



## Activity Report 2018

# Team STACK

## Software Stack for Massively Geo-Distributed Infrastructures

Inria teams are typically groups of researchers working on the definition of a common project, and objectives, with the goal to arrive at the creation of a project-team. Such project-teams may include other partners (universities or research institutions).

RESEARCH CENTER  
**Rennes - Bretagne-Atlantique**

THEME  
**Distributed programming and Software engineering**



## Table of contents

<b>1. Team, Visitors, External Collaborators</b> .....	<b>1</b>
<b>2. Overall Objectives</b> .....	<b>2</b>
2.1. STACK in a Nutshell .....	2
2.2. Toward a STACK for Geo-Distributed Infrastructures .....	3
<b>3. Research Program</b> .....	<b>3</b>
3.1. Overview .....	3
3.2. Resource Management .....	4
3.2.1. Performance Characterization of Low-Level Building Blocks .....	4
3.2.2. Geo-Distributed System Mechanisms .....	4
3.2.3. Capacity Planning and Placement Strategies .....	5
3.3. Programming Support .....	5
3.3.1. Programming Models and Languages Extensions .....	6
3.3.2. Deployment and Reconfiguration Challenges .....	6
3.4. Energy .....	8
3.5. Security .....	9
<b>4. Application Domains</b> .....	<b>10</b>
4.1. Overview .....	10
4.2. Industrial Internet .....	11
4.3. Internet of Skills .....	11
4.4. e-Health .....	11
4.5. Network Virtualization and Mobile Edge Services .....	12
<b>5. Highlights of the Year</b> .....	<b>12</b>
5.1.1. Remarkable results: research and third-party funding .....	12
5.1.2. Awards .....	12
<b>6. New Software and Platforms</b> .....	<b>12</b>
6.1. MAD .....	12
6.2. Nitro .....	13
6.3. VMPlaces .....	13
6.4. ENOS .....	14
6.5. Platforms .....	14
6.5.1. OpenStack .....	14
6.5.2. Grid'5000 .....	14
6.5.3. SeDuCe .....	15
6.5.4. SILECS .....	15
<b>7. New Results</b> .....	<b>15</b>
7.1. Resource Management .....	15
7.1.1. Geo-distributed Infrastructures .....	15
7.1.2. Cloud and HPC convergence .....	17
7.2. Programming Support .....	18
7.2.1. Deployment and reconfiguration in the Cloud .....	19
7.2.2. Autonomic computing and self-management .....	19
7.3. Energy-aware computing .....	20
7.4. Security and Privacy .....	21
7.4.1. Side-channel attacks and trusted Fog/Edge infrastructures .....	21
7.4.2. Secure and privacy-aware biomedical analyses .....	22
7.5. Experiment-Driven Research .....	22
<b>8. Bilateral Contracts and Grants with Industry</b> .....	<b>23</b>
<b>9. Partnerships and Cooperations</b> .....	<b>23</b>
9.1. Regional Initiatives .....	23

9.1.1.	CoMe4ACloud	23
9.1.2.	Oncoshare	23
9.1.3.	SyMeTRIC	23
9.1.4.	SysMics	24
9.2.	National Initiatives	24
9.2.1.	CominLabs laboratory of excellence	24
9.2.1.1.	PrivGen	24
9.2.1.2.	SeDuCe++	24
9.2.2.	ANR	24
9.2.2.1.	GRECO (ANR)	24
9.2.2.2.	KerStream (ANR)	25
9.2.3.	FSN	25
9.2.4.	CPER	25
9.2.5.	Inria Project Labs	25
9.2.6.	InriaHub	26
9.2.6.1.	MERCURY	26
9.2.6.2.	APOLLO/SOYUZ	26
9.2.7.	Fond d'amorçage IMT Industrie du Futur 2017	26
9.2.8.	Connect Talent	27
9.2.9.	Etoiles Montantes	27
9.3.	European Initiatives	27
9.4.	International Initiatives	28
9.5.	International Research Visitors	28
<b>10.</b>	<b>Dissemination</b> .....	<b>28</b>
10.1.	Promoting Scientific Activities	28
10.1.1.	Scientific Events Organisation	28
10.1.2.	Scientific Events Selection	29
10.1.2.1.	Chair of Conference Program Committees	29
10.1.2.2.	Member of the Conference Program Committees	29
10.1.3.	Journal	29
10.1.3.1.	Member of the Editorial Boards	29
10.1.3.2.	Reviewer - Reviewing Activities	29
10.1.4.	Invited Talks	29
10.1.5.	Scientific Expertise	30
10.1.6.	Research Administration	30
10.2.	Teaching - Supervision - Juries	30
10.2.1.	Supervision	30
10.2.2.	Juries	31
10.3.	Popularization	31
<b>11.</b>	<b>Bibliography</b> .....	<b>32</b>

## Team STACK

*Creation of the Team: 2017 November 01, updated into Project-Team: 2019 January 01*

*Project-team STACK was created in January 2019, in partnership with IMT Atlantique Bretagne-Pays de la Loire and in collaboration with the Laboratoire des Sciences du numérique de Nantes.*

### Keywords:

#### Computer Science and Digital Science:

- 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, Visitors, External Collaborators

Research Scientist

Shadi Ibrahim [Inria, Researcher]

#### **Faculty Members**

Hélène Coullon [IMT Atlantique, Inria Chair]  
Thomas Ledoux [IMT Atlantique, Associate Professor, HDR]  
Jean-Marc Menaud [IMT Atlantique, Professor, HDR]  
Jacques Noyé [IMT Atlantique, Associate Professor]  
Mario Südholt [IMT Atlantique, Professor, HDR]  
Adrien Lebre [Team leader, IMT Atlantique, Professor, HDR]

#### **Post-Doctoral Fellows**

Thomas Lambert [Inria, from Oct 2018]  
Jonathan Pastor [IMT Atlantique]  
Dimitri Pertin [Inria, until May 2018]  
Alexandre Van Kempen [Inria]

#### **PhD Students**

Mohamed Abderrahim [Orange Labs]  
Mohammad-Mahdi Bazm [Orange Labs]  
Maxime Belair [Orange Labs, from Oct 2018]  
Yewan Wang [Orange Labs]  
Fatima Zahra Boujdad [Univ de Nantes]  
Emile Cadorel [Armines]  
Maverick Chardet [Inria]  
Bastien Confais [CNRS, until Oct 2018]  
David Espinel [Orange Labs, from Feb 2018]  
Thuy Linh Nguyen [Inria]  
Dimitri Saingre [Armines, from Nov 2018]

#### **Technical staff**

Ronan-Alexandre Cherrueau [Inria]  
Marie Delavergne [Inria, from Oct 2018]  
Matthieu Simonin [Inria]

#### **Intern**

Hai Long Truong [Inria, from Apr 2018 until Jul 2018]

#### **Administrative Assistant**

Anne-Claire Binétruy [Inria]

#### **Visiting Scientist**

Thanh Quan Nguyen [University of Science, VNUHCM, Vietnam, from Oct 2018]

## **2. Overall Objectives**

### **2.1. STACK in a Nutshell**

The STACK team addresses challenges related to the management and advanced usages of Utility Computing infrastructures (i.e., Cloud, Fog, Edge, and beyond). More specifically, the team is interested in delivering appropriate system abstractions to operate and use massively geo-distributed ICT 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. Toward a STACK for Geo-Distributed Infrastructures

With the advent of Cloud Computing, modern applications have been developed on top of advanced software stacks composed of low-level system mechanisms, advanced middleware and software abstractions. While each of these layers has been designed to enable developers to efficiently use ICT resources without dealing with the burden of the underlying infrastructure aspects, the complexity of the resulting software stack has become a key challenge. As an example, Map/Reduce frameworks such as Hadoop have been developed to benefit from cpu/storage capacities of distinct servers. Running such frameworks on top of a virtualized cluster (*i.e.*, in a Cloud) can lead to critical situations if the resource management system decides to consolidate all the VMs on the same physical machine [105]. In other words, self-management decisions taken in isolation at one level (infrastructure, middleware, or application) may indirectly interfere with the decision taken by another layer, and globally affect the performance of the whole stack. Considering that geo-distributed ICT infrastructures significantly differ from the Cloud Computing ones regarding heterogeneity, resiliency, and the potential massive distribution of resources and networking environments [67], [100], we can expect that the complexity of the software stacks is going to increase. Such an assumption can be illustrated, for instance, by the software architecture proposed in 2016 by the ETSI Mobile edge computing Industry Specification Group [92]. This architecture is structured around a new layer in charge of orchestrating distinct independent cloud systems, *a.k.a.* Virtual Infrastructure Managers (VIMs) in their terminology. By reusing VIMs, ETSI targets an edge computing resource management that behaves in the same fashion as Cloud Computing ones. While mitigating development requirements, such a proposal hides all management decisions that might be taken in the VIM of one particular site and thus may lead to conflicting decisions and consequently to non-desired states overall.

Through the STACK team, we propose to investigate the software stack challenge as a whole. We claim it is the only way to limit as much as possible the complexity of the next generation software stack of geo-distributed ICT infrastructures. To reach our goal, we will identify major building blocks that should compose such a software stack, how they should be designed (*i.e.*, from the internal algorithms to the APIs they should expose), and finally how they should interact with each other.

Delivering such a software stack is an ambitious objective that goes beyond the activities of one research group. However, our expertise, our involvements in different consortiums (such as OpenStack) as well as our participation to different collaborative projects enable STACK members to contribute to this challenge in terms of architecture models, distributed system mechanisms and software artefacts, and finally, guideline reports on opportunities and constraints of geo-distributed ICT infrastructures.

## 3. Research Program

### 3.1. Overview

STACK research activities have been organized around four research topics. The two first 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 System/Middleware/Application layers, which generally composed a software stack, 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 two first axes, we identified them as independent research directions due to their critical aspects with respect to the societal challenges they represent. In the following, we detail the actions we plan to do in each research direction.

## 3.2. Resource Management

The challenge in this axis is to identify, design or revise mechanisms that are mandatory to operate and use a set of massively geo-distributed resources in an efficient manner [15]. This covers considering challenges at the scale of nodes, within one site (*i.e.*, one geographical location) and throughout the whole geo-distributed ICT infrastructure. It is noteworthy that the network community has been investigating similar challenges for the last few years [73]. To benefit from their expertise, in particular on how to deal with intermittent networks, STACK members have recently initiated exchanges and collaborative actions with some network research groups and telcos (see Sections 8.1 and 9.1). We emphasize, however, that we do not deliver contributions related to network equipments/protocols. The scientific and technical achievements we aim to deliver are related to the (distributed) system aspects.

### 3.2.1. Performance Characterization of Low-Level Building Blocks

Although Cloud Computing has enabled the consolidation of services and applications into a subset of servers, current operating system mechanisms do not provide appropriate abstractions to prevent (or at least control) the performance degradation that occurs when several workloads compete for the same resources [105]. Keeping in mind that server density is going to increase with physical machines composed of more and more cores and that applications will be more and more data intensive, it is mandatory to identify interferences that appear at a low level on each dimension (compute, memory, network, and storage) and propose countermeasures. In particular, previous studies [105], [65] on pros and cons of current technologies – virtual machines (VMs) [79], [87], containers and microservices – which are used to consolidate applications on the same server, should be extended: In addition to evaluating the performance we can expect from each of these technologies on a single node, it is important to investigate interferences that may result from cross-layer and remote communications [106]. We will consider in particular all interactions related to geo-distributed systems mechanisms/services that are mandatory to operate and use geo-distributed ICT infrastructures.

### 3.2.2. Geo-Distributed System Mechanisms

Although several studies have been highlighting the advantages of geo-distributed ICT infrastructures in various domains (see Section 3.), progress on how to operate and use such infrastructures is marginal. Current solutions [39] [40] are rather close to the initial Cisco Fog Computing proposal that only allows running domain-specific applications on edge resources and centralized Cloud platforms [47] (in other words, these solutions do not allow running stateful workloads in isolated environments such as containers or VMs). More recently, solutions leveraging the idea of federating VIMs (as the aforementioned ETSI MEC proposal [92]) have been proposed. ONAP [99], an industry-driven solution, enables the orchestration and automation of virtual network functions across distinct VIMs. From the academic side, FogBow [48] aims to support federations of Infrastructure-as-a-Service (IaaS) providers. Finally, NIST initiated a collaborative effort with IEEE to advance Federated Cloud platforms through the development of a conceptual architecture and a vocabulary<sup>1</sup>. Although all these projects provide valuable contributions, they face the aforementioned orchestration limitations (*i.e.*, they do not manage decisions taken in each VIM). Moreover, they all have been designed by only considering the developer/user's perspective. They provide abstractions to manage the life cycle of geo-distributed applications, but do not address administrative requirements.

To cope with specifics of Wide-Area networks while delivering most features that made Cloud Computing solutions successful also at the edge, our community should first identify limitations/drawbacks of current resource management system mechanisms with respect to the Fog/Edge requirements and propose revisions when needed [72], [85].

To achieve this aim, STACK members propose to conduct first a series of studies aiming at understanding the software architecture and footprint of major services that are mandatory for operating and using Fog/Edge infrastructures (storage backends, monitoring services, deployment/reconfiguration mechanisms, etc.). Leveraging these studies, we will investigate how these services should be deployed in order to deal with resources constraints, performance variability, and network split brains. We will rely on contributions that have been

<sup>1</sup><https://collaborate.nist.gov/twiki-cloud-computing/bin/view/CloudComputing/FederatedCloudPWGFC> (Dec 2018).



accomplished in distributed algorithms and self-\* approach for the last decade. In the short and medium term, we plan to evaluate the relevance of NewSQL systems [57] to store internal states of distributed system mechanisms in an edge context, and extend our proposals on new storage backends such as key/value stores [56], [98], and burst buffers [107]. We also plan to conduct new investigations on data-stream frameworks for Fog and Edge infrastructures [52]. These initial contributions should enable us to identify general rules to deliver other advanced system mechanisms that will be mandatory at the higher levels in particular for the deployment and reconfiguration manager in charge of orchestrating all resources.

### 3.2.3. Capacity Planning and Placement Strategies

An objective shared by users and providers of ICT infrastructures is to limit as much as possible the operational costs while providing the expected and requested quality of service (QoS). To optimize this cost while meeting QoS requirements, data and applications have to be placed in the best possible way onto physical resources according to data sources, data types (stream, graphs), application constraints (real-time requirements) and objective functions. Furthermore, the placement of applications must evolve through time to cope with the fluctuations in terms of application resource needs as well as the physical events that occur at the infrastructure level (resource creation/removals, hardware failures, etc.). This placement problem, *a.k.a.* the deployment and reconfiguration challenge as it will be described in Section 3.3, can be modeled in many different ways, most of the time by multi-dimensional and multi-objective bin-packing problems or by scheduling problems which are known to be difficult to solve. Many studies have been performed, for example, to optimize the placement of virtual machines onto ICT infrastructures [81]. STACK will inherit the knowledge acquired through previous activities in this domain, particularly its use of constraint programming strategies in autonomic managers [77], [76], relying on MAPE (monitor, analyze, plan, and execute) control loops. While constraint programming approaches are known to hardly scale, they enable the composition of various constraints without requiring to change heuristic algorithms each time a new constraint has to be considered [75]. We believe it is a strong advantage to deal with the diversity brought by geo-distributed ICT infrastructures. Moreover, we have shown in previous work that decentralized approaches can tackle the scalability issue while delivering placement decisions good enough and sometimes close to the optimal [91].

Leveraging this expertise, we propose, first, to identify new constraints raised by massively geo-distributed infrastructures (*e.g.*, data locality, energy, security, reliability and the heterogeneity and mobility of the underlying infrastructure). Based on this preliminary study, we will explore new placement strategies not only for computation sandboxes but for data (location, replication, streams, etc.) in order to benefit from the geo-distribution of resources and meet the required QoS. These investigations should lead to collaborations with operational research and optimization groups such as TASC, another research group from IMT Atlantique.

Second, we will leverage contributions made on the previous axis “Performance Characterization of Low-Level Building Blocks” to determine how the deployment of the different units (software components and data sets) should be executed in order to reduce as much as possible the time to reconfigure the system (*i.e.*, the *Execution* phase in the control loop). In some recent work [87], we have shown that the provisioning of a new virtual machine should be done carefully to mitigate boot penalties. More generally, proposing an efficient action plan for the *Execution* phase will be a major point as Wide-Area-Network specifics may lead to significant delays, in particular when the amount of data to be manipulated is important.

