource management in IoT networks, in [17], we present and evaluate an ultra-wideband (UWB) indoor processing architecture that allows the performing of simultaneous localizations of mobile tags, a network of low-power fixed anchors that provide forward-ranging measurements to a localization engine responsible for performing trilateration. The communications are scheduled globally (using UWB-TSCH), allowing deterministic channel access and low power consumption. We designed a centralized scheduler which organizes nodes into cells, allowing simultaneous



localizations and data transmissions, making indoor positioning system (IPS) more scalable and reducing deployment costs. We show through simulations that high positioning rates can be achieved, even in large (400-cell/400-tag) networks, with low schedule calculation times (less than 11s).

Finally, considering heterogeneity in emerging network traffic, traffic classification can enable a number of practical applications ranging from network monitoring to resource management. Numerous ML models have been applied to identify network applications. However, among the applied models so far, no model outperforms all the others. To solve these issues, we have proposed in [14] a novel DL-based approach that incorporates multiple Decision Tree based models. This approach employs a non-linear blending ensemble method by combining tree-based classifiers through DL in order to maximize generalization accuracy. This ensemble consists of two levels called base classifiers and meta-classifiers. In the first level, Decision Tree-based models are used as the base classifiers while in the second level, DL is used as a meta-model to combine the outputs of the base classifiers.

**Resource management for scientific workflows:** The goal of a workflow engine is to facilitate the writing, the deploying, and the execution of a scientific workflow (i.e., graph of coarse-grain and heterogeneous tasks) on distributed infrastructure. With the democratization of the Cloud paradigm, many workflow engines of the state-of-the-art offer a way to execute workflows on distant data centers by using the Infrastructure-as-a-Service (IaaS) or the Function-as-a-Service (FaaS) services of Cloud providers. Hence, workflow engines can take advantage of the (presumably) infinite resources and the economical model of the Cloud. However, two important limitations lie in this vision of Cloud-oriented workflow engines. First, by using existing services of Cloud providers, and by managing the workflows at the user side, the Cloud providers are unaware of both the workflows and their user needs, and cannot apply specific resource optimizations to their infrastructure. Second, for the same reasons, handling the heterogeneity of tasks (different operating systems) in workflows necessarily degrades either the transparency for the users (who must provision different types of resources), or the completion time performance of the workflows, because of the stacking of virtualization layers. In [24], we tackle these two limitations by presenting a new Cloud service dedicated to scientific workflows. Unlike existing workflow engines, this service is deployed and managed by the Cloud providers, and enables specific resource optimizations and offers a better control of the heterogeneity of the workflows. We evaluate our new service in comparison to Argo, a well-known workflow engine of the literature based on FaaS services. This evaluation was made on a real distributed experimental platform with a realistic and complex scenario.

## 8.2 Programming Support

**Participants:** Abdelghani Alidra, Geo Johns Antony, H el ene Coullon, Marie Delavergne, Mohamed Graiet, Thomas Ledoux, Adrien Lebre, Jolan Philippe.

**Fog Modeling:** In [11], 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.

**Model-Driven Engineering (MDE):** MDE is a software programming approach that raises the level of abstraction in traditional programming languages by using models with recurring design patterns. MDE simplifies the collaboration between development teams and more broadly promotes compatibility

between systems. MDE is increasingly used to help in the development of distributed systems, microservices architectures, and IoT systems but is also leveraged, for instance, in lowcode platforms such as Node-RED to write IoT applications.

In the thesis of Jolan Philippe (defended in December 2022) two important drawbacks of MDE have been studied. First, as most of the high-abstraction level programming models, when dealing with very big models most MDE frameworks become poorly efficient and slow, thus losing their initial advantage of speeding up the development process. Second, because of the underlying complexity of MDE approaches, the level of confidence in MDE frameworks is usually low. In the thesis the formal transformation engine (transformation being an important part of MDE approaches) CoqTL has been extended with a more scalable implementation that is proven equivalent to the initial one, and that is automatically dumped to a Apache Spark code. A prototype has been implemented on top of Spark, and has been evaluated. In addition to this, the thesis also explores how to automatically tune our Spark-based transformation engine according to the input models (size, nature etc.). This contribution notably presents a feature model of different configurations of the transformation engine: a feature model that could be used in future work to automatically configure and reconfigure the engine, thus opening a door to the second part of this section.

In [22], we address the issue of formally modeling the architecture of the IoT in order to be able to provide security guarantees and/or to respect real-time properties. In this work we proposed a novel correct-by-construction formal approach based on an Event-B method to describe the physical architecture of IoT layers. This formal approach inspects four layers: the physical layer, the gateway layer, the middleware layer, and the application layer. An Electrocardiogram (ECG) IoT system is applied in our model as a case study. Finally, we proved and validated the correctness of our formal model by using proof obligations and the model checking tool called Rodin.

**Software deployment/reconfiguration:** If speeding up and simplifying the development process of distributed systems is of prior importance because of the increasing complexity of geographically distributed infrastructures, speeding up and simplifying the deployment of these systems, and enabling and speeding up their dynamic evolution through time is another key to making these infrastructures usable. Indeed, because of the scale of both geo-distributed infrastructure and distributed systems (e.g., microservices architectures) manually handling deployments and reconfiguration procedures is an error-prone, tedious and complex task, probably impossible in practice. Hence, these procedures need to be automated, or even better they need to become completely autonomous.

