

RESEARCH CENTRE

Rennes - Bretagne Atlantique

IN PARTNERSHIP WITH:

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

2020

ACTIVITY REPORT

Project-Team

STACK

Software Stack for Massively Geo-Distributed Infrastructures

IN COLLABORATION WITH: Laboratoire des Sciences du numérique de
Nantes

DOMAIN

Networks, Systems and Services,
Distributed Computing

THEME

Distributed Systems and middleware

Contents

Project-Team STACK	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	3
2.1 STACK in a Nutshell	3
2.2 Toward a STACK for Geo-Distributed Infrastructures	3
3 Research program	4
3.1 Overview	4
3.2 Resource Management	4
3.3 Programming Support	6
3.4 Energy	8
3.5 Security	10
4 Application domains	11
4.1 Overview	11
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
5 Social and environmental responsibility	12
5.1 Footprint of research activities	12
5.2 Impact of research results	12
6 Highlights of the year	13
6.1 Awards	13
7 New software and platforms	13
7.1 New software	13
7.1.1 MAD	13
7.1.2 Nitro	14
7.1.3 VMPlaceS	14
7.1.4 ENOS	15
7.1.5 EnOSlib	15
7.1.6 Concerto	16
7.2 New platforms	16
7.2.1 OpenStack	16
7.2.2 Grid'5000	16
7.2.3 SeDuCe	16
7.2.4 SILECS/SLICES	17
8 New results	17
8.1 Resource Management	17
8.2 Programming Support	19
8.3 Energy-aware computing	20
8.4 Security and Privacy	21
9 Bilateral contracts and grants with industry	21
9.1 Bilateral contracts with industry	21

10 Partnerships and cooperations	22
10.1 International initiatives	22
10.1.1 Inria International Labs	22
10.2 International research visitors	22
10.2.1 Visits of international scientists	22
10.3 European initiatives	22
10.3.1 FP7 & H2020 Projects	22
10.4 National initiatives	23
10.4.1 Ademe	23
10.4.2 CominLabs laboratory of excellence	23
10.4.3 CNRS	23
10.4.4 ANR	23
10.4.5 FSN	24
10.4.6 CPER	25
10.4.7 Connect Talent	25
10.4.8 Etoiles Montantes	25
10.5 Regional initiatives	25
11 Dissemination	26
11.1 Promoting scientific activities	26
11.1.1 Scientific events: organisation	26
11.1.2 Scientific events: selection	27
11.1.3 Journal	27
11.1.4 Invited talks	27
11.1.5 Scientific expertise	27
11.1.6 Research administration	28
11.1.7 Member of the Selection Committee for Associate Professors	28
11.2 Teaching - Supervision - Juries	28
11.2.1 Teaching	28
11.2.2 Supervision	28
11.2.3 Juries	29
11.3 Popularization	30
11.3.1 Internal or external Inria responsibilities	30
11.3.2 Articles and contents	30
11.3.3 Education	30
12 Scientific production	30
12.1 Major publications	30
12.2 Publications of the year	30
12.3 Cited publications	33

Project-Team STACK

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

Keywords

Computer sciences and digital sciences

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

Other research topics and application domains

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

1 Team members, visitors, external collaborators

Research Scientist

- Shadi Ibrahim [Inria, Researcher, until Sep 2020]

Faculty Members

- Adrien Lebre [Team leader, IMT Atlantique, Professor, HDR]
- H el ene Coullon [IMT Atlantique, Chair]
- Remous Aris Koutsiamanis [IMT Atlantique, Associate Professor, from Sep 2020]
- Thomas Ledoux [IMT Atlantique, Associate Professor, Full Professor from Oct. 2020, HDR]
- Jean-Marc Menaud [IMT Atlantique, Professor, HDR]
- Jacques Noy e [IMT Atlantique, Associate Professor]
- Mario S udholt [IMT Atlantique, Professor, HDR]

Post-Doctoral Fellows

- David Guyon [Inria, until Sep 2020]
- Thomas Lambert [Inria, until Sep 2020]
- Jonathan Pastor [IMT Atlantique, until Jul 2020]
- Simon Robillard [IMT Atlantique]
- Hamza Sahli [Inria, until Apr 2020]
- Alexandre Van Kempen [IMT Atlantique, from Mar 2020 until Nov 2020]

PhD Students

- Maxime Belair [Orange Labs, CIFRE]
- Fatima-Zahra Boujdad [IMT Atlantique, until Mar 2020]
- Emile Cadorel [Armines, until Sep 2020]
- Maverick Chardet [Inria, until Sep 2020]
- Marie Delavergne [Inria]
- David Espinel [Orange Labs, CIFRE]
- Wilmer Edicson Garzon Alfonso [IMT Atlantique]
- Karim Manaouil [Inria, from Oct 2020 until Dec 2020]
- Jolan Philippe [IMT Atlantique]
- Dimitri Saingre [Armines]
- Sirine Sayadi [IMT Atlantique]

Technical Staff

- Ronan-Alexandre Cherrueau [Inria, Engineer]
- Jad Darrous [Inria, Engineer, until Aug 2020]
- Brice Nedelec [Sigma, Engineer]
- Rémy Pottier [IMT Atlantique, Engineer, until Oct 2020]
- Charlene Servantie [IMT Atlantique, Engineer, until Jun 2020]
- Matthieu Simonin [Inria, Engineer]

Interns and Apprentices

- Karim Manaouil [Inria, until Jul 2020]
- Nathan Rossard [Inria, from Mar 2020 until Aug 2020]

Administrative Assistant

- Anne-Claire Binétruy [Inria]

Visiting Scientist

- Twinkle Jain [Université Northeastern Boston - USA, from Feb 2020 until Aug 2020]

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 infrastructures, from the lowest (system) levels to the highest (application development) ones, and addressing cross-cutting dimensions such as energy or security.¹ 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 the 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 [99]. In other words, self-management decisions taken in isolation at one level (infrastructure, middleware, or application) may indirectly interfere with the decisions 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 [60, 93], we can expect that the complexity of the

¹Following the recruitment of Remous-Aris Koutsiamanis in October 2020, the team broadened its activities to the IoT area, covering now the continuum of Cloud→IoT devices under the resource management perspective.

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 [85]. 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 the 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 in 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, and IoT devices). They are transverse to the System/Middleware/Application layers, which generally compose 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 first two 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 [47]. 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 [66]. 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 9.1 and 10). 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.

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 [99]. 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 counter-measures. In particular, previous studies [99, 58] on pros and cons of current technologies – virtual machines (VMs) [72,

80], 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 [100]. 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.

Geo-Distributed System Mechanisms Although several studies have been highlighting the advantages of geo-distributed ICT infrastructures in various domains (see Section 4), progress on how to operate and use such infrastructures is marginal. Current solutions [29] [30] are rather close to the initial Cisco Fog Computing proposal that only allows running domain-specific applications on edge resources and centralized Cloud platforms [37] (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 [85]) have been proposed. ONAP [92], an industry-driven solution, enables the orchestration and automation of virtual network functions across distinct VIMs. From the academic side, FogBow [40] 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 vocabulary². 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 deliver means to administer the physical resources.

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 [65, 78].

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 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 [50] 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 [49, 91], and burst buffers [101]. We also plan to conduct new investigations on data-stream frameworks for Fog and Edge infrastructures [44]. 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.

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 [74]. STACK will inherit the knowledge acquired through previous activities in this domain, particularly its use of constraint programming strategies in autonomic managers [70, 69], relying on MAPE (monitor, analyze, plan, and execute) control loops.

²https://www.nist.gov/system/files/documents/2019/07/09/nist_cfra_20190709_draft_v1.0.pdf (Dec 2018).

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 [68]. 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 [84].

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 [80], 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 [57], 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.

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 [79], Hive [95], Hadoop [96], some of them are more general-purpose through the provision of several programming models, *e.g.*, Flink [43] and Spark [76]. Finally, some dedicated language

support has been considered for some models (*e.g.*, the language SPL underlying IBM Streams [71]) as well as core languages and calculi (*e.g.*, [39, 89]).

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 [52]. For example, providing a generic stream data processing model that can operate under both data stream [43] and operation stream [102] 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 [32, 90]. 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 [71]. 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.

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 [77, 82]. Many deployment tools have been proposed both in academia and industry [54]. For example, Ansible³, Chef⁴ and Puppet⁵ 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^{6,7}. 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 [90]. Thus, modeling the reconfiguration process consists in describing the life cycle of different components and their interactions. Most component-based approaches offer a fixed

³<https://www.ansible.com/>

⁴<https://www.chef.io/chef/>

⁵<https://puppet.com/>

⁶<https://kubernetes.io/>

⁷<https://jujucharms.com/>

life cycle, *i.e.*, identical for any component [59]. Two main contributions are able to customize life cycles, Fractal [42, 35] and its evolutions [32, 33, 56], and Aeolus [51]. 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 [86]. 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 [41]. 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 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 [67] 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 [88]. 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 [61]. Workload distribution for geo-distributed DCs is also another promising approach [63, 75, 97]. 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, etc.). 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) [55] and data (hard drives, SSD, in-memory storage, distributed file systems). For computations, in the continuity of our previous work [53, 70], 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 [48, 91]. 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 7.2), with the ultimate goal of reducing the CO₂ 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 [83].

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 [31]), for example to describe the maximum allowed CO₂ emissions of a Fog/Edge service. Unlike other approaches [64, 36, 62] 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 [67] 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 [87]. 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 [81, 98]. 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 [73, 45]. We provided preliminary results on such compositions [46] and have extended them to biomedical, notably genetic, analyses in the e-health domain [38]. 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 [34].

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 project⁸, see Section 10.

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-

⁸Privacy-preserving sharing and processing of genetic data, <https://privgen.cominlabs.u-bretagne-normandie.fr/fr>

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 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, etc.); the development of large "smart" applications handling large amounts of geo-distributed data and a large set of resources (video analytics, augmented reality, etc.); the advent of future medical prevention and treatment techniques thanks to the intensive use of ICT systems, etc. We expect our contributions 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 (aka. precision) medicine. The overall goal in these fields is the development of medication and treatment methods that are tailored towards small groups or even individual patients.

We 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. More generally, we are applying and developing corresponding techniques for the medical domains of genomics, immunobiology and transplantology in the international network SHLARC and the regional networks SysMics and Oncoshare (see Section 10): there, we investigate how to secure and preserve privacy if potentially sensitive personal data is moved and processed by distributed biomedical analyses.

