IN PARTNERSHIP WITH:
CNRS

Institut national des sciences appliquées de Rennes

Université Rennes 1

Activity Report 2016

Project-Team MYRIADS

Design and Implementation of Autonomous Distributed Systems

IN COLLABORATION WITH: Institut de recherche en informatique et systèmes aléatoires (IRISA)

RESEARCH CENTER
Rennes - Bretagne-Atlantique

THEME
Distributed Systems and middleware

Table of contents

Project-Team MYRIADS

Creation of the Team: 2010 January 01, updated into Project-Team: 2012 January 01

Keywords:

Computer Science and Digital Science:

1.1.2. - Hardware accelerators (GPGPU, FPGA, etc.)
1.1.4. - High performance computing
1.1.5. - Exascale
1.1.6. - Cloud
1.1.7. - Peer to peer
1.1.9. - Fault tolerant systems
1.1.13. - Virtualization
1.2. - Networks
1.2.4. - QoS, performance evaluation
1.2.5. - Internet of things
1.3. - Distributed Systems
1.6. - Green Computing
2.1.7. - Distributed programming
2.2.3. - Run-time systems
2.3.2. - Cyber-physical systems
2.4.2. - Model-checking
2.6. - Infrastructure software
2.6.1. - Operating systems
2.6.2. - Middleware
2.6.3. - Virtual machines
3.1.2. - Data management, quering and storage
3.1.3. - Distributed data
4.7. - Access control
4.9. - Security supervision
4.9.1. - Intrusion detection
4.9.3. - Reaction to attacks
5.6. - Virtual reality, augmented reality
6.1.3. - Discrete Modeling (multi-agent, people centered)
6.2.6. - Optimization
6.2.7. - High performance computing
7.1. - Parallel and distributed algorithms
7.3. - Optimization

Other Research Topics and Application Domains:

2.3. - Epidemiology
3.1. - Sustainable development
3.2. - Climate and meteorology
4.3. - Renewable energy production

4.4. - Energy delivery
4.4.1. - Smart grids
4.5. - Energy consumption
4.5.1. - Green computing
5.1. - Factory of the future
5.8. - Learning and training
6.1. - Software industry
6.1.1. - Software engineering
6.3. - Network functions
6.3.3. - Network Management
6.4. - Internet of things
6.5. - Information systems
6.6. - Embedded systems
8.1. - Smart building/home
8.2. - Connected city
8.5. - Smart society
9.1. - Education
9.1.1. - E-learning, MOOC
9.1.2. - Serious games
9.4.1. - Computer science
9.6. - Reproducibility
9.7. - Knowledge dissemination
9.7.1. - Open access
9.7.2. - Open data
9.8. - Privacy

1. Members

Research Scientists

Christine Morin [Team leader, Inria, Senior Researcher, HDR]
Deborah Agarwal [Inria, Senior Researcher, Inria International Chair]
Yvon Jégou [Inria, Researcher, until Sep. 2016]
Anne-Cécile Orgerie [CNRS, Researcher]

Faculty Members

Jean-Pierre Banâtre [Univ. Rennes I, Professor Emeritus]
Nikos Parlavantzas [INSA Rennes, Associate Professor]
Jean-Louis Pazat [INSA Rennes, Professor, HDR]
Guillaume Pierre [Univ. Rennes I, Professor]
Martin Quinson [ENS Rennes, Professor, HDR]
Cédric Tedeschi [Univ. Rennes I, Associate Professor]

Engineers

Anirvan Basu [Inria]
Julien Lefeuvre [Inria, granted by FP7 ECO2Clouds project, until Jul. 2016]
Manh Linh Pham [Inria, from Apr. 2016]
Sébastien Philippot [Inria]

Arnab Sinha [Inria, until Nov. 2016]
Alexandre Van Kempen [Inria]
David Margery [Inria, Research engineer]
Pascal Morillon [Univ. Rennes I]
Matthieu Simonin [Inria, Research engineer]

PhD Students

Arif Ahmed [Univ. Rennes I, from Oct. 2016]
Ismael Cuadrado Cordero [EMN, until Sep. 2016]
Anna Giannakou [Inria]
Baptiste Goupille-Lescar [Thales, granted by CIFRE]
David Guyon [Univ. Rennes I]
Sabbir Hasan Rochi [INSA Rennes]
Ancuta Iordache [Univ. Rennes I, until Jan. 2016]
Yunbo Li [EMN]
Edouard Outin [Inst. de Recherche Technologique B-COM, until Nov. 2016]
The Anh Pham [Inria, from Nov. 2016]
Genc Tato [Inria]
Amir Teshome Wonjiga [Inria]

Post-Doctoral Fellow

Benjamin Camus [Inria, from Oct. 2016]

Visiting Scientists

Carlos Ruiz Diaz [University of Guadalajara, Mexico, until Feb. 2016]
Palakiyem Wallah [University of Kara, Togo]

Administrative Assistant

Maryse Fouché [Inria]

Others

Simon Bihel [Inria, Intern, from May 2016 until Jul. 2016]
Jean Leon Cusinato [Inria, Intern, from Mar. 2016 until Aug. 2016]
Arnaud Fabre [Univ. Rennes I, Intern, from Feb. 2016 until Jul. 2016]
Yann Jegu [Inria, Intern, from May 2016 until Jul. 2016]
Fergal Martin-Tricot [Inria, Intern, from Jun. 2016 until Aug. 2016]
Louis Rilling [DGA, Researcher-Engineer, Associate Member]
Debaditya Roy [Univ. Rennes I, Intern, from May 2016 until Aug. 2016]
Kartik Sathyanarayanan [Inria, Intern, from May 2016 until Jul. 2016]
Simon Trebern [INSA Rennes, Intern, from Jun. 2016 until Jul. 2016]
Benjamin Trubert [Univ. Rennes I, Intern, from May 2016 until Aug. 2016]

2. Overall Objectives

2.1. General Objectives

MYRIADS is a joint team with INRIA, CNRS, UNIVERSITY RENNES 1, INSA RENNES and ENS RENNES. It is part of IRISA (D1 department on large scale systems) and INRIA RENNES – BRETAGNE ATLANTIQUE.

The objective of MYRIADS is to design and implement systems for autonomous service and resource management in interconnected and distributed clouds. The team tackles the challenges of dependable application execution and efficient resource management in highly distributed clouds.

2.2. Context

The MYRIADS team research activities are conducted in the context of the future of Internet.

Internet of Services. Myriads of applications are provided to more than one billion users [1] all over the world. Over time, these applications are becoming more and more sophisticated, a given application being a composition of services likely to be executed on various sites located in different geographical locations. The Internet of Services is spreading all domains: home, administration, business, industry and science. Everyone is involved in the Internet of Services: citizens, enterprises, scientists are application, service and resource consumers and/or providers over the Internet.

Outsourcing. Software is provided as a service over the Internet. Myriads of applications are available on-line to billions of users as, for instance, *GoogleApps* (Gmail). After decades in which companies used to host their entire IT infrastructures in-house, a major shift is occurring where these infrastructures are outsourced to external operators such as Data Centers and Computing Clouds. In the Internet of Services, not only software but also infrastructure are delivered as a service. Clouds turned computing and storage into a utility. Just like water or electricity, they are available in virtually infinite amounts and their consumption can be adapted within seconds like opening or closing a water tap. The main transition, however, is the change in business models. Companies or scientists do not need to buy and operate their own data centers anymore. Instead, the compute and storage resources are offered by companies on a “pay-as-you-go” basis. There is no more need for large hardware investments before starting a business. Even more, the new model allows users to adapt their resources within minutes, e.g., scale up to handle peak loads or rent large numbers of computers for a short experiment. The risk of wasting money by either under-utilization or undersized data centers is shifted from the user to the provider.

Sharing and Cooperation. Sharing information and cooperating over the Internet are also important user needs both in the private and the professional spheres. This is exemplified by various services that have been developed in the last decade. Peer-to-peer networks are extensively used by citizens in order to share musics and movies. A service like *Flickr* allowing individuals to share pictures is also very popular. Social networks such as *FaceBook* or *Linkedln* link millions of users who share various kinds of information within communities. Virtual organizations tightly connected to Grids allow scientists to share computing resources aggregated from different institutions (universities, computing centers...). The EGEE European Grid is an example of production Grid shared by thousands of scientists all over Europe.

2.3. Challenges

The term cloud was coined 10 years ago. Today cloud computing is widely adopted for a wide range of usage: information systems outsourcing, web service hosting, scientific computing, data analytics, back-end of mobile and IoT applications. There is a wide variety of cloud service providers (IaaS, PaaS, SaaS) resulting in difficulties for customers to select the services fitting their needs. Production clouds are powered by huge data centers that customers reach through the Internet. This current model raises a number of issues. Cloud computing generates a lot of traffic resulting in ISP providers needing to increase the network capacity. An increasing amount of always larger data centers consumes a lot of energy. Cloud customers experience poor quality of experience for highly interactive mobile applications as their requests are dealt with in data centers that are several hops away. The centralization of data in clouds also raises (i) security issues as clouds are a target of choice for attackers and (ii) privacy issues with data aggregation. Recently new cloud architectures have been proposed to overcome the scalability, latency, and energy issues of traditional centralized data centers. Various flavors of distributed cloud computing are emerging depending on the resources exploited: resources in the core network (distributed cloud), resources at the edge of the network (edge clouds) and even resources in the people swarms of devices (fog computing) enabling scalable cloud computing. These distributed clouds raise new challenges for resource and application management.

[1] According to World Stats, there are 3.67 billion Internet users i.e. more than half of the total world population in June 2016 http://www.internetworldstats.com/stats.htm.

The ultimate goal of Myriads team is making highly distributed clouds sustainable. By sustainability we mean green, efficient and secure clouds. We plan to study highly distributed clouds including edge clouds and fog computing. In this context, we will investigate novel techniques for greening clouds including the optimization of energy consumption in distributed clouds in the context of smart grids. As more and more critical information system are outsourced in the cloud and personal data captured by sensors embedded in smart objects and smartphones are stored in the cloud, we will investigate security and privacy issues in two directions: cloud security monitoring and personal data protection in cloud-based IoT applications.

System research requires experimental validation based on simulation and/or prototyping. Reproducible experimentation is essential. We will contribute to the design and implementation of simulators well suited to the study of distributed clouds (architecture, energy consumption) and of large scale experimentation platforms for distributed systems enabling reproducible experiments.

3. Research Program

3.1. Introduction

In this section, we present our research challenges along four work directions: resource and application management in distributed cloud architectures for scaling clouds in Section 3.2, energy management strategies for greening clouds in Section 3.3, security and data protection aspects for securing cloud-based information systems and applications in Section 3.4, and methods for experimenting with clouds in Section 3.5.

3.2. Scaling clouds

3.2.1. Resource management in hierarchical clouds

The next generation of utility computing appears to be an evolution from highly centralized clouds towards more decentralized platforms. Today, cloud computing platforms mostly rely on large data centers servicing a multitude of clients from the edge of the Internet. Servicing cloud clients in this manner suggests that locality patterns are ignored: wherever the client issues his/her request from, the request will have to go through the backbone of the Internet provider to the other side of the network where the data center relies. Besides this extra network traffic and this latency overhead that could be avoided, other common centralization drawbacks in this context stand in limitations in terms of security/legal issues and resilience.

At the same time, it appears that network backbones are over-provisioned for most of their usage. This advocates for placing computing resources directly within the backbone network. The general challenge of resource management for such clouds stands in trying to be locality-aware: for the needs of an application, several virtual machines may exchange data. Placing them *close* to each others can significantly improve the performance of the application they compose. More generally, building an overlay network which takes the hierarchical aspects of the platform without being a hierarchical overlay – which comes with load balancing and resilience issues is a challenge by itself.

The results of these works are planned to be integrated into the Discovery initiative [52] which aims at revisiting OpenStack to offer a cloud stack able to manage utility computing platforms where computing resources are located in small computing centers in the backbone's PoPs (Point of Presence) and interconnected through the backbone's internal links.