In [27], we tackle the planning problem for the reconfiguration of distributed systems in the component-based reconfiguration model Concerto. Specifically, given some tasks to execute and a desired final state of the system, we show how to compute a reconfiguration plan that guarantees satisfaction of inter-component dependencies and which is also optimized for parallel execution. Our technique relies on an SMT solver to compute the required dependencies between components and ultimately schedule the reconfiguration. We illustrate the use of this technique on a variety of synthetic examples as well as a real use case in the context of an OpenStack system.

#### **Reifying geo-distribution at the software level:**

One question to answer in the shift from the Cloud to the Edge computing paradigm is: how distributed applications developed for Cloud platforms can benefit from the opportunities of the Edge while dealing with inherent constraints of wide-area network links? Our solution to this question is to give the illusion of “single service images” spreading over the Edge infrastructure. Thanks to the modularity of micro-service based applications, one can deploy multiple instances of the same service (one per edge site) and deliver collaborations between them according to each request. This non-invasive approach is made possible by (i) a DSL that extends the application API and allows DevOps to program where/how the execution of each request should be executed, (ii) and its runtime, Cheops, a service that interprets and orchestrates each request in order to satisfy the geo-distribution parameters, allowing collaborations in a transparent manner for the underlying application. We demonstrate the relevance of our proposal by illustrating how Cheops can successfully geo-distribute the Kubernetes vanilla code [32, 25].

### 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).

### 8.4 Security and Privacy

**Participants:** Ons Aouedi, Wilmer Edicson Garzon Alfonso, Mohamed Graiet, Kan-daraj Piamrat, Sirine Sayadi, Mario Südholt.

This year we have provided multiple contributions related to the security and privacy properties of large-scale Cloud-based systems, notably for applications in the medical domain.

#### Secure and privacy-preserving distributed genetic analyses:

The amount of biological data collected and stored has grown significantly. By the amount of data collected, traditional scenarios are not feasible due to technical and legal restrictions moving data, especially on data of medical origin. This has motivated the scientific community to create global collaborations to analyze biomedical and share scientific findings responsibly. However, current tools do not support safely and efficiently these collaborations. In this paper, we review the state of the art identifying the needs to support these collaborations, which we characterize as Fully Distributed Collaborations (FDCs). We also investigate the technical and legal restrictions of sharing biomedical data. Additionally, we present a taxonomy of current tools to analyze biomedical data in distributed settings. The taxonomy considers three architectural key features to support FDC scenarios: data and computation placement, Privacy and Security, and, Performance and Scalability. The review reveals opportunities for multi-site analyses that encourage scientific collaborations while mitigating technical and legal constraints. Finally, we present these opportunities that serve as future work in this direction to promote collaborations around the analysis of biomedical data. [19]

The Allele Frequency Net Database (AFND) is a repository that allows users to research and analyze immune gene frequencies in different populations around the world. This database contains allele, haplotype and genotype format information. They are collected in a common database based on the results and the work of users. With the massive increase in medical data and the strengthening of data governance laws, the proposal for a new distributed and secure model for the historically centralized method in AFND has become important. In [21], we have developed Distributed AFND, an alternative distributed version of AFND that allows users to perform their research and analysis in cooperation with other remote sites without sharing their original data and monitoring data access.

In the context of Sirine Sayadi's PhD thesis [31], we have considered the development of distributed biomedical data analyses as well as their deployment in IT infrastructure and their integration into real medical tools. As part of a kidney transplant application, we have provided a new distributed contextualization solution that helps clinicians assess patients' kidney problems. For the same application, we have also proposed a new distributed version of Factor Analysis of Mixed Data (FAMD) for dimension reduction. In the context of analyses based on HLA genetic data, we have proposed a new distributed and secure model for the estimation of allele frequencies, haplotype frequencies and individual genotype frequencies. We also propose a new algorithm for the distributed estimation of HLA haplotype frequency using the expectation-maximization EM algorithm.

Despite the profitable usage of AI-based algorithms in the medical domain, these data-driven methods are facing issues such as the scarcity and privacy of user data, as well as the difficulty of institutions exchanging medical information. To solve these issues, Federated Learning (FL) appeared as a valuable

approach in the medical field, allowing patient data to stay where it is generated. However, FL is still unable to deliver all its promises and meets the more stringent requirements (e.g., latency, security) of a healthcare system based on multiple Internet of Medical Things (IoMT). In [16], we focused on the emerging deployment of FL, provide a broad overview of current approaches and existing challenges, and outline several directions of future work that are relevant to solving existing problems in federated healthcare, with a particular focus on security and privacy issues.

**Network security and privacy:** The rapid development of the Internet and smart devices trigger surges in network traffic making its infrastructure more complex and heterogeneous. The predominant usage of mobile phones, wearable devices and autonomous vehicles are examples of distributed networks, which generate a huge amount of data each and every day. The computational power of these devices have also seen steady progression which has created the need to transmit information, store data locally and drive network computations towards edge devices. Intrusion detection systems play a significant role in ensuring security and privacy of such devices. In this domain, Machine Learning and Deep Learning with Intrusion Detection Systems have gained great momentum due to their achievement of high classification accuracy.