Finally, we will investigate new approaches to decentralize the placement process while considering the geo-distributed context. Among the different challenges to address, we will study how a combination of autonomic managers, at both the infrastructure and application levels [63], could be proposed in a decentralized manner. Our first idea is to geo-distribute a fleet of small control loops over the whole infrastructure. By improving the locality of data collection and weakening combinatorics, these loops would allow the system to address responsiveness and quality expectations.

## 3.3. Programming Support

We pursue two main research directions relative to new programming support: first, developing new programming models with appropriate support in existing languages (libraries, embedded DSLs, etc.) and, second,

providing new means for deployment and reconfiguration in geo-distributed ICT environments, principally supporting the mapping of software onto the infrastructure. For both directions two levels of challenges are considered. On the one hand, the *generic* level refers to efforts on programming support that can be applied to any kind of distributed software, application or system. On this level, contributions could thus be applied to any of the three layers addressed by STACK (*i.e.*, system, middleware or application). On the other hand, the corresponding generic programming means may not be appropriate in practice (*e.g.*, requirements for more dedicated support, performance constraints, etc.), even if they may lead to interesting general properties. For this reason, a *specific* level is also considered. This level could be based on the generic one but addresses specific cases or domains.

### 3.3.1. Programming Models and Languages Extensions

The current landscape of programming support for cloud applications is fragmented. This fragmentation is based on apparently different needs for various kinds of applications, in particular, web-based, computation-based, focusing on the organization of the computation, and data-based applications, within the last case a quite strong dichotomy between applications considering data as sets or relations, close to traditional database applications and applications considering data as real-time streams. This has led to various programming models, in a loose sense, including for instance microservices, graph processing, dataflows, streams, etc. These programming models have mostly been offered to the application programmer in the guise of frameworks, each offering subtle variants of the programming models with various implementation decisions favoring particular application and infrastructure settings. Whereas most frameworks are dedicated to a given programming model, *e.g.*, basic Pregel [86], Hive [101], Hadoop [102], some of them are more general-purpose through the provision of several programming models, *e.g.*, Flink [51] and Spark [83]. Finally, some dedicated language support has been considered for some models (*e.g.*, the language SPL underlying IBM Streams [78]) as well as core languages and calculi (*e.g.*, [49], [96]).

This situation raises a number of challenges on its own, related to a better structuring of the landscape. It is necessary to better understand the various programming models and their possible relations, with the aim of facilitating, if not their complete integration, at least their composition, at the conceptual level but also with respect to their implementations, as specific languages and frameworks.

Switching to massively geo-distributed infrastructures adds to these challenges by leading to a new range of applications (*e.g.*, smart-\* applications) that, by nature, require mixing these various programming models, together with a much more dynamic management of their runtime.

In this context, STACK would like to explore two directions:

- First, we propose to contribute to generic programming models and languages to address composability of different programming models [58]. For example, providing a generic stream data processing model that can operate under both data stream [51] and operation stream [108] modes, thus streams can be processed in micro batches to favour high throughput or record by record to sustain low latency. Software engineering properties such as separation of concerns and composition should help address such challenges [42], [97]. They should also facilitate the software deployment and reconfiguration challenges discussed below.
- Second, we plan to revise relevant programming models, the associated specific languages, and their implementation according to the massive geo-distribution of the underlying infrastructure, the data sources, and application end-users. For example, although SPL is extensible and distributed, it has been designed to run on multi-cores and clusters [78]. It does not provide the level of dynamicity required by geo-distributed applications (*e.g.*, to handle topology changes, loss of connectivity at the edge, etc.). Moreover, as more network data transfers will happen within a massively geo-distributed infrastructure, correctness of data transfers should be guaranteed. This has potential impact from the programming models to their implementations.

### 3.3.2. Deployment and Reconfiguration Challenges

The second research direction deals with the complexity of deploying distributed software (whatever the layer, application, middleware or system) onto an underlying infrastructure. As both the deployed pieces

of software and the infrastructures addressed by STACK are large, massively distributed, heterogeneous and highly dynamic, the deployment process cannot be handled manually by developers or administrators. Furthermore, and as already mentioned in Section 3.2, the initial deployment of some distributed software will evolve through time because of the dynamicity of both the deployed software and the underlying infrastructures. When considering reconfiguration, which encompasses deployment as a specific case, the problem becomes more difficult for two main reasons: (1) the current state of both the deployed software and the infrastructure has to be taken into account when deciding on a reconfiguration plan, (2) as the software is already running the reconfiguration should minimize disruption time, while avoiding inconsistencies [84], [89]. Many deployment tools have been proposed both in academia and industry [60]. For example, Ansible<sup>2</sup>, Chef<sup>3</sup> and Puppet<sup>4</sup> are very well-known generic tools to automate the deployment process through a set of batch instructions organized in groups (*e.g.*, *playbooks* in Ansible). Some tools are specific to a given environment, like Kolla to deploy OpenStack, or the embedded deployment manager within Spark. Few reconfiguration capabilities are available in production tools such as *scaling* and *restart* after a fault<sup>5 6</sup>. Academia has contributed to generic deployment and reconfiguration models. Most of these contributions are component-based. Component models divide a distributed software as a set of component instances (or modules) and their assembly, where components are connected through well defined interfaces [97]. Thus, modeling the reconfiguration process consists in describing the life cycle of different components and their interactions. Most component-based approaches offer a fixed life cycle, *i.e.*, identical for any component [66]. Two main contributions are able to customize life cycles, Fractal [50], [45] and its evolutions [42], [43], [62], and Aeolus [64]. In Fractal, the *control* part of a component (*e.g.*, its life cycle) is modeled itself as a component assembly that is highly flexible. Aeolus, on the other hand, offers a finer control on both the evolution and the synchronization of the deployment process by modeling each component life cycle with a finite state machine.

A reconfiguration raises at least five questions, all of them are correlated: (1) *why software has to be reconfigured?* (monitoring, modeling and analysis) (2) *what should be reconfigured?* (software modeling and analysis), (3) *how should it be reconfigured?* (software modeling and planning decisions), (4) *where should it be reconfigured?* (infrastructure modeling and planning decisions), and (5) *when to reconfigure it?* (scheduling algorithms). STACK will contribute to all aspects of a reconfiguration process as described above. However, according to the expertise of STACK members, we will focus mainly on the three first questions: *why*, *what* and *how*, leaving questions *where* and *when* to collaborations with operational research and optimization teams.

First of all, we would like to investigate *why software has to be reconfigured?* Many reasons could be mentioned, such as hardware or software fault tolerance, mobile users, dynamicity of software services, etc. All those reasons are related somehow to the Quality of Service (QoS) or the Service Level Agreement (SLA) between the user and the Cloud provider. We first would like to explore the specificities of QoS and SLAs in the case of massively geo-distributed ICT environments [93]. By being able to formalize this question, analyzing the requirement of a reconfiguration will be facilitated.

Second, we think that four important properties should be enhanced when deploying and reconfiguring models in massively geo-distributed ICT environments. First, as low-latency applications and systems will be subject to deployment and reconfiguration, the performance and the ability to scale are important. Second, as many different kinds of deployments and reconfigurations will concurrently hold within the infrastructure, processes have to be reliable, which is facilitated by a fine-grained control of the process. Finally, as many different software elements will be subject to deployment and reconfiguration, common generic models and engines for deployment and reconfiguration should be designed [12]. For these reasons, we intend to go beyond Aeolus by: first, leveraging the expression of parallelism within the deployment process, which should lead to better performance; second, improving the separation of concerns between the component developer

---

<sup>2</sup><https://www.ansible.com/>

<sup>3</sup><https://www.chef.io/chef/>

<sup>4</sup><https://puppet.com/>

<sup>5</sup><https://kubernetes.io/>

<sup>6</sup><https://jujucharms.com/>

and the reconfiguration developer; third, enhancing the possibility to perform concurrent and decentralized reconfigurations.

Research challenges relative to programming support have been presented above. Many of these challenges are related, in different manners, to the resource management level of STACK or to crosscutting challenges, *i.e.*, energy and security. First, one can notice that any programming model or deployment and reconfiguration implementation should be based on mechanisms related to resource management challenges. For this reason, all challenges addressed within this section are linked with lower level building blocks presented in Section 3.2. Second, as detailed above, deployment and reconfiguration address at least five questions. The question *what?* is naturally related to programming support. However, questions *why*, *how?*, *where?* and *when?* are also related to Section 3.2, for example, to monitoring and capacity planning. Moreover, regarding the deployment and reconfiguration challenges, one can note that the same goals recursively happen when deploying the control building blocks themselves (bootstrap issue). This comforts the need to design generic deployment and reconfiguration models and frameworks. These low-level models should then be used as back-ends to higher-level solutions. Finally, as *energy* and *security* are crosscutting themes within the STACK project, many additional energy and security considerations could be added to the above challenges. For example, our deployment and reconfiguration frameworks and solutions could be used to guarantee the deployment of end-to-end security policies or to answer specific energy constraints [74] as detailed in the next section.

### 3.4. Energy

The overall electrical consumption of DCs grows according to the demand of Utility Computing. Considering that the latter has been continuously increasing since 2008, the energy footprint of Cloud services overall is nowadays critical with about 91 billion kilowatt-hours of electricity [95]. Besides the ecological impact, the energy consumption is a predominant criterion for providers since it determines a large part of the operational cost of their infrastructure. Among the different approaches that have been investigated to reduce the energy footprint, some studies have been investigating the use of renewable energy sources to power microDCs [68]. Workload distribution for geo-distributed DCs is also another promising approach [70], [82], [103]. Our research will extend these results with the ultimate goal of considering the different opportunities to control the energy footprint across the whole stack (hardware and software opportunities, renewable energy, thermal management ...). In particular, we identified several challenges that we will address in this context within the STACK framework:

- First, we propose to evaluate the energy efficiency of low-level building blocks, from the viewpoints of computation (VMs, containers, microkernel, microservices) [61] and data (hard drives, SSD, in-memory storage, distributed file systems). For computations, in the continuity of our previous work [59], [77], we will investigate workload placement policies according to energy (minimizing energy consumption, power capping, thermal load balancing, etc.). Regarding the data dimension, we will investigate, in particular, the trade-offs between energy consumption and data availability, durability and consistency [55], [98]. Our ambition is to propose an adaptive energy-aware data layout and replication scheme to ensure data availability with minimal energy consumption. It is noteworthy that these new activities will also consider our previous work on DCs partially powered by renewable energy (see the SeDuCe project, in Section 6.5), with the ultimate goal of reducing the CO<sub>2</sub> footprint.
- Second, we will complete current studies to understand pros and cons of massively geo-distributed infrastructures from the energy perspective. Addressing the energy challenge is a complex task that involves considering several dimensions such as the energy consumption due to the physical resources (CPU, memory, disk, network), the performance of the applications (from the computation and data viewpoints), and the thermal dissipation caused by air conditioning in each DC. Each of these aspects can be influenced by each level of the software stack (*i.e.*, low-level building blocks, coordination and autonomous loops, and finally application life cycle). In previous projects, we have studied and modeled the consumption of the main components, notably the network, as part of a single microDC. We plan to extend these models to deal with geo-distribution. The objective is

to propose models that will enable us to refine our placement algorithms as discussed in the next paragraph. These models should be able to consider the energy consumption induced by all WAN data exchanges, including site-to-site data movements as well as the end users' communications for accessing virtualized resources.

- Third, we expect to implement green-energy-aware balancing strategies, leveraging the aforementioned contributions. Although the infrastructures we envision increase complexity (because WAN aspects should also be taken into account), the geo-distribution of resources brings several opportunities from the energy viewpoint. For instance, it is possible to define several workload/data placement policies according to renewable energy availability. Moreover, a tightly-coupled software stack allows users to benefit from such a widely distributed infrastructure in a transparent way while enabling administrators to balance resources in order to benefit from green energy sources when available. An important difficulty, compared to centralized infrastructures, is related to data sharing between software instances. In particular, we will study issues raised by the distribution and replication of services across several microDCs. In this new context, many challenges must be addressed: where to place the data (Cloud, Edge) in order to mitigate data movements? What is the impact in terms of energy consumption, network and response time of these two approaches? How to manage the consistency of replicated data/services? All these aspects must be studied and integrated into our placement algorithms.
- Fourth, we will investigate the energy footprint of the current techniques that address failure and performance variability in large-scale systems. For instance, *stragglers* (*i.e.*, tasks that take a significantly longer time to finish than the normal execution time) are natural results of performance variability, they cause extra resource and energy consumption. Our goal is to understand the energy overhead of these techniques and introduce new handling techniques that take into consideration the energy efficiency of the platform [90].
- Finally, in order to answer specific energy constraints, we want to reify energy aspects at the application level and propose a metric related to the use of energy (Green SLA [41]), for example to describe the maximum allowed CO<sub>2</sub> emissions of a Fog/Edge service. Unlike other approaches [71], [46], [69] that attempt to identify the best trade-off, we want to offer to developers/end-users the opportunity to select the best choice between application performance, correctness and energy footprint. Such a capability will require reifying the energy dimension at the level of big-data and interactive applications. Besides, with the emergence of renewable energy (*e.g.*, solar panels for microDC), investigating the energy consumption *vs* performance trade-off [74] and the smart usage of green energy for ICT geo-distributed services seems promising. For example, we want to offer the opportunity to developers/end-users to control the scaling of the applications based on this trade-off instead of current approaches that only considered application load. Providing such a capability will also require appropriate software abstractions.

### 3.5. Security

Because of its large size and complex software structure, geo-distributed applications and infrastructures are particularly exposed to security and privacy issues [94]. They are subject to numerous security vulnerabilities that are frequently exploited by malicious attackers in order to exfiltrate personal, institutional or corporate data. Securing these systems require security and privacy models and corresponding techniques that are applicable at all software layers in order to guard interactions at each level but also between levels. However, very few security models exist for the lower layers of the software stack and no model enables the handling of interactions involving the complete software stack. Any modification to its implementation, deployment status, configuration, etc., may introduce new or trigger existing security and privacy issues. Finally, applications that execute on top of the software stack may introduce security issues or be affected by vulnerabilities of the stack. Overall, security and privacy issues are therefore interdependent with all other activities of the STACK team and constitute an important research topic for the team.



As part of the STACK activities, we consider principally security and privacy issues related to the vertical and horizontal compositions of software components forming the software stack and the distributed applications running on top of it. Modifications to the *vertical composition* of the software stack affect different software levels at once. As an example, side-channel attacks often target virtualized services (*i.e.*, services running within VMs); attackers may exploit insecure hardware caches at the system level to exfiltrate data from computations at the higher level of VM services [88], [104]. Security and privacy issues also affect *horizontal compositions*, that is, compositions of software abstractions on one level: most frequently horizontal compositions are considered on the level of applications/services but they are also relevant on the system level or the middleware level, such as compositions involving encryption and database fragmentation services.

The STACK members aim at addressing two main research issues: enabling full-stack (vertical) security and per-layer (horizontal) security. Both of these challenges are particularly hard in the context of large geo-distributed systems because they are often executed on heterogeneous infrastructures and are part of different administrative domains and governed by heterogeneous security and privacy policies. For these reasons they typically lack centralized control, are frequently subject to high latency and are prone to failures.

Concretely, we will consider two classes of security and privacy issues in this context. First, on a general level, we strive for a method for the programming and reasoning about compositions of security and privacy mechanisms including, but not limited to, encryption, database fragmentation and watermarking techniques. Currently, no such general method exists, compositions have only been devised for specific and limited cases, for example, compositions that support the commutation of specific encryption and watermarking techniques [80], [53]. We provided preliminary results on such compositions [54] and have extended them to biomedical, notably genetic, analyses in the e-health domain [11]. Second, on the level of security and privacy properties, we will focus on isolation properties that can be guaranteed through vertical and horizontal composition techniques. We have proposed first results in this context in form of a compositional notion of distributed side channel attacks that operate on the system and middleware levels [44].

It is noteworthy that the STACK members do not have to be experts on the individual security and privacy mechanisms, such as watermarking and database fragmentation. We are, however, well-versed in their main properties so that we can integrate them into our composition model. We also interact closely with experts in these techniques and the corresponding application domains, notably e-health for instance, in the context of the PRIVGEN project<sup>7</sup>, see Section 9.1.