3.2.2. Resource management in mobile edge clouds

Mobile edge cloud (MEC) infrastructures are composed of compute, storage and networking resources located at the edge of wide-area networks, in immediate proximity to the end users. Instead of treating the mobile operator's network as a high-latency dumb pipe between the end users and the external service providers, MEC platforms aim at deploying cloud functionalities *within* the mobile phone network, inside or close to the mobile access points. Doing so is expected to deliver added value to the content providers and the end users by enabling new types of applications ranging from Internet-of-Things applications to extremely interactive systems (e.g., augmented reality). Simultaneously, it will generate extra revenue streams for the mobile network operators, by allowing them to position themselves as cloud computing operators and to rent their already-deployed infrastructure to content and application providers.

Mobile edge clouds have very different geographical distribution compared to traditional clouds. While traditional clouds are composed of many reliable and powerful machines located in a very small number of data centers and interconnected by very high-speed networks, mobile edge cloud are composed of a very large number of points-of-presence with a couple of weak and potentially unreliable servers, interconnected with each other by commodity long-distance networks. This creates new demands for the organization of a scalable mobile edge computing infrastructure, and opens new directions for research.

The main challenges that we plan to address are:

- How should an edge cloud infrastructure be designed such that it remains scalable, fault-tolerant, controllable, energy-efficient, etc.?
- How should applications making use of edge clouds be organized? One promising direction is to explore the extent to which stream-data processing platforms such as Apache Spark and Apache Flink can be adapted to become one of the main application programming paradigms in such environments.

3.2.3. Self-optimizing applications in multi-cloud environments

As the use of cloud computing becomes pervasive, the ability to deploy an application on a multi-cloud infrastructure becomes increasingly important. Potential benefits include avoiding dependence on a single vendor, taking advantage of lower resource prices or resource proximity, and enhancing application availability. Supporting multi-cloud application management involves two tasks. First, it involves selecting an initial multi-cloud application deployment that best satisfies application objectives and optimizes performance and cost. Second, it involves dynamically adapting the application deployment in order to react to changes in execution conditions, application objectives, cloud provider offerings, or resource prices. Handling price changes in particular is becoming increasingly complex. The reason is the growing trend of providers offering sophisticated, dynamic pricing models that allow buying and selling resources of finer granularities for shorter time durations with varying prices.

Although multi-cloud platforms are starting to emerge, these platforms impose a considerable amount of effort on developers and operations engineers, provide no support for dynamic pricing, and lack the responsiveness and scalability necessary for handling highly-distributed, dynamic applications with strict quality requirements. The goal of this work is to develop techniques and mechanisms for automating application management, enabling applications to cope with and take advantage of the dynamic, diverse, multi-cloud environment in which they operate.

The main challenges arising in this context are:

- selecting effective decision-making approaches for application adaptation,
- supporting scalable monitoring and adaptation across multiple clouds,
- performing adaptation actions in a cost-efficient and safe manner.

3.3. Greening clouds

ICT (Information and Communications Technologies) ecosystem now approaches 5% of world electricity consumption and this ICT energy use will continue grow fast because of the information appetite of Big Data, big networks and big infrastructures as Clouds that unavoidably leads to big power.

3.3.1. Smart grids and clouds

We propose exploiting Smart Grid technologies to come to the rescue of energy-hungry Clouds. Unlike in traditional electrical distribution networks, where power can only be moved and scheduled in very limited ways, Smart Grids dynamically and effectively adapt supply to demand and limit electricity losses (currently 10% of produced energy is lost during transmission and distribution).

For instance, when a user submits a Cloud request (such as a Google search for instance), it is routed to a data center that processes it, computes the answer and sends it back to the user. Google owns several data centers spread across the world and for performance reasons, the center answering the user's request is more likely to be the one closest to the user. However, this data center may be less energy efficient. This request may have consumed less energy, or a different kind of energy (renewable or not), if it had been sent to this further data center. In this case, the response time would have been increased but maybe not noticeably: a different trade-off between quality of service (QoS) and energy-efficiency could have been adopted.

While Clouds come naturally to the rescue of Smart Grids for dealing with this big data issue, little attention has been paid to the benefits that Smart Grids could bring to distributed Clouds. To our knowledge, no previous work has exploited the Smart Grids potential to obtain and control the energy consumption of entire Cloud infrastructures from underlying facilities such as air conditioning equipment (which accounts for 30% to 50% of a data center's electricity bill) to network resources (which are often operated by several actors) and to computing resources (with their heterogeneity and distribution across multiple data centers). We aim at taking advantage of the opportunity brought by the Smart Grids to exploit renewable energy availability and to optimize energy management in distributed Clouds.

3.3.2. Energy cost models

Cloud computing allows users to outsource the computer resources required for their applications instead of using a local installation. It offers on-demand access to the resources through the Internet with a pay-as-you-go pricing model.However, this model hides the electricity cost of running these infrastructures.

The costs of current data centers are mostly driven by their energy consumption (specifically by the air conditioning, computing and networking infrastructure). Yet, current pricing models are usually static and rarely consider the facilities' energy consumption per user. The challenge is to provide a fair and predictable model to attribute the overall energy costs per virtual machine and to increase energy-awareness of users.

Another goal consists in better understanding the energy consumption of computing and networking resources of Clouds in order to provide energy cost models for the entire infrastructure including incentivizing cost models for both Cloud providers and energy suppliers. These models will be based on experimental measurement campaigns on heterogeneous devices. Inferring a cost model from energy measurements is an arduous task since simple models are not convincing, as shown in our previous work. We aim at proposing and validating energy cost models for the heterogeneous Cloud infrastructures in one hand, and the energy distribution grid on the other hand. These models will be integrated into simulation frameworks in order to validate our energy-efficient algorithms at larger scale.

3.3.3. Energy-aware users

In a Cloud moderately loaded, some servers may be turned off when not used for energy saving purpose. Cloud providers can apply resource management strategies to favor idle servers. Some of the existing solutions propose mechanisms to optimize VM scheduling in the Cloud. A common solution is to consolidate the mapping of the VMs in the Cloud by grouping them in a fewer number of servers. The unused servers can then be turned off in order to lower the global electricity consumption.

Indeed, current work focuses on possible levers at the virtual machine suppliers and/or services. However, users are not involved in the choice of using these levers while significant energy savings could be achieved with their help. For example, they might agree to delay slightly the calculation of the response to their applications on the Cloud or accept that it is supported by a remote data center, to save energy or wait for the availability of renewable energy. The VMs are black boxes from the Cloud provider point of view. So, the user is the only one to know the applications running on her VMs.

We plan to explore possible collaborations between virtual machine suppliers, service providers and users of Clouds in order to provide users with ways of participating in the reduction of the Clouds energy consumption. This work will follow two directions: 1) to investigate compromises between power and performance/service quality that cloud providers can offer to their users and to propose them a variety of options adapted to their workload; and 2) to develop mechanisms for each layer of the Cloud software stack to provide users with a quantification of the energy consumed by each of their options as an incentive to become greener.

3.4. Securing clouds

3.4.1. Security monitoring SLO

While the trend for companies to outsource their information system in clouds is confirmed, the problem of securing an information system becomes more difficult. Indeed, in the case of infrastructure clouds, physical resources are shared between companies (also called tenants) but each tenant controls only parts of the shared resources, and, thanks to virtualization, the information system can be dynamically and automatically reconfigured with added or removed resources (for example starting or stopping virtual machines), or even moved between physical resources (for example using virtual machine migration). Partial control of shared resources brings new classes of attacks between tenants, and security monitoring mechanisms to detect such attacks are better placed out of the tenant-controlled virtual information systems, that is under control of the cloud provider. Dynamic and automatic reconfigurations of the information system make it unfeasible for a tenant's security administrator to setup the security monitoring components to detect attacks, and thus an automated self-adaptable security monitoring service is required.

Combining the two previous statements, there is a need for a dependable, automatic security monitoring service provided to tenants by the cloud provider. Our goal is to address the following challenges to design such a security monitoring service:

1. to define relevant Service-Level Objectives (SLOs) of a security monitoring service, that can figure in the Service-Level Agreement (SLA) signed between a cloud provider and a tenant;
2. to design heuristics to automatically configure provider-controlled security monitoring software components and devices so that SLOs are reached, even during automatic reconfigurations of tenants' information systems;
3. to design evaluation methods for tenants to check that SLOs are reached.

Moreover in challenges 2 and 3 the following sub-challenges must be addressed:

- although SLAs are bi-lateral contracts between the provider and each tenant, the implementation of the contracts is based on shared resources, and thus we must study methods to combine the SLOs;
- the designed methods should have a minimal impact on performance.

3.4.2. Data Protection in Cloud-based IoT Services

The Internet of Things is becoming a reality. Individuals have their own swarm of connected devices (e.g. smartphone, wearables, and home connected objects) continually collecting personal data. A novel generation of services is emerging exploiting data streams produced by the devices' sensors. People are deprived of control of their personal data as they don't know precisely what data are collected by service providers operating on Internet (oISP), for which purpose they could be used, for how long they are stored, and to whom they are disclosed. In response to privacy concerns the European Union has introduced, with the Global Data Protection Regulation (GDPR), new rules aimed at enforcing the people's rights to personal data protection. The GDPR also gives strong incentives to oISPs to comply. However, today, oISPs can't make their systems GDPR-compliant since they don't have the required technologies. We argue that a new generation of system is mandatory for enabling oISPs to conform to the GDPR. We plan to to design an open source distributed operating system for native implementation of new GDPR rules and ease the programming of compliant cloud-based IoT services. Among the new rules, transparency, right of erasure, and accountability are the most challenging ones to be implemented in IoT environments but could fundamentally increase people's confidence in oISPs. Deployed on individuals' swarms of devices and oISPs' cloud-hosted servers, it will enforce detailed data protection agreements and accountability of oISPs' data processing activities. Ultimately we will show to what extend the new GDPR rules can be implemented for cloud-based IoT services.

3.5. Experimenting with Clouds

Cloud platforms are challenging to evaluate and study with a sound scientific methodology. As with any distributed platform, it is very difficult to gather a global and precise view of the system state. Experiments are not reproducible by default since these systems are shared between several stakeholder. This is even worsen by the fact that microscopic differences in the experimental conditions can lead to drastic changes since typical Cloud applications continuously adapt their behavior to the system conditions.

3.5.1. Experimentation methodologies for clouds

We propose to combine two complementary experimental approaches: direct execution on testbeds such as Grid'5000, that are eminently believable but rather labor intensive, and simulations (using *e.g.* SimGrid) that are much more light-weighted, but requires are careful assessment. One specificity of the Myriads team is that we are working on these experimental methodologies *per se*, raising the standards of *good experiments* in our community.

We plan to make SimGrid widely usable beyond research laboratories, in order to evaluate industrial systems and to teach the future generations of cloud practitioners. This requires to frame the specific concepts of Cloud systems and platforms in actionable interfaces. The challenge is to make the framework both easy to use for simple studies in educational settings while modular and extensible to suit the specific needs of every advanced industrial-class users.

We aim at leveraging the convergence opportunities between methodologies by further bridging simulation and real testbeds. The predictions obtained from the simulator should be validated against some real-world experiments obtained on the target production platform, or on a similar platform. This (in)validation of the predicted results often improves the understanding of the modeled system. On the other side, it may even happen that the measured discrepancies are due to some mis-configuration of the real platform that would have been undetected without this (in)validation study. In that sense, the simulator constitutes a precious tool for the quality assurance of real testbeds such as Grid'5000.

Scientists need more help to make there Cloud experiments fully reproducible, in the sprit of Open Science exemplified by the HAL Open Archive, actively backed by Inria. Users still need practical solutions to archive, share and compare the whole experimental settings, including the raw data production (particularly in the case of real testbeds) and their statistical analysis. This is a long lasting task to which we plan to collaborate through the research communities gathered around the Grid'5000 and SimGrid scientific instruments.

Finally, since correction and performance can constitute contradictory goals, it is particularly important to study them jointly. To that extend, we want to bridge the performance studies, that constitute our main scientific heritage, to correction studies leveraging formal techniques. SimGrid already includes to exhaustively explore the possible executions. We plan to continue this work to ease the use of the relevant formal methods to the experimenter studying Cloud systems.