However, the privacy and security aspects potentially get jeopardized due to the need of storing and communicating data to a centralized server. On the contrary, federated learning (FL) fits in appropriately as a privacy-preserving decentralized learning technique that does not transfer data but trains models locally and transfers the parameters to the centralized server. In [9], we aimed to present an extensive and exhaustive review on the use of FL in intrusion detection system. In order to establish the need for FL, various types of IDS, relevant ML approaches and its associated issues are discussed. The paper presents detailed overview of the implementation of FL in various aspects of anomaly detection. The allied challenges of FL implementations are also identified which provides an idea on the scope of future direction of research. The paper finally presents the plausible solutions associated with the identified challenges in FL based intrusion detection system implementation acting as a baseline for prospective research.

Security has also become a critical issue for Industry 4.0 due to different emerging cyber-security threats. Despite the huge amount of data generated by the Internet of Things (IoT) devices in Industry 4.0, it is difficult to get labeled data, because data labeling is costly and time-consuming. This poses many challenges for several DL approaches, which require labeled data. To handle these issues, we proposed in [13] a novel federated semi-supervised learning scheme, that takes advantage of both unlabeled and labeled data in a federated way. First, an AutoEncoder (AE) is trained on each device (using unlabeled local/private data) to learn the representative and low-dimensional features. Then, a cloud server aggregates these models into a global AE using Federated Learning (FL). Finally, the cloud server composes a supervised neural network, by adding fully connected layers (FCN) to the global encoder (the first part of the global AE) and trains the resulting model using publicly available labeled data.

Moreover, with the increasing development of beyond 5G and network softwarization, there is more more flexibility and agility in the network than ever. This can be exploited to integrate intelligence into the network, for example, for conceiving intelligent Intrusion detection systems (IDS). However, traditional approaches in this domain require all data (and their associated labels) to be centralized. Such approaches lead to: (i) a large bandwidth overhead, as raw data needs to be transmitted to the server, (ii) low incentives for devices to send their private data, and (iii) large computing and storage resources needed on the server side to label and treat all this data. To cope with the above limitations, we have proposed in [23], a semi-supervised federated learning model for IDS. Moreover, we use network softwarisation for automation and deployment. Our model combines Federated Learning and Semi-Supervised Learning where the clients train unsupervised models (using unlabeled data) to learn the representative and low-dimensional features and the server conducts a supervised model (using labeled data).

Finally, in [28] we have addressed the challenge of the regulation of traffic violations in intelligent transportation system networks. We have proposed using a blockchain smart contract-based method to perform vehicle speed detection and vehicle information collection through vehicle re-identification

using smart contracts. The smart contract is a conditioned filter that follows regulatory rules from reporting violation points and fines for each violation to penalties. We have implemented the proposed algorithm and we have presented an evaluation of the proposed method leading to satisfactory results in comparison to other known methods.

## 9 Bilateral contracts and grants with industry

### 9.1 Bilateral contracts with industry

#### Kelio (formely Bodet Software)

**Participants:** Thomas Ledoux.

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

In 2022, several activities were conducted, such as a hackathon with 70 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.

#### OVHcloud

**Participants:** 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.

#### Orange

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

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

Dan Freeman Mahoro initiated a new collaboration on the topic of the “digital twin”, a dynamic virtual copy of a physical asset, process, system or environment that looks and behaves identically to its real-world counterpart. This cooperation has resulted in a joint doctoral dissertation entitled "Digital twins of complex cyber-physical systems," which began in April 2022, under a co-supervision with Orange team in Rennes.

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 to international teams

##### Research stays abroad

###### Kandaraj Piamrat

**Visited institution:** National Institute of Informatics (NII)

**Country:** Japan

**Dates:** 24/10/2022-5/11/2022

**Context of the visit:** Initiating collaboration on intelligent network management

**Mobility program/type of mobility:** NII MOU Grant, research stay

###### Ons Aouedi

**Visited institution:** National Institute of Informatics (NII)

**Country:** Japan

**Dates:** 1/09/2022-5/11/2022

**Context of the visit :** NII internship

**Mobility program/type of mobility:** Internship



## 10.2 European initiatives

### 10.2.1 H2020 projects

#### SLICES-PP

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

SLICES-RI (Research Infrastructure), which was recently included in the 2021 ESFRI roadmap, aims to answer these problems by 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 (44K€) but the group is strongly involved at the national level taking part to different discussions related to SLICES-FR, the French node of SLICES.

## 10.3 National initiatives

### 10.3.1 ANR

#### SeMaFoR (Self-Management of Fog Resources)

**Participants:** Thomas Ledoux (*coordinator*), H el ene Coullon, Abdelghani Alidra.

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 e (Delys), Alter way/Smile (SME). The Stack team supervises the project.

The main result of the year 2022 was the survey "A feature-based survey of Fog modeling languages" [11].

SeMaFoR is running for 42 months (starting in March 2021 with an allocated budget of 506k€, 230K€ for Stack). See [Semafor](#) web site for more information.