More generally, we highlight that security issues in distributed systems are very closely related to the other STACK challenges, dimensions and research directions. Guaranteeing security properties across the software stack and throughout software layers in highly volatile and heterogeneous geo-distributed systems is expected to harness and contribute results to the self-management capabilities investigated as part of the team's resource management challenges. Furthermore, security and privacy properties are crosscutting concerns that are intimately related to the challenges of application life cycle management. Similarly, the security issues are also closely related to the team's work on programming support. This includes new means for programming, notably in terms of event and stream programming, but also the deployment and reconfiguration challenges, notably concerning automated deployment. As a crosscutting functionality, the security challenges introduced above must be met in an integrated fashion when designing, constructing, executing and adapting distributed applications as well as managing distributed resources.

## 4. Application Domains

### 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

<sup>7</sup>Privacy-preserving sharing and processing of genetic data, <https://privgen.cominlabs.u-bretagne-normandie.fr/fr>

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 infrastructures 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 ...); the development of large “smart” applications handling large amounts of geo-distributed data and a large set of resources (video analytics, augmented reality ...); the advent of future medical prevention and treatment techniques thanks to the intensive use of ICT systems ... We expect our contributions will favor the rise of efficient, correct and sustainable massively geo-distributed infrastructures 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 Utility Computing infrastructures can cope with this requirement, it is important to determine how the quality of service of such applications should be defined and how latency and bandwidth constraints can be guaranteed at the infrastructure level.

## 4.4. e-Health

The e-Health domain constitutes an important societal application domain of the two previous areas. The STACK teams is investigating distribution, security and privacy issues in the fields of systems and personalized 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 are working, as part of the ongoing PrivGen CominLabs collaborative project on new means for the sharing of genetic data and applications in the Cloud. We are applying and developing such techniques in the regional networks SysMics and Oncoshare: there, we investigate how to secure and preserve privacy if potentially sensitive personal data is moved and processed by distributed biomedical analyses.

We are also involved in the SyMeTRIC regional initiative where preliminary studies have been conducted in order to build a common System Medicine computing infrastructure to accelerate the discovery and validation of bio-markers in the fields of oncology, transplantation, and chronic cardiovascular diseases. The challenges were related to the need of being able to perform analyses on data that cannot be moved between distinct locations.

The STACK team will continue to contribute to the e-Health domain by harnessing advanced architectures, applications and infrastructures 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 infrastructures, in particular through working groups such as Mobile Edge Computing at the European Telecommunication Standards Institute<sup>8</sup>. The initial reason is that geo-distributed infrastructures will 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. Highlights of the Year

### 5.1. Highlights of the Year

#### 5.1.1. Remarkable results: research and third-party funding

Regarding scientific results, the team has produced a number of outstanding results on resource and data managements in large-scale infrastructures, notably on how to place VMs in Clouds [5], and on how to manage VM images in geo-distributed clouds [18]. On the software side, the team has proposed a new model-based Architecture to design and implement autonomic and heterogeneous Cloud Systems [12]. Finally on the energy side, the team has deployed the SeDuce platform that allows researchers to investigate energy concerns in data-centers thanks to a numerous of energy sensors deployed across the dedicated facility [19], [20], [34].

Concerning third-party funding, 2018 has seen the acceptance of the VERDI “Etoiles Montantes” project. “Etoiles Montantes” is a highly-competitive call with the goal of bootstrapping ERC submissions.

#### 5.1.2. Awards

In 2018, the team has received two best paper awards and one individual award:

- **Programme Jeunes Talents France Chine 2018** Shadi Ibrahim was selected for the “Programme Jeunes Talents France Chine” award.

#### BEST PAPERS AWARDS:

[12]

H. BRUNELIERE, Z. AL-SHARA, F. ALVARES, J. LEJEUNE, T. LEDOUX. *A Model-based Architecture for Autonomic and Heterogeneous Cloud Systems*, in "CLOSER 2018 - 8th International Conference on Cloud Computing and Services Science", Funchal, Portugal, March 2018, vol. 1, pp. 201-212, Best Paper Award [DOI : 10.5220/0006773002010212], <https://hal.archives-ouvertes.fr/hal-01705248>

[19]

J. PASTOR, J.-M. MENAUD. *SeDuCe: a Testbed for Research on Thermal and Power Management in Datacenters*, in "GREEN 2018 - Third International Conference on Green Communications, Computing and Technologies", Venise, Italy, September 2018, pp. 1-6, <https://hal.inria.fr/hal-01855452>

## 6. New Software and Platforms

### 6.1. MAD

*Madeus Application Deployer*

<sup>8</sup><http://www.etsi.org/news-events/news/1078-2016-04-etsi-mobile-edge-computing-publishes-foundation-specifications>.



KEYWORDS: Automatic deployment - Distributed Software - Component models - Cloud computing

SCIENTIFIC DESCRIPTION: MAD is a Python implementation of the Madeus deployment model for multi-component distributed software. Precisely, it allows to: 1. describe the deployment process and the dependencies of distributed software components in accordance with the Madeus model, 2. describe an assembly of components, resulting in a functional distributed software, 3. automatically deploy the component assembly of distributed software following the operational semantics of Madeus.

RELEASE FUNCTIONAL DESCRIPTION: Initial submission with basic functionalities of MAD

NEWS OF THE YEAR: Operational prototype.

- Participants: Christian Pérez, Dimitri Pertin, H el ene Coullon and Maverick Chardet
- Partners: IMT Atlantique - LS2N - LIP
- Contact: H el ene Coullon
- Publications: [Madeus: A formal deployment model - Behavioral interfaces for reconfiguration of component models](#)

## 6.2. Nitro

KEYWORDS: Cloud storage - Virtual Machine Image - Geo-distribution

SCIENTIFIC DESCRIPTION: Nitro is a storage system that is designed to work in geo-distributed cloud environments (i.e., over WAN) to efficiently manage Virtual Machine Images (VMIs).

Nitro employs fixed-size deduplication to store VMIs. This technique contributes to minimizing the network cost. Also, Nitro incorporates a network-aware scheduling algorithm (based on max flow algorithm) to determine which chunks should be pulled from which site in order to reconstruct the corresponding image on the destination site, with minimal (provisioning) time.

FUNCTIONAL DESCRIPTION: Geo-distributed Storage System to optimize Images (VM, containers, ...) management, in terms of cost and time, in geographically distributed cloud environment (i.e. data centers are connected over WAN).

- Authors: Jad Darrous, Shadi Ibrahim and Christian P erez
- Contact: Shadi Ibrahim
- URL: <https://gitlab.inria.fr/jdarrous/nitro>

## 6.3. VMPlaces

KEYWORDS: Simulation - Virtualization - Scheduling

FUNCTIONAL DESCRIPTION: VMPlaces is a dedicated framework to evaluate and compare VM placement algorithms. This framework is composed of two major components: the injector and the VM placement algorithm. The injector is the generic part of the framework (i.e. the one you can directly use) while the VM placement algorithm is the part you want to study (or compare with available algorithms). Currently, the VMPlaceS is released with three algorithms:

Entropy, a centralized approach using a constraint programming approach to solve the placement/reconfiguration VM problem

Snooze, a hierarchical approach where each manager of a group invokes Entropy to solve the placement/reconfiguration VM problem. Note that in the original implementation of Snooze, it is using a specific heuristic to solve the placement/reconfiguration VM problem. As the sake of simplicity, we have simply reused the entropy scheduling code.

DVMS, a distributed approach that dynamically partitions the system and invokes Entropy on each partition.

- Participants: Adrien L ebre, Flavien Quesnel, Jonathan Pastor, Mario S udholt and Takahiro Hirofuchi
- Contact: Adrien L ebre
- URL: <http://beyondtheclouds.github.io/VMPlaceS/>

## 6.4. ENOS

*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.

- Partner: Orange Labs
- Contact: Adrien Lèbre
- URL: <http://enos.readthedocs.io/en/stable/>

## 6.5. Platforms

### 6.5.1. OpenStack

OpenStack is the de facto open-source management system to operate and use Cloud Computing infrastructures. 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 of this community. We created and animated between 2016 and 2018 the Fog/Edge/Massively Distributed (FEMDC) Special Interest Group <sup>9</sup> 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 [100], 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.

### 6.5.2. Grid'5000

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. Researchers can experiment with a fully customized software stack thanks to bare-metal deployment features, and can isolate their experiment at the networking layer advanced monitoring and measurement features for traces collection of networking and power consumption, providing a deep understanding of experiments designed to support Open Science and reproducible research, with full traceability of infrastructure and software changes on the testbed. 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.

<sup>9</sup>[https://wiki.openstack.org/wiki/Fog\\_Edge\\_Massively\\_Distributed\\_Clouds](https://wiki.openstack.org/wiki/Fog_Edge_Massively_Distributed_Clouds)

### 6.5.3. SeDuCe

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 should soon be connected to renewqble energy sources (solar panels). This “green” datacenter will enable 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).

### 6.5.4. SILECS

STACK Members are involved in the definition and bootstrap of the SILECS infrastructure (IR ministère). 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, ...).

## 7. New Results

### 7.1. Resource Management

**Participants:** Mohamed Abderrahim, Ronan-Alexandre Cherrueau, Bastien Confais, Jad Darrous, Shadi Ibrahim, Adrien Lebre, Matthieu Simonin, Emile Cadorel, H  l  ne Coullon, Jean-Marc Menaud.

Our contributions regarding resource management can be divided into two main topics described below: contributions related to (i) geo-distributed cloud infrastructures (*e.g.*, Fog and Edge computing) and (ii) the convergence of Cloud and HPC infrastructures.

#### 7.1.1. Geo-distributed Infrastructures

In [15], we provide reflections regarding how fog/edge infrastructures can be operated. While it is clear that edge infrastructures are required for emerging use-cases related to IoT, VR or NFV, there is currently no resource management system able to deliver all features for the edge that made cloud computing successful (*e.g.*, an OpenStack for the edge). Since building a system from scratch is seen by many as impractical, our community should investigate different appraoches. This study, which has been achieved with Ericsson colleagues, provides a list of the features required to operate and use edge computing resources, and investigate how an existing IaaS manager (*i.e.*, OpenStack) satisfies these requirements. Finally, we identify from this study two approaches to design an edge infrastructure manager that fulfils our requirements, and discuss their pros and cons.

In [18], we propose a new novel VMI management system for distributed cloud infrastructures. Most large cloud providers, like Amazon and Microsoft, replicate their Virtual Machine Images (VMIs) on multiple geographically distributed data centers to offer fast service provisioning. Provisioning a service may require to transfer a VMI over the wide-area network (WAN) and therefore is dictated by the distribution of VMIs and the network bandwidth in-between sites. Nevertheless, existing methods to facilitate VMI management (*ie*, retrieving VMIs) overlook network heterogeneity in geo-distributed clouds. To deal with such a limitation, we design, implement and evaluate Nitro, a novel VMI management system that helps to minimize the transfer time of VMIs over a heterogeneous WAN. To achieve this goal, Nitro incorporates two complementary features. First, it makes use of deduplication to reduce the amount of data which will be transferred due to the high similarities within an image and in-between images. Second, Nitro is equipped with a network-aware data transfer strategy to effectively exploit links with high bandwidth when acquiring data and thus expedites the provisioning time. Experimental results show that our network-aware data transfer strategy offers the optimal solution when acquiring VMIs while introducing minimal overhead. Moreover, Nitro outperforms state-of-the-art VMI storage systems (*eg*, OpenStack Swift) by up to 77%.

In [22] we perform a performance evaluation of two communication bus mechanisms available in the open-stack eco-system. Cloud computing depends on communication mechanisms implying location transparency. Transparency is tied to the cost of ensuring scalability and an acceptable request responses associated to the locality. Current implementations, as in the case of OpenStack, mostly follow a centralized paradigm but they lack the required service agility that can be obtained in decentralized approaches. In an edge scenario, the communicating entities of an application can be dispersed. In this context, we perform a study on the inter-process communication of OpenStack when its agents are geo-distributed. More precisely, we are interested in the different Remote Procedure Calls (OARPCs) implementations of OpenStack and their behaviours with regards to three classical communication patterns: anycast, unicast and multicast. We discuss how the communication middleware can align with the geo-distribution of the RPC agents regarding two key factors: scalability and locality. We reached up to ten thousands communicating agents, and results show that a router-based deployment offers a better trade-off between locality and load-balancing. Broker-based suffers from its centralized model which impact the achieved locality and scalability.

In [5], we give a complete overview of VMPlaceS, a dedicated framework we have been implementing since 2015 in order to evaluate and compare VM placement algorithms. Most current infrastructures for cloud computing leverage static and greedy policies for the placement of virtual machines. Such policies impede the optimal allocation of resources from the infrastructure provider viewpoint. Over the last decade, more dynamic and often more efficient policies based, *e.g.*, on consolidation and load balancing techniques, have been developed. Due to the underlying complexity of cloud infrastructures, these policies are evaluated either using limited scale testbeds/in-vivo experiments or ad-hoc simulators. These validation methodologies are unsatisfactory for two important reasons: they (i) do not model precisely enough real production platforms (size, workload variations, failure, etc.) and (ii) do not enable the fair comparison of different approaches. More generally, new placement algorithms are thus continuously being proposed without actually identifying their benefits with respect to the state of the art. In this article, we present and discuss most of the features provided by VMPlaceS, a dedicated simulation framework that enables researchers (i) to study and compare VM placement algorithms from the infrastructure perspective, (ii) to detect possible limitations at large scale and (iii) to easily investigate different design choices. Built on top of the SimGrid simulation platform, VMPlaceS provides programming support to ease the implementation of placement algorithms and runtime support dedicated to load injection and execution trace analysis. To illustrate the relevance of VMPlaceS, we first discuss a few experiments that enabled us to study in details three well known VM placement strategies. Diving into details, we also identify several modifications that can significantly increase their performance in terms of reactivity. Second, we complete this overall presentation of VMPlaceS by focusing on the energy efficiency of the well-know FFD strategy. We believe that VMPlaceS will allow researchers to validate the benefits of new placement algorithms, thus accelerating placement research and favouring the transfer of results to IaaS production platforms.

In [27], we present different heuristics that address the placement challenge in Fog/Edge infrastructures. As Fog Computing brings processing and storage resources to the edge of the network, there is an increasing need of automated placement (*i.e.*, host selection) to deploy distributed applications. Such a placement must conform to applications' resource requirements in a heterogeneous Fog infrastructure, and deal with the complexity brought by Internet of Things (IoT) applications tied to sensors and actuators. In this study, we present and evaluate four heuristics to address the problem of placing distributed IoT applications in the fog. By combining proposed heuristics, our approach is able to deal with large scale problems, and to efficiently make placement decisions fitting the objective: minimizing placed applications' average response time. The proposed approach has been validated through comparative simulation of different heuristic combinations with varying sizes of infrastructures and applications.

In [35], we introduce the premises of monitoring function chaining concepts with the ultimate goal of delivering an holistic monitoring system for Fog/Edge infrastructures. By relying on small sized and massively distributed infrastructures, the Edge computing paradigm aims at supporting the low latency and high bandwidth requirements of the next generation services that will leverage IoT devices (*e.g.*, video cameras, sensors). To favor the advent of this paradigm, management services, similar to the ones that made the success of Cloud computing platforms, should be proposed. However, they should be designed in order to cope with the

limited capabilities of the resources that are located at the edge. In that sense, they should mitigate as much as possible their footprint. Among the different management services that need to be revisited, we investigate in this study the monitoring one. Monitoring functions tend to become compute-, storage- and network-intensive, in particular because they will be used by a large part of applications that rely on real-time data. To reduce as much as possible the footprint of the whole monitoring service, we propose to mutualize identical processing functions among different tenants while ensuring their quality-of-service (QoS) expectations. We formalize our approach as a constraint satisfaction problem and show through micro-benchmarks its relevance to mitigate compute and network footprints.

In [17], we discuss the limitations of meta-data management in Fog/Edge infrastructures. A few storage systems have been proposed to store data in those infrastructures. Most of them are relying on a Distributed Hash Table (DHT) to store the location of objects which is not efficient because the node storing the location of the data may be placed far away from the object replicas. In this paper, we propose to replace the DHT by a tree-based approach mapping the physical topology. Servers look for the location of an object by requesting successively their ancestors in the tree. Location records are also relocated close to the object replicas not only to limit the network traffic when requesting an object, but also to avoid an overload of the root node. We also propose to modify the Dijkstra's algorithm to compute the tree used. Finally, we evaluate our approach using the object store InterPlanetary File System (IPFS) on Grid'5000 using both a micro experiment with a simple network topology and a macro experiment using the topology of the French National Research and Education Network (RENATER). We show that the time to locate an object in our approach is less than 15 ms on average which is around 20% better than using a DHT.