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 Institute⁹. 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 Social and environmental responsibility

5.1 Footprint of research activities

In addition to the international travels¹⁰, the environmental footprint of our research activities is linked to our intensive use of large-scale testbeds such as Grid'5000 (STACK members are often in the top10 list of the largest consumers). Although the access to such facilities is critical to move forward in our research roadmap, it is important to recognize that they have a strong environmental impact as described in the next paragraph.

5.2 Impact of research results

The environmental impact of digital technology is a major scientific and societal challenge. Even though the software remains virtual objects, it is executed on very real hardware contributing to the carbon footprint. This impact materializes during the manufacture / destruction of hardware infrastructure (estimated at 45% of digital consumption in 2018 by the The Shift Project) and during the software use phase via terminals, networks and data centers. data (estimated at 55%). Stack members have been studying various approaches for several years to reduce the energy footprint of digital infrastructures during the use phase. The work carried out revolves around two main axes: (i) reducing the energy footprint of infrastructures and (ii) adapting the software applications hosted by these infrastructures according to the energy available. More precisely, this second axe investigates possible improvements

⁹<http://www.etsi.org/news-events/news/1078-2016-04-etsi-mobile-edge-computing-publishes-foundation-specifications>.

¹⁰Due to the COVID pandemia, STACK members did not travel this year overall.

that could be made by the end-users of the software themselves. At scale, involving end-users in decision-making processes concerning energy consumption would lead to more frugal Cloud computing. For instance, in the GL4MA project (cf. Section 10), we propose that the end-users customize the software in SaaS mode to contribute to reducing their carbon footprint by using either fewer resources or resources powered directly by renewable energy.

6 Highlights of the year

Regarding scientific results, the team has produced a number of outstanding results on the management of resources and data in large-scale infrastructures. In particular, the team published a survey in the ACM Computing Journal on the service placement problem in Fog and Edge Computing [1]. The team obtained also three publications in the 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGrid 2020), a major conference in the area of Cloud Computing systems [4, 2, 3].

On the software side, the team has pursued its efforts on the development of the EnosLib library and the resulting artifacts to help researchers perform experiment campaigns: <https://discovery.gitlabpages.inria.fr/enoslib/theyuseit.html>. In particular, we would like to point out the recent extensions toward the Kubernetes framework.

On the platform side, we continued our effort and took part to the different actions around the SILECS initiative¹¹ (and its European part, SLICES¹²)

6.1 Awards

In 2020, the team has received one individual award:

- Adrien Lebre received the IEEE TCC Editorial Excellence and Eminence Award to recognize his exceptional contributions to the TCC Editorial Board.

We would like also to highlight two elements that underline the visibility and recognition of the team nationally and internationally. First, at the national level, the team has obtained two grants from the *Appel ANR générique 2020* for the SeMaFoR and PICNIC projects. The former's objective is to model, design and develop a generic and decentralised solution for the self-management of Fog resources using a Fog Architecture Description Language, a collaborative/consensual decision-making process and an automatic reconfiguration coordination mechanism. The latter addresses the problem of transferring large data-sets between two geographically distant sites. More precisely, the objective is to develop network mechanisms in the linux kernel to sending in parallel a file by breaking the VM bottleneck by using the property that the file is (all or part) duplicated on several disks.

Second, at the international level, Helene Coullon obtained a 20% Adjunct Professor position at the Arctic University of Norway in Tromsø (UiT)¹³. This position has been offered to Helene after an invited talk (at the end of 2019), and discussions on collaborations with the team involved in the DAO project¹⁴ led by Otto Anshus. The position has started in september 2020 and will end in august 2022. Helene has the responsibility to give a few courses to the UiT students (Masters and PhD) in advanced distributed systems. Furthermore, this position facilitates interactions for research collaborations on the DAO project in which Helene brings an expertise in dynamic software configuration and reconfiguration.

7 New software and platforms

7.1 New software

7.1.1 MAD

Name: Madeus Application Deployer

¹¹<https://www.silecs.net>

¹²<http://slices-ri.eu>

¹³<https://en.uit.no/startside>

¹⁴<https://site.uit.no/dao/>

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.

Functional 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 Contributions: Initial submission with basic functionalities of MAD

News of the Year: Operational prototype.

Publications: [hal-01858150](#), [hal-01897803](#)

Authors: Maverick Chardet, H el ene Coullon, Christian P erez, Dimitri Pertin

Contacts: Maverick Chardet, Christian P erez, Dimitri Pertin, H el ene Coullon

Participants: Christian P erez, Dimitri Pertin, H el ene Coullon, Maverick Chardet

Partners: IMT Atlantique, LS2N, LIP

7.1.2 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).

URL: <https://gitlab.inria.fr/jdarrous/nitro>

Authors: Jad Darrous, Shadi Ibrahim, Christian P erez

Contacts: Jad Darrous, Shadi Ibrahim, Christian P erez

7.1.3 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.

URL: <http://beyondtheclouds.github.io/VMPlaceS/>

Contact: Adrien Lèbre

Participants: Adrien Lèbre, Jonathan Pastor, Mario Südholt

7.1.4 ENOS

Name: Experimental eNvironment for OpenStack

Keywords: OpenStack, Experimentation, Reproducibility

Functional Description: Enos workflow :

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

URL: <http://enos.readthedocs.io/en/stable/>

Contacts: Adrien Lèbre, Matthieu Simonin

Partner: Orange Labs

7.1.5 EnOSlib

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

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

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

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

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

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

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

Contact: Matthieu Simonin

7.1.6 Concerto

Name: Concerto

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

Functional Description: Concerto is an implementation of the formal model Concerto written in Python. Concerto allows to : 1. describe the life-cycle and the dependencies of software components, 2. describe a components assembly that forms the overall life-cycle of a distributed software, 3. automatically reconfigure a Concerto assembly of components by using a set of reconfiguration instructions as well as a formal operational semantics.

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

Contacts: Maverick Chardet, H el ene Coullon, Christian P erez

Partners: IMT Atlantique, LS2N, LIP

7.2 New platforms

7.2.1 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 Group¹⁵ 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 [93], the major result from the academic view point is the aforementioned EnOS solution, a holistic framework to conduct performance evaluations of OpenStack (control and data plane). In May 2018, the FEMDC SiG turned into a larger group under the control of the OpenStack foundation. This group gathers large companies such as Verizon, ATT, etc. Although our involvement has been less important in 2020, our participation is still significant. For instance, we co-signed the second white paper delivered by the edge working group in 2020 [94].

7.2.2 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. STACK members are strongly involved into the management and the supervision of the testbed, notably through the steering committee or the SeDuCe testbed described hereafter.

7.2.3 SeDuCe

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 research to cover a full

¹⁵https://wiki.openstack.org/wiki/Fog_Edge_Massively_Distributed_Clouds

spectrum of the energetical aspect of datacenters, such as cooling and power consumption depending of experimental conditions.

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

In 2020, we developed a first step for a deployment and reservation system for the Edge Computing (Raspberry Pi Cluster). In November 2020, we also began the study of a device to cloud deployment system, funded by the CNRS, in the Kabuto project, in connection with the SILECS project.

7.2.4 SILECS/SLICES

STACK Members are involved in the definition and bootstrap of the SILECS infrastructure. This infrastructure can be seen as a merge of the Grid’5000 and FIT testbeds with the goal of providing a common platform for experimental computer Science (Next Generation Internet, Internet of things, clouds, HPC, big data, etc.). In 2020, STACK members took part to the PIA3 SILECS proposal submitted to the national PIA-3 call ¹⁶ and the submission to the European EFRI SLICES proposal ¹⁷.

8 New results

8.1 Resource Management

Participants Adwait Jitendra Bauskar, Ronan-Alexandre Cherrueau, Bastien Confais, Marie Delavergne, David Espinel, David Guyon, Shadi Ibrahim, Remous Aris Koutsiamanis, Thomas Lambert, Adrien Lebre, Karim Manaouil, Javier Rojas Balderrama, Matthieu Simonin, Alexandre Van Kempen.

In 2020, we had multiple contributions in the field of resource management of geo-distributed cloud infrastructures, with a particular focus on Fog and Edge computing. First, we present contributions regarding the resource placement problem, for both traditional and data streaming applications. We then address issues regarding the network complications in Fog/Edge computing regarding the network itself and storage. To improve reproducibility and experimentation in resource management we also introduce a new simulation tool for edge computing use-cases. Following our multi-year work on this subject we then re-evaluate past suggested best practices and the readiness of modern resource management systems (Kubernetes) for the edge-computing case. Finally, we also start addressing the IoT part of the Cloud ↔ IoT continuum by presenting a scheduler for wireless IoT network resources.

In [5], we presented a survey of current research conducted on Service Placement Problem (SPP) in the Fog/Edge Computing. To support the large and various applications generated by the Internet of Things (IoT), Fog Computing has been introduced to complement the Cloud Computing and offer Cloud-like services at the edge of the network with low latency and real-time responses. Large-scale, geographical distribution and heterogeneity of edge computational nodes make service placement in such infrastructure a challenging issue. Diversity of user expectations and IoT devices characteristics also complexify the deployment problem. Based on a new classification scheme, we analysed current proposals and discussed issues and challenges our community should deal with.

In [21], we propose a model to evaluate Maximum Sustainable Throughput (MST) for operators placement in the Edge, with a strong focus on the heterogeneous nature of network. We then introduce an optimal operators placement for stream data applications in the Edge using constraint programming. Through simulations, we show how existing placement strategies that target overall communications reduction often fail to keep up with the rate of data streams. Importantly, the constraint programming-based operators placement is able to sustain up to 5x increased data ingestion compared to baseline strategies including resource-aware placements and graph partitioning-based placement.

¹⁶<https://anr.fr/en/investments-for-the-future/investments-for-the-future/>

¹⁷<http://slices-ri.eu>

In [13], we introduce DIMINET, a distributed module to interconnect independent networking resources across multiple locations in an automatized and transparent manner. It is nowadays accepted that the deployment of a geo-distributed cloud infrastructure, leveraging for instance Point-of-Presences at the edge of the network, could better fit the requirements of Network Function Virtualization services and Internet of Things applications. The envisioned architecture to operate such a widely distributed infrastructure relies on executing one instance of a Virtual Infrastructure Manager (VIM) per location and implement appropriate code to enable collaborations between them when needed. However, delivering the mechanisms that allow the collaborations is complex and error prone task. This is particularly true for the one in charge of establishing connectivity among VIM instances on-demand. Besides the re-configuration of the network equipment, the main challenge is to design a mechanism that can offer usual network virtualization operations to the users while dealing with scalability and intermittent network properties of geo-distributed infrastructures. To deal with these challenges, we designed DIMINET, a DIstributed Module for Inter-site NETworking services. DIMINET relies on a decentralized architecture where each agent communicates with others only if needed. Moreover, there is no global view of all networking resources but each agent is in charge of interconnecting resources that have been created locally. This approach enables us to mitigate management traffic and keep each site operational in case of network partitions. A promising approach to make other cloud-services collaborative on-demand.