#### PicNic (Transfert de grands volumes de donn ees 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 Parallélisme, Institut de Cancérologie de l'Ouest / Informatique, Institut de Recherche en Informatique de Toulouse, Laboratoire des Sciences du Numérique 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*), Alexis Bitailou, Hélène Coulon, Remous-Aris Koutsiamanis, Thomas Ledoux, Jean-Marc Menaud, Jacques Noyé, Kandaraj Piamrat, Eloi Perdereau, Mario Südholt.

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 (Agileo Automation, Mydatamodels, Dupliprint, 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 autoNoMoUs 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 recherche that is part of the Nantes Excellence Initiative in Medecine and Engineering. Its main objective is the development of new methods for precision medecine, 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:** Hélène Coullon, Thomas Ledoux, Pierre Jacquet.

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

The PhD thesis of Pierre Jacquet is positioned in the 1st axis and addresses the following question: fostering the Frugal Design of Cloud Native Applications.

Hélène Coullon is positioned in the 2nd axis and addresses the following question: pulling the energy cost of a reconfiguration execution within reconfiguration decisions.

#### 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

**Participants:** Hélène Coullon, Baptiste Jonglez, Remous-Aris Koutsiamanis, Adrien Lebre, Thomas Ledoux, Jacques Noyé, Jean-Marc Menaud, Kandaraj Piamrat, Mario Südholt.

## 11.1 Promoting scientific activities

### 11.1.1 Scientific events: organisation

#### Member of the organizing committees

- Hélène Coullon: Publicity co-chair CCGrid 2022.
- Adrien Lebre: Member of the steering committee of the international conference of Fog and Edge Computing (ICFEC).
- Remous-Aris Koutsiamanis: Member of the organisation committee of the "Journées du GDR Réseaux et Systèmes distribués", (27-28 Apr. 2022).

### 11.1.2 Scientific events: selection

#### Chair of conference program committees

- Kandaraj Piamrat : Co-chair of CoReS 2022

#### Member of the conference program committees

- Hélène Coullon: ICCS 2022, SBAC-PAD 2022, ICE 2022, CompPAS 2022
- Adrien Lebre: UCC 2022, CloudCom 2022, CloudNet 2022, ICFEC 2022.
- Kandaraj Piamrat: IEEE ICC (CSM symposium) 2022, IEEE CCNC (IoT: From Sensors to Vertical Applications Track) 2022, IEEE IWQoS 2022, IEEE WiMob 2022, IEEE ICC Workshop on Data Driven Intelligence for Networks and Systems (DDINS) 2022
- Jean-Marc Menaud: MODERN SYSTEMS'22, SMARTGREENS'22, SDS'22.
- Mario Südholt: CloudCom 2022

#### Reviewer

- Kandaraj Piamrat : IEEE CloudNet, IEEE ICC, IEEE CCNC, IEEE IWQoS, IEEE WiMob, IEEE ICC workshop DDINS
- Remous-Aris Koutsiamanis : IEEE CSCN 2022, IEEE WiMob 2022, CoReS 2022, IEEE ISCC 2022, ICCS 2022, IEEE ICC 2022

### 11.1.3 Journal

#### Member of the editorial boards

- Adrien Lebre : Associate Editor of the IEEE Transactions on Cloud Computing.

#### Member of steering committees

- Mario Südholt : Member of the steering committee of the internal journal and conference of Programming.

#### Reviewer - reviewing activities

- Adrien Lebre: ACM Transactions on Computer.
- Kandaraj Piamrat : Computer Networks (Elsevier), IEEE Transactions on cognitive communications and networking
- Remous-Aris Koutsiamanis : Wireless Networks (Springer), IEEE Internet of Things Journal, Ad Hoc Networks (Elsevier)

#### 11.1.4 Invited talks

- Adrien Lebre : keynote speaker at [CIOT 2022](#) and [SSS 2022](#), invited speaker at "journée Infrastructures pour la souveraineté numérique" (CNAM Paris).

#### 11.1.5 Scientific expertise

- Adrien Lebre : Member of the steering committee of the PEPR Cloud proposal (IMT representative).
- Kandaraj Piamrat : Reviewer for Information and Communication Technology Call 2022 at the Vienna Science and Technology Fund (WWTF)

#### 11.1.6 Research administration

- Adrien Lebre : Member of the executive committee of the Grid'5000 GIS (Groupement d'intérêt scientifique), Co-director of the [I/O Lab](#), a joint lab between Inria and Orange Labs.
- Jean-Marc Menaud : Deputy Director of the Laboratory of Digital Sciences of Nantes (LS2N) UMR CNRS 6004.

### 11.2 Teaching - Supervision - Juries

#### 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.

- 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.
- Hélène Coullon : Responsible for the LOGIN training in computer science at IMT Atlantique (last year)
- Thomas Ledoux : Head of the Filière 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 numérique](#).
- Jacques Noyé : Deputy head of the Automation, Production and Computer Sciences Department of IMT Atlantique.
- Mario Südholt : Representative for MSc-level and PhD-level studies of the API department of IMT Atlantique.
- 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, Sept 2019, March 2023, Director: A. Lebre.
- PhD: Geo Johns Anthony, Resource management in geo-distributed infrastructures: the Kubernetes case, April 2021, Sept 2024, Director: A. Lebre.