### 7.1.2. Cloud and HPC convergence

Geo-distribution of Cloud Infrastructures is not the only current trend of utility computing. Another important challenge is to favor the convergence of Cloud and HPC infrastructures, in other words on-demand HPC. Among challenges of this convergence is, for example, how to exploit HPC systems to execute data-intensive workflows effectively, as well as how to schedule tasks and jobs in Cloud, HPC, or hybrid HPC/Cloud infrastructures to meet data volatility and the ever-growing heterogeneity in the computation demands of workflows.

With the growing needs of users and size of data, commodity-based infrastructure will strain under the heavy weight of Big Data. On the other hand, HPC systems offer a rich set of opportunities for Big Data processing. As first steps toward Big Data processing on HPC systems, several research efforts have been devoted to understanding the performance of Big Data applications on these systems. Yet the HPC specific performance considerations have not been fully investigated. In [28], we conduct an experimental campaign to provide a clearer understanding of the performance of Spark, the de facto in-memory data processing framework, on HPC systems. We ran Spark using representative Big Data workloads on Grid'5000 testbed to evaluate how the latency, contention and file system's configuration can influence the application performance. We discuss the implications of our findings and draw attention to new ways (e.g., burst buffers) to improve the performance of Spark on HPC systems.

Motivated by our work [28], we extend Eley [107], a burst buffer solution that aims to accelerate the performance of Big Data applications, to be interference-aware. Specifically, while data prefetching reduce the response time of Big data applications as data inputs will be stored on a low-latency device close to computing nodes, it may come at a high cost for the HPC applications: the continuous interaction with the parallel file system (i.e., I/O read requests) may introduce a huge interference at the parallel file system level and thus end up with a degraded and unpredictable performance for HPC applications. In [7], we introduce interference and performance models for both HPC and Big Data applications in order to identify the performance gain and the interference cost of the prefetching technique of Eley; and demonstrate how Eley chooses the best action to optimize the prefetching while guaranteeing the pre-defined QoS requirement of HPC applications. For example, with 5% QoS requirement of the HPC application, Eley reduces the execution time of Big Data applications by up to 30% compared to the Naive burst buffer solution (NaiveBB) while guaranteeing the QoS requirement. On the other hand, the NaiveBB violates the QoS requirement by up to 58%.



Besides Clouds, Data Stream Processing (DSP) applications are widely deployed in HPC systems, especially the ones which require timely responses. DSP applications are often modelled as a directed acyclic graph: operators with data streams among them. Inter-operator communications can have a significant impact on the latency of DSP applications, accounting for 86% of the total latency. Despite their impact, there has been relatively little work on optimizing inter-operator communications, focusing on reducing inter-node traffic but not considering inter-process communication (IPC) inside a node, which often generates high latency due to the multiple memory-copy operations. In [26], we introduce a new DSP system designed specifically to address the high latency caused by inter-operator communications, called TurboStream. To achieve this goal, we introduce (1) an improved IPC framework with OSRBuffer, a DSP-oriented buffer, to reduce memory-copy operations and waiting time of each single message when transmitting messages between the operators inside one node, and (2) a coarse-grained scheduler that consolidates operator instances and assigns them to nodes to diminish the inter-node IPC traffic. Using a prototype implementation, we show that our improved IPC framework reduces the end-to-end latency of intra-node IPC by 45.64% to 99.30%. Moreover, TurboStream reduces the latency of DSP by 83.23% compared to JStorm.

Current data stream or operation stream paradigms cannot handle data burst efficiently, which probably results in noticeable performance degradation. In [25], we introduce a dual-paradigm stream processing, called DO (Data and Operation) that can adapt to stream data volatility. It enables data to be processed in micro-batches (ie, operation stream) when data burst occurs to achieve high throughput, while data is processed record by record (ie, data stream) in the remaining time to sustain low latency. DO embraces a method to detect data bursts, identify the main operations affected by the data burst and switch paradigms accordingly. Our insight behind DO's design is that the trade-off between latency and throughput of stream processing frameworks can be dynamically achieved according to data communication among operations in a fine-grained manner (ie, operation level) instead of framework level. We implement a prototype stream processing framework that adopts DO. Our experimental results show that our framework with DO can achieve 5x speedup over operation stream under low data stream sizes, and outperforms data stream on throughput by 2.1 x to 3.2 x under data burst.

In the context of the Hydda project, where hybrid HPC/Cloud infrastructures are studied, heterogeneous dataflows, composed of coarse-grain tasks interconnected through data dependencies, are scheduled. Indeed, in heterogeneous dataflows, genomics dataflows for instance, some tasks may need HPC infrastructures (*e.g.*, simulation) while other are suited for Cloud infrastructures (*e.g.*, Big Data). Different quality of services are also expected from one task to the other. In [31] the scheduling of heterogeneous scientific dataflows is studied while minimizing the Cloud provider operational costs, by introducing a deadline-aware algorithm. Scheduling in a Cloud environment is a difficult optimization problem. Usually, works around the scheduling of scientific dataflows focus on public Clouds where the management of the infrastructure is an unknown black box. Thus, many works offer scheduling algorithms built to choose the best set of virtual machines through time such that the cost of the enduser is minimized. This paper presents a new algorithm based on HEFT that aims at minimizing the number of machines used by the Cloud provider, by taking deadlines into account.

## 7.2. Programming Support

**Participants:** Zakaria Al-Shara, Frederico Alvares, Maverick Chardet, H el ene Coullon, Thomas Ledoux, Jacques Noy e, Dimitri Pertin.

Our contributions regarding the programming support are divided in two topics. First, we have contributed to automated deployment and reconfiguration with three publications. Second, we have contributed to autonomic computing and self-management in the Cloud with two publications. While these topics are strongly related (*i.e.*, a reconfiguration system is an autonomic controller), we have decided to distinguish two different levels of contributions, one being based on deployment and reconfiguration execution, or software commissioning (low level system commands), while the other uses model-driven software engineering techniques to build common self-management models for the Cloud (high level abstractions).

### 7.2.1. *Deployment and reconfiguration in the Cloud*

Distributed software architecture is composed of multiple interacting modules, or components. Deploying such software consists in installing them on a given infrastructure and leading them to a functional state. However, since each module has its own life cycle and might have various dependencies with other modules, deploying such software is a very tedious task, particularly on massively distributed and heterogeneous infrastructures. To address this problem, many solutions have been designed to automate the deployment process. In [14], we introduce Madeus, a component-based deployment model for complex distributed software. Madeus accurately describes the life cycle of each component by a Petri net structure, and is able to finely express the dependencies between components. The overall dependency graph it produces is then used to reduce deployment time by parallelizing deployment actions. While this increases the precision and performance of the model, it also increases its complexity. For this reason, the operational semantics needs to be clearly defined to prove results such as the termination of a deployment. In this paper, we formally describe the operational semantics of Madeus, and show how it can be used in a use-case: the deployment of OpenSatsck, a real and large distributed software.

Distributed software and infrastructures also become more and more dynamic. Therefore, there is a need for models assisting their management, including their reconfiguration. We focus on three properties for reconfigurations. First, we think that the efficiency of a reconfiguration is of first importance as a running service should not be interrupted for a long period of time (downtime minimization). Second, we think that it is important to offer generic reconfiguration models to help developers building complex reconfigurations. Such models offer safety properties and a clear expressivity to guide the developer. Third, multiple actors are involved in reconfigurations. On one side, developers of components are responsible for describing components life cycles, while on the other side, different developers or IT administrators could be responsible for the reconfiguration design of a complete distributed software composed of multiple connected components. To be able to simplify the reconfiguration design, it is important to offer the good abstraction level to each actor by guaranteeing a separation of concerns. Existing reconfiguration models are either specific to a subset of reconfigurations or are unable to provide both good performance and high separation of concerns between the actors interacting with them. In [32], we present an extension that could be applied both to Aeolus (an existing reconfiguration model) and Madeus (our deployment model). This extension introduces reconfiguration to Madeus, and enhances the separation of concerns compared to Aeolus. To this purpose, we introduce the behavior concept such that more elaborated life-cycles can be handled by Madeus. The obtained life-cycle defined by the component developer is complex and not adapted to the reconfiguration designer. Thus, we also introduce a minimal view of each life-cycle, namely behavioral interfaces, such that the reconfiguration is still possible but hides intricate details of each component life-cycle.

In [13] we present our complete plans to extend Madeus to support reconfiguration and to provide a good separation of concerns.

### 7.2.2. *Autonomic computing and self-management*

A Cloud needs autonomic controllers to be handled efficiently. Such controllers mostly follow a loop with four steps: monitor the system or the infrastructure, analyze the situation according to the monitoring and a set of models, plan and execute actions in consequences. In the Cloud management, multiple autonomic controllers have to be designed at each level of service (*e.g.*, IaaS, PaaS, SaaS etc.). Moreover, each autonomic controller is connected to the others. In the context of massively geo-distributed infrastructures such as Fog computing, autonomic controllers will also be decentralized, thus increasing the need for generic models of autonomic controllers and their coordination.

In the CoMe4ACloud project [3], [12], we propose a generic model-based architecture for autonomic management of Cloud systems. We derive a generic unique Autonomic Manager (AM) capable of managing any Cloud service, regardless of its XaaS layer. This AM is based on a constraint solver which aims at finding the optimal configuration for the modeled XaaS, *i.e.* the best balance between costs and revenues while meeting the constraints established by the SLA between the producer and the consumer of the Cloud service. In [12], we introduce the designed model-based architecture, and notably its core generic XaaS modeling language. We

present as well the interoperability with a Cloud standard (TOSCA). In [3], we evaluate our approach in two different ways. Firstly, we analyze qualitatively the impact of the AM behavior on the system configuration when a given series of events occurs. We show that the AM takes decisions in less than 10 s for several hundred nodes simulating virtual/physical machines. Secondly, we demonstrate the feasibility of the integration with real Cloud systems, such as OpenStack, while still remaining generic.

### 7.3. Energy-aware computing

**Participants:** Jean-Marc Menaud, Shadi Ibrahim, Thomas Ledoux, Emile Cadorel, Yewan Wang, Jonathan Pastor.

Energy consumption is one of the major challenges of modern datacenters and supercomputers. Our works in Energy-aware computing can be categorized into two subdomains: Software level (SaaS, PaaS) and Infrastructure level (IaaS).

At Software level, we worked on the general Cloud applications architecture and HPC applications.

In particular, in his habilitation thesis [2], Thomas Ledoux shows that dynamic reconfiguration in Cloud computing can provide an answer to an important societal challenge, namely digital and energetic transitions. Unlike current work providing solutions in the lower layers of the Cloud to improve the energy efficiency of data centers, Thomas Ledoux advocates a software eco-elasticity approach on the high layers of the Cloud. Inspired by both the concept of frugal innovation (Jugaad) and the mechanism of energy brownout, he proposes a number of original artifacts – such as Cloud SLA, eco-elasticity in the SaaS layers, virtualization of energy or green energy-aware SaaS applications, etc. – to reduce the carbon footprint of Cloud architectures.

However, by applying Green Programming techniques, developers have to iteratively implement and test new versions of their software, thus evaluating the impact of each code version on their energy, power and performance objectives. This approach is manual and can be long, challenging and complicated, especially for High Performance Computing applications. In [21], we formally introduce the definition of the Code Version Variability (CVV) leverage and present a first approach to automate Green Programming (i.e., CVV usage) by studying the specific use-case of an HPC stencil-based numerical code, used in production. This approach is based on the automatic generation of code versions thanks to a Domain Specific Language (DSL), and on the automatic choice of code version through a set of actors. Moreover, a real case study is introduced and evaluated through a set of benchmarks to show that several trade-offs are introduced by CVV. Finally, different kinds of production scenarios are evaluated through simulation to illustrate possible benefits of applying various actors on top of the CVV automation. While this work takes HPC applications as a use-case the presented automated green programming technique could be applied to any kind of production application onto any kind of infrastructures.

In general, many Big Data processing applications nowadays run on large-scale multi-tenant clusters. Due to hardware heterogeneity and resource contentions, straggler problem has become the norm rather than the exception in such clusters. To handle the straggler problem, speculative execution has emerged as one of the most widely used straggler mitigation techniques. Although a number of speculative execution mechanisms have been proposed, as we have observed from real-world traces, the questions of “when” and “where” to launch speculative copies have not been fully discussed and hence cause inefficiencies on the performance and energy of Big Data applications. In [29], we propose a performance model and an energy consumption model to reveal the performance and energy variations with different speculative execution solutions. We further propose a window-based dynamic resource reservation and a heterogeneity-aware copy allocation technique to answer the “when” and “where” questions for speculative executions. Evaluations using real-world traces show that our proposed technique can improve the performance of Big Data applications by up to 30% and reduce the overall energy consumption by up to 34%.

At infrastructure level, we worked on power and thermal management from server to datacenter. In fact, with the advent of Cloud Computing, the size of datacenters is ever increasing and the management of servers and their power consumption and heat production have become challenges. The management of the heat produced by servers has been experimentally less explored than the management of their power consumption. It can



be partly explained by the lack of a public testbed that provides reliable access to both thermal and power metrics of server rooms. In [34], [20], [19] we had describe SeDuCe, a testbed that targets research on power and thermal management of servers, by providing public access to precise data about the power consumption and the thermal dissipation of 48 servers integrated in Grid'5000 as the new ecotype cluster. We presented the chosen software and hardware architecture for the SeDuCe testbed. Future work will focus on two areas: adding renewable energy capabilities to the SeDuCe testbed, and improving the precision of temperature sensors.

If SeDuCe testbed is focused on the management of the power consumption and heat produced by servers at room level, we realized in [30], [24], [23], studies on power consumption (and heat impact) of physical servers. First, we characterized some potential factors on the power variation of the servers, such as: original fabrication, position in the rack, voltage variation and temperature of components on motherboard. The results show that certain factors, such as original fabrication, ambient temperature and CPU temperature, have noticeable effects on the power consumption of servers. The experimental results emphasize the importance of adding these external factors into the metric, so as to build an energy predictive model adaptable in real situations.

## 7.4. Security and Privacy

**Participants:** Mario Südholt, Mohammad Mahdi Bazm, Fatima-Zahra Boujdad, Jean-Marc Menaud.

This year the team has provided two major contributions on security and privacy challenges in distributed systems. First, we have developed our models and techniques for the detection and mitigation of side-channel attacks. Second, we have provided a first model and implementation techniques for secure and privacy-preserving distributed biomedical analyses, notably genomic ones.

### 7.4.1. Side-channel attacks and trusted Fog/Edge infrastructures

In [4], we investigate Cloud computing infrastructures, which are based on the sharing of hardware resources among different clients. The infrastructures leverage virtualization to share physical resources among several self-contained execution environments like virtual machines and Linux containers. Isolation is a core security challenge for such a paradigm. It may be threatened through side-channels, created due to the sharing of physical resources like caches of the processor or by mechanisms implemented in the virtualization layer. Side-channel attacks (SCAs) exploit and use such leaky channels to obtain sensitive data like kernel information. We clarify the nature of this threat for cloud infrastructures. Current SCAs are done locally and exploit isolation challenges of virtualized environments to retrieve sensitive information. We also introduce the concept of distributed side-channel attack (DSCA). We explore how such attacks can threaten isolation of any virtualized environments. Finally, we study a set of different applicable countermeasures for attack mitigation in cloud infrastructures.

In [9], we investigate Fog and Edge computing for the provision of large pools of resources at the edge of the network that may be used for distributed computing. Fog infrastructure heterogeneity also results in complex configuration of distributed applications on computing nodes. Linux containers are a mainstream technique allowing to run packaged applications and micro services. However, running applications on remote hosts owned by third parties is challenging because of untrusted operating systems and hardware maintained by third parties. To meet such challenges, we may leverage trusted execution mechanisms. In this work, we propose a model for distributed computing on Fog infrastructures using Linux containers secured by Intel's Software Guard Extensions (SGX) technology. We implement our model on a Docker and OpenSGX platform. The result is a secure and flexible approach for distributed computing on Fog infrastructures.

In [10], we contribute to the research on cache-based side-channel attacks and show the security impact of these attacks on cloud computing. The detection of cache-based side-channel attacks has received more attention in IaaS cloud infrastructures because of improvements in the attack techniques. However, detection of such attacks requires high resolution information, and it is also a challenging task because of the fine-granularity of the attacks. In this paper, we present an approach to detect cross-VM cache-based side-channel attacks through using hardware fine-grained information provided by Intel Cache Monitoring Technology

(CMT) and Hardware Performance Counters (HPCs) following the Gaussian anomaly detection method. The approach shows a high detection rate with a 2% performance overhead on the computing platform.