In [22], we describe in detail the storage solution we designed between 2017 and 2019 for Fog/Edge infrastructures. Our initial proposal consisted in coupling an object store system and a scale-out NAS (Network Attached Storage) allowing both scalability and performance. This architecture has been extended with a new protocol inspired from the Domain Name System (DNS) to manage replicas in a context of mobility. The different versions have been evaluated over Grid'5000, presenting each time promising performance.

In [12], we introduce a simulation engine dedicated to the evaluation and comparison of scheduling and data movement policies for edge computing use-cases. Although several strategies have been proposed, they have been evaluated through ad-hoc simulator extensions that are, when available, usually not maintained. This is a critical problem because it prevents researchers to-easily-perform fair comparisons between different proposals. We proposed to address this limitation by developing an extension to the Batsim/SimGrid toolkit. Our tool includes a plug-in system that allows researchers to add new models in order to cope with the diversity of edge computing devices. Moreover, it includes an injector that allows the simulator to replay a series of events captured in real infrastructures. We demonstrated the relevance of such a simulation toolkit by studying two scheduling strategies with four data movement policies on top of a simulated version of the Qarnot Computing platform, a production edge infrastructure based on smart heaters. We chose this use-case as it illustrates the heterogeneity as well as the uncertainties of edge infrastructures. Our ultimate goal is to gather industry and academics around a common simulator so that efforts made by one group can be factorised by others.

In [26], we present a break through approach in the design of resource management systems (RMSes) for the edge. Two years ago, at HotEdge'18, we alerted our community of the importance of delivering a RMS to favor the advent of the edge computing paradigm. While new initiatives have been proposed, they are far from offering the expected features to administrate and use geo-distributed infrastructures. However, we claim our community can move forward by focusing on the collaborations of multiple instances of the same RMS instead of developing a new solution from scratch. Our proposal leverages service concepts, dynamic composition as well as programming software abstractions. Beyond the development of a resource management system for edge infrastructures, our proposal may lead to a new way of distributing applications where intermittent network is the norm.

In [28], we provide reflections regarding how Kubernetes, the well-known container orchestration platform can cope with edge infrastructure specifics. With the advent of Edge Computing era, DevOps expect to find features that made the success of containerized applications in the cloud, also at the edge. However, orchestration systems have not been designed to deal with resources geo-distribution aspects such as latency, intermittent networks, locality-awareness, etc. In other words, it is unclear whether they could be directly used on top of such massively distributed infrastructures or whether they must be revised. To provide the first elements of answers, we conducted an experimental campaign to analyze the impact of WAN links on the vanilla Kubernetes. This study raised several challenges our community should consider in order to better address geo-distribution concerns in orchestration platforms.

In [19], we propose and evaluate a flexible centralized controller for scheduling wireless networks. The

context of this work encompasses wireless networks within the wider Internet of Things (IoT) field and in particular addresses the requirements and limitations within the narrower Industrial Internet of Things (IIoT) sub-field. The overall aim of this work is to produce wireless networking solutions for industrial applications. The challenges include providing high reliability and low latency guarantees, comparable to existing wired solutions, within a noisy wireless medium and using generally computationally- and energy-restrained network nodes. We describe the development of a centralized controller for Wireless Industrial Networks, currently aimed at IEEE Std 802.15.4-2015 Time Slotted Channel Hopping protocol. Our controller takes a high-level network-centric problem description as input, translates it to a low-level representation and uses that to retrieve a solution from a Satisfiability Modulo Theories (SMT) solver, translating the solution back to a higher-level network-centric representation. The advantages of our solution are the ability to gain the added flexibility, higher ease of deployment, and lower deployment cost offered by wireless networks by generating configurable and flexible schedules for these applications.

8.2 Programming Support

Participants Maverick Chardet, H el ene Coullon, Jolan Philippe.

In 2020, our contributions on programming support are related to two aspects, at first glance without any link. Yet, they both are related to the autonomic adaptation of systems. One can note that making distributed systems autonomous in their deployment and management is of high importance at the scale of Fog and Edge computing, to avoid human errors and guarantee safety properties. It is usual when speaking about autonomic to introduce the MAPE-K loop, where (M) is the monitoring of systems and environment, (A) is the analysis of the situation and the decision of new adaptations, (P) is the planning phase of the adaptation and (E) is the execution of the plan. These 4 steps share a knowledge (K). The first contribution detailed below answers the (A) phase, in the specific case of low-code platforms, by computing the most adapted parallelism and configuration strategy according to the current state (*i.e.*, application, environment). The second publication contribute to building a generic model and tool to build efficient and safe execution of distributed systems adaptation (E), *i.e.*, reconfiguration. By formalizing its semantics and offering performance prediction tools, this second contribution also open the door to automated computation of a reconfiguration plan (P).

In [14] we discuss programming support to automate the choice of the most adequate parallelism strategy and distributed configuration for Low-code development platforms. Low-code development platforms are taking an important place in the model-driven engineering ecosystem, raising new challenges, among which transparent efficiency or scalability. Indeed, the increasing size of models leads to difficulties in interacting with them efficiently. To tackle this scalability issue, some tools are built upon specific computational strategies exploiting reactivity, or parallelism. However, their performances may vary depending on the specific nature of their usage. Choosing the most suitable computational strategy for a given usage is a difficult task which should be automated. Besides, the most efficient solutions may be obtained by the use of several strategies at the same time. This paper motivates the need for a transparent multi-strategy execution mode for model-management operations. We present an overview of the different computational strategies used in the model-driven engineering ecosystem, and use a running example to introduce the benefits of mixing strategies for performing a single computation. This example helps us present our design ideas for a multi-strategy model-management system. The code-related and DevOps challenges that emerged from this analysis are also presented.

In [11] the dynamic reconfiguration of distributed software systems is tackled. Reconfiguration is nowadays gaining interest because of the emergence of dynamic IoT and smart applications as well as large scale dynamic infrastructures (e.g., Fog and Edge computing). When quality of service and experience is of prime importance, efficient reconfiguration is necessary, as well as performance predictability to decide when a reconfiguration should occur. This paper tackles the problem of efficient execution of a reconfiguration plan and its predictability with Concerto, a reconfiguration model supporting a high level of parallelism in reconfiguration operations. Evaluation performed on synthetic cases and on two real production scenarios show that Concerto provides better performance than state-of-the-art systems with accurate time estimation.

8.3 Energy-aware computing

Participants Emile Cadorel, H el ene Coullon, Remous Aris Koutsiamanis, Thomas Ledoux, Jean-Marc Menaud, Jonathan Pastor, Dimitri Saingre, Yewan Wang.

In 2020, we achieved three contributions in the field of Energy-aware computing. First, we have addressed blockchain-based solutions and proposed a framework to benchmark different blockchain technologies (e.g., CPU consumption). Then, tackling the problem of scheduling multiuser workflows from the Cloud provider point of view, we have proposed an algorithm that optimizes two metrics: the energy consumption and the fairness between users. Finally, we have addressed the issue of network energy consumption in the context of IoT devices.

Over the last years, research activities on blockchain technologies have fairly increased. Firstly introduced with Bitcoin, some projects have since emerged to create or improve blockchain features like privacy while others propose to overcome technical limitations such as scalability and energy consumption. New proposals are often evaluated with ad hoc tools and experimental environments. Reproducibility and comparison of these new contributions with the state of the art of the blockchain technologies are therefore complicated. To the best of our knowledge, only a few tools partially address the design of a generic benchmarking of blockchain technologies (e.g., load generation). In [15], we introduce BCTMark, a generic framework for benchmarking blockchain technologies on an emulated network in a reproducible way. To illustrate the portability of experiments using BCTMark, we have conducted some experiments on two different testbeds: a cluster of Dell PowerEdge R630 servers (Grid'5000) and one of Raspberry Pi 3+. Experiments have been conducted on three different blockchain systems (Ethereum Clique/Ethash and Hyperledger Fabric) to measure their CPU consumption and energy footprint for different numbers of clients.

In [10], we deal with the problem of scheduling multiuser scientific workflows with unpredictable random arrivals and uncertain task execution times in a Cloud environment from the Cloud provider point of view. The solution consists in a deadline sensitive online algorithm, named NEARDEADLINE, that optimizes two metrics: the energy consumption and the fairness between users. Scheduling workflows in a private Cloud environment is a difficult optimization problem as capacity constraints must be fulfilled additionally to dependencies constraints between tasks of the workflows. Furthermore, NEARDEADLINE is built upon a new workflow execution platform. As far as we know no existing work tries to combine both energy consumption and fairness metrics in their optimization problem. The experiments conducted on a real infrastructure (clusters of Grid'5000) demonstrate that the NEARDEADLINE algorithm offers real benefits in reducing energy consumption, and enhancing user fairness.

In [18], we propose a set of extensions to a popular routing protocol for Industrial IoT devices to improve the energy consumption and reduce the waiting time for devices joining the network. The IPv6 Routing Protocol for Low-Power and Lossy Networks (RPL) is the de facto routing protocol for Low Power and Lossy Networks (LLNs). It is a proactive and link-layer agnostic routing protocol standardized as RFC 6550 by the Internet Engineering Task Force (IETF). Based on the distance-vector technique, RPL builds a Destination Oriented Directed Acyclic Graph (DODAG) topology. To establish and maintain the routes, RPL uses DODAG Information Object (DIO) control packets, that are transmitted in broadcast, for RPL nodes to propagate the DODAG related information, while its transmission frequency depends on Trickle timer algorithm, i.e., the less stable the network the more DIOs are transmitted. Thus, when a new node intends to join the RPL DODAG, it listens for a DIO message from nearby nodes, which may take a very long time if the network is in a stable state. Therefore, RFC 6550 is equipped with the DODAG Informational Solicitation (DIS) message to solicit DIOs from nearby RPL nodes, similar to the Router Solicitation in IPv6 Neighbor Discovery. However, the solicitation procedure is not the most efficient one, since it resets the Trickle timers in the nodes that receive the DIS message and, thus, they transmit an unnecessarily large number of DIOs that congest the network and consume energy in the nodes. In this paper, we propose to augment RFC 6550, the RPL routing protocol, with additional DIS flags and options that allow a RPL node to better control how the nearby RPL nodes will respond to its solicitation for DIOs. Our performance evaluation in Contiki-NG & COOJA demonstrates that we can reduce the control packets in the network by up to 45.5%.