- PhD: Jolan Philippe, Contribution to the analysis of the design-space of a distributed transformation engine, Setp 2019, Dec 2022, Director: G. Sunye (NaoMod Nantes), Advisors: H. Coullon, M. Tisi (NaoMod).
- PhD: Wilmer Garzon, Secure distributed workflows for biomedical data analytics, Sept 2019, Jan 2023, Director: M. Südholt, Advisor: D. Benavides.
- PhD: Sirine Sayadi, Distributed architectures and secure software containers for multi-site medical cooperation, Sept. 2019, Dec 2022, Director: M. Südholt.
- PhD: 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: Abdou Seck, Parallel transfer service for the exchange of large volumes of data between Datacenters, June 2022, June 2025, Director: J.-M. Menaud, Advisor: R.-A. Koutsiamanis.
- PhD: Pierre Jacquet, Fostering the Frugal Design of Cloud Native Applications, Oct 2021, Oct 2024, Director R. Rouvoy (SPIRALS Lille), Advisor: T. Ledoux.
- PhD: Hiba Awad, A Model-based Approach for Multi-Scale and Dynamic Distributed Systems, Nov. 2021 - March 2025. Director: T. Ledoux.
- PhD: Ons Aouedi, Machine learning-Enabled Network Traffic Analysis, Sept. 2019, Nov. 2022, Director: B. Parrein (STR, Nantes) Advisor: K. Piamrat.
- PhD: 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: Dan Freeman Mahoro, Digital twins of complex cyber-physical systems, Apr. 2022 - March 2025, Director: T. Ledoux (since Apr. 2022).
- Post-doc: Abdelghani Alidra, Modeling Fog-based systems, Sept 2021, Oct 2022, Advisor: T. Ledoux.
- 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.
- Resarch engineer: Alexis Bitailou, Development support within the OTPaaS project, Sept 2022, Jan 2023, Advisor: A. Lebre.

### 11.2.3 Juries

- Hélène Coullon was a member of the selection committee for two associate professor positions in Rennes.
- Adrien Lebre was a member of the selection committee for two associate professor positions (Evry, Toulouse) and one professor position (Grenoble).
- Thomas Ledoux was the chairman of the PhD thesis jury of Yuwei Wang, "Evolution of Microservice-based Applications: Modelling and Safe Dynamic Updating", Télécom SudParis, Oct. 27, 2022.
- Thomas Ledoux was the chairman of the PhD thesis jury of Jolan Philippe, "Contribution to the Analysis of the Design-Space of a Distributed Transformation Engine", IMT Atlantique, Dec. 19, 2022.
- Jean-Marc Menaud was a member of the selection committee for a professor position in Grenoble.
- Jean-Marc Menaud was an examiner of Safuriyawu Ahmed Thesis - Resilient IoT-based Monitoring System for the Nigerian Oil and Gas Industry. INSA Lyon, 12/22.

- Jean-Marc Menaud was a reviewer of Humberto Alvarez - An energy saving perspective for distributed environments: Deployment, scheduling and simulation with multidimensional entities for Software and Hardware - Université de Pau et Des Pays de l'Adour, 06/22.
- Jean-Marc Menaud was a reviewer of Brice EKANE APAH - Optimisation des entrées-sorties dans les architectures multi-tiers Input-output optimization in multi-tiers architectures Université Grenoble Alpes - 12/22.
- Jean-Marc Menaud was a reviewer of Houssam KANSO - Contributing to the Energy Efficiency of Smart Homes : An Automated Management Framework - Université de Pau et Des Pays de l'Adour, 12/22.
- Remous-Aris Koutsiamanis was an examiner of the thesis of Charlier Maximilien entitled "High density and large scale Ultra Wideband indoor positioning infrastructure", University of Mons (UMons), Belgium, May 3, 2022.
- Remous-Aris Koutsiamanis was an examiner of the thesis of David Hauweele entitled "Optimization of Radio Duty Cycling protocols in Wireless Sensor Networks", University of Mons (UMons), Belgium, December 16, 2022.

## 11.3 Popularization

### 11.3.1 Articles and contents

- Thomas Ledoux , "[Tribune] Pour une informatique moins gourmande en énergie, formons les ingénieurs", L'Usine nouvelle, Jan. 2022.
- Jean-Marc Menaud, Énergie. Les data centers consomment-ils plus d'électricité que la SNCF ? Ouest-France Nov. 2022.
- Jean-Marc Menaud, Électricité : le numérique trop gourmand. Ouest-France Oct. 2022.
- Jean-Marc Menaud, La transition énergétique passe aussi par la formation et la recherche, Le Cortex Oct. 2022.

### 11.3.2 Education

Thomas Ledoux is co-leader of the Ecolog initiative. This national project is the result of the 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 is to create a body of training courses available in open source and dedicated to green computing.

### 11.3.3 Interventions

Thomas Ledoux took part in three round tables related to digital sustainability/green computing: "Before Nantes Digital Week 2022" organized by Nantes Métropole (Sept. 15, 2022); Technoférence "Le numérique face à ses responsabilités" organized by Pôle de compétitivité Images & Réseaux (Oct. 06, 2022); conference GreenTech Forum 2022 organized by Planet Tech'Care - Numeum (Dec 1st, 2022).