#### 7.4.2. *Secure and privacy-aware biomedical analyses*

In [11], we study the need for the sharing of genetic data, for instance, in genome-wide association studies, which is incessantly growing. In parallel, serious privacy concerns rise from a multi-party access to genetic information. Several techniques, such as encryption, have been proposed as solutions for the privacy-preserving sharing of genomes. However, existing programming means do not support guarantees for privacy properties and the performance optimization of genetic applications involving shared data. We propose two contributions in this context. First, we present new cloud-based architectures for cloud-based genetic applications that are motivated by the needs of geneticists. Second, we propose a model and implementation for the composition of watermarking with encryption, fragmentation, and client-side computations for the secure and privacy-preserving sharing of genetic data in the cloud.

### 7.5. Experiment-Driven Research

**Participants:** Adrien Lebre, Bastien Confais, Ronan-Alexandre Cherrueau, Matthieu Simonin, Thuy-Linh Nguyen.

Because STACK members have to perform a significant number of evaluations of complex software stack at large scale, the team contributes to the recent area of software-defined experiments and reproducible research.

In [36], [16] we propose a new approach to ensure reproducibility and repeatability of scientific experiments. Similar to the LAMP stack that considerably eased the web developers life, we advocate the need of an analogous software stack to help the experimenters making reproducible research. In 2018, we propose the EnosStack, an open source software stack especially designed for reproducible scientific experiments. EnosStack enables to easily describe experimental workflows meant to be re-used, while abstracting the underlying infrastructure running them. Being able to switch experiments from a local to a real testbed deployment greatly lower code development and validation time. In this paper, we describe the abstractions that have driven its design, before presenting a real experiment we deployed on Grid'5000 to illustrate its usefulness. We also provide all the experiment code, data and results to the community.

Similar to the previous work, we discuss in [37] a large experimental campaign that allows us to understand in details the boot duration of both virtualization techniques under various storage devices and resources contentions. While many studies have been focusing on reducing the time to manipulate Virtual Machine/Container images in order to optimize provisioning operations in a Cloud infrastructure, only a few studies have considered the time required to boot these systems. Some previous researches showed that the whole boot process can last from a few seconds to few minutes depending on co-located workloads and the number of concurrent deployed machines. The paper explains how we analyzed thoroughly the boot time of VMs, Docker containers on top of bare-metal servers, and Docker containers inside VMs. We discuss a methodology that enables us to perform fully-automatized and reproducible experimental campaigns on a scientific testbed. Thanks to this methodology, we conducted more than 14.400 experiments on Grid'5000 testbed for a bit more than 500 hours. The results we collected provide an important information related to the boot time behavior of these two virtualization technologies.

In [33], we presented the first experiment that has been done, as far as we know, on top of the Grid'5000 and FIT testbeds. More precisely, we discuss how we evaluated a new storage service for edge/IoT scenarios. Our proof-of-concept relies on the Interplanetary Object Store (IPFS), a Scale-Out NAS deployed on each site and a tree-based approach for the meta-data management [17]. This proposal enables (i) IoT devices to write locally on their closest site and (ii) to relocate automatically the objects on the sites they are requested, leading to low access times. The contribution of this work is a discussion of our attempt of using the two platforms simultaneously as well as the problems we encountered to interconnect them. Our ultimate goal is to give guidelines on how can researchers perform evaluations in a realistic environment : IoT devices comes from the FIT/IoT-lab and Fog nodes from Grid'5000.

## 8. Bilateral Contracts and Grants with Industry

### 8.1. Bilateral Contracts with Industry

**Participants:** Adrien Lebre [Contact point], Ronan-Alexandre Cherrueau, Marie Delavergne, Alexandre Van Kempen.

During 2017, we agreed with Orange Labs (Lannion) to conduct a dedicated study on the evaluation of AMQP message bus alternatives within the OpenStack ecosystem. This bilateral contract (“Contrat de Recherche Externalisé”) officially started in Sept 2017 for one year. With the allocated budget ( 100K), we hired a new research engineer, Alexandre Van Kempen. Alexandre Van Kempen works with Ronan-Alexandre Cherrueau (Temporary Resarch Engineer, hired in the context of the MERCURY InriHub) and Matthieu Simonin (Permanent Research Engineer from the Rennes Bretagne Atlantique Center) on conducting this analysis. In addition to extending the EnOS framework previously presented, they are performing several experiments with the support of the OpenStack open-source community (in particulat RedHat). The goal of the study is to identify major drawbacks of the default RabbitMQ solution with respect to the Fog/Edge requirements and evaluate whether some alternatives are available in the open-source ecosytem.

## 9. Partnerships and Cooperations

### 9.1. Regional Initiatives

#### 9.1.1. CoMe4ACloud

**Participant:** Thomas Ledoux [coordinator].

The project, started in October 2016, was completed in March 2018. CoMe4ACloud was an Atlanstic 2020 funded project and supported a one year post-doc position. The project was led by STACK research team and involved also AtlanModels and TASC, all of them from the LS2N and situated at IMT Atlantique.

The high-level objective of the CoMe4ACloud (Constraints and Model Engineering for Autonomic Clouds) project was to provide an end-to-end solution for autonomic Cloud services. To that end, we relied on techniques of Constraint Programming so as a decision-making tool and Model-driven Engineering to ease the automatic generation of the so-called autonomic managers as well as their synchronization with the managed system (i.e., the Cloud layers).

This year, we have focus on the dissemination of the results. We got the best paper award of CLOSER 2018 (the 8th International Conference on Cloud Computing and Services Science) [12] and published in the journal FGCS [3]. We also gave a pitch in the annual Atlanstic 2020 meeting in November.

See <https://come4acloud.github.io> for more information.

#### 9.1.2. Oncoshare

**Participant:** Mario Südholt [coordinator].

The ONCOSHARE project (ONCOlogy big data SHAring for Research) will demonstrate, through a multi-disciplinary cooperation within the Western CANCEROPOLE network, the feasibility and the added value of a Cancer Patient Centered Information Common for in-silico research. The STACK team will work on challenges to the security and the privacy of user data in this context.

This project is financed by three French regions from 2018-2021.

#### 9.1.3. SyMeTRIC

**Participant:** Jean-Marc Menaud [coordinator].

SyMeTRIC is a regional federated project in Systems Medicine funded by the Pays de la Loire french region. Systems Medicine approaches can be compared to Systems Biology. They aim at integrating several information sources to design and validate bio-models and biomarkers to anticipate and enhance patients follow-up (diagnosis, treatment response prediction, prognosis).

This project is ending in 2018.

#### 9.1.4. SysMics

**Participants:** Jean-Marc Menaud, Mario Südholt [coordinator].

The SysMics project aims at federating the NExT scientific community toward a common objective: anticipate the emergence of systems medicine by co-developing 3 approaches in population-scale genomics: genotyping by sequencing, cell-by-cell profiling and microbiome analysis. STACK investigates new means for secure and privacy-aware computations in the context of personalized medicine, notably genetic analyses.

This project is financed by the Nantes excellency initiative in Medecine and Informatics (NExT) from 2018-22.

## 9.2. National Initiatives

### 9.2.1. CominLabs laboratory of excellence

#### 9.2.1.1. PrivGen

**Participants:** Fatima-Zahra Boujdad, Mario Südholt [coordinator].

PrivGen (“Privacy-preserving sharing and processing of genetic data”) is a three-year project that has been started in Oct. 2016 and is conducted by three partners: a team of computer scientists from the LATIM Inserm institute in Brest mainly working on data watermarking techniques, a team of geneticists from an Inserm institute in Rennes working on the gathering and interpretation of genetic data, and the STACK team. The project provides funding of 330 KEUR altogether with an STACK share of 120 KEUR.

The project considers challenges related to the outsourcing of genetic data that is in the Cloud by different stakeholders (researchers, organizations, providers, etc.). It tackles several limitations of current security solutions in the cloud, notably the lack of support for different security and privacy properties at once and computations executed at different sites that are executed on behalf of multiple stakeholders.

The partners are working on three main challenges:

- Mechanisms for a continuous digital content protection
- Composition of security and privacy-protection mechanisms
- Distributed processing and sharing of genetic data

The Ascola team is mainly involved in providing solutions for the second and third challenges.

#### 9.2.1.2. SeDuCe++

**Participants:** Jonathan Pastor, Jean-Marc Menaud [coordinator].

SeDuCe++ is an extended version of the SeDuCe project. Funded by the LS2N (CNRS), an allocated budget of 10KEuros for one year, it aims at studying the energy footprint of extreme edge infrastructure.

### 9.2.2. ANR

#### 9.2.2.1. GRECO (ANR)

**Participants:** Adrien Lebre [Contact point], Alexandre Van Kempen.

The GRECO project (Resource manager for cloud of Things) is an ANR project (ANR-16-CE25-0016) running for 42 months (starting in January 2017 with an allocated budget of 522KEuros, 90KEuro for ASCOLA).

The consortium is composed of 4 partners: Qarnot Computing (coordinator) and 3 academic research group (DATAMOVE and AMA from the LIG in Grenoble and ASCOLA from Inria Rennes Bretagne Atlantique).

The goal of the GRECO project (<https://anr-greco.net>) is to design a manager for cloud of things. The manager should act at the IaaS, PaaS and SaaS layer of the cloud. One of the principal challenges will consist in handling the execution context of the environment in which the cloud of things operates. Indeed, unlike classical resource managers, connected devices imply to consider new types of networks, execution supports, sensors and new constraints like human interactions. The great mobility and variability of these contexts complexify the modelling of the quality of service. To face this challenge, we intend to innovate in designing scheduling and data management systems that will use machine learning techniques to automatically adapt their behaviour to the execution context. Adaptation here requires a modelling of the recurrent cloud of things usages, the modelling of the dynamics of physical cloud architecture.

#### 9.2.2.2. *KerStream* (ANR)

**Participant:** Shadi Ibrahim [Coordinator].

The KerStream project (Big Data Processing: Beyond Hadoop!) is an ANR JCJC (Young Researcher) project (ANR-16-CE25-0014-1) running for 48 months (starting in January 2017 with an allocated budget of 238KEuros).

The goal of the KerStream project is to address the limitations of Hadoop when running Big Data stream applications on large-scale clouds and do a step beyond Hadoop by proposing a new approach, called KerStream, for scalable and resilient Big Data stream processing on clouds. The KerStream project can be seen as the first step towards developing the first French middleware that handles Stream Data processing at Scale.

### 9.2.3. *FSN*

#### 9.2.3.1. *Hydda* (FSN)

**Participants:** Helene Coullon, Jean-Marc Menaud [coordinator].

The HYDDA project aims to develop a software solution allowing the deployment of Big Data applications (with hybrid design (HPC/Cloud)) on heterogeneous platforms (cluster, Grid, private Cloud) and orchestrators (Task scheduler like Slurm, Virtual orchestrator (like Nova for OpenStack or Swarm for Docker). The main questions we are investigating are :

- How to propose an easy-to-use service to host (from deployment to elimination) application components that are both typed Cloud and HPC?
- How propose a service that unifies the HPCaaS (HPC as a service) and the Infrastructure as a Service (IaaS) in order to offer resources on demand and to take into account the specificities of scientific applications?
- How optimize resources usage of these platforms (CPU, RAM, Disk, Energy, etc.) in order to propose solutions at the least cost?

### 9.2.4. *CPER*

#### 9.2.4.1. *SeDuCe*

**Participants:** Adrien Lebre, Jean-Marc Menaud [coordinator], Jonathan Pastor.

The SeDuCe project (Sustainable Data Centers: Bring Sun, Wind and Cloud Back Together), aims to design an experimental infrastructure dedicated to the study of data centers with low energy footprint. This innovative data center will be the first experimental data center in the world for studying the energy impact of cloud computing and the contribution of renewable energy (solar panels, wind turbines) from the scientific, technological and economic viewpoints. This project is integrated in the national context of grid computing (Grid'5000), and the Constellation project, which will be an inter-node (Pays de la Loire, Brittany).

### 9.2.5. *Inria Project Labs*

#### 9.2.5.1. *DISCOVERY*

**Participants:** Helene Coullon, Marie Delavergne, Shadi Ibrahim, Adrien Lebre [coordinator], Dimitri Pertin [Until May 2018], Ronan Alexandre Rcherreau, Matthieu Simonin, Alexandre Van Kempen.

To accommodate the ever-increasing demand for Utility Computing (UC) resources, while taking into account both energy and economical issues, the current trend consists in building larger and larger Data Centers in a few strategic locations. Although such an approach enables UC providers to cope with the actual demand while continuing to operate UC resources through centralized software system, it is far from delivering sustainable and efficient UC infrastructures for future needs.

The DISCOVERY initiative <sup>10</sup> aims at exploring a new way of operating Utility Computing (UC) resources by leveraging any facilities available through the Internet in order to deliver widely distributed platforms that can better match the geographical dispersal of users as well as the ever increasing demand. Critical to the emergence of such locality-based UC (also referred as Fog/Edge Computing) platforms is the availability of appropriate operating mechanisms. The main objective of DISCOVERY is to design, implement, demonstrate and promote a new kind of Cloud Operating System (OS) that will enable the management of such a large-scale and widely distributed infrastructure in an unified and friendly manner.

The consortium is composed of experts in the following research areas: large-scale infrastructure management systems, networking and P2P algorithms. Moreover, two key network operators, namely Orange and RENATER, are involved in the project.

By deploying and using a Fog/Edge OS on backbones, our ultimate vision is to enable large parts of the Internet to be hosted and operated by its internal structure itself: a scalable set of resources delivered by any computing facilities forming the Internet, starting from the larger hubs operated by ISPs, governments and academic institutions, to any idle resources that may be provided by end users.

STACK leads the DISCOVERY IPL and contributes mainly around two axes: VM life cycle management and deployment/reconfiguration concerns.

## 9.2.6. InriaHub

### 9.2.6.1. MERCURY

**Participants:** Ronan-Alexandre Cherrueau, Adrien Lebre [coordinator], Matthieu Simonin.

STACK, in particular within the framework of the DISCOVERY initiative has been working on the massively distributed use case since 2013. With the development of several proof-of-concepts around OpenStack, the team has had the opportunity to start an InriaHub action. Named MERCURY, the goal of this action is twofold: (i) support the research development made within the context of DISCOVERY and (ii) favor the transfer toward the OpenStack community.

Further information available at: <http://beyondtheClouds.github.io>.

### 9.2.6.2. APOLLO/SOYUZ

**Participants:** Ronan-Alexandre Cherrueau, Adrien Lebre [coordinator], Matthieu Simonin.

The Apollo/Soyuz is the second InriaHub action attached the DISCOVERY IPL. While MERCURY aims mainly at supporting development efforts within the DISCOVERY IPL, the APOLLO/SOYUZ is focusing on the animation and the dissemination of the DISCOVERY activities within the different open-source ecosystem (*i.e.*, OpenStack, OPNFV, etc.). One additional engineer will join the current team in January 2019.

Further information available at: <http://beyondtheClouds.github.io>.

## 9.2.7. Fond d'amorçage IMT Industrie du Futur 2017

### 9.2.7.1. aLIFE

**Participants:** Hélène Coullon [coordinator], Jacques Noyé.

We have organized, in partnership with colleagues from IMT Atlantique, the aLIFE workshop between industry and academia, which took place in Nantes during two days on 30-31 January.

---

<sup>10</sup><http://beyondtheclouds.github.io>

The objective was to share experience and success stories, as well as open challenges related to the contribution of software-related research to Factories of the Future, in French *Apport de l'industrie du Logiciel à l'Industrie du Futur Européenne (aLIFE)*. 86 people registered to the workshop, organized around plenary sessions and discussion panels, with speakers from Airbus, Baldwin Partners, Comau (Italy), Dassault Systèmes, e.l.m. Leblanc, La Poste, Naval Group, Predict, Fraunhofer (Germany), KTH (Sweden), Polytechnique Montréal (Canada), and TUM (Germany).

### 9.2.8. Connect Talent

#### 9.2.8.1. Apollo (Connect Talent)

**Participant:** Shadi Ibrahim [Coordinator].

The Apollo project (Fast, efficient and privacy-aware Workflow executions in massively distributed Data-centers) is an individual research project “Connect Talent” running for 36 months (starting in November 2017 with an allocated budget of 201KEuros).