In [20], we propose and evaluate three different ways of implementing multi-path routing in Industrial IoT devices to improve reliability, latency and energy consumption. The IPv6 Routing Protocol for Low-Power and Lossy Networks (RPL) networks is designed for Internet of Things (IoT) networks to generate routes between devices with minimal processing. This protocol creates a DODAG (Destination Oriented Directed Acyclic Graph) network topology through the use of DODAG Information Object (DIO) control packets. The DODAG routes the data packets upstream to the destination device. In order to obtain a reliable network, we implement Packet Replication and Elimination (PRE) to perform multi-path data transmission via multiple parent devices. However, there is no standard way to select an alternative path. This document presents three types of Alternative Parent (AP) selection following a braided model. We focus on analyzing its performance in terms of delay and compromise between network traffic and reliability.

Finally, Yewan Wang defended in March her PhD [23] on the evaluation and the modelisation of the energy impacts of data centers, in terms of hardware / software architecture and associated environment. The main contribution is a global power modeling to estimate the overall energy consumption of a cluster using informations such as ambient temperature, cooling system configurations and server loads.

8.4 Security and Privacy

Participants Fatima Zahra Boujdad, Wilmer Edicson Garzon Alfonso, Sirine Sayadi, Mario Südholt.

This year the team has continued work on the protection of sensitive data and the security of distributed biomedical analyses in the context of the PhD theses of Fatima-zahra Boujdad (in cooperation with partners informaticiens and geneticiens working at INSERM laboratories in Brest), Wilmer Garzon (in cooperation with a Colombian technical university) and Sirine Sayadi (in cooperation with a medical group from the university hospital Nantes).

We have provided solutions to the problem of preserving privacy of sensitive data in distributed biomedical analyses, notably in the field of precision medicine, that is, analyses and treatments tailored to small groups of patients or even individual patients. [16]

An important aspect of precision medicine consists in patient-centered contextualization analyses that are used as part of biomedical interactive tools. Such analyses often harness data of large populations of patients from different research centers and can often benefit from a distributed implementation. However, performance and the security/privacy concerns of sharing sensitive biomedical data can become a major issue.

We have investigated these issues in the context of the Kidney Transplantation Application (KITAPP). We presented a motivation for distributed implementations in this context, notably for computing percentiles for contextualization. We motivated privacy and performance issues and presented a novel system architecture as well as a distributed implementation to tackle them. Its evaluation in a realistic multi-site environment has shown that our approach reduces data sharing to a large extent, and thus enables imposing strong data privacy guarantees on distributed biomedical analyses.

9 Bilateral contracts and grants with industry

9.1 Bilateral contracts with industry

Participants Ronan-Alexandre Cherrueau, Marie Delavergne, Adrien Lebre (*Contact point*), Karim Manaouil, Matthieu Simonin.

Following the ENOS bilateral contract (“Contrat de Recherche Externalisé”) between Orange and Inria (Sep 2017-Oct 2018), we agreed with Orange Labs to pursue this collaboration around a second contrat. This new contrat, which is going to last 18 months for a budget of 150K€, targets the following objectives:

- Strengthen the Enos framework and the resulting EnosLib solution (see Section 6.4 and Section 6.5).

- Define an experimental protocol allowing the automatized and reproducible evaluation of resource management/orchestrations systems in a WANWide context.
- Develop a DSL to reify location aspects at the CLI level in order to create new resources (image, VM, etc.) through a set of OpenStack instances while guaranteeing a notion of master copy.

10 Partnerships and cooperations

10.1 International initiatives

10.1.1 Inria International Labs

Inria@SiliconValley

Associate Teams involved in the Inria International Lab:

Hermes

Title: *Accelerating the Performance of Multi-Site Scientific applications through Coordinated Data management*

Duration: 2019 - 2021

Coordinator: Shadi Ibrahim

Partners:

- Scientific Data Management Group, Lawrence Berkeley National Laboratory (United States)

Inria contact: Shadi Ibrahim

Summary: Advances in computing, experimental, and observational facilities are enabling scientists to generate and analyze unprecedented volumes of data. A critical challenge facing scientists in this era of data deluge is storing, moving, sharing, retrieving, and gaining insight from massive collections of data efficiently. Existing data management and I/O solutions on high-performance computing (HPC) systems require significant enhancements to handle the three V's of Big Data (volume, velocity, and variety) in order to improve productivity of scientists. Even more challenging, many scientific Big Data and machine learning applications require data to be shared, exchanged, and transferred among multiple HPC sites. Towards overcoming these challenges, in this project, we aim at accelerating scientific Big Data application performance through coordinated data management that addresses performance limitations of managing data across multiple sites. In particular, we focus on challenges related to the management of data and metadata across sites, distributed burst buffers, and online data analysis across sites.

10.2 International research visitors

10.2.1 Visits of international scientists

Twinkle Jain: PhD student at Northeastern university visited the STACK team from February to August 2020.

10.3 European initiatives

10.3.1 FP7 & H2020 Projects

Lowcomote

Participants H  l  ne Coullon, Jolan Philippe.

H  l  ne Coullon is a member of the advisory board of the Lowcomote ITN project (H2020), on the subject of low code platforms. In particular, H  l  ne brings her expertise in distributed and high performance computing applied to model-driven engineering. She supervises Jolan Philippe, one PhD student of the project and a member of the team.

SLICES-DS

Participants Adrien Lebre.

Adrien Lebre is a member of the SLICES Design Study project (ESFRI) that targets a Europe-wide test-platform designed to support large-scale, experimental research. It will provide advanced compute, storage and network components, interconnected by dedicated high-speed links. Pushing forward, the project's main goal is to strengthen the research excellence and innovation capacity of European researchers and scientists in the design and operation of digital infrastructures.

10.4 National initiatives

10.4.1 Ademe

GLAMA

Participants Brice N  delec, Thomas Ledoux (*coordinator*).

The Green Label for Microservices Architecture (GLAMA) project aims to design and develop a technological platform (tools, framework, dedicated languages) for the self management of eco-responsible micro-service architectures for the Cloud. The experiments will be carried out through case studies provided by Sigma Informatique and the presence of renewable energy will initially be simulated. At the end of the project, the technological platform will be deployed as part of the CPER SeDuCe platform.

This project is funded by the Ademe (call Perfecto) running for 18 months (starting in September 2019) with an allocated budget of 116 480   (the majority of aid is dedicated to the R&D engineer salary).

10.4.2 CominLabs laboratory of excellence

10.4.3 CNRS

Kabuto

Participants Remy Pottier, Jean-Marc Menaud (*coordinator*).

The Kabuto project aims to develop a software solution allowing the reservation and deployment IoT and Edge devices to HPC nodes. Strated in November 2020 for one year, Kabuto finances a post doc over one year (70 K  ).

10.4.4 ANR

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 522K€, 90K€ for STACK).

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

The goal of the GRECO project ¹⁸ is to design a manager for cloud of things. The manager should act at the IaaS, PaaS and SaaS layer of the cloud. To move forward to this objective, we have been designing a simulator to innovate in designing scheduling and data management systems. This simulator leverage the Simgrid/PyBATSIM solution

KerStream (ANR)

Participants 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 238K€).

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.

10.4.5 FSN

Hydda (FSN)

Participants Emile Cadorel, Dimitri Saingre, Rémy Pottier, Hélène 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 to 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 to optimize resources usage of these platforms (CPU, RAM, Disk, Energy, etc.) in order to propose solutions at the least cost?

Started in 2017, (with an allocated budget of 4000K€, 380K€ for STACK), the project ended in October 2020.

¹⁸<https://anr-greco.net>

10.4.6 CPER

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).

10.4.7 Connect Talent

Apollo (Connect Talent)

Participants 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 201K€).

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.

10.4.8 Etoiles Montantes

VeRDi

Participants Emile Cadorel, Hélène Coullon (*Coordinator*), Dimitri Pertin, Simon Robillard, Charlène Servantie.

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 June 2021 (extended due to Covid-19) 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.

10.5 Regional initiatives

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 medicine systems medicine by co-developing three 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) with a global envelope of 150 K€ (10 K€ for our team) from 2018-22.

SHLARC

Participants Mario Südholt (*coordinator*), Sirine Sayadi.

The SHLARC project is an international network involving more than 20 partners from more than 15 countries located on four continents. The network aims at improving HLA imputation techniques in the domain of immunobiology, notably by investigation better computational methods for the corresponding biomedical analyses.

The ambition of the SHLARC is to bring together international expertise to solve essential questions on immune-related pathologies through innovative algorithms and powerful computation tool development. To achieve this goal, we determined 3 main objectives

- **Data.** By bringing together scientists from around the world, we will collectively increase the amount of SNP+HLA data available, both in terms of quantity and diversity.
- **Applied mathematical and computer sciences.** We will further optimize SNP-HLA imputation methods using the attribute-bagging HIBAG tool, and particularly for genetically diverse and admixed populations.
- **Accessibility and service to the scientific community.** Following the Haplotype Reference Consortium (HRC) initiative, the network envisions building a free, user-friendly webserver where researchers can access improved imputation protocols by simply uploading their data and obtaining the best possible HLA imputation for their dataset.

In this context, the STACK team is working on improved analysis techniques that harness distributed infrastructures.

This project is financed by the Nantes excellency initiative in Medecine and Informatics (NExT) from 2019-22 with a global envelope of 100 K€ (and 5 K€ for our team)

Oncoshare

Participants Mario Südholt (*coordinator*).

The ONCOSHARE project (ONCOlogy big data SHARING for Research) will demonstrate, through a multidisciplinary cooperation within the Western CANCERPOLE 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 with a global envelope of 150 K€ (and 10 K€ for our team).

11 Dissemination

11.1 Promoting scientific activities

11.1.1 Scientific events: organisation

- A. Lebre co-organized the GdR RSD yearly meeting (100 participants), Nantes, January 2020.

Member of the organizing committees

- Shadi Ibrahim: workshop co-chair for 40th IEEE International Conference on Distributed Computing Systems (ICDCS 2020), November 29 - December 1, 2020, Singapore.
- Adrien Lebre is a member of the steering committee of the international conference of Fog and Edge Computing (ICFEC).

11.1.2 Scientific events: selection

Chair of conference program committees

- H el ene Coullon: track vice-chair of CCGrid 2020, track “Programming models and runtime systems”
- Shadi Ibrahim: program vice chair for the Architecture, Networking, Data Centers track of the 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGrid 2020), May 11 – 14, 2020.
- Shadi Ibrahim: track co-chair for the Data, Storage and Visualization Track of ACM HPCAsia 2020, Fukuoka, Japan, January, 2020.