## 12 Scientific production

### 12.1 Major publications

- [1] S. Agrawal, S. Sarkar, O. Aouedi, G. Yenduri, K. Piamrat, M. Alazab, S. Bhattacharya, P. K. R. Maddikunta and T. R. Gadekallu. 'Federated Learning for intrusion detection system: Concepts, challenges and future directions'. In: *Computer Communications* (Sept. 2022), pp. 1–19. DOI: [10.1016/j.comcom.2022.09.012](https://doi.org/10.1016/j.comcom.2022.09.012). URL: <https://hal.archives-ouvertes.fr/hal-03779486>.

- [2] E. Ahvar, A.-C. Orgerie and A. Lebre. ‘Estimating Energy Consumption of Cloud, Fog and Edge Computing Infrastructures’. In: *IEEE Transactions on Sustainable Computing* 7.2 (Apr. 2022), pp. 277–288. DOI: [10.1109/TSUSC.2019.2905900](https://doi.org/10.1109/TSUSC.2019.2905900). URL: <https://hal.archives-ouvertes.fr/hal-02083080>.
- [3] 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.archives-ouvertes.fr/hal-03759010>.
- [4] O. Aouedi, K. Piamrat and D. Bagadthey. ‘Handling partially labeled network data: a semi-supervised approach using stacked sparse autoencoder’. In: *Computer Networks* 207 (22nd Apr. 2022), pp. 1–12. DOI: [10.1016/j.comnet.2021.108742](https://doi.org/10.1016/j.comnet.2021.108742). URL: <https://hal.archives-ouvertes.fr/hal-03524935>.
- [5] 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* (9th Mar. 2022), pp. 1–10. DOI: [10.1109/TII.2022.3156642](https://doi.org/10.1109/TII.2022.3156642). URL: <https://hal.archives-ouvertes.fr/hal-03602788>.
- [6] O. Aouedi, K. Piamrat and B. Parrein. ‘Ensemble-based Deep Learning model for network traffic classification’. In: *IEEE Transactions on Network and Service Management* (22nd July 2022), pp. 1–12. DOI: [10.1109/TNSM.2022.3193748](https://doi.org/10.1109/TNSM.2022.3193748). URL: <https://hal.archives-ouvertes.fr/hal-03736603>.
- [7] O. Aouedi, A. Sacco, K. Piamrat and G. Marchetto. ‘Handling Privacy-Sensitive Medical Data With Federated Learning: Challenges and Future Directions’. In: *IEEE Journal of Biomedical and Health Informatics* (23rd June 2022), pp. 1–14. DOI: [10.1109/JBHI.2022.3185673](https://doi.org/10.1109/JBHI.2022.3185673). URL: <https://hal.archives-ouvertes.fr/hal-03703925>.
- [8] W. Garzón, L. D. Benavides Navarro, A. Gaignard, R. Redon and M. Südholt. ‘A taxonomy of tools and approaches for distributed genomic analyses’. In: *Informatics in Medicine Unlocked* 32 (2022), pp. 1–17. DOI: [10.1016/j.imu.2022.101024](https://doi.org/10.1016/j.imu.2022.101024). URL: <https://hal.archives-ouvertes.fr/hal-03748752>.

## 12.2 Publications of the year

### International journals

- [9] S. Agrawal, S. Sarkar, O. Aouedi, G. Yenduri, K. Piamrat, M. Alazab, S. Bhattacharya, P. K. R. Maddikunta and T. R. Gadekallu. ‘Federated Learning for intrusion detection system: Concepts, challenges and future directions’. In: *Computer Communications* (Sept. 2022), pp. 1–19. DOI: [10.1016/j.comcom.2022.09.012](https://doi.org/10.1016/j.comcom.2022.09.012). URL: <https://hal.science/hal-03779486>.
- [10] E. Ahvar, A.-C. Orgerie and A. Lebre. ‘Estimating Energy Consumption of Cloud, Fog and Edge Computing Infrastructures’. In: *IEEE Transactions on Sustainable Computing* 7.2 (Apr. 2022), pp. 277–288. DOI: [10.1109/TSUSC.2019.2905900](https://doi.org/10.1109/TSUSC.2019.2905900). URL: <https://hal.science/hal-02083080>.
- [11] 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>.
- [12] O. Aouedi, K. Piamrat and D. Bagadthey. ‘Handling partially labeled network data: a semi-supervised approach using stacked sparse autoencoder’. In: *Computer Networks* 207 (22nd Apr. 2022), pp. 1–12. DOI: [10.1016/j.comnet.2021.108742](https://doi.org/10.1016/j.comnet.2021.108742). URL: <https://hal.science/hal-03524935>.
- [13] 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* (9th Mar. 2022), pp. 1–10. DOI: [10.1109/TII.2022.3156642](https://doi.org/10.1109/TII.2022.3156642). URL: <https://hal.science/hal-03602788>.
- [14] O. Aouedi, K. Piamrat and B. Parrein. ‘Ensemble-based Deep Learning model for network traffic classification’. In: *IEEE Transactions on Network and Service Management* (22nd July 2022), pp. 1–12. DOI: [10.1109/TNSM.2022.3193748](https://doi.org/10.1109/TNSM.2022.3193748). URL: <https://hal.science/hal-03736603>.