3.5.2. Use cases

In system research it is important to work on real-world use cases from which we extract requirements inspiring new research directions and with which we can validate the system services and mechanisms we propose. In the framework of our close collaboration with the Data Science Technology department of the LBNL, we will investigate cloud usage for scientific data management. Next-generation scientific discoveries are at the boundaries of datasets, e.g., across multiple science disciplines, institutions and spatial and temporal scales. Today, data integration processes and methods are largely adhoc or manual. A generalized resource infrastructure that integrates knowledge of the data and the processing tasks being performed by the user in the context of the data and resource lifecycle is needed. Clouds provide an important infrastructure platform that can be leveraged by including knowledge for distributed data integration.

4. Application Domains

4.1. Application Domains

The Myriads team investigates the design and implementation of system services. Thus its research activities address a broad range of application domains. We validate our research results with selected use cases in the following application domains:

- Web services, Service oriented applications,
- Business applications,
- Bio-informatics applications,
- Computational science applications,
- Data science applications,
- Numerical simulations,
- Energy and sustainable development,
- Smart cities.

5. Highlights of the Year

5.1. Highlights of the Year

- The PaaSage European project was successfully completed in November 2016 with an excellent rating from the reviewers. The PaaSage project developed a model-based, cross-cloud development and deployment platform that overcomes platform heterogeneity while enabling dynamic, fully-automated application scaling and cloud bursting. The main Myriads contribution is the Adapter subsystem, responsible for supporting dynamic, cross-cloud application adaptation.

5.1.1. Awards

- Baptiste Goupille-Lescar won the prize of the organizing committee of MMS Challenge 2016 (INSA Science Day).
- Anna Giannakou won the "Most Promising Experiment" award at the Grid'5000 winter school in February 2016 for her work "Towards Self Adaptable Security Monitoring in IaaS Clouds".

6. New Software and Platforms

6.1. ConPaaS

KEYWORDS: Cloud computing - PaaS
SCIENTIFIC DESCRIPTION ConPaaS [60] is a runtime environment for hosting applications in the cloud. It aims at offering the full power of the cloud to application developers while shielding them from the associated complexity of the cloud. ConPaaS is designed to host both high-performance scientific applications and online Web applications. It automates the entire life-cycle of an application, including collaborative development, deployment, performance monitoring, and automatic scaling. This allows developers to focus their attention on application-specific concerns rather than on cloud-specific details.
FUNCTIONAL DESCRIPTION

ConPaaS is a runtime environment for hosting applications in the cloud. It aims at offering the full power of the cloud to application developers while shielding them from the associated complexity of the cloud. ConPaaS is designed to host both high-performance scientific applications and online Web applications. It automates the entire life-cycle of an application, including collaborative development, deployment, performance monitoring, and automatic scaling. This allows developers to focus their attention on application-specific concerns rather than on cloud-specific details.

- Participants: Guillaume Pierre, Eliya Buyukkaya, Ancuta Iordache, Morteza Neishaboori, Dzenan Softic, Genc Tato and Teodor Crivat
- Contact: Guillaume Pierre
- URL: http://www.conpaas.eu/

6.2. GinFlow

KEYWORDS: Workflow - Distributed computing - Distributed - Distributed Applications - Framework
FUNCTIONAL DESCRIPTION GinFlow decentralizes the coordination of the execution of workflow-based applications. GinFlow relyies on an architecture where multiple service agents (SA) coordinate each others through a shared space containing the workflow description and current status. GinFlow allows the user to define several variants of a workflow and to switch from one to the other during run time.

- Participants: Matthieu Simonin, Cédric Tedeschi, Hector Fernandez, Javier Rojas Balderrama and Thierry Priol
- Partner: Université de Rennes 1
- Contact: Cédric Tedeschi
- URL: http://ginflow.inria.fr

6.3. Merkat

KEYWORDS: Resource management - Cloud - Elastic scaling - Market mechanisms - Service Level Objectives - HPC
FUNCTIONAL DESCRIPTION

Merkat is a platform that allows users of an organization to automatically manage and scale their applications while maximizing the infrastructure's utilization [12]. Merkat is generic and extensible, allowing users to automate the application deployment and management process. Users have the flexibility to control how many resources are allocated to their applications and to define their own resource demand adaptation policies. Merkat applies an unique approach to multiplex the infrastructure capacity between the applications, by implementing a proportional-share market and allowing applications to adapt autonomously to resource price and their given performance objectives. The price of the acquired resources acts as a control mechanism to ensure that resources are distributed to applications according to the user's value for them. Merkat was evaluated on Grid'5000 with several scientific applications.

- Participants: Stefania Costache, Christine Morin and Nikos Parlavantzas
- Contact: Nikos Parlavantzas
- URL: http://www.irisa.fr/myriads/software/Merkat/

6.4. Meryn

KEYWORDS: Resource management - PaaS - Cloud - Market mechanisms - Service Level Agreements
FUNCTIONAL DESCRIPTION

Meryn is an open, SLA-driven PaaS architecture that supports cloud bursting and allows hosting an extensible set of application types. Meryn relies on a decentralized optimization policy that aims at maximizing the overall provider profit, taking into account the penalties incurred when quality guarantees are unsatisfied. Meryn builds on the Snooze VM manager, and currently supports batch and MapReduce applications.

- Participants: Nikos Parlavantzas, Djawida Dib and Christine Morin
- Contact: Nikos Parlavantzas
- URL: http://www.irisa.fr/myriads/software/Meryn/

6.5. PaaSage Adapter

KEYWORDS: Cloud computing - Dynamic adaptation - Cloud applications management
FUNCTIONAL DESCRIPTION

The purpose of the Adapter is to transform the current configuration of a cloud application into a target configuration in an efficient and safe way. The Adapter is part of PaaSage, an open-source platform for modeling, deploying and executing applications on different clouds in an optimal manner. The Adapter has the following responsibilities: (1) validating reconfiguration plans, (2) applying the plans to the running system, and (3) maintaining an up-to-date representation of the current system state.

- Participants: Nikos Parlavantzas, Arnab Sinha, Manh Linh Pham, and Christine Morin
- Contact: Nikos Parlavantzas
- URL: https://team.inria.fr/myriads/software-and-platforms/paasage-adapter/

6.6. Resilin

KEYWORDS: Map Reduce - Parallel processing - Hadoop - Cloud - PaaS
FUNCTIONAL DESCRIPTION

Resilin is an open-source system for creating and managing MapReduce execution platforms over clouds. Resilin is compatible with the Amazon Elastic MapReduce (EMR) API, but it goes beyond Amazon's proprietary EMR solution in allowing users (e.g. companies, scientists) to leverage resources from one or more public and/or private clouds. This enables performing MapReduce computations over a large number of geographically-distributed and diverse resources. Resilin can be deployed across most of the open-source and commercial IaaS cloud management systems (e.g., OpenStack, OpenNebula, Amazon EC2). Once deployed, Resilin takes care of provisioning Hadoop clusters and submitting MapReduce jobs, allowing users to focus on writing their MapReduce applications rather than managing cloud resources. Resilin is implemented in the Python language and uses the Apache Libcloud library to interact with IaaS clouds. Resilin has been evaluated on multiple clusters of the Grid'5000 experimentation testbed. The results show that Resilin enables the use of geographically distributed resources with a limited impact on MapReduce job execution time.

- Participants: Ancuta Iodache, Christine Morin, Pierre Riteau, Nikos Parlavantzas and Matthieu Simonin
- Contact: Christine Morin
- URL: http://resilin.inria.fr

6.7. SAIDS

KEYWORDS: Cloud - Security
FUNCTIONAL DESCRIPTION SAIDS is a self-adaptable intrusion detection system for IaaS clouds. To maintain an effective level of intrusion detection, SAIDS monitors changes in the virtual infrastructure of a Cloud environment and reconfigures its components (security probes) accordingly. SAIDS can also reconfigure probes in the case of a change in the list of running services.

- Participants: Anna Giannakou, Jean-Léon Cusinato, Christine Morin, Jean-Louis Pazat, Louis Rilling and Fergal Martin-Tricot,
- Contact: Christine Morin
- URL: https://bil.inria.fr

6.8. SimGrid

KEYWORDS: Large-scale Emulators - Grid Computing - Distributed Applications
SCIENTIFIC DESCRIPTION SimGrid is a toolkit that provides core functionalities for the simulation of distributed applications in heterogeneous distributed environments. The simulation engine uses algorithmic and implementation techniques toward the fast simulation of large systems on a single machine. The models are theoretically grounded and experimentally validated. The results are reproducible, enabling better scientific practices.

Its models of networks, cpus and disks are adapted to (Data)Grids, P2P, Clouds, Clusters and HPC, allowing multi-domain studies. It can be used either to simulate algorithms and prototypes of applications, or to emulate real MPI applications through the virtualization of their communication, or to formally assess algorithms and applications that can run in the framework.

The formal verification module explores all possible message interleavings in the application, searching for states violating the provided properties. We recently added the ability to assess liveness properties over arbitrary and legacy codes, thanks to a system-level introspection tool that provides a finely detailed view of the running application to the model checker. This can for example be leveraged to verify both safety or liveness properties, on arbitrary MPI code written in C/C++/Fortran.

- Participants: Martin Quinson,
- Partners: Frédéric Suter, Arnaud Legrand, Adrien Lèbre, Luka Stanisic, Augustin Degomme.
- Contact: Martin Quinson
- URL: http://simgrid.gforge.inria.fr/

6.9. Snooze

KEYWORDS: Energy management - Fault-tolerance - Self-organization - Self-healing - Cloud computing - Consolidation - Virtualization
SCIENTIFIC DESCRIPTION

Snooze is a scalable, resilient and energy-aware virtual machine management framework for clouds. It is the result of Eugen Feller's PhD thesis which has been funded by the ANR EcoGrappe project (2008 - 2012).

The objectives of the Snooze ADT are threefold: (i) to distribute Snooze system as an open source software and to provide support to the user community (ii) to implement additional features to make it more user-friendly (iii) to integrate it with other open source software stacks and in public testbeds to favour its dissemination.

Snooze is a highly modular system for IaaS clouds. For the scientific community, it constitutes a unique framework for the experimentation of resource management policies in a real system. More generally, it allows any organization to operate a large-scale cluster as a resilient and energy-aware computing infrastructure enabling on demand provisioning of virtual clusters.
FUNCTIONAL DESCRIPTION

Snooze is a self-organizing and energy aware Cloud management framework.

Snooze is a novel Infrastructure-as-a-Service (IaaS) cloud-management system, which is designed to scale across many thousands of servers and virtual machines (VMs) while being easy to configure, highly available, and energy efficient. For scalability, Snooze performs distributed VM management based on a hierarchical architecture. To support ease of configuration and high availability Snooze implements self-configuring and self-healing features. Finally, for energy efficiency, Snooze integrates a holistic energy management approach via VM resource (i.e. CPU, memory, network) utilization monitoring, underload/overload detection and mitigation, VM consolidation (by implementing a modified version of the Sercon algorithm [59]), and power management to transition idle servers into a power saving mode. Snooze is a highly modular piece of software. It has been extensively evaluated on the Grid'5000 testbed using realistic applications.

- Participants: Eugen Feller, Christine Morin, Jiajun Cao, Gene Cooperman, Yvon Jégou, David Margery and Matthieu Simonin

- Contact: Christine Morin
- URL: http://snooze.inria.fr/

6.10. VEP

KEYWORDS: Cloud - Security - Computing - IaaS - Standards - OVF - CIMI - SLA
SCIENTIFIC DESCRIPTION

Virtual Execution Platform (VEP) is a Contrail service that sits just above IaaS layer at the service provider end of the Contrail cloud federation. The VEP service provides a uniform interface for managing the whole lifecycle of elastic applications on the cloud and hides the details of the IaaS layer to the user. VEP applications are described in OVF (Open Virtualization Format) standard format. Resource usage is controlled by CEE (Constrained Execution Environment) rules which can be derived from SLAs (Service Level Agreement). The VEP service integrates a monitoring system where the major events about the application, mainly resource usage, are made available to the user.