Member of the conference program committees

- H el ene Coullon: Euro-Par 2020, ICCS 2020, ISCC 2020
- Shadi Ibrahim: SC 2020 (regular paper and poster tracks), ICDCS 2020, HiPC 2020, ICAPDS 2020, ICA3PP 2020, HPDBC@IPDPS 2020, COMPAS 2020, HPSC 2020, BDPS@ICDCS 2020.
- Adrien Lebre: ICC 2020, Europar 2020, IC2E 2020, ICDCS 2020, CloudCom 2020, UCC 2020.
- Jean-Marc Menaud: SMARTGREENS’20, SOFTCOM 20.

11.1.3 Journal

Member of the editorial boards

- Shadi Ibrahim: Young Associate Editor of the Springer Frontiers of Computer Science journal.
- Adrien Lebre: Associate Editor of the IEEE Transactions on Cloud Computing.

Reviewer - reviewing activities

- H el ene Coullon: Transactions on Cloud Computing, Transactions on Parallel and Distributed Systems
- Shadi Ibrahim: IEEE Transactions on Computers, IEEE Transactions on Network and Service Management, ACM Transactions on Modeling and Performance Evaluation of Computing Systems.
- Remous Aris Koutsiamanis: IEEE Transactions on Cloud Computing, IEEE Internet of Things Journal, Springer Wireless Networks.
- Jacques Noy e: PeerJ Computer Science.

11.1.4 Invited talks

- T. Ledoux has given a talk to the Technof erence P ole Images & r eseaux « Voyage dans le XaaS », June 2020.

11.1.5 Scientific expertise

- T. Ledoux contributed for the Allistene alliance to the national inter-alliances seminar "Energie d ecarbon ee, changement climatique, sant e environnementale et biodiversit e : les impacts de nos choix sur les nouvelles voies de recherche interdisciplinaires", Dec. 2020.
- A. Lebre contributed to the Allistene working group dedicated to the definition of the Infrastructure roadmap (“Instrument de recherche”), Dec 2020.

11.1.6 Research administration

- A. Lebre has been co-chairing of the transversal action “Virtualization” of the GdR RSD between 2015 and November 2020.
- A. Lebre has been 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 between 2015 and November 2020.
- A. Lebre is a member of the executive and architect committees of the Grid’5000 GIS (Groupement d’intérêt scientifique).
- T. Ledoux has been vice-head of the cross-cutting theme "Gestion de l'énergie et maîtrise des impacts environnementaux" at LS2N
- Jean-Marc Menaud has been head of the scientific theme "Science du logiciel et systèmes distribués" at LS2N

11.1.7 Member of the Selection Committee for Associate Professors

- H. Coullon was a member of a Selection Committee (CS) at the University of Rennes 1
- H. Coullon was a member of a Selection Committee (CS) at the University Nice Sophia Antipolis
- T. Ledoux was member of a Selection Committee (CS) at INSA Rennes

11.2 Teaching - Supervision - Juries

11.2.1 Teaching

- S. Ibrahim is the co-coordinator of the international Master’s program in Cloud Computing and Services at University of Rennes 1.
- 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.
- H. Coullon is responsible for the Computer Science domain of the new apprenticeship program in Industry 4.0 (FIT) of IMT Atlantique. This 3-year program leads to the award of a Master degree in Industry 4.0 from the IMT Atlantique.

11.2.2 Supervision

- Phd: Marie Delavergne, director: A. Lebre.
- Phd: David Espinel, director: A. Lebre.
- PhD: Dimitri Saingre, advisor: T. Ledoux, director: J.-M. Menaud.
- PhD: Maverick Chardet, advisor: H. Coullon, director: C. Perez (Avalon).
- PhD: Emile Cadorel, advisor: H. Coullon, director: J.-M. Menaud.
- PhD: Jolan Philippe, advisors: H. Coullon, M. Tisi (NaoMod), director: G. Sunye (NaoMod).
- PhD: Yewan Wang, advisor: J.-M. Menaud.
- PhD: Maxime Belair, advisor: J.-M. Menaud.
- PhD: Fatima-zahra Boujdad, advisor: M. Südholt.

- PhD: Wilmer Garzon, advisor: M. Südholt.
- PhD: Sirine Sayadi, advisor: M. Südholt.
- Postdoc: David Guyon, advisor: S. Ibrahim, until September 2020.
- Postdoc: Thomas Lambert, advisor: S. Ibrahim, until October 2020.
- Postdoc: Jonathan Pastor, advisor: J.-M. Menaud, until June 2020?
- Postdoc: Rémy Pottier, advisor: J.-M. Menaud.
- Postdoc: Simon Robillard, advisor: H. Coullon.
- Postdoc: Hamza Sahli, advisor: T. Ledoux (until April 2020).
- Postdoc: Alexandre Van Kempen: A. Lebre (until November 2020).
- Engineer: Ronan-Alexandre Cherrueau, advisor: A. Lebre.
- Engineer: Brice Nédelec, advisor: T. Ledoux.
- Engineer: Charène Servantie, advisor: H. Coullon (until June 2020).
- Engineer: Jad Darrous, advisor: S. Ibrahim. (until Sept 2020).

11.2.3 Juries

- H. Coullon was a member of the PhD committee of Salwa Swaf, "Formal Methods meet Security in a Cost Aware Cloud Brokerage Solution", Univ. Orléans, Dec. 15, 2020.
- S. Ibrahim was a reviewer of the PhD thesis of Muhammad Abdullah "Performance-aware Cloud Resource Management for Microservices-based Applications". Punjab University, Pakistan, Sep, 2020.
- S. Ibrahim was was a member of the PhD committee of Arif Ahmed "Efficient Cloud Application Deployment in Distributed Fog Infrastructures". Universite de Rennes 1, France, Jan. 20, 2020.
- A. Lebre was a member of the Phd Committee of Adrien Faure, "Advanced Simulation for Resource Management", Grenoble Alpes University, Dec 2, 2020.
- A. Lebre was a reviewer of the Phd Committee of Arthur Chevalier, "Optimisation du placement des licences logicielles dans le Cloud pour un déploiement économique et efficient", Ecole Normale Supérieure de Lyon, Nov, 24, 2020.
- A. Lebre was a reviewer of the Phd Committee of Bruno Donassolo, "Orchestration des applications IoT dans le Fog", Grenoble Alpes University, Nov, 4, 2020.
- A. Lebre was a member of the Phd Committee of Jean Emile Dartois, "Leveraging Cloud unused heterogeneous resources for applications with SLA guarantees", Rennes 1 University, Sept 4, 2020.
- A. Lebre was a member of the Phd Committee of Gregoire Todeschi, "Optimisation des caches de fichiers dans les environnements virtualisés", Institut National Polytechnique de Toulouse (France), June 8, 2020.
- T. Ledoux was a member of the PhD committee of Francis Laniel, "Vers une consolidation mémoire pour les conteneurs grâce à un retour applicatif", Sorbonne Univ., Nov. 9, 2020.
- T. Ledoux was a reviewer of the PhD committee of Neil Ayeb, "Autonomic and decentralized device management for the Internet of Things", Univ. Grenoble Alpes, Nov. 25, 2020.
- J.-M. Menaud was a reviewer of the PhD committee of Mathieu Bacou "Performance et gestion de ressources dans un cloud multi-virtualisé" Univ Toulouse, May 12, 2020.
- J.-M. Menaud was a reviewer of the PhD committee of Marc Platini "Apprentissage machine appliqué à l'analyse et à la prédiction des défaillances dans les systèmes HPC", Univ Grenoble, May 20, 2020.

11.3 Popularization

11.3.1 Internal or external Inria responsibilities

- A. Lebre is a co-director of the <I/O> Lab, a joint lab between Inria and Orange Labs.
- A. Lebre is a member of the scientific committee of the joint lab between Inria and Nokia Bell Labs.

11.3.2 Articles and contents

- Adrien Lebre is co-author of the second white paper of the Edge Computing working group of the OpenStack foundation [94]

11.3.3 Education

- T. Ledoux is head of the "Filière informatique nantaise" since Sept. 2020. This entity, created by the University of Nantes, the Ecole Centrale de Nantes and IMT Atlantique, aims to bring together the main players in IT training in Nantes to ensure a coherent and ambitious training offer that meets the present and future challenges of the IT discipline. It is organised around a Council made up of representatives from the academic and socio-economic worlds.

12 Scientific production

12.1 Major publications

- [1] F. Ait Salaht, F. Desprez and A. Lebre. 'An overview of service placement problem in Fog and Edge Computing'. In: *ACM Computing Surveys* 53.3 (June 2020), Article 65, 35 pages. DOI: [10.1145/3391196](https://doi.org/10.1145/3391196). URL: <https://hal.inria.fr/hal-02596419>.
- [2] E. Cadorel, H. Coullon and J.-M. Menaud. 'Online Multi-User Workflow Scheduling Algorithm for Fairness and Energy Optimization'. In: *CCGrid2020 : 20th International Symposium on Cluster, Cloud and Internet Computing*. Melbourne, Australia, Nov. 2020. DOI: [10.1109/CCGrid49817.2020.00-36](https://doi.org/10.1109/CCGrid49817.2020.00-36). URL: <https://hal.archives-ouvertes.fr/hal-02551733>.
- [3] M. Chardet, H. Coullon and C. Pérez. 'Predictable Efficiency for Reconfiguration of Service-Oriented Systems with Concerto'. In: *CCGrid 2020 : 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing*. Melbourne, Australia: IEEE, May 2020. DOI: [10.1109/CCGrid49817.2020.00-59](https://doi.org/10.1109/CCGrid49817.2020.00-59). URL: <https://hal.inria.fr/hal-02535077>.
- [4] D. Sarmiento, A. Lebre, L. Nussbaum and A. Chari. 'Multi-site Connectivity for Edge Infrastructures DIMINET: Distributed Module for Inter-site NETWORKing'. In: *CCGRID 2020 - 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing*. IEEE and The University of Melbourne. Melbourne, Australia, May 2020, pp. 1–10. URL: <https://hal.archives-ouvertes.fr/hal-02573638>.