The goal of the Apollo project is to investigate novel scheduling policies and mechanisms for fast, efficient and privacy-aware data-intensive workflow executions in massively distributed data-centers.

### 9.2.9. Etoiles Montantes

#### 9.2.9.1. VeRDi

**Participant:** Hélène Coullon [Coordinator].

VeRDi is an acronym for Verified Reconfiguration Driven by execution. The VeRDi project is funded by the French region Pays De La Loire where Nantes is located. The project starts in November 2018 and ends on December 2020 with an allocated budget of 172800€.

It aims at addressing distributed software reconfiguration in an efficient and verified way. The aim of the VeRDi project is to build an argued disruptive view of the problem. To do so we want to validate the work already performed on the deployment in the team and extend it to reconfiguration.

## 9.3. European Initiatives

### 9.3.1. H2020 Projects

#### 9.3.1.1. BigStorage

Title: BigStorage: Storage-based Convergence between HPC and Cloud to handle Big Data

Programm: H2020

Duration: January 2015 - December 2018

Coordinator: Universidad politecnica de Madrid

Partners:

Barcelona Supercomputing Center - Centro Nacional de Supercomputacion (Spain)

Ca Technologies Development Spain (Spain)

Commissariat A L Energie Atomique et Aux Energies Alternatives (France)

Deutsches Klimarechenzentrum (Germany)

Foundation for Research and Technology Hellas (Greece)

Fujitsu Technology Solutions (Germany)

Johannes Gutenberg Universitaet Mainz (Germany)

Universidad Politecnica de Madrid (Spain)

Seagate Systems Uk (United Kingdom)

Inria contact: G. Antoniu & A. Lebre

The consortium of this European Training Network (ETN) 'BigStorage: Storage-based Convergence between HPC and Cloud to handle Big Data' will train future data scientists in order to enable them and us to apply holistic and interdisciplinary approaches for taking advantage of a data-overwhelmed world, which requires HPC and Cloud infrastructures with a redefinition of storage architectures underpinning them - focusing on meeting highly ambitious performance and energy usage objectives. There has been an explosion of digital data, which is changing our knowledge about the world. This huge data collection, which cannot be managed by current data management systems, is known as Big Data. Techniques to address it are gradually combining with what has been traditionally known as High Performance Computing. Therefore, this ETN will focus on the convergence of Big Data, HPC, and Cloud data storage, its management and analysis. To gain value from Big Data it must be addressed from many different angles: (i) applications, which can exploit this data, (ii) middleware, operating in the cloud and HPC environments, and (iii) infrastructure, which provides the Storage, and Computing capable of handling it. Big Data can only be effectively exploited if techniques and algorithms are available, which help to understand its content, so that it can be processed by decision-making models. This is the main goal of Data Science. We claim that this ETN project will be the ideal means to educate new researchers on the different facets of Data Science (across storage hardware and software architectures, large-scale distributed systems, data management services, data analysis, machine learning, decision making). Such a multifaceted expertise is mandatory to enable researchers to propose appropriate answers to applications requirements, while leveraging advanced data storage solutions unifying cloud and HPC storage facilities.'

## 9.4. International Initiatives

### 9.4.1. Inria International Partners

#### 9.4.1.1. Informal International Partners

Huazhong university of Science and Technology (HUST): We collaborate on resource management and task scheduling for stream data applications in the cloud.

National University of Singapore (NUS): We collaborate on mitigating stragglers for Big Data applications in clouds and optimizing graph processing in geo-distributed data-centers.

ShenZhen University: We collaborate on data management in HPC systems, mitigating stragglers for Big Data applications in clouds and optimizing graph processing in geo-distributed data-centers.

## 9.5. International Research Visitors

### 9.5.1. Visits to International Teams

#### 9.5.1.1. Research Stays Abroad

HUST, China: From October 20 to November 5, S. Ibrahim visited the Services Computing Technology and System Lab at Huazhong university of Science and Technology.

Flagstaff, Northern Arizona University, USA: From September 10 to September 16, H. Coullon visited the School of Informatics, Computing and Cyber Systems, Northern Arizona University, Flagstaff, USA.

## 10. Dissemination

### 10.1. Promoting Scientific Activities

#### 10.1.1. Scientific Events Organisation

##### 10.1.1.1. Member of the Organizing Committees



- A. Lebre has co-organized the 1st edition of the SILECS (Grid'5000/FIT) School, Nice, April 2018 (50 persons).

### **10.1.2. Scientific Events Selection**

#### *10.1.2.1. Chair of Conference Program Committees*

- H. Coullon was the chair of the IEEE Big Data Congress 2018 "Quality of Big Data Services" track.
- S. Ibrahim was the program chair of the IEEE Big Data Congress 2018.
- S. Ibrahim was the program co-chair of the International workshop on the Convergence of Extreme Scale Computing and Big Data Analysis (CEBDA 2018), co-located with IPDPS'18, Vancouver, Canada, May 2018.
- S. Ibrahim was the track co-chair for the "Big data management and scalable storage technologies track" of the 6th International Conference on Emerging Internet, Data and Web Technologies (EIDWT 2018).
- A. Lebre was the chair of the System Software track at SBAC-PAD'18.

#### *10.1.2.2. Member of the Conference Program Committees*

- H. Coullon was a member of the program committees of CloudCom'18, ICCS'18, ScalCom'18, CIOt'18.
- S. Ibrahim was a member of the program committees of SC'18, CCGrid'18, ICA3PP'18, IPCCC'18, SCA'18, CloudCom'18, Trust-Com'18, Innovate-Data'18, PDSW-DISCS@SC'18, WAC-18@HPDC'18, HPBDC@IPDPS'18.
- A. Lebre was a member of the program committees of IC2E'18, CCGRID'18, IC FEC'18, CloudCom'18, and the Vancouver OpenStack summit.
- T. Ledoux was a member of the program committees of the conference Compas'18 and Cross-Cloud'18@EuroSys, Greens'18@ICSE workshops.
- M. Südholt was a member of the program committees of CloudCom'18, ProWeb'18, ICICS'18, SBAC-PAD'18, and COP'18.

### **10.1.3. Journal**

#### *10.1.3.1. Member of the Editorial Boards*

- A. Lebre is an Associate Editor of the IEEE Transactions on BigData
- S. Ibrahim is a Guest Editor of Springer FGCS Journal – Special Issue on the Convergence of Extreme Scale Computing and Big Data Analysis.
- M. Südholt is an Associate Editor of the journal Programming (Springer).

#### *10.1.3.2. Reviewer - Reviewing Activities*

- H. Coullon has been a reviewer for the Future Generation Computer Systems journal
- H. Coullon has been a reviewer for the following conferences: IPDPS'18, CCGrid'18.
- A. Lebre has been a reviewer for the Future Generation Computer Systems journal.
- T. Ledoux has been a reviewer for the following journals: Journal of Systems Architecture - Elsevier; Future Generation Computer Systems.
- S. Ibrahim has been a reviewer for the following journals: Journal of Parallel and Distributed Computing, Future Generation Computer Systems, IEEE Internet of Things, Computer Networks, and Big Data Research.
- J.-M. Menaud has been a reviewer for IEEE Transactions on Parallel and Distributed Systems

### **10.1.4. Invited Talks**

- H. Coullon has been invited to present a talk at the Northern Arizona University, USA, September 2018: “Performance, Software engineering and verification: a winning combination for HPC and Utility computing?”.
- H. Coullon has been invited to present a talk at IriSa 68NQRT seminar, Rennes, France: “Toward efficient and safe deployment and reconfiguration of distributed software”.
- H. Coullon has been invited to present a talk at the Languages, Compilation, and Semantics LIP Seminar in Lyon, France: “Toward efficient and safe deployment and reconfiguration of distributed software”.
- R.A. Cherrueau has been invited to present a talk at the “Opérations France Grille” workshop on “Edge Computing Infrastructure with OpenStack”
- M. Südholt has been invited to present a talk at the French Institute of Bioinformatics (IFB) on “Compositional security and privacy for biomedical analyses using shared genetic data.”
- M. Südholt has been invited to present a talk at the IMT Cybersecurity day on “Privacy and sharing of genomic data.”
- J.-M. Menaud has been invited to present a talk at the IMT Energy and IoT day on “Energy Management on Cloud/Edge Computing”

#### 10.1.5. Scientific Expertise

- A. Lebre has been the chair of the OpenStack “Fog/Edge/Massively Distributed Clouds” Special Interest Group until May 2018 (further information at: [https://wiki.openstack.org/wiki/Massively\\_Distributed\\_Clouds](https://wiki.openstack.org/wiki/Massively_Distributed_Clouds)).
- A. Lebre is a member of the scientific committee of the Inria - Nokia/Bell Labs.

#### 10.1.6. Research Administration

- S. Ibrahim is the co-coordinator of the international Master’s program in Cloud Computing and Services at University of Rennes 1.
- S. Ibrahim was a member of Grid’5000 Sites Committee – Responsible for the Rennes site till April 2018.
- A. Lebre is a member of the executive committee of the GDR CNRS RSD “Réseau et Système distribué” and Co-leader of the transversal action Virtualization and Clouds of this GDR since 2015.
- A. Lebre is a member of the executive and architect committees of the Grid’5000 GIS (Groupement d’intérêt scientifique).
- A. Lebre is a member of the executive committee of the <I/O> Lab, a joint lab between Inria and Orange Labs.
- T. Ledoux is the head of the apprenticeship program in Software Engineering FIL (<http://www.imt-atlantique.fr/formation/ingenieur-par-apprentissage/ingenieurs-specialite-ingenierie-logicielle>). This 3-year program leads to the award of a Master degree in Software Engineering from the IMT Atlantique.
- J.-M. Menaud is the organizer of "Pôle Science du Logiciel et des Systèmes Distribués" in Laboratoire des Sciences du Numérique à Nantes (LS2N) since June 2015.
- J. Noyé is deputy head of the Automation, Production and Computer Sciences department of IMT Atlantique.

## 10.2. Teaching - Supervision - Juries

### 10.2.1. Supervision

- HdR : Thomas Ledoux, “Reconfiguration dynamique d’architectures logicielles : des métaclasse aux « nuages verts », Université de Nantes, defended on July 2018, the 17th.

- PhD: Bastien Confais, co-director: A. Lebre, director: B. Parrein (RIO, LS2N Nantes), “Conception d’un système de partage de données adapté à un environnement de Fog Computing”, defended on July 2018, the 10th.
- Phd: Mohamed Abderrahim (Orange CIFRE), director: A. Lebre, “Conception d’un système de supervision programmable et reconfigurable pour une infrastructure informatique et réseau répartie”, defended on Dec 2018, the 19th.
- PhD: Mohammad Mahdi Bazm: codirectors: Mario Südholt, J-M. Menaud.
- PhD: Maxime Belair, director: J-M. Menaud.
- PhD: Emile Cadorel, director: advisor: H. Coullon, director: J-M. Menaud.
- PhD: Fatima-zahra Boujdad, director: Mario Südholt.
- PhD: Maverick Chardet, advisor: H. Coullon, director: A. Lebre.
- PhD: David Espinel (Orange CIFRE), director: A. Lebre.
- PhD: Yewan Wang, director: J-M. Menaud.
- PhD: Thuy-Linh NGuyen, director: A. Lebre
- PhD: Dimitri Saingre, advisor: T. Ledoux, director: J-M. Menaud.
- PhD: Jad Darrous, advisor: S. Ibrahim, director: C. Perez (Avalon)
- Postdoc: Thomas Lambert, advisor: S. Ibrahim.
- Postdoc: Jonathan Pastor, advisor: J-M. Menaud.
- Postdoc: Alexandre Van Kempen, advisor: A. Lebre.

### 10.2.2. Juries

- H. Coullon was a reviewer of the PhD of Stéphanie Challita, “Inferring Models from Cloud APIs and Reasoning over Them: A Tooled and Formal Approach”, University of Lille, Inria Lille-Nord Europe, France, Dec. 21 2018.
- H. Coullon was a reviewer of the PhD of Gustavo Sousa, “A Software Product Lines-Based Approach for the Setup and Adaptation of Multi-Cloud Environments”, University of Lille, France, June 5 2018.
- J.-M. Menaud was a member of the PhD committee of David Guyon (Dec. 7, 2018) "Supporting Energy-awareness for Cloud Users", UBL, Dec 2018.
- S. Ibrahim was reviewer of the PhD of Paul Hermann Lensing, “Direct Lookup and Hash-Based Metadata Placement - Impact on Architecture, Performance and Scalability of Local and Distributed File Systems”, Universitat Politecnica de Catalunya Barcelona, Spain, Oct. 10 2018.
- A. Lebre was a member of the PhD Committee of Xia Yé, “Combining Heuristics for Optimizing and Scaling the Placement of IoT Applications in the Fog”, University of Grenoble, Dec 2018.
- T. Ledoux was reviewer of the PhD of Yahya Al-Dhuraibi, "Flexible Framework for Elasticity in Cloud Computing", Univ. Lille, December 2018

## 10.3. Popularization

- J.-M. Menaud, Online publication <http://www.lemondeinformatique.fr>: *L’IMT Atlantique veut réduire le coût des calculs lourds.*
- J.-M. Menaud, Online publication <http://www.mtom-mag.com>: *Le projet HYDDA, développer une plateforme innovante, associant Cloud et HPC.*
- J.-M. Menaud, Online publication <http://www.agence-api.fr>: *L’Institut Mines-Telecom Atlantique monte une plateforme cloud et HPC.*

# 11. Bibliography

## Publications of the year

### Doctoral Dissertations and Habilitation Theses

- [1] B. CONFAIS. *A sharing data system adapted to a Fog Computing environment*, Université de Nantes, July 2018, <https://hal.archives-ouvertes.fr/tel-01868308>
- [2] T. LEDOUX. *Dynamic reconfiguration of software architectures: from metaclasses to "green Clouds"*, Université de Nantes, Faculté des sciences et des techniques, July 2018, Habilitation à diriger des recherches, <https://tel.archives-ouvertes.fr/tel-01864344>

### Articles in International Peer-Reviewed Journals

- [3] Z. AL-SHARA, F. ALVARES, H. BRUNELIERE, J. LEJEUNE, C. PRUD'HOMME, T. LEDOUX. *CoMe4ACloud: An End-to-End Framework for Autonomic Cloud Systems*, in "Future Generation Computer Systems", September 2018, vol. 86, pp. 339-354 [DOI : 10.1016/J.FUTURE.2018.03.039], <https://hal.archives-ouvertes.fr/hal-01762716>
- [4] M.-M. BAZM, M. LACOSTE, M. SÜDHOLT, J.-M. MENAUD. *Isolation in Cloud Computing Infrastructures: New Security Challenges*, in "Annals of Telecommunications, Springer", 2018, <https://hal.inria.fr/hal-01874206>
- [5] A. LEBRE, J. PASTOR, A. SIMONET, M. SÜDHOLT. *Putting the Next 500 VM Placement Algorithms to the Acid Test: The Infrastructure Provider Viewpoint*, in "IEEE Transactions on Parallel and Distributed Systems", July 2018, pp. 1-14 [DOI : 10.1109/TPDS.2018.2855158], <https://hal.inria.fr/hal-01816248>
- [6] Y. LI, A.-C. ORGERIE, I. RODERO, B. LEMMA AMERSHO, M. PARASHAR, J.-M. MENAUD. *End-to-end Energy Models for Edge Cloud-based IoT Platforms: Application to Data Stream Analysis in IoT*, in "Future Generation Computer Systems", October 2018, vol. 87, pp. 667-678 [DOI : 10.1016/J.FUTURE.2017.12.048], <https://hal.archives-ouvertes.fr/hal-01673501>
- [7] O. YILDIZ, A. C. ZHOU, S. IBRAHIM. *Improving the Effectiveness of Burst Buffers for Big Data Processing in HPC Systems with Eley*, in "Future Generation Computer Systems", March 2018, vol. 86, pp. 308-318 [DOI : 10.1016/J.FUTURE.2018.03.029], <https://hal.archives-ouvertes.fr/hal-01742090>

### Invited Conferences

- [8] M. ABDERRAHIM, M. OUZZIF, K. GUILLOUARD, J. FRANÇOIS, A. LEBRE, C. PRUD'HOMME, X. LORCA. *Efficient Resource Allocation for Multi-tenant Monitoring of Edge Infrastructures*, in "PDP 2019 - 27th Euromicro International Conference on Parallel, Distributed and Network-Based Processing", Pavie, Italy, February 2019, pp. 1-8, <https://hal.inria.fr/hal-01987946>