The VEP service provides a RESTful interface and can be exploited directly by users on top of the provider IaaS. OpenNebula and OpenStack IaaS frameworks were initially supported. During the VEP-S EIT ICT Labs activity in 2014, VEP was extended with a new OCCI IaaS driver which allows to control any IaaS framework providing a standard OCCI API. Support for the new OCCI SLA proposition from OGF has also been added and allows to represent the VEP CEEs in a standard format. Finally, during this activity, the Zabbix open source distributed monitoring system was integrated to VEP.
FUNCTIONAL DESCRIPTION VEP is a management tool for IaaS clouds with a REST interface and simple GUI for administrator. It is an extensible and reusable software for easy deployment of distributed applications. It provides advance reservation, pro-active fault tolerance. It is SLA aware and manages elasticity.

- Participants: Yvon Jégou, Roberto Gioacchino Cascella, Florian Dudouet, Filippo Gaudenzi, Christine Morin and Arnab Sinha
- Contact: Christine Morin
- URL: https://project.inria.fr/vep/

7. New Results

7.1. Scaling Clouds

7.1.1. Heterogeneous Resource Management

Participants: Baptiste Goupille-Lescar, Ancuta Iordache, Christine Morin, Manh Linh Pham, Nikos Parlavantzas, Guillaume Pierre, Arnab Sinha.

7.1.1.1. High performance in the cloud with FPGA virtualization

Participants: Ancuta Iordache, Guillaume Pierre.

Cloud platforms are becoming increasingly heterogeneous, with the availability of large numbers of virtual machine instance types as well as accelerator devices such as GPUs. In collaboration with Maxeler technologies, we have proposed a technique to virtualize FPGAs and make them available as first-class high-performance computation devices in the cloud [24]. The increasing variety of computation, storage and networking resources in the cloud is an opportunity for adjusting the provisioned resources to the individual needs of each application, but making an informed choice is extremely difficult. We therefore proposed application profiling techniques which can automatically identify the configuration which provides the best performance/cost tradeoff [49]. These two results were developed as part of the HARNESS European project, and they constitute Anca Iordache's PhD thesis [50]. FPGA virtualization is being further developed by Maxeler technologies toward commercial exploitation, and application profiling has been integrated in the open-source ConPaaS platform.

7.1.1.2. Multi-cloud application execution

Participants: Manh Linh Pham, Nikos Parlavantzas, Arnab Sinha.

Within the PaaSage European project, we improved and extended the Adapter subsystem, the part of the PaaSage platform that dynamically adapts the application deployment to changes in current runtime conditions [45]. Specifically, we added full support for causal connection between the running system and the runtime model and extended the plan validation functionality to use historical reconfiguration information. Moreover, we assisted industrial PaaSage partners with applying the PaaSage platform in diverse business scenarios.

7.1.1.3. Adaptive resource management for high-performance, multi-sensor systems

Participants: Baptiste Goupille-Lescar, Christine Morin, Nikos Parlavantzas.

In the context of our collaboration with Thales Research and Technology, we are applying cloud resource management techniques to high-performance, multi-sensor, embedded systems with real-time constraints. The objective is to increase the flexibility and efficiency of resource allocation in such systems, enabling the execution of dynamic sets of applications with strict QoS requirements. In 2016, we focused on characterising the targeted applications and platforms and developing a simulator in order to explore relevant resource management solutions. This work is performed in the context of Baptiste Goupille-Lescar's PhD work.

7.1.2. Distributed Cloud Computing

Participants: Nikos Parlavantzas, Jean-Louis Pazat, Guillaume Pierre, Genc Tato, Cédric Tedeschi, Alexandre Van Kempen.

7.1.2.1. Application self-optimization in multi-cloud environments

Participant: Nikos Parlavantzas.

Current approaches to application adaptation in multi-cloud environments are typically static, platform dependent, complex, and error prone. To address these limitations, we are combining the use of software product lines (SPLs) with models@run-time techniques. This work is performed in the context of the thesis of Carlos Ruiz Diaz, a PhD student at the University of Guadalajara, co-advised by Nikos Parlavantzas. The work focuses on the development of an SPL-based framework supporting initial cloud configuration as well as proactive, dynamic adaptation in a systematic, platform-independent way. The evaluation of this framework is currently in progress.

7.1.2.2. Edge clouds

Participants: Guillaume Pierre, Genc Tato, Cédric Tedeschi, Alexandre Van Kempen.

Mobile edge cloud computing aims to deploy cloud resources even closer to the end users, typically within mobile network access points. This is useful for hyper-interactive applications such as augmented reality which demand ultra-low network latencies (2-5 ms) between the end-user device and the cloud instances serving it. In contrast, current mobile networks exhibit network latencies in the order of 50-150 ms between the device and any cloud. We extended the ConPaaS open-source cloud platform to support the deployment of cloud applications in a distributed set of Raspberry Pi machines: instead of reaching the cloud through a wide-area network, in this setup each cloud node is also equipped with a wifi hotspot which allows local users to access it directly [53]. This work is ongoing, and a paper on this topic is currently being reviewed.

Getting closer to the edge user can be done through provisioning computing resources in Points of Presence (PoPs) within the telco's backbone network. The Discovery project [52] aims at revisiting the OpenStack Cloud stack to allow to disperse several smaller cloud facilities and connect them together to make them appear as a single Cloud entity. Genc Tato's PhD aims at proposing the building blocks on top of such an infrastructure to abstract out the network, route queries, store and retrieve objects (VMs and data). We have devised an overlay network to support such functionalities keeping in mind to maximise the laziness of the maintenance protocol to avoid any useless cost. A paper is being written on the subject.

7.1.2.3. Community Clouds

Participant: Jean-Louis Pazat.

Hosting services on an edge infrastructure based on devices owned and operated by end-users may be interesting for serving a community of users. However, these devices (such as internet boxes, disks or small computers have heterogeneous capabilities and no guaranteed availability. It is therefore challenging to ensure to the guest application a minimal hosting service level, like availability or Quality of Service. The management of the hosting service should adapt to the characteristics of the infrastructure. We are designing an architecture for a middleware capable of adapting the deployment of services on edge devices to ensure a given Quality of Service to access the service. While the middleware requires a minimal knowledge of the underlying infrastructure, its adaptation decisions are based on the feedbacks of users of the deployed service, like measured network latency. The environment relies on the use of micro-services which are composed to build the end-user services. This allows many adaptation strategies to adapt the system during run-time.

7.1.3. Scaling workflows with GinFlow

Participants: Matthieu Simonin, Cédric Tedeschi.

In 2016, we deployed GinFlow over 800 cores of the Grid'5000 platform, running Montage workflows comprising 118 tasks, and artificial workflows made of more than 3000 tasks. The ability of GinFlow to support adaptation and versioning of workflow with seamless transitions between workflow alternatives at runtime has been validated experimentally and presented on the Inria booth at SuperComputing in November 2016. These results have been presented at the IPDPS conference [32], and have been submitted to a journal special issue on workflows.

7.2. Greening Clouds

7.2.1. Energy Models

Participants: Yvon Jégou, Anne-Cécile Orgerie, Edouard Outin, Jean-Louis Pazat, Martin Quinson.

Simulating the impact of DVFS within SimGrid Simulation is a a popular approach for studying the performance of HPC applications in a variety of scenarios. However, simulators do not typically provide insights on the energy consumption of the simulated platforms. The goal of this ongoing work is to enable energy-aware experimentation within the SimGrid simulation toolkit, by introducing a model of energy consumption for computing applications making use of Dynamic Voltage and Frequency Scaling (DVFS) techniques.

Simulating Energy Consumption of Wired Networks In this work, we aim at simulating the energy consumption of wired networks which receive little attention in the Cloud computing community even though they represent key elements of these distributed architectures. To this end, we are contributing to the well-known open-source simulator ns3 by developing an energy consumption module named ECOFEN. This simulator embeds green levers: low power idle (IEEE 802.3az) and adaptive link rate. An article is currently under review on this topic.

Multicriteria scheduling for large-scale HPC environments Energy consumption is one of the main limiting factor for the design and deployment of large scale numerical infrastructures. The road towards "Sustainable Exascale" is a challenge with a target of 50 Gflops per watt. As platforms become more and more heterogeneous (co-processors, GPUs, low power processors...), an efficient scheduling of applications and services at large scale remains a challenge. In this context, we explore a multicriteria scheduling model and framework for large scale HPC systems. This work is done in collaboration with ROMA and Avalon teams from LIP in Lyon [29], [37].

Dynamic resource management for energy-efficiency The B-Com project, a joint private/public focusing on transfer, targets the design and the implementation of Watcher, a software module used to optimize an OpenStack cloud (in terms of performance, storage optimization or energy savings). This Software module is in the "Big Tent" software development process of OpenStack. In cooperation with Olivier Barais (Diverse Inria Team), we focus on dynamic management of cloud resources for energy-efficiency. Our approach relies on machine learning techniques, models@run-time and dynamic adaptation, and is intended to be included in Watcher. At regular intervals of time, we optimize the use of cloud resources by checking if a better placement of Virtual Machines on physical resources can be achieved, taking into account the migration cost. To achieve this, we have an energy model of the resources which is regularly updated using machine learning techniques that helps optimization algorithms to check if a better configuration can be reached energy-wise. This year we worked on the evaluation of the energy model [28].

7.2.2. Involving users in Energy Saving

Participants: Deborah Agarwal, Ismael Cuadrado Cordero, David Guyon, Christine Morin, Anne-Cécile Orgerie.

Energy-efficient cloud elasticity for data-driven applications Data centers hosting cloud systems consume enormous amounts of energy. Reducing this consumption becomes an urgent challenge with the rapid growth of cloud utilization. An existing solution to lower this consumption is to turn off as many servers as possible, but these solutions do not involve the user as a main lever to save energy. We introduce a system that proposes to the user to run her application with degraded performance in order to promote a better consolidation and thus to turn off more servers. Experimentation results using the Montage workflow show promising outcomes [47], [48]. We also performed a simulation-based evaluation on how much an energy-aware cloud system could save in energy consumed depending on the proportion of users selecting a green execution mode. These results based on the simulation of two typical daily uses of a data center running 3 real scientific applications will be published in Euromicro PDP 2017.

Energy-efficient and network-aware resource allocation in Cloud infrastructures The ever-growing appetite of new applications for network resources leads to an unprecedented electricity bill, and for these bandwidth-hungry applications, networks can become a significant bottleneck. Towards this end, we proposed microclouds, a fully autonomous energy-efficient subnetwork of clients of the same service, designed to keep the greenest path between its node. This semi-decentralized PaaS architecture for real-time multiple-users applications geographically distributes the computation among the clients of the cloud, moving the computation away from the datacenter to save energy - by shutting down or downgrading non utilized resources such as routers and switches, servers, etc. - and provides lower latencies for users. In this work, we have also analyzed the use of incentives for Mobile Clouds, and proposed a new auction system adapted to the high dynamism and heterogeneity of these systems [20], [19] [46].

7.2.3. Exploiting Renewable Energy in Datacenters

Participants: Sabbir Hasan Rochi, Yunbo Li, Anne-Cécile Orgerie, Jean-Louis Pazat.

Resource allocation in a Cloud partially powered by renewable energy sources We propose here to design a disruptive approach to Cloud resource management which takes advantage of renewable energy availability to perform opportunistic tasks. This Cloud receives a fixed amount of power from the regular electric Grid. This power allows it to run usual tasks. In addition, this Cloud is also connected to renewable energy sources (such as windmills or solar cells) and when these sources produce electricity, the Cloud can use it to run more tasks. The proposed resource management system integrates a prediction model to be able to forecast these extra-power periods of time in order to schedule more work during these periods. This work is done in collaboration with Ascola team from LINA in Nantes [44], [51][9].

Creating green-energy adaptivity awareness in SaaS application In addition to “green” resource allocation at the IaaS level in Datacenters, we think that users should be involved in “greening” their energy use (SaaS level). We propose that applications should have multiple “modes” of execution, each mode using a different level of energy and providing a different service level. For example, a B2C application may provide more or less recommandations . If this applcation can be dynamically swiched between these modes depending on the availability of green energy, the IaaS can optimize resource allocation better. To enforce his, we have designed green energy aware controllers.

This work is done in collaboration with Ascola team [23], [9].

7.3. Securing Clouds

7.3.1. Security monitoring in clouds

Participants: Jean Leon Cusinato, Anna Giannakou, Fergal Martin-Tricot, Christine Morin, Jean-Louis Pazat, Louis Rilling, Amir Teshome Wonjiga.