12.2 Publications of the year

International journals

- [5] F. Ait Salaht, F. Desprez and A. Lebre. 'An overview of service placement problem in Fog and Edge Computing'. In: *ACM Computing Surveys* 53.3 (June 2020), Article 65, 35 pages. DOI: [10.1145/3391196](https://doi.org/10.1145/3391196). URL: <https://hal.inria.fr/hal-02596419>.
- [6] O. Beaumont, T. Lambert, L. Marchal and B. Thomas. 'Performance Analysis and Optimality Results for Data-Locality Aware Tasks Scheduling with Replicated Inputs'. In: *Future Generation Computer Systems* 111 (Oct. 2020), pp. 582–598. DOI: [10.1016/j.future.2019.08.024](https://doi.org/10.1016/j.future.2019.08.024). URL: <https://hal.inria.fr/hal-02275473>.
- [7] T. Lagos Jenschke, R.-A. Koutsiamanis, G. Papadopoulos and N. Montavont. 'ODESe: On-Demand Selection for multi-path RPL networks'. In: *Ad Hoc Networks* 114 (2021), p. 102431. DOI: [10.1016/j.adhoc.2021.102431](https://doi.org/10.1016/j.adhoc.2021.102431). URL: <https://hal.archives-ouvertes.fr/hal-03122968>.

- [8] D. E. Sarmiento, A. Lebre, L. Nussbaum and A. Chari. ‘Decentralized SDN Control Plane for a Distributed Cloud-Edge Infrastructure: A Survey’. In: *Communications Surveys and Tutorials, IEEE Communications Society*. IEEE Communications Surveys & Tutorials 23.1 (25th Feb. 2021), pp. 256–281. DOI: [10.1109/COMST.2021.3050297](https://doi.org/10.1109/COMST.2021.3050297). URL: <https://hal.archives-ouvertes.fr/hal-03119901>.

International peer-reviewed conferences

- [9] M. Belair, S. Laniece and J.-M. Menaud. ‘SNAPPY: Programmable Kernel-Level Policies for Containers’. In: SAC 2021 : 36th ACM/SIGAPP Symposium On Applied Computing. Gwangju / Virtual, South Korea: <https://www.sigapp.org/sac/sac2021/>, 22nd Mar. 2021. URL: <https://hal.inria.fr/hal-03108231>.
- [10] E. Cadorel, H. Coullon and J.-M. Menaud. ‘Online Multi-User Workflow Scheduling Algorithm for Fairness and Energy Optimization’. In: CCGrid2020 : 20th International Symposium on Cluster, Cloud and Internet Computing. Melbourne, Australia, 3rd Nov. 2020. DOI: [10.1109/CCGrid49817.2020.00-36](https://doi.org/10.1109/CCGrid49817.2020.00-36). URL: <https://hal.archives-ouvertes.fr/hal-02551733>.
- [11] M. Chardet, H. Coullon and C. Pérez. ‘Predictable Efficiency for Reconfiguration of Service-Oriented Systems with Concerto’. In: CCGrid 2020 : 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing. Melbourne, Australia, 4th June 2020. DOI: [10.1109/CCGrid49817.2020.00-59](https://doi.org/10.1109/CCGrid49817.2020.00-59). URL: <https://hal.inria.fr/hal-02535077>.
- [12] A. A. Da Silva, C. Mommessin, P. Neyron, D. Trystram, A. Bauskar, A. Lebre, A. V. Kempen, Y. Ngoko and Y. Ricordel. ‘Evaluating Computation and Data Placements in Edge Infrastructures through a Common Simulator’. In: SBAC-PAD 2020 - IEEE 32nd International Symposium on Computer Architecture and High Performance Computing. Porto, Portugal, 8th Sept. 2020. URL: <https://hal.inria.fr/hal-02915346>.
- [13] D. Espinel Sarmiento, A. Lebre, L. Nussbaum and A. Chari. ‘Multi-site Connectivity for Edge Infrastructures DIMINET:DIstributed Module for Inter-site NETworking’. In: CCGRID 2020 - 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing. Melbourne, Australia: <http://cloudbus.org/ccgrid2020/>, 2020, pp. 1–10. URL: <https://hal.archives-ouvertes.fr/hal-02573638>.
- [14] J. Philippe, M. Tisi, H. Coullon and G. Sunyé. ‘Towards Transparent Combination of Model Management Execution Strategies for Low-Code Development Platforms’. In: *ACM/IEEE 23rd International Conference on Model Driven Engineering Languages and Systems*, ACM/IEEE 23rd International Conference on Model Driven Engineering Languages and Systems. (Virtual Conference), Canada: <https://lowcode-workshop.github.io/>, 19th Oct. 2020. DOI: [10.1145/3417990.3420206](https://doi.org/10.1145/3417990.3420206). URL: <https://hal.archives-ouvertes.fr/hal-02952952>.
- [15] D. Saingre, T. Ledoux and J.-M. Menaud. ‘BCTMark: a framework for benchmarking blockchain technologies’. In: AICCSA 2020 - 17th IEEE/ACS International Conference on Computer Systems and Applications. 2020 IEEE/ACS 17th International Conference on Computer Systems and Applications (AICCSA). Antalya, Turkey, 2nd Nov. 2020, pp. 1–8. DOI: [10.1109/AICCSA50499.2020.9316536](https://doi.org/10.1109/AICCSA50499.2020.9316536). URL: <https://hal.archives-ouvertes.fr/hal-02923038>.
- [16] S. Sayadi, E. Geffard, M. Südholt, N. Vince and P.-A. Gourraud. ‘Distributed Contextualization of Biomedical Data: a case study in precision medicine’. In: AICCSA 2020 - 17th IEEE/ACS International Conference on Computer Systems and Applications. Antalya, Turkey, 2nd Nov. 2020, pp. 1–6. URL: <https://hal.inria.fr/hal-02922930>.
- [17] S. Sayadi, E. Geffard, M. Südholt, N. Vince and P.-A. Gourraud. ‘Secure distribution of Factor Analysis of Mixed Data (FAMD) and its application to personalized medicine of transplanted patients’. In: AINA 2021 : 35th International Conference on Advanced Information Networking and Applications. Toronto, Canada: <https://link.springer.com/conference/aina>, 12th May 2021. URL: <https://hal.archives-ouvertes.fr/hal-03141653>.

- [18] D. Sourailidis, R.-A. Koutsiamanis, G. Papadopoulos, D. Barthel and N. Montavont. ‘RFC 6550: On Minimizing the Control Plane Traffic of RPL-based Industrial Networks’. In: 2020 IEEE 21st International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoW-MoM). 2020 IEEE 21st International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoW-MoM). Cork, Ireland, 9th Oct. 2020, pp. 439–444. DOI: [10.1109/WoWMoM49955.2020.00080](https://doi.org/10.1109/WoWMoM49955.2020.00080). URL: <https://hal.archives-ouvertes.fr/hal-02749045>.

Conferences without proceedings

- [19] R.-A. Koutsiamanis, G. Papadopoulos, B. Quoitin and N. Montavont. ‘A Centralized Controller for Reliable and Available Wireless Schedules in Industrial Networks’. In: MSN 2020 - 16th International Conference on Mobility, Sensing and Networking. Virtual, Japan: <https://conference.csis.cityu.edu.hk/msn2020/>, 17th Dec. 2020, pp. 1–9. URL: <https://hal.archives-ouvertes.fr/hal-02987638>.
- [20] T. Lagos Jenschke, G. Papadopoulos, R.-A. Koutsiamanis and N. Montavont. ‘Stratégies d’ancêtre commun pour les réseaux RPL multi-chemins’. In: CORES2020. Lyon, France, 28th Sept. 2020. URL: <https://hal.archives-ouvertes.fr/hal-02651685>.
- [21] T. Lambert, D. Guyon and S. Ibrahim. ‘Rethinking Operators Placement of Stream Data Application in the Edge’. In: CIKM 2020- 29th ACM International Conference on Information and Knowledge Management. Dublin, Ireland, 19th Oct. 2020. URL: <https://hal.archives-ouvertes.fr/hal-02942759>.

Scientific book chapters

- [22] B. Confais, A. Lebre and B. Parrein. ‘A Fog storage software architecture for the Internet of Things’. In: *Advances in Edge Computing: Massive Parallel Processing and Applications*. Advances in Parallel Computing 35. <http://ebooks.iospress.nl/volumearticle/53824>, 2020, pp. 61–105. DOI: [10.3233/APC200004](https://doi.org/10.3233/APC200004). URL: <https://hal.archives-ouvertes.fr/hal-02496105>.

Doctoral dissertations and habilitation theses

- [23] Y. Wang. ‘Evaluating and modeling the energy impacts of data centers, in terms of hardware / software architecture and associated environment’. Ecole nationale supérieure Mines-Télécom Atlantique, 9th Mar. 2020. URL: <https://tel.archives-ouvertes.fr/tel-02948725>.

Reports & preprints

- [24] A. Bauskar, A. da Silva, A. Lebre, C. Mommessin, P. Neyron, Y. Ngoko, Y. Ricordel, D. Trystram and A. Van Kempen. *Investigating Placement Challenges in Edge Infrastructures through a Common Simulator (extended version)*. INRIA, 25th May 2020, p. 21. URL: <https://hal.inria.fr/hal-02153203>.
- [25] M. Chardet, H. Coullon, C. Pérez, D. Pertin, C. Servantie and S. Robillard. *Enhancing Separation of Concerns, Parallelism, and Formalism in Distributed Software Deployment with Madeus*. 2nd June 2020. URL: <https://hal.inria.fr/hal-02737859>.
- [26] R.-A. Cherrueau, M. Delavergne, A. Lebre, J. Rojas Balderrama and M. Simonin. *Edge Computing Resource Management System: Two Years Later!* Inria Rennes Bretagne Atlantique, 1st Apr. 2020. URL: <https://hal.inria.fr/hal-02527366>.
- [27] D. Espinel Sarmiento, A. Lebre, L. Nussbaum and A. Chari. *Decentralized SDN Control Plane for a Distributed Cloud-Edge Infrastructure: A Survey*. INRIA, 15th June 2020, p. 40. URL: <https://hal.inria.fr/hal-02868984>.
- [28] K. Manaouil and A. Lebre. *Kubernetes and the Edge?* Inria Rennes - Bretagne Atlantique, 20th Oct. 2020, p. 19. URL: <https://hal.inria.fr/hal-02972686>.