### International Conferences with Proceedings

- [9] M.-M. BAZM, M. LACOSTE, M. SÜDHOLT, J.-M. MENAUD. *Secure Distributed Computing on Untrusted Fog Infrastructures Using Trusted Linux Containers*, in "CloudCom 2018 - 10th IEEE International Conference on Cloud Computing Technology and Science", Nicosia, Cyprus, December 2018, pp. 1-7, <https://hal.inria.fr/hal-01875777>

- [10] M.-M. BAZM, T. SAUTEREAU, M. LACOSTE, M. SÜDHOLT, J.-M. MENAUD. *Cache-Based Side-Channel Attacks Detection through Intel Cache Monitoring Technology and Hardware Performance Counters*, in "FMEC 2018 - Third IEEE International Conference on Fog and Mobile Edge Computing", Barcelona, Spain, IEEE, April 2018, pp. 1-6 [DOI : 10.1109/FMEC.2018.8364038], <https://hal.inria.fr/hal-01762803>
- [11] F.-Z. BOUJAD, M. SÜDHOLT. *Constructive Privacy for Shared Genetic Data*, in "CLOSER 2018 - 8th International Conference on Cloud Computing and Services Science", Funchal, Madeira, Portugal, Proceedings of CLOSER 2018, March 2018, pp. 1-8, <https://hal.archives-ouvertes.fr/hal-01692620>
- [12] *Best Paper*  
H. BRUNELIERE, Z. AL-SHARA, F. ALVARES, J. LEJEUNE, T. LEDOUX. *A Model-based Architecture for Autonomic and Heterogeneous Cloud Systems*, in "CLOSER 2018 - 8th International Conference on Cloud Computing and Services Science", Funchal, Portugal, March 2018, vol. 1, pp. 201-212, Best Paper Award [DOI : 10.5220/0006773002010212], <https://hal.archives-ouvertes.fr/hal-01705248>.
- [13] M. CHARDET. *Combining Separation of Concerns and Performance in Distributed Software Reconfiguration*, in "Doctoral symposium @ Middleware 2018 - ACM/IFIP International Middleware Conference", Rennes, France, December 2018, pp. 1-6, <https://hal.inria.fr/hal-01939853>
- [14] M. CHARDET, H. COULLON, D. PERTIN, C. PÉREZ. *Madeus: A formal deployment model*, in "4PAD 2018 - 5th International Symposium on Formal Approaches to Parallel and Distributed Systems (hosted at HPCS 2018)", Orléans, France, July 2018, pp. 1-8, <https://hal.inria.fr/hal-01858150>
- [15] R.-A. CHERRUEAU, A. LEBRE, D. PERTIN, F. WUHI, J. M. SOARES. *Edge Computing Resource Management System: a Critical Building Block! Initiating the debate via OpenStack*, in "HotEdge 2018 - USENIX Workshop on Hot Topics in Edge Computing", Boston, MA, United States, July 2018, pp. 1-6, <https://hal.archives-ouvertes.fr/hal-01812747>
- [16] R.-A. CHERRUEAU, M. SIMONIN, A. VAN KEMPEN. *EnosStack: A LAMP-like stack for the experimenter*, in "INFOCOM WKSHPs 2018 - IEEE International Conference on Computer Communications", Honolulu, United States, IEEE, April 2018, pp. 336-341 [DOI : 10.1109/INFOCOMW.2018.8407024], <https://hal.inria.fr/hal-01916484>
- [17] B. CONFAIS, A. LEBRE, B. PARREIN. *A Tree-Based Approach to locate Object Replicas in a Fog Storage Infrastructure*, in "GLOBECOM 2018 - IEEE Global Communications Conference", Abu Dhabi, United Arab Emirates, IEEE GlobeCom 2018 Conference Proceedings, IEEE, December 2018, pp. 1-6, <https://hal.archives-ouvertes.fr/hal-01946365>
- [18] J. DARROUS, S. IBRAHIM, A. C. ZHOU, C. PÉREZ. *Nitro: Network-Aware Virtual Machine Image Management in Geo-Distributed Clouds*, in "CCGrid 2018 - 18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing", Washington D.C., United States, IEEE, May 2018, pp. 553-562 [DOI : 10.1109/CCGRID.2018.00082], <https://hal.inria.fr/hal-01745405>

- [19] *Best Paper*  
J. PASTOR, J.-M. MENAUD. *SeDuCe: a Testbed for Research on Thermal and Power Management in Datacenters*, in "GREEN 2018 - Third International Conference on Green Communications, Computing and Technologies", Venice, Italy, September 2018, pp. 1-6, <https://hal.inria.fr/hal-01855452>.
- [20] J. PASTOR, J.-M. MENAUD. *SeDuCe: Toward a testbed for research on thermal and power management in datacenters*, in "E2DC 2018 - 7th International Workshop on Energy-Efficient Data Centres", Karlsruhe, Germany, ACM, June 2018, vol. 18, pp. 513-518 [DOI : 10.1145/3208903.3213523], <https://hal.inria.fr/hal-01793521>
- [21] I. RAÏS, H. COULLON, L. LEFÈVRE, C. PÉREZ. *Automatic Energy Efficient HPC Programming: A Case Study*, in "ISPA 2018", Melbourne, Australia, December 2018, <https://hal.inria.fr/hal-01922301>
- [22] J. ROJAS BALDERRAMA, M. SIMONIN. *Scalability and Locality Awareness of Remote Procedure Calls: An Experimental Study in Edge Infrastructures*, in "CloudCom", Nicosia, Cyprus, December 2018, <https://hal.inria.fr/hal-01891567>
- [23] Y. WANG, D. NÖRTERSCHÄUSER, S. LE MASSON, J.-M. MENAUD. *Etude de l'influence des aspects thermiques sur la consommation et l'efficacité énergétique des serveurs*, in "SFT 2018 - 26ème Congrès Français de Thermique", Pau, France, May 2018, pp. 1-8, <https://hal.inria.fr/hal-01763524>
- [24] Y. WANG, D. NÖRTERSCHÄUSER, S. LE MASSON, J.-M. MENAUD. *Potential effects on server power metering and modeling*, in "CloudComp 2018 - 8th EAI International Conference on Cloud Computing", Exeter, United Kingdom, September 2018, pp. 1-12, <https://hal.inria.fr/hal-01869705>
- [25] S. WU, Z. LIU, S. IBRAHIM, L. GU, H. JIN, F. CHEN. *Dual-Paradigm Stream Processing*, in "ICPP 2018 - 47th International Conference on Parallel Processing", Eugene, United States, ACM, August 2018, Article No. 83 [DOI : 10.1145/3225058.3225120], <https://hal.inria.fr/hal-01834668>
- [26] S. WU, M. LIU, S. IBRAHIM, H. JIN, L. GU, F. CHEN, Z. LIU. *TurboStream: Towards Low-Latency Data Stream Processing*, in "ICDCS 2018 - 38th IEEE International Conference on Distributed Computing Systems", Vienna, Austria, IEEE, July 2018, pp. 1-11 [DOI : 10.1109/ICDCS.2018.00099], <https://hal.inria.fr/hal-01807302>
- [27] Y. XIA, X. ETCHEVERS, L. LETONDEUR, A. LEBRE, T. COUPAYE, F. DESPREZ. *Combining Heuristics to Optimize and Scale the Placement of IoT Applications in the Fog*, in "UCC 2018 - 11th IEEE/ACM Conference on Utility and Cloud Computing", Zurich, Switzerland, December 2018, pp. 1-11, <https://hal.inria.fr/hal-01942097>
- [28] O. YILDIZ, S. IBRAHIM. *On the Performance of Spark on HPC Systems: Towards a Complete Picture*, in "SCA 2018 - SupercomputingAsia", Singapore, Singapore, Springer, March 2018, pp. 70-89 [DOI : 10.1007/978-3-319-69953-0\_5], <https://hal.archives-ouvertes.fr/hal-01742016>
- [29] A. C. ZHOU, T.-D. PHAN, S. IBRAHIM, B. HE. *Energy-Efficient Speculative Execution using Advanced Reservation for Heterogeneous Clusters*, in "ICPP 2018 - 47th International Conference on Parallel Process-

ing", Eugene, United States, ACM, August 2018, article n°8 [DOI : 10.1145/3225058.3225084], <https://hal.inria.fr/hal-01807496>

### National Conferences with Proceedings

- [30] Y. WANG, D. NÖRTERSCHÄUSER, S. MASSON, J.-M. MENAUD. *Etude de l'influence de la température du processeur sur la consommation des serveurs*, in "ComPAS 2018 - Conférence d'informatique en Parallélisme, Architecture et Système", Toulouse, France, July 2018, pp. 1-8, <https://hal.inria.fr/hal-01807805>

### Conferences without Proceedings

- [31] E. CADOREL, H. COULLON, J.-M. MENAUD. *Ordonnancement multi-objectifs de workflows dans un Cloud privé*, in "ComPAS 2018 - Conférence d'informatique en Parallélisme, Architecture et Système", Toulouse, France, July 2018, pp. 1-8, <https://hal.archives-ouvertes.fr/hal-01801453>
- [32] M. CHARDET, H. COULLON, C. PÉREZ. *Behavioral interfaces for reconfiguration of component models*, in "Compas 2018 - Conférence d'informatique en Parallélisme, Architecture et Système", Toulouse, France, July 2018, pp. 1-8, <https://hal.inria.fr/hal-01897803>
- [33] B. CONFAIS, A. LEBRE, B. PARREIN. *Improving locality of an object store working in a Fog environment*, in "1st Grid'5000-FIT school", Nice, France, April 2018, pp. 1-2, <https://hal.archives-ouvertes.fr/hal-01759998>
- [34] J. PASTOR, J.-M. MENAUD. *SeDuCe: a Testbed for Research on Thermal and Power Management in Datacenters*, in "SoftCom 2018 - 26th International Conference on Software, Telecommunications and Computer Networks", Split, Croatia, September 2018, pp. 1-6, <https://hal.archives-ouvertes.fr/hal-01855459>

### Research Reports

- [35] M. ABDERRAHIM, M. OUZZIF, K. GUILLOUARD, J. FRANÇOIS, X. LORCA, C. PRUD'HOMME, A. LEBRE. *A Framework for Edge Infrastructures Monitoring*, Orange Labs ; Inria Nancy - Grand Est ; IMT-Atlantique, October 2018, n° RR-9215, pp. 1-14, <https://hal.inria.fr/hal-01897570>
- [36] R.-A. CHERRUEAU, M. SIMONIN, A. VAN KEMPEN. *EnosStack: A LAMP-like stack for the experimenter*, Inria Rennes - Bretagne Atlantique, January 2018, n° RR-9146, pp. 1-14, <https://hal.inria.fr/hal-01689726>
- [37] T. L. NGUYEN, A. LEBRE. *Conducting Thousands of Experiments to Analyze VMs, Dockers and Nested Dockers Boot Time*, Inria, November 2018, n° RR-9221, pp. 1-13, <https://hal.inria.fr/hal-01919354>
- [38] T. L. NGUYEN, R. NOU, A. LEBRE. *YOLO: Speeding up VM Boot Time by reducing I/O operations*, Inria, January 2019, n° RR-9245, pp. 1-18, <https://hal.inria.fr/hal-01983626>

### References in notes

- [39] *Akamai Cloudlets*, 2018, (Accessed: 2018-03-08), <http://cloudlets.akamai.com>
- [40] *Amazon Lambda@Edge*, 2018, (Accessed: 2018-03-08), <https://aws.amazon.com/lambda/edge/>
- [41] C. ATKINSON, T. SCHULZE, S. KLINGERT. *Facilitating Greener IT through Green Specifications*, in "IEEE Software", May 2014, vol. 31, n° 3, pp. 56-63, <http://dx.doi.org/10.1109/MS.2014.19>



- [42] F. BAUDE, D. CAROMEL, C. DALMASSO, M. DANELUTTO, V. GETOV, L. HENRIO, C. PÉREZ. *GCM: a grid extension to Fractal for autonomous distributed components*, in "annals of telecommunications", 2009, vol. 64, n<sup>o</sup> 1-2, pp. 5-24, <http://dx.doi.org/10.1007/s12243-008-0068-8>
- [43] F. BAUDE, L. HENRIO, C. RUZ. *Programming distributed and adaptable autonomous components—the GCM/ProActive framework*, in "Software: Practice and Experience", May 2014, <https://hal.inria.fr/hal-01001043>
- [44] M.-M. BAZM, M. LACOSTE, M. SÜDHOLT, J.-M. MENAUD. *Side-Channels Beyond the Cloud Edge : New Isolation Threats and Solutions*, in "IEEE International Conference on Cyber Security in Networking (CSNet) 2017", Rio de Janeiro, Brazil, October 2017, <https://hal.inria.fr/hal-01593144>
- [45] G. BLAIR, T. COUPAYE, J.-B. STEFANI. *Component-based architecture: the Fractal initiative*, in "Annals of telecommunications", February 2009, vol. 64, n<sup>o</sup> 1, pp. 1–4, <https://doi.org/10.1007/s12243-009-0086-1>
- [46] P. BODIK, R. GRIFFITH, C. SUTTON, A. FOX, M. I. JORDAN, D. A. PATTERSON. *Automatic Exploration of Datacenter Performance Regimes*, in "Proceedings of the 1st Workshop on Automated Control for Datacenters and Clouds", New York, NY, USA, ACDC '09, ACM, 2009, pp. 1–6, <http://doi.acm.org/10.1145/1555271.1555273>
- [47] F. BONOMI, R. MILITO, J. ZHU, S. ADDEPALLI. *Fog computing and its role in the internet of things*, in "Proceedings of the first edition of the MCC workshop on Mobile cloud computing", ACM, 2012, pp. 13–16
- [48] F. BRASILEIRO, G. SILVA, F. ARAÚJO, M. NÓBREGA, I. SILVA, G. ROCHA. *Fogbow: A Middleware for the Federation of IaaS Clouds*, in "The 16th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)", IEEE, 2016, pp. 531–534
- [49] O. BRAČEVAC, S. ERDWEG, G. SALVANESCHI, M. MEZINI. *CPL: A Core Language for Cloud Computing*, in "Proceedings of the 13th International Conference on Modularity (MODULARITY 2016)", ACM Press, 2016, pp. 94–105, <http://doi.acm.org/10.1145/2889443.2889452>
- [50] E. BRUNETON, T. COUPAYE, M. LECLERCQ, V. QUEMA, J.-B. STEFANI. *An Open Component Model and Its Support in Java*, in "Component-Based Software Engineering", Berlin, Heidelberg, I. CRNKOVIC, J. A. STAFFORD, H. W. SCHMIDT, K. WALLNAU (editors), Springer Berlin Heidelberg, 2004, pp. 7–22
- [51] P. CARBONE, A. KATSIFODIMOS, S. EWEN, V. MARKL, S. HARIDI, K. TZOUMAS. *Apache Flink?: Stream and Batch Processing in a Single Engine*, in "IEEE Data Eng. Bull.", 2015, vol. 38, pp. 28-38
- [52] V. CARDELLINI, F. L. PRESTI, M. NARDELLI, G. R. RUSSO. *Towards Hierarchical Autonomous Control for Elastic Data Stream Processing in the Fog*, in "Euro-Par 2017: Parallel Processing Workshops: Euro-Par 2017 International Workshops, Santiago de Compostela, Spain, August 28-29, 2017, Revised Selected Papers", Springer, 2018, vol. 10659, pp. 106-117
- [53] F.-C. CHANG, H.-C. HUANG, H.-M. HANG. *Combined Encryption and Watermarking Approaches for Scalable Multimedia Coding*, in "Advances in Multimedia Information Processing - PCM 2004: 5th Pacific Rim Conference on Multimedia, Tokyo, Japan, November 30 - December 3, 2004. Proceedings, Part III", K. AIZAWA, Y. NAKAMURA, S. SATOH (editors), Springer Berlin Heidelberg, Berlin, Heidelberg, 2005, pp. 356–363