In the INDIC project we aim at making security monitoring a dependable service for IaaS cloud customers. To this end, we study three topics:

- defining relevant SLA terms for security monitoring,
- enforcing and verifying SLA terms,
- making the SLA terms enforcement mechanisms self-adaptable to cope with the dynamic nature of clouds.

The considered enforcement and verification mechanisms should have a minimal impact on performance.

In 2016 we improved the SAIDS approach, that we proposed in 2015, and that makes a network intrusion detection system (NIDS) deployed in a cloud operator infrastructure self-adaptable. In particular, we validated that the approach is generic enough to handle signature-based NIDSs (support for Snort and Suricata was implemented) as well as event-based NIDSs (support for Bro was implemented). An experimental evaluation of SAIDS has also been started in order to submit a full paper for publication in 2017. Jean-Léon Cusinato contributed to this work during his master internship.

We also improved the AL-SAFE approach, that we proposed in 2015, and that secures an application-level firewall by isolating it from the customer virtual machine and makes it self-adaptable [36], [35]. In particular, we validated that the self-adaptation architecture introduced for SAIDS could be reused to address firewalls, and the prototype was improved to implement stateful filtering. Fergal Martin-Tricot contributed to this work during his master internship. We also evaluated AL-SAFE experimentally on the prototype as well as analytically regarding the security correctness. The design and the evaluation of AL-SAFE were published in the CloudCom 2016 conference [21].

Regarding SLA definition and enforcement, in 2016 we have studied a verification method to enable a Cloud customer to verify that an NIDS located in the operator infrastructure is configured correctly according to the Service-Level Objectives (SLO) figuring in the SLA. A simple example of SLO is being used for this study, and further work should address more complete SLO regarding NIDSs. A prototype of the proposed verification method was implemented on OpenStack and Open vSwitch, and the NIDS software used is Snort. An evaluation of the verification method has been started and will include both experiments on the Grid'5000 platform and a correctness analysis. The design and evaluation of the verification method will be submitted in a full paper for publication in 2017.

7.3.2. Risk assessment in clouds

Participant: Christine Morin.

Attack graphs are leveraged in networks to exhibit the various scenarios available to compromise the system. They allow to uncover vulnerabilities chains exploitable by attackers based on network connectivity and vulnerabilities pre-requisites. In physical infrastructures, the acquisition of the topology has been vastly addressed in existing works with either passive or active discovery methods. Considering the Cloud context, in which virtualization attacks and virtual infrastructure dynamism are introduced, new methods need to be developed. We have designed a topology builder able to keep the topology and connectivity up to date in cloud environments. Based on the use of an IaaS cloud management system and a SDN (Software-Defined Networking) controller, our approach encompasses two steps: (i) when plugged into a running system, the topology builder retrieves the current topology and builds the associated connectivity: this represents the static topology and connectivity retrieval, in which we assume the network configuration to be fixed ; (ii) the topology builder listens to change events generated inside the infrastructure and within the SDN controller in order to update the topology and connectivity previously built: this represents the dynamic topology and connectivity retrieval. A prototype has been developed based on OpenStack cloud management system and

ONOS SDN open source technologies. This work is carried out in the context of Pernelle Mensah's PhD thesis and in collaboration with Nokia and CIDRE Inria project-team.

7.4. Experimenting with Clouds

7.4.1. Simulation

Participants: Simon Bihel, Martin Quinson.

Providing better interfaces to the users for Cloud Studies. Aware that the current user interface is a impediment to the adoption of our framework by the scientific community, we tried to propose a new, simplified API through the internship of Simon Bihel this summer. We identified several use cases and usage scenario that relevant to our context, and started implementing the new interface that we will provide. This work is still under progress.

Production-ready simulator of large-scale distributed systems. We are currently involved in a complete reorganization of the SimGrid implementation. The goal is two-fold: first we want reduce the tool's learning curve to help beginners. At the same time, we want to normalize the tool's internals so that power users can modify it and/or script the kernel behavior easily. Eventually, we are targeting usages in production and teaching contexts. This long term overhaul is still underway.

7.4.2. Experimentation Testbed

Participants: Anirvan Basu, Julien Lefeuvre, David Margery, Pascal Morillon.

Providing ready to use scripts to deploy popular and complex stacks. The study of complex software stacks on Grid'5000 has always been possible due to the reconfigurability properties of the testbed. Nevertheless, for newcomers with little background in system administration, automating the deployment of these stacks on Grid'5000 has always proved difficult. In 2016, we have provided scripts, that users can fork on github to customise to their needs, to deploy OpenStack, Ceph, Hadoop over Ceph or Sparkle. These have been presented to users during the 2016 winter school.

7.4.3. Use cases

Participants: Deborah Agarwal, Yvon Jégou, Nikos Parlavantzas, Manh Linh Pham, Christine Morin, Kartik Sathyanarayanan, Arnab Sinha.

7.4.3.1. Experimental Evaluation of Data Stream Processing Frameworks

We worked on evaluating data stream processing environments deployed in clouds. We compared the throughput, latency and energy consumption of Spark Streaming, Storm and Heron real-time data processing environments executed on top of Linux clusters and on top of virtual clusters deployed on top of the OpenStack IaaS cloud. The preliminary evaluation was conducted using the word count application on the twitter data stream. All experiments were conducted on Grid'5000 experimentation platform. The experimental results are described in a technical report to be published in 2017. This work was carried out by Kartik Sathyanarayanan, a student intern in Myriads team in the framework of DALHIS associate team.

7.4.3.2. Simulation framework for studying between-herd pathogen spread in a region

In our collaboration with Inra in the context of the Mihmes project, we worked on the design of decision tools to evaluate the epidemio-economic effectiveness of disease prevention and control strategies at the scales of the herd, the region and the supply chain. We developed a generic service-based framework to efficiently execute models of infection dynamics in a metapopulation of cattle herds on large-scale computing infrastructures. Our framework has been designed to execute complex regional models combining within-herds epidemiological models. The framework automatically distributes the simulation runs on multiple servers in a cluster and exploits the parallelism of the multicore servers. It relies on OpenMP for parallelizing simulation loops and deals with server heterogeneity and failures. We leveraged PaaSage software stack to deploy the framework on several IaaS clouds.

7.4.3.3. Mobile application for reliable collection of field data for Fluxnet

Critical to the interpretation of Fluxnet carbon flux data is the ancillary information and measurements taken at the tower sites. The submission and update of this data using excel sheets is difficult and error prone. In partnership with ICOS in the framework of DALHIS associate team, we are innovating the data submission and organization method through a responsive web User Interface able to run on desktop, mobile etc.; thus easing the data lookup and entry process from anywhere including the field sites. Continuing with our initial usability feedback experiences gathered last year on the application interface designs, we decided on the mobile application workflow for implementation. We developed a first prototype based on the PhoneGap [2] platform which provided the advantage of the same development code generating mobile application for IOS, Android and Windows platform simultaneously. The main functionality realized in the application prototype is that the user can download all the site data required by logging in through the application; and then view/edit them at the tower site (even in offline mode). The next logical step would be developing the synchronization and validation of data held locally in the application with the servers.

8. Bilateral Contracts and Grants with Industry

8.1. Bilateral Contracts with Industry

8.1.1. Technicolor (2016-2017)

Participant: Guillaume Pierre.

Our collaboration with Technicolor focuses on the design of a scalable and elastic virtual vistumer premises equipment based on Network Function Virtualization, Software-Defined Networking and Cloud technologies. In 2016 we completed the system design and an engineer from Technicolor started implementing the system. We expect to conduct further experiments and write a joint publication on this topic in 2017.

8.2. Bilateral Grants with Industry

8.2.1. Thales Research and Technology (2016-2018)

Participants: Baptiste Goupille-Lescar, Christine Morin, Nikos Parlavantzas.

Our collaboration with Thales Research and Technology focuses on the development of distributed Cyber-Physical Systems, such as those developed by Thales to monitor and react to changing physical environments. These systems need to be highly adaptable in order to cope with the dynamism and diversity of their operating environments. Notably, they require distributed, parallel architectures that support dynamic sets of applications, not known in advance, while providing strong QoS guarantees. The objective of this collaboration is to explore adaptive resource management mechanisms for such systems that can adapt to changes in the requirements and in the availability of resources. This contract funds Baptiste Goupille-Lescar's PhD grant.

8.2.2. Nokia (2015-2018)

Participant: Christine Morin.

Together with CIDRE Inria project-team we are involved in a collaboration with Nokia on security policy adaptation driven by risk evaluation in modern communication infrastructures. To address the need for efficient security supervision mechanisms, approaches such as attack graphs generation, coupled to a risk-based assessment have been used to provide an insight into a system's threat exposure. In comparison to static infrastructures, clouds exhibit a dynamic nature and are exposed to new attack scenarios due to virtualization. The goal of this collaboration is thus to revisit existing methods in the context of clouds. This contract funds Pernelle Mensah's PhD grant. Pernelle is a member of CIDRE project-team.

[2] http://phonegap.com/

9. Partnerships and Cooperations

9.1. Regional Initiatives

9.1.1. IRT B-Com

Participants: Yvon Jégou, Edouard Outin, Jean-Louis Pazat.

Yvon Jégou and Jean-Louis Pazat are at IRT B-Com [3] one day per week. With Édouard Outin, B-com PhD student, they contribute to the B-Com *Indeed* project, which aims at developing a distributed cloud software stack with a high degree of adaptability.

In the last period, they were involved in the elaboration of new projects in the Cloud Computing lab of B-Com.

9.1.2. CominLabs EPOC project (2013-2017)

Participants: Sabbir Hasan Rochi, Yunbo Li, Anne-Cécile Orgerie, Jean-Louis Pazat.

In this project, partners aim at focusing on energy-aware task execution from the hardware to application's components in the context of a mono-site data center (all resources are in the same physical location) which is connected to the regular electric Grid and to renewable energy sources (such as windmills or solar cells). In this context, we tackle three major challenges:

- Optimizing the energy consumption of distributed infrastructures and service compositions in the presence of ever more dynamic service applications and ever more stringent availability requirements for services.
- Designing a clever cloud's resource management which takes advantage of renewable energy availability to perform opportunistic tasks, then exploring the trade-off between energy saving and performance aspects in large-scale distributed systems.
- Investigating energy-aware optical ultra high-speed interconnection networks to exchange large volumes of data (VM memory and storage) over very short periods of time.

9.1.3. INDIC - Cybersecurity Pole of Excellence (2014-2018)

Participants: Anna Giannakou, Christine Morin, Jean-Louis Pazat, Louis Rilling, Amir Teshome Wonjiga.

Our study carried out in the framework of a collaboration with DGA-MI aims at defining and enforcing SLA for security monitoring of virtualized information systems. To this aim we study three topics:

- defining relevant SLA terms for security monitoring,
- enforcing and evaluating SLA terms,
- making the SLA terms enforcement mechanisms self-adaptable to cope with the dynamic nature of clouds.

The considered enforcement and evaluation mechanisms should have a minimal impact on performance. The funding from DGA funds two PhD students: Anna Giannakou and Amir Teshome Wonjiga.

9.2. National Initiatives

9.2.1. Inria ADT GinFlow (2014-2016)

Participants: Christine Morin, Matthieu Simonin, Cédric Tedeschi.

The GinFlow technological development action funded by INRIA targets the development of a fully-operational workflow management system based on the HOCL-TS software prototype developed during the PhD thesis of Hector Fernandez between 2009 and 2012. Also, it allows the integration of this software with the TIGRES workflow engine developed at the Lawrence Berkeley National Lab so as to make the workflows submitted using the TIGRES programming model run in a decentralized fashion. These developments led to the release of the GinFlow software and its deposit at the APP (Agence de Protection des Programmes).

[3] http://b-com.org/wp/

9.2.2. Inria ADT SaaP (2016-2018)

Participant: Martin Quinson.

The SaaP technological development action (SimGrid As A Platform) funded by INRIA targets the refactoring of SimGrid to make it ready to use in production and teaching contexts. Our ultimate goal is to sustain the development of the framework by involving 5 to 10 companies that are using it internally. Our target of the teaching context is thus an intermediate goal, as we think that the best solution to ensure the adoption of our tool by the industrial engineers is that they discovered the tool during their studies.