12.3 Cited publications

- [29] Akamai Cloudlets. (Accessed: 2018-03-08). 2018. URL: <http://cloudlets.akamai.com>.
- [30] Amazon Lambda@Edge. (Accessed: 2018-03-08). 2018. URL: <https://aws.amazon.com/lambda/edge/>.
- [31] C. Atkinson, T. Schulze and S. Klingert. 'Facilitating Greener IT through Green Specifications'. In: *IEEE Software* 31.3 (May 2014), pp. 56–63. DOI: [10.1109/MS.2014.19](https://doi.org/10.1109/MS.2014.19). URL: <http://dx.doi.org/10.1109/MS.2014.19>.
- [32] F. Baude, D. Caromel, C. Dalmaso, M. Danelutto, V. Getov, L. Henrio and C. Pérez. 'GCM: a grid extension to Fractal for autonomous distributed components'. English. In: *annals of telecommunications* 64.1-2 (2009), pp. 5–24. DOI: [10.1007/s12243-008-0068-8](https://doi.org/10.1007/s12243-008-0068-8). URL: <http://dx.doi.org/10.1007/s12243-008-0068-8>.
- [33] F. Baude, L. Henrio and C. Ruz. 'Programming distributed and adaptable autonomous components—the GCM/ProActive framework'. In: *Software: Practice and Experience* (May 2014). URL: <https://hal.inria.fr/hal-01001043>.
- [34] M.-M. Bazm, M. Lacoste, M. Südholt and 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, Oct. 2017. URL: <https://hal.inria.fr/hal-01593144>.
- [35] G. Blair, T. Coupaye and J.-B. Stefani. 'Component-based architecture: the Fractal initiative'. In: *Annals of telecommunications* 64.1 (Feb. 2009), pp. 1–4. DOI: [10.1007/s12243-009-0086-1](https://doi.org/10.1007/s12243-009-0086-1). URL: <https://doi.org/10.1007/s12243-009-0086-1>.
- [36] P. Bodik, R. Griffith, C. Sutton, A. Fox, M. I. Jordan and D. A. Patterson. 'Automatic Exploration of Datacenter Performance Regimes'. In: *Proceedings of the 1st Workshop on Automated Control for Datacenters and Clouds*. ACDC '09. Barcelona, Spain: ACM, 2009, pp. 1–6. DOI: [10.1145/1555271.1555273](https://doi.org/10.1145/1555271.1555273). URL: <http://doi.acm.org/10.1145/1555271.1555273>.
- [37] F. Bonomi, R. Milito, J. Zhu and 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.
- [38] F.-z. Boujdad and M. Südholt. 'Constructive Privacy for Shared Genetic Data'. In: *CLOSER 2018 - 8th International Conference on Cloud Computing and Services Science*. Proceedings of CLOSER 2018. Funchal, Madeira, Portugal, Mar. 2018, pp. 1–8. URL: <https://hal.archives-ouvertes.fr/hal-01692620>.
- [39] O. Bračevac, S. Erdweg, G. Salvaneschi and M. Mezini. 'CPL: A Core Language for Cloud Computing'. In: *Proceedings of the 13th International Conference on Modularity (MODULARITY 2016)*. Málaga, Spain: ACM Press, 2016, pp. 94–105. URL: <http://doi.acm.org/10.1145/2889443.2889452>.
- [40] F. Brasileiro, G. Silva, F. Araújo, M. Nóbrega, I. Silva and 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.
- [41] H. Bruneliere, Z. Al-Shara, F. Alvares, J. Lejeune and T. Ledoux. 'A Model-based Architecture for Autonomic and Heterogeneous Cloud Systems'. In: *CLOSER 2018 - 8th International Conference on Cloud Computing and Services Science*. Vol. 1. Best Paper Award. Funchal, Portugal, Mar. 2018, pp. 201–212. DOI: [10.5220/0006773002010212](https://doi.org/10.5220/0006773002010212). URL: <https://hal.archives-ouvertes.fr/hal-01705248>.
- [42] E. Bruneton, T. Coupaye, M. Leclercq, V. Quema and J.-B. Stefani. 'An Open Component Model and Its Support in Java'. In: *Component-Based Software Engineering*. Ed. by I. Crnkovic, J. A. Stafford, H. W. Schmidt and K. Wallnau. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 7–22.
- [43] P. Carbone, A. Katsifodimos, S. Ewen, V. Markl, S. Haridi and K. Tzoumas. 'Apache Flink?: Stream and Batch Processing in a Single Engine'. In: *IEEE Data Eng. Bull.* 38 (2015), pp. 28–38.

- [44] V. Cardellini, F. L. Presti, M. Nardelli and 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*. Vol. 10659. Springer. 2018, pp. 106–117.
- [45] F.-C. Chang, H.-C. Huang and 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*. Ed. by K. Aizawa, Y. Nakamura and S. Satoh. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 356–363.
- [46] R.-A. Cherrueau, R. Douence and 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, Aug. 2015. URL: <https://hal.inria.fr/hal-01168303>.
- [47] R.-A. Cherrueau, A. Lebre, D. Pertin, F. Wuhib and J. Soares. 'Edge Computing Resource Management System: a Critical Building Block! Initiating the debate via OpenStack'. In: *The USENIX Workshop on Hot Topics in Edge Computing (HotEdge'18)*. July 2018.
- [48] H. E. Chihoub, S. Ibrahim, Y. Li, G. Antoniu, M. S. Perez and 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)*. Oct. 2015, pp. 146–153.
- [49] B. Confais, A. Lebre and 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.
- [50] J. C. Corbett, J. Dean, M. Epstein, A. Fikes, C. Frost, J. J. Furman, S. Ghemawat, A. Gubarev, C. Heiser and P. Hochschild. 'Spanner: Google's globally distributed database'. In: *ACM Transactions on Computer Systems (TOCS)* 31.3 (2013), p. 8.
- [51] R. D. Cosmo, J. Mauro, S. Zacchiroli and G. Zavattaro. 'Aeolus: A component model for the cloud'. In: *Information and Computation* 239.Supplement C (2014), pp. 100–121. DOI: <https://doi.org/10.1016/j.ic.2014.11.002>. URL: <http://www.sciencedirect.com/science/article/pii/S0890540114001424>.
- [52] H. Coullon, J. Bigot and C. Pérez. 'Extensibility and Composability of a Multi-Stencil Domain Specific Framework'. In: *International Journal of Parallel Programming* (Nov. 2017). DOI: [10.1007/s10766-017-0539-5](https://doi.org/10.1007/s10766-017-0539-5). URL: <https://doi.org/10.1007/s10766-017-0539-5>.
- [53] H. Coullon, G. Le Louët and J.-M. Menaud. 'Virtual Machine Placement for Hybrid Cloud using Constraint Programming'. In: *ICPADS 2017*. 2017.
- [54] H. Coullon, D. Pertin and 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, Aug. 2017. URL: <https://hal.inria.fr/hal-01532489>.
- [55] I. Cuadrado-Cordero, A.-C. Orgerie and 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.
- [56] P.-C. David, T. Ledoux, M. Léger and T. Coupaye. 'FPath and FScript: Language support for navigation and reliable reconfiguration of Fractal architectures'. In: *annals of telecommunications - annales des télécommunications* 64.1 (Feb. 2009), pp. 45–63. DOI: [10.1007/s12243-008-0073-y](https://doi.org/10.1007/s12243-008-0073-y). URL: <https://doi.org/10.1007/s12243-008-0073-y>.
- [57] F. A. De Oliveira, T. Ledoux and R. Sharrock. 'A framework for the coordination of multiple autonomic managers in cloud environments'. In: *Self-Adaptive and Self-Organizing Systems (SASO), 2013 IEEE 7th International Conference on*. IEEE. 2013, pp. 179–188.
- [58] W. Felter, A. Ferreira, R. Rajamony and 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.

- [59] A. Flissi, J. Dubus, N. Dolet and 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://doi.org/10.1109/CCGRID.2008.59). URL: <https://hal.archives-ouvertes.fr/hal-00259836>.
- [60] P. Garcia Lopez, A. Montresor, D. Epema, A. Datta, T. Higashino, A. Iamnitchi, M. Barcellos, P. Felber and E. Riviere. 'Edge-centric Computing: Vision and Challenges'. In: *SIGCOMM Comput. Commun. Rev.* 45.5 (Sept. 2015), pp. 37–42. DOI: [10.1145/2831347.2831354](https://doi.org/10.1145/2831347.2831354). URL: <http://doi.acm.org/10.1145/2831347.2831354>.
- [61] Í. Goiri, W. Katsak, K. Le, T. D. Nguyen and R. Bianchini. 'Parasol and GreenSwitch: Managing Datacenters Powered by Renewable Energy'. In: *SIGARCH Comput. Archit. News* 41.1 (Mar. 2013), pp. 51–64.
- [62] Í. Goiri, K. Le, T. D. Nguyen, J. Guitart, J. Torres and R. Bianchini. 'GreenHadoop: Leveraging Green Energy in Data-processing Frameworks'. In: *Proceedings of the 7th ACM European Conference on Computer Systems*. EuroSys '12. Bern, Switzerland: ACM, 2012, pp. 57–70. DOI: [10.1145/2168836.2168843](https://doi.org/10.1145/2168836.2168843). URL: <http://doi.acm.org/10.1145/2168836.2168843>.
- [63] H. Goudarzi and M. Pedram. 'Geographical Load Balancing for Online Service Applications in Distributed Datacenters'. In: *IEEE international conference on cloud computing (CLOUD 2013)*. 2013.
- [64] L. Gu, D. Zeng, A. Barnawi, S. Guo and I. Stojmenovic. 'Optimal Task Placement with QoS Constraints in Geo-Distributed Data Centers Using DVFS'. In: *IEEE Transactions on Computers* 64.7 (July 2015), pp. 2049–2059. DOI: [10.1109/TC.2014.2349510](https://doi.org/10.1109/TC.2014.2349510). URL: <http://dx.doi.org/10.1109/TC.2014.2349510>.
- [65] K. Ha, Y. Abe, T. Eiszler, Z. Chen, W. Hu, B. Amos, R. Upadhyaya, P. Pillai and 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*. SEC '17. San Jose, California: ACM, 2017. DOI: [10.1145/3132211.3134453](https://doi.org/10.1145/3132211.3134453). URL: <http://doi.acm.org/10.1145/3132211.3134453>.
- [66] B. Han, V. Gopalakrishnan, L. Ji and S. Lee. 'Network function virtualization: Challenges and opportunities for innovations'. In: *IEEE Communications Magazine* 53.2 (Feb. 2015), pp. 90–97. DOI: [10.1109/MCOM.2015.7045396](https://doi.org/10.1109/MCOM.2015.7045396). URL: <http://dx.doi.org/10.1109/MCOM.2015.7045396>.
- [67] M. S. Hasan, F. Alvares, T. Ledoux and J. L. Pazat. 'Investigating Energy Consumption and Performance Trade-Off for Interactive Cloud Application'. In: *IEEE Transactions on Sustainable Computing* 2.2 (Apr. 2017), pp. 113–126. DOI: [10.1109/TSUSC.2017.2714959](https://doi.org/10.1109/TSUSC.2017.2714959). URL: <http://dx.doi.org/10.1109/TSUSC.2017.2714959>.
- [68] F. Hermenier, J. Lawall and G. Muller. 'Btrplace: A flexible consolidation manager for highly available applications'. In: *IEEE Transactions on dependable and Secure Computing* 10.5 (2013), pp. 273–286.
- [69] F. Hermenier, A. Lebre and 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)*. HPDC '10. Chicago, Illinois: ACM, 2010, pp. 658–666. DOI: [10.1145/1851476.1851574](https://doi.org/10.1145/1851476.1851574). URL: <http://doi.acm.org/10.1145/1851476.1851574>.
- [70] F. Hermenier, X. Lorca, J.-M. Menaud, G. Muller and J. Lawall. 'Entropy: a consolidation manager for clusters'. In: *Proceedings of the 2009 ACM SIGPLAN/SIGOPS international conference on Virtual execution environments*. ACM. 2009, pp. 41–50.
- [71] M. Hirzel, S. Schneider and B. Gedik. 'SPL: An Extensible Language for Distributed Stream Processing'. In: *toplas* 39.1 (Mar. 2017). URL: <http://doi.acm.org/10.1145/3039207>.
- [72] S. Ibrahim, B. He and H. Jin. 'Towards Pay-As-You-Consume Cloud Computing'. In: *2011 IEEE International Conference on Services Computing*. July 2011, pp. 370–377. DOI: [10.1109/SCC.2011.38](https://doi.org/10.1109/SCC.2011.38). URL: <http://dx.doi.org/10.1109/SCC.2011.38>.
- [73] S. Iftikhar, S. Khan, Z. Anwar et al. 'GenInfoGuard: A Robust and Distortion-Free Watermarking Technique for Genetic Data'. In: *PLOS ONE* 10.2 (Feb. 2015), pp. 1–22. DOI: [10.1371/journal.pone.0117717](https://doi.org/10.1371/journal.pone.0117717). URL: <https://doi.org/10.1371/journal.pone.0117717>.