- [54] R.-A. CHERRUEAU, R. DOUENCE, M. SÜDHOLT. *A Language for the Composition of Privacy-Enforcement Techniques*, in "IEEE RATSP 2015, The 2015 IEEE International Symposium on Recent Advances of Trust, Security and Privacy in Computing and Communications ", Helsinki, Finland, August 2015, <https://hal.inria.fr/hal-01168303>
- [55] H. E. CHIHOUB, S. IBRAHIM, Y. LI, G. ANTONIU, M. S. PEREZ, L. BOUGE. *Exploring Energy-Consistency Trade-Offs in Cassandra Cloud Storage System*, in "27th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)", October 2015, pp. 146-153
- [56] B. CONFAIS, A. LEBRE, B. PARREIN. *An Object Store Service for a Fog/Edge Computing Infrastructure based on IPFS and Scale-out NAS*, in "1st IEEE International Conference on Fog and Edge Computing-ICFEC'2017", 2017
- [57] J. C. CORBETT, J. DEAN, M. EPSTEIN, A. FIKES, C. FROST, J. J. FURMAN, S. GHEMAWAT, A. GUBAREV, C. HEISER, P. HOCHSCHILD, ET AL. . *Spanner: Google's globally distributed database*, in "ACM Transactions on Computer Systems (TOCS)", 2013, vol. 31, n<sup>o</sup> 3, 8 p.
- [58] H. COULLON, J. BIGOT, C. PÉREZ. *Extensibility and Composability of a Multi-Stencil Domain Specific Framework*, in "International Journal of Parallel Programming", November 2017, <https://doi.org/10.1007/s10766-017-0539-5>
- [59] H. COULLON, G. LE LOUËT, J.-M. MENAUD. *Virtual Machine Placement for Hybrid Cloud using Constraint Programming*, in "ICPADS 2017", 2017
- [60] H. COULLON, D. PERTIN, C. PÉREZ. *Production Deployment Tools for IaaS: an Overall Model and Survey*, in "IEEE International Conference on Future Internet of Things and Cloud (FiCloud) 2017", Prague, Czech Republic, August 2017, <https://hal.inria.fr/hal-01532489>
- [61] I. CUADRADO-CORDERO, A.-C. ORGERIE, J.-M. MENAUD. *Comparative Experimental Analysis of the Quality-of-Service and Energy-Efficiency of VMs and Containers' Consolidation for Cloud Applications*, in "SoftCOM: International Conference on Software, Telecommunications and Computer Networks", 2017
- [62] P.-C. DAVID, T. LEDOUX, M. LÉGER, T. COUPAYE. *FPath and FScript: Language support for navigation and reliable reconfiguration of Fractal architectures*, in "annals of telecommunications - annales des télécommunications", February 2009, vol. 64, n<sup>o</sup> 1, pp. 45–63, <https://doi.org/10.1007/s12243-008-0073-y>
- [63] F. A. DE OLIVEIRA, T. LEDOUX, R. SHARROCK. *A framework for the coordination of multiple autonomic managers in cloud environments*, in "Self-Adaptive and Self-Organizing Systems (SASO), 2013 IEEE 7th International Conference on", IEEE, 2013, pp. 179–188
- [64] R. DI COSMO, J. MAURO, S. ZACCHIROLI, G. ZAVATTARO. *Aeolus: A component model for the cloud*, in "Information and Computation", 2014, vol. 239, n<sup>o</sup> Supplement C, pp. 100 - 121 [DOI : 10.1016/J.IC.2014.11.002], <http://www.sciencedirect.com/science/article/pii/S0890540114001424>
- [65] W. FELTER, A. FERREIRA, R. RAJAMONY, J. RUBIO. *An updated performance comparison of virtual machines and linux containers*, in "Performance Analysis of Systems and Software (ISPASS), 2015 IEEE International Symposium On", IEEE, 2015, pp. 171–172

- [66] A. FLISSI, J. DUBUS, N. DOLET, P. MERLE. *Deploying on the Grid with DeployWare*, in "Eighth IEEE International Symposium on Cluster Computing and the Grid", France, May 2008, pp. 177-184 [DOI : 10.1109/CCGRID.2008.59], <https://hal.archives-ouvertes.fr/hal-00259836>
- [67] P. GARCIA LOPEZ, A. MONTRESOR, D. EPEMA, A. DATTA, T. HIGASHINO, A. IAMNITCHI, M. BARCELLOS, P. FELBER, E. RIVIERE. *Edge-centric Computing: Vision and Challenges*, in "SIGCOMM Comput. Commun. Rev.", September 2015, vol. 45, n<sup>o</sup> 5, pp. 37-42, <http://doi.acm.org/10.1145/2831347.2831354>
- [68] Í. GOIRI, W. KATSAK, K. LE, T. D. NGUYEN, R. BIANCHINI. *Parasol and GreenSwitch: Managing Datacenters Powered by Renewable Energy*, in "SIGARCH Comput. Archit. News", March 2013, vol. 41, n<sup>o</sup> 1, pp. 51-64
- [69] Í. GOIRI, K. LE, T. D. NGUYEN, J. GUITART, J. TORRES, R. BIANCHINI. *GreenHadoop: Leveraging Green Energy in Data-processing Frameworks*, in "Proceedings of the 7th ACM European Conference on Computer Systems", New York, NY, USA, EuroSys '12, ACM, 2012, pp. 57-70, <http://doi.acm.org/10.1145/2168836.2168843>
- [70] H. GOUDARZI, M. PEDRAM. *Geographical Load Balancing for Online Service Applications in Distributed Datacenters*, in "in IEEE international conference on cloud computing (CLOUD 2013)", 2013
- [71] L. GU, D. ZENG, A. BARNAWI, S. GUO, I. STOJMENOVIC. *Optimal Task Placement with QoS Constraints in Geo-Distributed Data Centers Using DVFS*, in "IEEE Transactions on Computers", July 2015, vol. 64, n<sup>o</sup> 7, pp. 2049-2059, <http://dx.doi.org/10.1109/TC.2014.2349510>
- [72] K. HA, Y. ABE, T. EISZLER, Z. CHEN, W. HU, B. AMOS, R. UPADHYAYA, P. PILLAI, M. SATYANARAYANAN. *You Can Teach Elephants to Dance: Agile VM Handoff for Edge Computing*, in "Proceedings of the Second ACM/IEEE Symposium on Edge Computing", New York, NY, USA, SEC '17, ACM, 2017, pp. 12:1-12:14, <http://doi.acm.org/10.1145/3132211.3134453>
- [73] B. HAN, V. GOPALAKRISHNAN, L. JI, S. LEE. *Network function virtualization: Challenges and opportunities for innovations*, in "IEEE Communications Magazine", February 2015, vol. 53, n<sup>o</sup> 2, pp. 90-97, <http://dx.doi.org/10.1109/MCOM.2015.7045396>
- [74] M. S. HASAN, F. ALVARES, T. LEDOUX, J. L. PAZAT. *Investigating Energy Consumption and Performance Trade-Off for Interactive Cloud Application*, in "IEEE Transactions on Sustainable Computing", April 2017, vol. 2, n<sup>o</sup> 2, pp. 113-126, <http://dx.doi.org/10.1109/TSUSC.2017.2714959>
- [75] F. HERMENIER, J. LAWALL, G. MULLER. *Btrplace: A flexible consolidation manager for highly available applications*, in "IEEE Transactions on dependable and Secure Computing", 2013, vol. 10, n<sup>o</sup> 5, pp. 273-286
- [76] F. HERMENIER, A. LEBRE, J.-M. MENAUD. *Cluster-wide Context Switch of Virtualized Jobs*, in "Proceedings of the Virtualization Technologies in Distributed Computing Workshop (co-located with ACM HPDC'10)", New York, NY, USA, HPDC '10, ACM, 2010, pp. 658-666, <http://doi.acm.org/10.1145/1851476.1851574>
- [77] F. HERMENIER, X. LORCA, J.-M. MENAUD, G. MULLER, 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

- [78] M. HIRZEL, S. SCHNEIDER, B. GEDIK. *SPL: An Extensible Language for Distributed Stream Processing*, in "toplas", March 2017, vol. 39, n<sup>o</sup> 1, pp. 5:1–5:39, <http://doi.acm.org/10.1145/3039207>
- [79] S. IBRAHIM, B. HE, H. JIN. *Towards Pay-As-You-Consume Cloud Computing*, in "2011 IEEE International Conference on Services Computing", July 2011, pp. 370-377, <http://dx.doi.org/10.1109/SCC.2011.38>
- [80] S. IFTIKHAR, S. KHAN, Z. ANWAR, ET AL. . *GenInfoGuard: A Robust and Distortion-Free Watermarking Technique for Genetic Data*, in "PLOS ONE", February 2015, vol. 10, n<sup>o</sup> 2, pp. 1-22, <https://doi.org/10.1371/journal.pone.0117717>
- [81] B. JENNINGS, R. STADLER. *Resource Management in Clouds: Survey and Research Challenges*, in "Journal of Network and Systems Management", July 2015, vol. 23, n<sup>o</sup> 3, pp. 567–619, <https://doi.org/10.1007/s10922-014-9307-7>
- [82] E. JONARDI, M. A. OXLEY, S. PASRICHA, A. A. MACIEJEWSKI, H. J. SIEGEL. *Energy cost optimization for geographically distributed heterogeneous data centers*, in "2015 Sixth International Green and Sustainable Computing Conference (IGSC)", December 2015, pp. 1-6, <http://dx.doi.org/10.1109/IGCC.2015.7393677>
- [83] H. KARAU, A. KONWINSKI, P. WENDELL, M. ZAHARIA. *Learning Spark*, O'Reilly Media, February 2015
- [84] J. KRAMER, J. MAGEE. *The evolving philosophers problem: dynamic change management*, in "IEEE Transactions on Software Engineering", November 1990, vol. 16, n<sup>o</sup> 11, pp. 1293-1306, <http://dx.doi.org/10.1109/32.60317>
- [85] A. LEBRE, J. PASTOR, A. SIMONET, F. DESPREZ. *Revising OpenStack to Operate Fog/Edge Computing Infrastructures*, in "The IEEE International Conference on Cloud Engineering (IC2E)", April 2017, pp. 138-148, <http://dx.doi.org/10.1109/IC2E.2017.35>
- [86] G. MALEWICZ, M. H. AUSTERN, A. J. BIK, J. C. DEHNERT, I. HORN, N. LEISER, G. CZAJKOWSKI. *Pregel: A System for Large-scale Graph Processing*, in "Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data", New York, NY, USA, SIGMOD '10, ACM, 2010, pp. 135–146, <http://doi.acm.org/10.1145/1807167.1807184>
- [87] T. L. NGUYEN, A. LEBRE. *Virtual Machine Boot Time Model*, in "Parallel, Distributed and Network-based Processing (PDP), 2017 25th Euromicro International Conference on", IEEE, 2017, pp. 430–437
- [88] K. OKAMURA, Y. OYAMA. *Load-based covert channels between Xen virtual machines*, in "Proceedings of the 2010 ACM Symposium on Applied Computing", ACM, 2010, pp. 173–180
- [89] P. OREIZY, N. MEDVIDOVIC, R. N. TAYLOR. *Runtime Software Adaptation: Framework, Approaches, and Styles*, in "Companion of the 30th International Conference on Software Engineering", New York, NY, USA, ICSE Companion '08, ACM, 2008, pp. 899–910, <http://doi.acm.org/10.1145/1370175.1370181>
- [90] T. D. PHAN, S. IBRAHIM, G. ANTONIU, L. BOUGE. *On Understanding the Energy Impact of Speculative Execution in Hadoop*, in "2015 IEEE International Conference on Data Science and Data Intensive Systems", December 2015, pp. 396-403

- [91] F. QUESNEL, A. LEBRE, M. SÜDHOLT. *Cooperative and reactive scheduling in large-scale virtualized platforms with DVMS*, in "Concurrency and Computation: Practice and Experience", 2013, vol. 25, n<sup>o</sup> 12, pp. 1643–1655
- [92] D. SABELLA, A. VAILLANT, P. KUURE, U. RAUSCHENBACH, F. GIUST. *Mobile-Edge Computing Architecture: The role of MEC in the Internet of Things*, in "IEEE Consumer Electronics Magazine", October 2016, vol. 5, n<sup>o</sup> 4, pp. 84-91, <http://dx.doi.org/10.1109/MCE.2016.2590118>
- [93] D. SERRANO, S. BOUCHENAK, Y. KOUKI, F. ALVARES, T. LEDOUX, J. LEJEUNE, J. SOPENA, L. ARANTES, P. SENS. *SLA guarantees for cloud services*, in "Future Generation Computer Systems", 2016, vol. 54, n<sup>o</sup> Supplement C, pp. 233 - 246 [DOI: 10.1016/j.future.2015.03.018], <http://www.sciencedirect.com/science/article/pii/S0167739X15000801>
- [94] Q. SHEN, X. LIANG, X. S. SHEN, X. LIN, H. Y. LUO. *Exploiting Geo-Distributed Clouds for a E-Health Monitoring System With Minimum Service Delay and Privacy Preservation*, in "IEEE Journal of Biomedical and Health Informatics", March 2014, vol. 18, n<sup>o</sup> 2, pp. 430-439, <http://dx.doi.org/10.1109/JBHI.2013.2292829>
- [95] W. SHI, T. F. WENISCH. *Energy-Efficient Data Centers*, in "IEEE Internet Computing", 2017, vol. 21, n<sup>o</sup> 4, pp. 6-7, <http://dx.doi.org/10.1109/MIC.2017.2911429>
- [96] J. STEFANI. *Components as Location Graphs*, in "Formal Aspects of Component Software - 11th International Symposium, FACS 2014, Bertinoro, Italy, September 10-12, 2014, Revised Selected Papers", 2014, pp. 3–23, [https://doi.org/10.1007/978-3-319-15317-9\\_1](https://doi.org/10.1007/978-3-319-15317-9_1)
- [97] C. SZYPERSKI. *Component Software: Beyond Object-Oriented Programming*, 2nd, Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2002
- [98] Y. TALEB, S. IBRAHIM, G. ANTONIU, T. CORTES. *Characterizing Performance and Energy-Efficiency of the RAMCloud Storage System*, in "2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)", June 2017, pp. 1488-1498, <http://dx.doi.org/10.1109/ICDCS.2017.51>
- [99] THE LINUX FOUNDATION. *Open Network Automation Platform*, 2018, (Accessed: 2018-03-08), <https://www.onap.org/>
- [100] THE OPENSTACK FOUNDATION. *Cloud Edge Computing: Beyond the Data Center (White Paper)*, January 2018, (Accessed: 2018-03-08), <https://www.openstack.org/assets/edge/OpenStack-EdgeWhitepaper-v3-online.pdf>
- [101] A. THUSOO, Z. SHAO, S. ANTHONY, D. BORTHAKUR, N. JAIN, J. SEN SARMA, R. MURTHY, H. LIU. *Data Warehousing and Analytics Infrastructure at Facebook*, in "Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data", SIGMOD '10, ACM Press, 2010, pp. 1013–1020, <http://doi.acm.org/10.1145/1807167.1807278>
- [102] T. WHITE. *Hadoop: The Definitive Guide*, 4th, O'Reilly Media, April 2015
- [103] A. WIERMAN, Z. LIU, I. LIU, H. MOHSENIAN-RAD. *Opportunities and challenges for data center demand response*, in "International Green Computing Conference", November 2014, pp. 1-10, <http://dx.doi.org/10.1109/IGCC.2014.7039172>

- 
- [104] J. XIAO, Z. XU, H. HUANG, H. WANG. *Security implications of memory deduplication in a virtualized environment*, in "2013 43rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)", IEEE, 2013, pp. 1–12
- [105] F. XU, F. LIU, H. JIN, A. V. VASILAKOS. *Managing performance overhead of virtual machines in cloud computing: A survey, state of the art, and future directions*, in "Proceedings of the IEEE", 2014, vol. 102, n<sup>o</sup> 1, pp. 11–31
- [106] O. YILDIZ, M. DORIER, S. IBRAHIM, R. ROSS, G. ANTONIU. *On the Root Causes of Cross-Application I/O Interference in HPC Storage Systems*, in "2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS)", May 2016, pp. 750-759, <http://dx.doi.org/10.1109/IPDPS.2016.50>
- [107] O. YILDIZ, A. C. ZHOU, S. IBRAHIM. *Eley: On the Effectiveness of Burst Buffers for Big Data Processing in HPC Systems*, in "2017 IEEE International Conference on Cluster Computing (CLUSTER)", September 2017, pp. 87-91, <http://dx.doi.org/10.1109/CLUSTER.2017.73>
- [108] M. ZAHARIA, T. DAS, H. LI, T. HUNTER, S. SHENKER, I. STOICA. *Discretized Streams: Fault-tolerant Streaming Computation at Scale*, in "Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles", New York, NY, USA, SOSP '13, ACM, 2013, pp. 423–438, <http://doi.acm.org/10.1145/2517349.2522737>