The technical actions envisioned for this ADT are the complete rearchitecturation of the software (to make it easier to script a new model within the tool kernel) and a reorganization of the interfaces (for a better integration in the Java and python language). This work is lead in collaboration with the whole SimGrid community, which provide valuable feedback.

9.2.3. Inria IPL Discovery (2015-2019)

Participants: Anne-Cécile Orgerie, Matthieu Simonin, Cédric Tedeschi.

The Inria IPL Discovery officially started in September 2015. It targets the design, development and deployment of a distributed Cloud infrastructure within the network's backbone. It will be based upon a set of building blocks whose design will take locality as a primary constraint, so as to minimize distant communications and consequently achieve better network traffic, partition management and improved availability.

Its developments are planned to get integrated within the OpenStack framework. An energy/cost benefit analysis of the fully distributed Discovery architecture will also be performed to show the energy efficiency of the chosen approach.

9.2.4. Inria IPL CityLab (2015-2018)

Participant: Christine Morin.

The Inria Project Lab (IPL) CityLab@Inria (https://citylab.inria.fr) studies ICT solutions toward smart cities that promote both social and environmental sustainability. A strong emphasis of the Lab is on the undertaking of a multi-disciplinary research program through the integration of relevant scientific and technology studies, from sensing up to analytics and advanced applications, so as to actually enact the foreseen smart city Systems of Systems. City-scale experiments of the proposed platforms and services are planned in cities in California and France, thereby learning lessons from diverse setups.

Myriads investigates advanced cloud solutions for the Future Internet, which are critical for the processing of urban data. It leverages its experience in cloud computing and Internet of services while expanding its research activities to the design and implementation of cloud services to support crowd-Xing applications and mobile social applications.

In 2016, Christine Morin was involved in the MOOC entitled *Villes Intelligentes : défis technologiques et sociétaux* (Smart cities : technological and social challenges) run on the FUN platform from January to March 2016. She prepared eight sequences on urban data management in clouds. In 2016, we also conducted a comparative experimental evaluation of data stream processing environments executed on clusters and in a cloud. We compared the performance and energy consumption of Heron, Storm and SparkStreaming frameworks.

9.2.5. Inria IPL Hac Specis (2016-2020)

Participants: Anne-Cécile Orgerie, Martin Quinson.

The goal of the HAC SPECIS (High-performance Application and Computers: Studying PErformance and Correctness In Simulation) project (http://hacspecis.gforge.inria.fr/) is to answer methodological needs of HPC application and runtime developers and to allow to study real HPC systems both from the correctness and performance point of view. To this end, we gather experts from the HPC, formal verification and performance evaluation community.

The Anh Pham started a PhD thesis in November 2016, co-advised by Thierry Jéron (team SUMO, formal methods) and Martin Quinson. The envisionned work will pursue the previous efforts to formally assess distributed applications within the SimGrid framework.

9.2.6. COSMIC PRE (2016 - 2018)

Participants: Benjamin Camus, Anne-Cécile Orgerie.

The distributed nature of Cloud infrastructures involves that their components are spread across wide areas, interconnected through different networks, and powered by diverse energy sources and providers, making overall energy monitoring and optimization challenging. The COSMIC project aims at taking advantage of the opportunity brought by the Smart Grids to exploit renewable energy availability and to optimize energy management in distributed Clouds. This PRE, led by Anne-Cécile Orgerie also involves Fanny Dufossé from Dolphin team (Inria Lille) and Benjamin Camus, who has started a 18 months post-doc in October 2016 in the context of this project.

9.2.7. MIHMES ANR Investissements d'Avenir (2012 - 2018)

Participants: Yvon Jégou, Christine Morin, Manh Linh Pham, Nikos Parlavantzas.

The MIMHES project (http://www.inra.fr/mihmes) led by INRA/BioEpAR aims at producing scientific knowledge and methods for the management of endemic infectious animal diseases and veterinary public health risks. Myriads team provides software tools to efficiently manage and ease the use of a distributed computing infrastructure for the execution of different simulation applications.

In 2016, we further developed a distributed framework which allows to exploit multiple compute servers in parallel. Parallelism is exploited both at server level using OpenMP and at data-center level using this framework. To facilitate the deployment of the workloads on heterogeneous environments, this framework limits the requirements concerning the server configurations. They need not share any file system, the workloads can be programmed in differing programming language. These servers need only the capability to communicate through the network. The system allows to dynamically add and stop servers. To some extend, it is tolerant to server failures. The framework had being repackaged to facilitate its reuse for new workloads. We also worked on the automated deployment of the framework on top of one or multiple IaaS clouds.

9.2.8. PIA ELCI (2015-2018)

Participant: Anne-Cécile Orgerie.

The PIA ELCI project deals with software environment for computation-intensive applications. It is leaded by BULL. In the context of this project, we collaborate with ROMA and Avalon teams from Lyon: we co-supervise a PhD student (Issam Rais) funded by this project with these teams on multicriteria scheduling for large-scale HPC environments.

9.2.9. CNRS PEPS EcoSmart (2016)

Participant: Anne-Cécile Orgerie.

Smart Grids are connected to telecommunication networks and can thus optimize the production, distribution and consumption of electricity. Virtualized distributed systems (Clouds) are the major players in providing services over the Internet. The success of these on-demand services makes the energy consumption of these systems worrying. This project aims to optimize the energy consumption of these large consumers, namely virtualized distributed Clouds consisting of computing, storage and communication resources. The objective is to exploit the capabilities offered by smart grids to control the consumption of these systems and be able to influence it according to the availability or the nature of the electricity used.

9.2.10. CNRS GDS EcoInfo

Participant: Anne-Cécile Orgerie.

The EcoInfo group deals with reducing environmental and societal impacts of Information and Communications Technologies from hardware to software aspects. This group aims at providing critical studies, lifecycle analyses and best practices in order to improve the energy efficiency of printers, servers, data centers, and any ICT equipment in use in public research organizations.

9.3. European Initiatives

9.3.1. FP7 & H2020 Projects

9.3.1.1. PaaSage

Title: PaaSage: Model Based Cloud Platform Upperware

Programm: FP7

Duration: October 2012 - September 2016

Coordinator: ERCIM

Partners:

Akademia Gorniczo-Hutnicza Im. Stanislawa Staszica W Krakowie (Poland)

Automotive Simulation Center Stuttgart Ev (Germany)

Be.Wan Sprl (Belgium)

Centred'Excellence en Technologies de l'Information et de la Communication (Belgium)

Geie Ercim (France)

Evry As (Norway)

Flexiant (United Kingdom)

Foundation for Research and Technology Hellas (Greece)

Gesellschaft Fur Wissenschaftliche Datenverarbeitung Mbh Gottingen (Germany)

Ibsac - Intelligent Business Solutions Ltd (Cyprus)

Inria (France)

Lufthansa Systems (Germany)

Stiftelsen Sintef (Norway)

Science and Technology Facilities Council (United Kingdom)

University of Cyprus (Cyprus)

Universitaet Stuttgart (Germany)

Universitaet Ulm (Germany)

Universitetet I Oslo (Norway)

Inria contact: Christine Morin

PaaSage (2012-2016) (http://www.paasage.eu) is an FP7 collaborative project that develops an open-source cloud platform, with an accompanying methodology and language, which enables developers to access cloud services in a technology-neutral approach while guiding them to configure their applications for best performance. PaaSage facilitates application deployment on multiple clouds while enhancing the flexibility, adaptivity and scalability of applications. Myriads develops the Adapter subsystem that supports dynamic, cross-cloud application adaptation. In 2016, we improved the Adapter implementation and evaluated its use within the business scenarios of PaaSage partners.

9.3.1.2. Fed4Fire

Title: Federation for FIRE

Programm: FP7

Duration: October 2012 - September 2016

Coordinator: Interdisciplinary institute for broadband technology (iMinds, Belgium)

Partners:

University of Southampton (It Innovation, United Kingdom)
Universite Pierre et Marie Curie - paris6 (UPMC, France)
Fraunhofer-Gesellschaft zur Foerderung der Angewandten Forschung e.v (Fraunhofer, Germany)
Technische Universitat Berlin (TUB, Germany)
The University of Edinburgh (UEDIN, United Kingdom)
National Ict Australia Limited (NICTA, Australia)
Atos Spain SA (Atos, Spain)
Panepistimio Thessalias (University of Thessaly) (UTH, Greece)
National Technical University of Athens (NTUA, Greece)
University of Bristol (UNIVBRIS, United Kingdom)
Fundacio Privada i2cat, Internet I Innovacio Digital a Catalunya (i2cat, Spain)
Eurescom-European Institute for Research and Strategic Studies in Telecommunications (EUR, Gmbh Germany)
Delivery of Advanced Network Technology to Europe limited (DANTE limited, United Kingdom)
Universidad de Cantabria (UC, Spain)
National Information Society agency (NIA, Korea (republic of))

Inria contact: Walid Dabbous

Fed4FIRE is an FP7 Integrated Project project running between October 2012 and September 2016 (http://www.fed4fire.eu), extended to December 2016. In Fed4FIRE, we investigate the means by which our experimental platforms (BonFIRE, and in a secondary way Grid'5000) could be made interoperable with a wider eco-system of experimental platforms in Europe and beyond. 2016 is the sustainability year of Fed4FIRE, and as usage from experimenters is not bringing any revenue, we closed the BonFIRE platform as it was become unmaintainable without significant effort.

9.3.2. Collaborations in European Programs, Except FP7 & H2020

9.3.2.1. NESUS

Participant: Anne-Cécile Orgerie.

Program: ICT COST
Project acronym: NESUS
Project title: Network for Sustainable Ultrascale Computing (ICT COST Action IC1305)
Duration: 2014 - 2018
Coordinator: Prof. Jesus Carretero, University Carlos III of Madrid, Spain, http://www.nesus.eu
Other partners: 33 COST countries and 11 non-COST countries

Abstract: Ultrascale systems are envisioned as large-scale complex systems joining parallel and distributed computing systems that will be two to three orders of magnitude larger that today's systems. The EU is already funding large scale computing systems research, but it is not coordinated across researchers, leading to duplications and inefficiencies. The goal of the NESUS Action is to establish an open European research network targeting sustainable solutions for ultrascale computing aiming at cross fertilization among HPC, large scale distributed systems, and big data management. The network will contribute to gluing disparate researchers working across different areas and provide a meeting ground for researchers in these separate areas to exchange ideas, to identify synergies, and to pursue common activities in research topics such as sustainable software solutions (applications and system software stack), data management, energy efficiency, and resilience. Some of the most active research groups of the world in this area are members of this proposal. This Action will increase the value of these groups at the European-level by reducing duplication of efforts and providing a more holistic view to all researchers, it will promote the leadership of Europe, and it will increase their impact on science, economy, and society. Anne-Cécile Orgerie is co-responsible of the focus group on metrics, monitoring, instrumentation and profiling in the Working Group 5 on Energy Efficiency.

9.3.3. Collaborations with Major European Organizations

Partner 1: EPFL, Network architecture lab (Switzerland)

We collaborate with Katerina Argyraki's research group on the integration of networking and cloud computing technologies in order to support placement constraints between cloud resources.

Partner 2: VU University Amsterdam, dept. of Computer Science (the Netherlands)

We collaborate with Thilo Kielmann's research group at VU University Amsterdam on research and development around the ConPaaS system.

Partner 3: University of Neuchâtel, dept. of Computer Science (Switzerland)

We collaborate with Pascal Felber's research group on energy efficiency in Clouds and in particular on the design of energy cost models for virtual machines.