- [15] O. Aouedi, K. Piamrat and B. Parrein. ‘Intelligent Traffic Management in Next-Generation Networks’. In: *Future internet* (28th Jan. 2022), pp. 1–35. DOI: [10.3390/fi14020044](https://doi.org/10.3390/fi14020044). URL: <https://hal.science/hal-03546653>.
- [16] O. Aouedi, A. Sacco, K. Piamrat and G. Marchetto. ‘Handling Privacy-Sensitive Medical Data With Federated Learning: Challenges and Future Directions’. In: *IEEE Journal of Biomedical and Health Informatics* (23rd June 2022), pp. 1–14. DOI: [10.1109/JBHI.2022.3185673](https://doi.org/10.1109/JBHI.2022.3185673). URL: <https://hal.science/hal-03703925>.
- [17] 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* (2nd Mar. 2022), pp. 219–258. DOI: [10.3390/iot3010013](https://doi.org/10.3390/iot3010013). URL: <https://hal.inria.fr/hal-03919553>.
- [18] 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 (1st June 2022), pp. 1464–1477. DOI: [10.1109/TPDS.2021.3111159](https://doi.org/10.1109/TPDS.2021.3111159). URL: <https://hal.inria.fr/hal-03324177>.
- [19] W. Garzón, L. D. Benavides Navarro, A. Gaignard, R. Redon and M. Südholt. ‘A taxonomy of tools and approaches for distributed genomic analyses’. In: *Informatics in Medicine Unlocked* 32 (2022), pp. 1–17. DOI: [10.1016/j.imu.2022.101024](https://doi.org/10.1016/j.imu.2022.101024). URL: <https://hal.science/hal-03748752>.
- [20] P. Kumar Reddy Maddikunta, Q.-V. Pham, D. C. Nguyen, T. Huynh-The, O. Aouedi, G. Yenduri, S. Bhattacharya and T. Reddy Gadekallu. ‘Incentive techniques for the Internet of Things: A survey’. In: *Journal of Network and Computer Applications (JNCA)* (July 2022), pp. 1–25. DOI: [10.1016/j.jnca.2022.103464](https://doi.org/10.1016/j.jnca.2022.103464). URL: <https://hal.science/hal-03737297>.
- [21] S. Sayadi, V. Douillard, N. Vince, M. Südholt and P.-A. Gourraud. ‘Distributing human leukocyte antigen HLA database in histocompatibility: a shift in HLA data governance’. In: *Exploration of Immunology 2* (1st Nov. 2022), pp. 749–759. DOI: [10.37349/ei.2022.00080](https://doi.org/10.37349/ei.2022.00080). URL: <https://hal.science/hal-03747555>.
- [22] Z. H. Toman, L. Hamel, S. H. Toman and M. Graiet. ‘Correct-by-Construction Approach for Formal Verification of IoT Architecture’. In: *Procedia Computer Science* 207 (2022), pp. 2598–2609. DOI: [10.1016/j.procs.2022.09.318](https://doi.org/10.1016/j.procs.2022.09.318). URL: <https://hal.science/hal-03938808>.

#### International peer-reviewed conferences

- [23] O. Aouedi, K. Piamrat, G. Muller and K. Singh. ‘Intrusion detection for Softwarized Networks with Semi-supervised Federated Learning’. In: *ICC 2022 - IEEE International Conference on Communications*. ICC 2022 - IEEE International Conference on Communications. Seoul, South Korea: IEEE, 16th June 2022. URL: <https://hal.science/hal-03544099>.
- [24] E. Cadorel, H. Coullon and J.-M. Menaud. ‘Handling heterogeneous workflows in the Cloud while enhancing optimizations and performance’. In: *CLOUD 2022 - IEEE 15th International Conference on Cloud Computing*. Barcelona, Spain: IEEE, 10th July 2022, pp. 49–58. DOI: [10.1109/CLOUD55607.2022.00021](https://doi.org/10.1109/CLOUD55607.2022.00021). URL: <https://hal.science/hal-03922772>.
- [25] M. Delavergne, G. J. Antony and A. Lebre. ‘Cheops, a Service to Blow Away Cloud Applications to the Edge’. In: *Lecture Notes in Computer Science*. ICSOC 2022 - 20th International Conference on Service-Oriented Computing. Vol. LNCS-13740. Service-Oriented Computing. Sevilla, Spain: Springer Nature Switzerland, 22nd Nov. 2022, pp. 530–539. DOI: [10.1007/978-3-031-20984-0\\_37](https://doi.org/10.1007/978-3-031-20984-0_37). URL: <https://hal.inria.fr/hal-03926688>.
- [26] A. Lebre, B. Nédelec and A. van Kempen. ‘AS-cast: Lock Down the Traffic of Decentralized Content Indexing at the Edge’. In: *Meng, W., Lu, R., Min, G., Vaidya, J. (eds) Algorithms and Architectures for Parallel Processing, Lecture Notes in Computer Science book series (LNCS)*. ICA3PP 2022 - 22nd International Conference on Algorithms and Architectures for Parallel Processing. Vol. 13777. Copenhagen, Denmark, 2022, pp. 433–454. DOI: [10.1007/978-3-031-22677-9\\_23](https://doi.org/10.1007/978-3-031-22677-9_23). URL: <https://hal.inria.fr/hal-03931226>.