- [74] B. Jennings and R. Stadler. ‘Resource Management in Clouds: Survey and Research Challenges’. In: *Journal of Network and Systems Management* 23.3 (July 2015), pp. 567–619. DOI: [10.1007/s10922-014-9307-7](https://doi.org/10.1007/s10922-014-9307-7). URL: <https://doi.org/10.1007/s10922-014-9307-7>.
- [75] E. Jonardi, M. A. Oxley, S. Pasricha, A. A. Maciejewski and H. J. Siegel. ‘Energy cost optimization for geographically distributed heterogeneous data centers’. In: *2015 Sixth International Green and Sustainable Computing Conference (IGSC)*. Dec. 2015, pp. 1–6. DOI: [10.1109/IGCC.2015.7393677](https://doi.org/10.1109/IGCC.2015.7393677). URL: <http://dx.doi.org/10.1109/IGCC.2015.7393677>.
- [76] H. Karau, A. Konwinski, P. Wendell and M. Zaharia. *Learning Spark*. O’Reilly Media, Feb. 2015.
- [77] J. Kramer and J. Magee. ‘The evolving philosophers problem: dynamic change management’. In: *IEEE Transactions on Software Engineering* 16.11 (Nov. 1990), pp. 1293–1306. DOI: [10.1109/32.60317](https://doi.org/10.1109/32.60317). URL: <http://dx.doi.org/10.1109/32.60317>.
- [78] A. Lebre, J. Pastor, A. Simonet and F. Desprez. ‘Revising OpenStack to Operate Fog/Edge Computing Infrastructures’. In: *The IEEE International Conference on Cloud Engineering (IC2E)*. Apr. 2017, pp. 138–148. DOI: [10.1109/IC2E.2017.35](https://doi.org/10.1109/IC2E.2017.35). URL: <http://dx.doi.org/10.1109/IC2E.2017.35>.
- [79] G. Malewicz, M. H. Austern, A. J. Bik, J. C. Dehnert, I. Horn, N. Leiser and G. Czajkowski. ‘Pregel: A System for Large-scale Graph Processing’. In: *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data*. SIGMOD ’10. Indianapolis, Indiana, USA: ACM, 2010, pp. 135–146. URL: [http://doi.acm.org/10.1145/1807167.1807184](https://doi.acm.org/10.1145/1807167.1807184).
- [80] T. L. Nguyen and 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.
- [81] K. Okamura and 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.
- [82] P. Oreizy, N. Medvidovic and R. N. Taylor. ‘Runtime Software Adaptation: Framework, Approaches, and Styles’. In: *Companion of the 30th International Conference on Software Engineering*. ICSE Companion ’08. Leipzig, Germany: ACM, 2008, pp. 899–910. DOI: [10.1145/1370175.1370181](https://doi.org/10.1145/1370175.1370181). URL: <http://doi.acm.org/10.1145/1370175.1370181>.
- [83] T. D. Phan, S. Ibrahim, G. Antoniu and L. Bouge. ‘On Understanding the Energy Impact of Speculative Execution in Hadoop’. In: *2015 IEEE International Conference on Data Science and Data Intensive Systems*. Dec. 2015, pp. 396–403.
- [84] F. Quesnel, A. Lebre and M. Südholt. ‘Cooperative and reactive scheduling in large-scale virtualized platforms with DVMS’. In: *Concurrency and Computation: Practice and Experience* 25.12 (2013), pp. 1643–1655.
- [85] D. Sabella, A. Vaillant, P. Kuure, U. Rauschenbach and F. Giust. ‘Mobile-Edge Computing Architecture: The role of MEC in the Internet of Things’. In: *IEEE Consumer Electronics Magazine* 5.4 (Oct. 2016), pp. 84–91. DOI: [10.1109/MCE.2016.2590118](https://doi.org/10.1109/MCE.2016.2590118). URL: <http://dx.doi.org/10.1109/MCE.2016.2590118>.
- [86] D. Serrano, S. Bouchenak, Y. Kouki, F. A. de Oliveira Jr., T. Ledoux, J. Lejeune, J. Sopena, L. Arantes and P. Sens. ‘SLA guarantees for cloud services’. In: *Future Generation Computer Systems* 54.Supplement C (2016), pp. 233–246. DOI: <https://doi.org/10.1016/j.future.2015.03.018>. URL: <http://www.sciencedirect.com/science/article/pii/S0167739X15000801>.
- [87] Q. Shen, X. Liang, X. S. Shen, X. Lin and 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* 18.2 (Mar. 2014), pp. 430–439. DOI: [10.1109/JBHI.2013.2292829](https://doi.org/10.1109/JBHI.2013.2292829). URL: <http://dx.doi.org/10.1109/JBHI.2013.2292829>.
- [88] W. Shi and T. F. Wenisch. ‘Energy-Efficient Data Centers’. In: *IEEE Internet Computing* 21.4 (2017), pp. 6–7. DOI: [10.1109/MIC.2017.2911429](https://doi.org/10.1109/MIC.2017.2911429). URL: <http://dx.doi.org/10.1109/MIC.2017.2911429>.

- [89] J.-B. 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. URL: https://doi.org/10.1007/978-3-319-15317-9_1.
- [90] C. Szyperski. *Component Software: Beyond Object-Oriented Programming*. 2nd. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 2002.
- [91] Y. Taleb, S. Ibrahim, G. Antoniu and 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. DOI: [10.1109/ICDCS.2017.51](https://doi.org/10.1109/ICDCS.2017.51). URL: <http://dx.doi.org/10.1109/ICDCS.2017.51>.
- [92] The Linux Foundation. *Open Network Automation Platform*. (Accessed: 2018-03-08). 2018. URL: <https://www.onap.org/>.
- [93] The OpenStack Foundation. *Cloud Edge Computing: Beyond the Data Center (White Paper)*. (Accessed: 2020-02-08). Jan. 2018. URL: <https://www.openstack.org/assets/edge/OpenStack-EdgeWhitepaper-v3-online.pdf>.
- [94] The OpenStack Foundation. *Edge Computing: Next Steps in Architecture, Design and Testing*. (Accessed: 2020-02-08). Jan. 2020. URL: <https://www.openstack.org/edge-computing/edge-computing-next-steps-in-architecture-design-and-testing>.
- [95] A. Thusoo, Z. Shao, S. Anthony, D. Borthakur, N. Jain, J. Sen Sarma, R. Murthy and H. Liu. ‘Data Warehousing and Analytics Infrastructure at Facebook’. In: *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data*. SIGMOD ’10. Indianapolis, Indiana, USA: ACM Press, 2010, pp. 1013–1020. URL: <http://doi.acm.org/10.1145/1807167.1807278>.
- [96] T. White. *Hadoop: The Definitive Guide*. 4th. O’Reilly Media, Apr. 2015.
- [97] A. Wierman, Z. Liu, I. Liu and H. Mohsenian-Rad. ‘Opportunities and challenges for data center demand response’. In: *International Green Computing Conference*. Nov. 2014, pp. 1–10. DOI: [10.1109/IGCC.2014.7039172](https://doi.org/10.1109/IGCC.2014.7039172). URL: <http://dx.doi.org/10.1109/IGCC.2014.7039172>.
- [98] J. Xiao, Z. Xu, H. Huang and 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.
- [99] F. Xu, F. Liu, H. Jin and 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 102.1* (2014), pp. 11–31.
- [100] O. Yildiz, M. Dorier, S. Ibrahim, R. Ross and 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. DOI: [10.1109/IPDPS.2016.50](https://doi.org/10.1109/IPDPS.2016.50). URL: <http://dx.doi.org/10.1109/IPDPS.2016.50>.
- [101] O. Yildiz, A. C. Zhou and 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)*. Sept. 2017, pp. 87–91. DOI: [10.1109/CLUSTER.2017.73](https://doi.org/10.1109/CLUSTER.2017.73). URL: <http://dx.doi.org/10.1109/CLUSTER.2017.73>.
- [102] M. Zaharia, T. Das, H. Li, T. Hunter, S. Shenker and I. Stoica. ‘Discretized Streams: Fault-tolerant Streaming Computation at Scale’. In: *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*. SOSP ’13. Farmington, Pennsylvania: ACM, 2013, pp. 423–438. DOI: [10.1145/2517349.2522737](https://doi.org/10.1145/2517349.2522737). URL: <http://doi.acm.org/10.1145/2517349.2522737>.