9.4. International Initiatives

9.4.1. Inria International Labs

9.4.1.1. DALHIS

Title: Data Analysis on Large Heterogeneous Infrastructures for Science

International Partner (Institution - Laboratory - Researcher):

> Lawrence Berkeley National Laboratory (United States) - Data Science and Technology department - Deb Agarwal

Start year: 2016

See also: https://project.inria.fr/dalhis/

Data produced by scientific instruments (large facilities like telescopes or field data), large-scale experiments, and high-fidelity simulations are increasing in magnitude and complexity. Existing data analysis methods, tools and infrastructure are often difficult to use and unable to provide the complete data management, collaboration, and curation environment needed to manage these complex, dynamic, and large-scale data analysis environments. The goal of the Inria-LBL DALHIS associate team involving the Myriads (PI) and Avalon Inria project-teams and the Data Science and Technology (DST) department at Lawrence Berkeley National Laboratory (LBL) is to create a collaborative distributed software ecosystem to manage data lifecycle and enable data analytics on distributed data sets and resources. Specifically, our goal is to build a dynamic software stack that is user-friendly, scalable, energy-efficient and fault tolerant. Our research will determine appropriate execution environments that allow users to seamlessly execute their end-to-end dynamic data analysis workflows in various resource environments and scales while meeting energy-efficiency, performance and fault tolerance goals. We will engage in deep partnerships with scientific teams (Fluxnet in environmental science and SNFactory and LSST experiences in cosmology) and use a mix of user research with system software R&D to address specific challenges that these communities face. Our experience will in turn inform future research directions.

9.4.2. Inria International Partners

9.4.2.1. Informal International Partners

Partner: Rutgers University, dept. of Computer Science (New Jersey, United States)

We collaborate with Manish Parashar's research group on energy efficiency in edge Clouds and in particular on the design of energy cost models for such environments.

Partner: Northeastern University, dept. of Computer Science (Massachusetts, United States)

We collaborate with Gene Cooperman 's research group on virtualization technologies for the study of large-scale distributed systems.

Partner: University of Guadalajara (Mexico)

We collaborate with the team of Prof. H√©ctor Duran-Limon on application and resource management in the cloud. In 2016, we produced a joint journal publication [14]. Nikos Parlavantzas is co-advising a PhD student enrolled in the University of Guadalajara (Carlos Ruiz Diaz).

Partner: Tlemcen University (Algeria)

We collaborate with Djawida Dib on energy-efficient fault-tolerant resource and application management in containerized clouds. Christine Morin will co-advise a PhD student enrolled in the University of Tlemcen (Yasmina Bouizem) from December 2016.

9.5. International Research Visitors

9.5.1. Visits of International Scientists

Deb Agarwal, senior scientist at Lawrence Berkeley National Laboratory, who has been awarded an Inria International Chair for the 2015-2019 period, visited Myriads team during two months from May 1st to June 30th, 2016.

Christian Heinrich, PhD student in the Polaris team at Inria Grenoble, spent one month in October in the Myriads team to improve how large-scala distributed systems are declared in the SimGrid framework.

Professor Gene Cooperman, Northeastern University, Boston, USA, visited the Myriads team for one week in June to revive our collaboration on the virtualization of large-scale distributed systems.

Professor Peter Hubwieser, Technic University of Munchen, Germany, visited the Myriads team for two weeks in November to start a collaboration on the didactics of computer science with Martin Quinson.

Carlos Ruiz Diaz, PhD student in the University of Guadalajara, Mexico, spent 6 months in the Myriads team (from September 2015 to February 2016) to advance his work on adapting cloud configurations.

9.5.1.1. Internships

Benjamin Trubert

Date: May-August 2016

Institution: University of Rennes 1

Supervisor: Guillaume Pierre

Kartik Sathyanarayanan

Date: May-July 2016

Institution: Birla Institute of Technology & Science, Pilani (India)

Supervisor: Christine Morin

9.5.2. Visits to International Teams

9.5.2.1. Research Stays Abroad

- Anna Giannakou did a 3-month research internship in the Data Science and Technology department of the Lawrence Berkeley National Laboratory from August to October 2016. She worked with Sean Peisert, staff scientist, on building a workflow for anomaly Detection in HPC environments using statistical data.
- Yunbo Li did a 2 month research internship in the Computer Science department of Rutgers University from August to September 2016. He worked with Prof. Manish Parashar on building an energy cost model for edge cloud applications.

10. Dissemination

10.1. Promoting Scientific Activities

10.1.1. Scientific Events Organisation

10.1.1.1. Member of the Organizing Committees

- Anne-Cécile Orgerie was the main organizer of the research school E3-RSD on energy-efficiency in networks and distributed systems held in Dinard in June 2016 (http://e3rsd.irisa.fr).
- Louis Rilling was the main organizer of the SEC2 2016 workshop on security in Clouds held in Lorient in July 2016 in conjunction with the Compas 2016 conference (https://sec2-2016.inria.fr/).
- Guillaume Pierre was co-workshop/tutorials chair of the ACM/IFIP/USENIX Middleware 2016 conference.
- Martin Quinson was the main organizer of the research school SUD gathering the SimGrid user community held in Fréjus in January.
- Nikos Parlavantzas was a member of the organizing committee for MMS 2016 (INSA Science Day) in April 2016

10.1.2. Scientific Events Selection

10.1.2.1. Chair of Conference Program Committees

- Guillaume Pierre was the program chair of the Systems track of the Compas 2016 national conference.
- Guillaume Pierre was a program co-chair of CloudDP 2016 (6th International Workshop on Cloud Data and Platforms)

10.1.2.2. Member of the Conference Program Committees

- Christine Morin was a member of the program committee of the ACM/IEEE, SC'16, IEEE ICDCS 2016, ACM/IEEE CC-Grid 2016, IEEE CloudCom 2016, IEEE NFV-SDN 2016, Euromicro PDP 2016, ICCCN 2016 conferences and Resilience'16, ScienceCloud'16, CrossCloud'16 workshops.
- Anne-Cécile Orgerie was a member of the program committee of the PASA 2016 workshop and the Compas 2016 conference.
- Guillaume Pierre was a member of the program committees of IEEE IC2E 2016, ICPP 2016, DADS track of ACM SAC 2016, CloudCom 2016, and IEEE ScalCom 2016.
- Martin Quinson was a member of the program committees of SimulTech 2016.
- Nikos Parlavantzas was a member of the program committees of VHPC 16, and CLOSER 2016.
- Cédric Tedeschi was a member of the program committee of ICWS 2016, ICCS 2016, CloSer 2016 and Compas 2016 conferences. He was a member of the program committee of the ERROR 2016 workshop.

10.1.2.3. Reviewer

- Jean-Louis Pazat was a reviewer for the Compas Conference.
- Martin Quinson was a reviewer for the Compas Conference.

10.1.3. Journal

10.1.3.1. Member of the Editorial Boards

Christine Morin is associate editor in the IEEE's Transactions on Parallel and Distributed Systems' Editorial Board.

10.1.3.2. Reviewer - Reviewing Activities

- Cédric Tedeschi was a reviewer for the TPDS journal.

10.1.4. Invited Talks

- “Developing an Understanding of Ecosystems and Energy-Water Interactions Through Data“, keynote talk, Inria@Silicon Valley BIS workshop, Paris, June 8th, 2016.
- "Infection dynamics in a metapopulation of cattle herds: models and their execution in large-scale computing infrastructure"; Christine Morin (together with Pauline Ezanno), talk for Rencontres INRA-Inria, Pont Royal en Provence, October 3rd, 2016.
- “Le coût énergétique du stockage dans le Cloud”, Anne-Cécile Orgerie, talk for the scientific days of the University of Nantes, Nantes, June 10, 2016.
- “Green Computing and Sustainability”, Anne-Cécile Orgerie, Keynote at the Annual Workshop of the French section of URSI (International Union of Radio Science), Rennes, France, March 2016.
- “L’axe Green IT à l’IRISA”, Anne-Cécile Orgerie, talk for the HCERES evaluation of IRISA, Rennes, January 21, 2016.
- “Real-time stream data processing for hyper-interactive applications in fog computing platforms”, Guillaume Pierre, invited presentation at the US-EU Invited Workshop on the Next Generation Internet of Things, Los Angeles (USA), March 31st 2016.
- “Toward hyper-interactive applications in fog computing platforms”, Guillaume Pierre, invited presentation at the 9th Cloud control workshop, Friiberghs Herrgård (Sweden), 27-29 June 2016.
- “Experimental methodologies for large-scale distributed systems”, Martin Quinson, research school E3-RSD on energy-efficiency in networks and distributed systems held in Dinard in June 2016.

10.1.5. Leadership within the Scientific Community

- Christine Morin was invited to participate in the CRA-W career mentoring workshop held in Washington DC on November 19-20, 2016. Together with AJ Brush she organized a panel on "Growing Research/Surviving First 2 Years" and together with Lori Pollock she organized a panel on "Networking and Finding Advocates". She also did 1-1 mentoring for junior computer scientists (postdocs, assistant-professors, junior scientists and junior R&D engineers).
- Anne-Cécile Orgerie is co-responsible for the Green axis of the CNRS GDR RSD (Network and Distributed Systems working group).
- Anne-Cécile Orgerie is secretary of the ASF: the French chapter of ACM SIGOPS.
- Cédric Tedeschi is a member of the steering committee of the Compas conference.

10.1.6. Scientific Expertise

Yvon Jégou is a member of the **Comité de Sélection et de Validation** (CSV) of the *Images & Réseaux* cluster (until June 2016).

Christine Morin was a member of the HCERES evaluation committee of the LORIA (visit in Nancy December 13-15, 2016). She was a remote reviewer for the Call 2015 of CHIST-ERA on the topic User-Centric Security, Privacy and Trust in the Internet of Things. She was a member of the Jury for the Inria - Académie des Sciences prizes 2016.

Jean-Louis Pazat is the coordinator of experts in Information Technology for the evaluation of international bilateral collaborations at the ministry of research and education.

Martin Quinson is a member of the scientific committee of the Blaise Pascal foundation for the scientific dissemination of Maths and Computer Science, and also of the 2017 forum “Vivid Maths” (mathématiques vivantes) organized by the CFEM (French commission for the teaching of maths). He acted as an expert for a project of the UNIT foundation, the technical and engineering numerical university.

Nikos Parlavantzas acted as an expert reviewer for ANR and ANRT CIFRE projects.

10.1.7. Research Administration

- Yvon Jégou is a member of the Grid 5000 executive committee (until June 2016).
- Christine Morin is a member of the board of the Project-Team Committee of Inria Rennes Bretagne Atlantique.
- Christine Morin was the scientific coordinator for the Inria evaluation seminar of the "distributed system and middleware" theme, held in Rungis in October 2016.
- Christine Morin has been a member of the University of Rennes 1 board of directors and of the International Affairs Commission and its board since March 2016.
- Anne-Cécile Orgerie is officer (chargée de mission) for the IRISA cross-cutting axis on Green IT.
- Jean-Louis Pazat was the leader of the “Large Scale Systems” department of IRISA until October 2016, after which this responsibility was taken over by Martin Quinson.
- Jean-Louis Pazat was the leader of the IRISA-INSA computing lab until October 2016

10.2. Teaching - Supervision - Juries

10.2.1. Teaching

Christine Morin is responsible for the Internet of Services: Programming Models & Things and Clouds (ISI) teaching unit of the Master in research in Computer Science of the University of Rennes 1 and of the EIT Digital Master School at the University of Rennes 1.

Christine Morin:

- Master 1: FreeRoom project (spanning the two semesters) co-supervised with Cédric Tedeschi and Paul Couderc, University of Rennes 1, France.
- Master 2: 4.5 hours ETD, Energy management in Large Scale Distributed Systems, Centrale-Supelec, Rennes, France.
- Master 2: Internet of Services: Programming Models & Things and Clouds (ISI), 12 hours ETD, EIT Digital Master School, University of Rennes 1, France.