- [27] S. Robillard and H. Coullon. ‘SMT-Based Planning Synthesis for Distributed System Reconfigurations’. In: FASE 2022 : 25th International Conference on Fundamental Approaches to Software Engineering. Vol. 13241. Lecture Notes in Computer Science. Munich, Germany, 2022, pp. 268–287. DOI: [10.1007/978-3-030-99429-7\\_15](https://doi.org/10.1007/978-3-030-99429-7_15). URL: <https://hal.inria.fr/hal-03536643>.
- [28] N. Yousfi, M. Kmimech, I. Abbassi, H. Hamdi and M. Graiet. ‘ITS Traffic Violation Regulation Based on Blockchain Smart Contracts’. In: *Communications in Computer and Information Science*. ICCCI 2022 - 14th International Conference on Computational Collective Intelligence. Vol. 1653. Advances in Computational Collective Intelligence. Hammamet, Tunisia: Springer International Publishing, 21st Sept. 2022, pp. 459–471. DOI: [10.1007/978-3-031-16210-7\\_38](https://doi.org/10.1007/978-3-031-16210-7_38). URL: <https://hal.science/hal-03938813>.

#### Conferences without proceedings

- [29] 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. URL: <https://hal.inria.fr/hal-03946638>.

#### Doctoral dissertations and habilitation theses

- [30] O. Aouedi. ‘Machine Learning-Enabled Network Traffic Analysis’. Nantes Université, 2nd Dec. 2022. URL: <https://hal.science/tel-03931930>.
- [31] S. Sayadi. ‘Distributed architectures and secure software containers for multi-site medical cooperation’. Ecole nationale supérieure Mines-Télécom Atlantique, 5th Dec. 2022. URL: <https://theses.hal.science/tel-03922210>.

#### Reports & preprints

- [32] M. Delavergne, G. J. Antony and A. Lebre. *Cheops, a service to blow away Cloud applications to the Edge*. RR-9486. Inria Rennes - Bretagne Atlantique, 6th Sept. 2022, pp. 1–16. URL: <https://hal.inria.fr/hal-03770492>.

#### Other scientific publications

- [33] O. Aouedi, K. Piamrat, G. Muller and K. Singh. ‘Détection d’intrusions fédérée et semi-supervisée pour l’IoT’. In: Plate-Forme Intelligence Artificielle PFIA 2022. Saint-Etienne, France, 27th June 2022. URL: <https://hal-emse.ccsd.cnrs.fr/emse-03922716>.

### 12.3 Cited publications

- [34] O. Aouedi, K. Piamrat and D. Bagadthey. ‘Handling partially labeled network data: a semi-supervised approach using stacked sparse autoencoder’. In: *Computer Networks* (2022).
- [35] O. Aouedi, K. Piamrat, S. Hamma and J. Perera. ‘Network Traffic Analysis using Machine Learning: an unsupervised approach to understand and slice your network’. In: *Annals of Telecommunications* (2021), pp. 1–13.
- [36] 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).
- [37] O. Aouedi, K. Piamrat and B. Parrein. ‘Decision tree-based blending method using deep-learning for network management’. In: *IEEE/IFIP Network Operations and Management Symposium*. 2021.
- [38] N. Ayeub, 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>.



- [39] 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>.
- [40] 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](https://doi.org/10.1007/978-3-030-62327-2_5).
- [41] 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.
- [42] M. Belair, S. Laniece 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>.
- [43] 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>.
- [44] 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>.
- [45] 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.
- [46] 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.
- [47] 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://hal.inria.fr/hal-03103714). URL: <https://hal.inria.fr/hal-03103714>.
- [48] 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.
- [49] 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://hal.inria.fr/hal-03212421). URL: <https://hal.inria.fr/hal-03212421>.
- [50] 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.
- [51] 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>.
- [52] T. Coupaye, S. Bolle, S. Derrien, P. Folz, P. Meyé, 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.

- [53] T. Coupaye and J.-B. 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](https://doi.org/10.1007/978-3-540-71774-4_13).
- [54] 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>.
- [55] 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.
- [56] 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.-W. Kim, C.-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>.
- [57] 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>.
- [58] 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>.
- [59] 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://hal.inria.fr/hal-03137952). URL: <https://hal.inria.fr/hal-03137952>.
- [60] 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://hal.inria.fr/hal-03513456). URL: <https://hal.inria.fr/hal-03513456>.
- [61] 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.
- [62] 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.
- [63] 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.
- [64] 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).
- [65] 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.
- [66] 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>.
- [67] 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](http://dx.doi.org/10.1007/978-3-642-13238-4_5).

- [68] 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).
- [69] 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/ICCAC.2014.29](https://doi.org/10.1109/ICCAC.2014.29). URL: <https://doi.org/10.1109/ICCAC.2014.29>.
- [70] 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](https://doi.org/10.1007/978-3-030-85910-7_25).
- [71] 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.
- [72] 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.
- [73] 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>.
- [74] 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.
- [75] 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.
- [76] 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).
- [77] 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>.
- [78] 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.
- [79] 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>.
- [80] 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>.
- [81] 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>.

- 
- [82] 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.
- [83] 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).