Anne-Cécile Orgerie:

- Licence 3: ARCSYS2: Architecture and System - 24 hours of lecture at ENS Rennes
- Master 1: Cloud & Big Data - 12 hours of lecture and 12 hours of practical sessions at ENS Rennes
- Master 2: Green ICT - 5 hours of invited lecture at Telecom SudParis
- PhD students: Energy issues in HPC and Clouds - 1.5 jours of invited lecture at the research school of the ETN BigStorage

Guillaume Pierre (at the University of Rennes 1):

- License 3: Systèmes (25 hours ETD)
- License 3: Organisation et utilisation des systèmes d’exploitation 2 (67 hours ETD)
- Master 2: Techniques de développement logiciel dans le Cloud (39 hours ETD)
- Master 1: Service Technologies (24 hours ETD)
- Master 2: Approche algorithmique des applications et systèmes répartis (32 hours ETD)
- DU D√©veloppeur-Concepteur Logiciel: Systèmes (42 hours ETD)

Martin Quinson (at ENS Rennes):

- Licence 3: Programming and Software Engineering (30 hours ETD); ARCSYS2 - architecture et syst√®mes 2 (60 hours ETD); Pedagogy (15 hours ETD).
- Agregation Science Industrielle: Programming and Software Engineering (20 hours ETD); Operating Systems and C programming (20 hours ETD); Networking (20 hours ETD).
- Master 2: Pedagogy and Scientific Mediation for Computer Science (30 hours EDT)

Jean-Louis Pazat (at INSA Rennes):

- Licence 3: Parallel and Multicore Programming (36 hours ETD)
- Master 1: Parallel and Distributed Programming (36 hours ETD)
- Master 1: Scientific Parallel Programming (36 hours ETD)
- Licence 3: Networks and SOA (20 hours ETD)

Nikos Parlavantzas (at INSA Rennes):

- Master 1: Performance Evaluation (32 hours ETD)
- Master 1: Operating Systems (36 hours ETD)
- Master 1: Big Data and Applications (15 hours ETD)
- Master 1: Networks and SOA (20 hours ETD)
- Master 1: Advanced Operating Systems (20 hours ETD)
- Master 1: Parallel programming (12 hours ETD)
- Master 1: Software Development Project (30 hours ETD)
- Licence 3: Networking (16 hours ETD)
- Master 2: Component-based Software Engineering (16 hours ETD)

Cédric Tedeschi (220 hours ETD at Univ. Rennes 1):

- Master 1: Cooperation and concurrency in Systems and Networks
- Master 1: Projects (in charge of the *UE*)
- Master 1: Parallel programming
- Master 2: Internet of Services and Infrastructures

10.2.2. Supervision

PhD in progress: Arif Ahmed, “Scalable Decentralized Edge Cloud Infrastructures”, started in October 2016, Guillaume Pierre.

PhD in progress: Ismael Cuadrado Cordero, Energy-efficient and network-aware resource allocation in Cloud infrastructures, started in October 2013, Christine Morin, Anne-Cécile Orgerie.

PhD in progress: Anna Giannakou, Self-adaptable security monitoring in Clouds, started in March 2014, Louis Rilling, Christine Morin, Jean-Louis Pazat.

PhD in progress: Baptiste Goupille-Lescar, Designing agile, distributed cyber-physical systems with advanced collaboration capabilities, started in January 2016, Michel Barreteau (Thales) Eric Lenormand (Thales), Christine Morin, Nikos Parlavantzas.

PhD in progress: David Guyon, Supporting energy-awareness for cloud users, started in September 2015, Anne-Cécile Orgerie, Christine Morin.

PhD in progress: Sabbir Hasan, SLA Driven Cloud autoscaling for optimizing energy footprint, started in December 2013, Thomas Ledoux (Ascola), Jean-Louis Pazat.

PhD: Ancuta Iordache, “Performance-cost tradeoffs in heterogeneous clouds”, University of Rennes 1, successfully defended on September 9th 2016. Guillaume Pierre.

PhD in progress: Yunbo Li, Resource allocation in a Cloud partially powered by renewable energy sources, started in October 2013, Anne-Cécile Orgerie, Jean-Marc Menaud (Ascola).

PhD in progress: Pernelle Mensah, Security policy adaptation driven by risk evaluation in modern communication infrastructures, started in December 2015, Samuel Dubus (Alcatel-Lucent), Christine Morin, Guillaume Piolle (Cidre), Eric Totel (Cidre).

PhD in progress: Édouard Outin, A multi-objective adaptation system for the management of a Distributed Cloud, started in October 2013, Olivier Barais (Triskell), Yvon Jégou, Jean-Louis Pazat.

PhD in progress: The Anh Pham, Dynamic Formal Verification of High Performance Runtimes and Applications, started in November 2016, Martin Quinson, Thierry Jéron.

PhD in progress: Issam Rais, Multi criteria scheduling for large scale High Performance Computing environments, started in October 2015, Anne-Cécile Orgerie, Anne Benoit (ROMA), Laurent Lefèvre (Avalon).

PhD in progress: Bruno Stevant, Resource allocation strategies for service distribution at the Internet edge to optimize end-to-end latency, started in December 2014, Jean-Louis Pazat.

PhD in progress: Genc Tato, Locality-aware Lazy Overlay Networks for WANS, started in December 2015, Marin Bertier, Cédric Tedeschi, Christine Morin.

PhD in progress: Amir Teshome, Definition and enforcement of Service-Level Agreements for Cloud security monitoring, started in October 2015, Louis Rilling, Christine Morin.

10.2.3. Juries

- Christine Morin was a reviewer in the PhD committee of Maxime Lorrillere, university Pierre et Marie Curie, February 4th, 2016.
- Christine Morin was a reviewer in the PhD committee of Richard Relaza, university of Toulouse, February 12th, 2016.
- Christine Morin was a member in the PhD committee of Anca Iordache, university of Rennes 1, September 9th, 2016.
- Christine Morin chaired the PhD committee of Ji Liu, university of Montpellier, November 3rd, 2016.
- Christine Morin was a member in the PhD committee of Daniel Balouek-Thomert, ENS Lyon, December 5th, 2016.
- Christine Morin chaired the PhD committee of Violaine Villebonnet, ENS Lyon, December 6th, 2016.
- Christine Morin was a reviewer in the PhD committee of Vincente Kherbache, university of Nice, December 7th, 2016.
- Anne-Cécile Orgerie was a member in the PhD committee of Frédéric Dumont, Ecole des Mines de Nantes, 21st of September 2016.
- Anne-Cécile Orgerie was a member in the PhD committee of Maxime Colmant, University of Lille, 24th of November 2016.
- Jean-Louis Pazat chaired the PhD committee of Pierre Obame Meye, University of Rennes 1, December 1st, 2016.
- Jean-Louis Pazat chaired the PhD committee of Fran√ßois Lehericey, INSA Rennes, September 20, 2016.
- Guillaume Pierre was a reviewer in the PhD committee of Luc André, University of Lorraine, May 13th 2016.
- Guillaume Pierre was a reviewer in the HDR committee of Christoph Neumann, University of Rennes 1, November 23rd 2016.
- Guillaume Pierre was a reviewer in the PhD committee of Vincent Kherbache, university of Nice, December 7th, 2016.
- Louis Rilling was a member in the PhD committee of Allan Blanchard, University of Orléans, 6th of December 2016.
- Martin Quinson was a reviewer in the PhD committee of Dimitri Pertin, University of Nantes, April 22th, 2016.
- Martin Quinson was a member in the PhD committee of David Beniamine, University of Grenoble, December 5th, 2016.

- Cédric Tedeschi was a member of the PhD committee of Sajith Kalathingal, University of Rennes 1, December 13th, 2016.

10.3. Popularization

- Martin Quinson is a scientific expert in a teaching manual of Computer Science for primary schools, authored within the "La main à la pâte" popularization network. This manual was released this summer.
- Martin Quinson co-founded a working group of the Société Informatique de France toward the creation of new unplugged activities to teach the informatics to pupils. This group met twice in Lyon this year, in April and November.
- Martin Quinson participated and co-organized to several workshops and open events where high-school students were invited to discover Computer Science (March at ENS Rennes, November on the Beaulieu campus). He was a scientific expert to the MathC2+ event at ENS-Rennes in June.
- Martin Quinson participated to the full day meeting “Splash Education” targeting at determining the fundamental programming concepts that should be taught to every pupils, co-located with the SPLASH conference http://2016.splashcon.org/track/splash-2016-splash-e. He organized the 3-days event on instrumented teaching of computer science, co-located with the Orphee-RDV event https://apprentissageinstrumentdelinformatique.wordpress.com/.
- Martin Quinson continued the development of the PLM web platform, which is an exerciser to teach programming to beginners. He also submitted several project applications to pursue this work in the future. Unfortunately, none of these applications have been accepted so far. Prof Peter Hubwieser (Technical University of Munchen, chair of didactics of Computer Science) visited us for two weeks in November. Developing the PLM and exploiting the data already gathered were central elements of this work meeting. A joint publication is currently prepared, targeting the ItiCSE'17 conference.

11. Bibliography

Major publications by the team in recent years

[1] S. Costache, S. Kortas, C. Morin, N. Parlavantzas. *Market-based Autonomous Resource and Application Management in Private Clouds*, in "Journal of Parallel and Distributed Computing", February 2017, vol. 100, pp. 85-102 [*DOI :* 10.1016/j.jpdc.2016.10.003], https://hal.archives-ouvertes.fr/hal-01378536

[2] E. Feller, L. Ramakrishnan, C. Morin. *Performance and energy efficiency of big data applications in cloud environments: A Hadoop case study*, in "Journal of Parallel and Distributed Computing", 2015, vol. 79-80, pp. 80-89 [*DOI :* 10.1016/j.jpdc.2015.01.001], https://hal.archives-ouvertes.fr/hal-01271141

[3] E. Feller, L. Rilling, C. Morin. *Energy-Aware Ant Colony Based Workload Placement in Clouds*, in "The 12th IEEE/ACM International Conference on Grid Computing (GRID-2011)", Lyon, France, September 2011, https://hal.inria.fr/inria-00626042

[4] A. Iordache, C. Morin, N. Parlavantzas, P. Riteau. *Resilin: Elastic MapReduce over Multiple Clouds*, Inria, October 2012, n^o RR-8081, https://hal.inria.fr/hal-00737030

[5] A. Iordache, G. Pierre, P. Sanders, J. G. d. F. Coutinho, M. Stillwell. *High Performance in the Cloud with FPGA Groups*, in "9th IEEE/ACM International Conference on Utility and Cloud Computing (UCC 2016)", Shanghai, China, December 2016, https://hal.inria.fr/hal-01356998

[6] A.-C. ORGERIE, M. DIAS DE ASUNCAO, L. LEFÈVRE. *A Survey on Techniques for Improving the Energy Efficiency of Large Scale Distributed Systems*, in "ACM Computing Surveys", December 2014, vol. 46, n° 4, To appear, https://hal.inria.fr/hal-00767582

[7] J. ROJAS BALDERRAMA, M. SIMONIN, C. TEDESCHI. *GinFlow: A Decentralised Adaptive Workflow Execution Manager*, in "30th IEEE International Parallel & Distributed Processing Symposium", Chicago, United States, Proceedings of the 30th IEEE International Parallel & Distributed Processing Symposium, May 2016, https://hal.archives-ouvertes.fr/hal-01292135

Publications of the year

Doctoral Dissertations and Habilitation Theses

[8] A. IORDACHE. *Performance-cost trade-offs in heterogeneous clouds*, Université Rennes 1, September 2016, https://tel.archives-ouvertes.fr/tel-01419975

Articles in International Peer-Reviewed Journals

[9] N. BELDICEANU, B. DUMAS FERIS, P. GRAVEY, S. HASAN, C. JARD, T. LEDOUX, Y. LI, D. LIME, G. MADI-WAMBA, J.-M. MENAUD, P. MOREL, M. MORVAN, M.-L. MOULINARD, A.-C. ORGERIE, J.-L. PAZAT, O. H. ROUX, A. SHARAIHA. *Towards energy-proportional Clouds partially powered by renewable energy*, in "Computing", January 2017, vol. 99, n° 1, 20 p. [*DOI :* 10.1007/S00607-016-0503-Z], https://hal.inria.fr/hal-01340318

[10] M. BERTIER, M. PERRIN, C. TEDESCHI. *On the Complexity of Concurrent Multiset Rewriting*, in "International Journal of Foundations of Computer Science", 2016, vol. 27, n° 1 [*DOI :* 10.1142/S0129054116500052], https://hal.inria.fr/hal-01326849

[11] E. CARON, A. K. DATTA, C. TEDESCHI, F. PETIT. *Self-Stabilizing Prefix Tree Based Overlay Networks*, in "International Journal of Foundations of Computer Science", 2016, vol. 27, n° 5, pp. 607–630 [*DOI :* 10.1142/S0129054116500192], http://hal.upmc.fr/hal-01347457

[12] S. COSTACHE, S. KORTAS, C. MORIN, N. PARLAVANTZAS. *Market-based Autonomous Resource and Application Management in Private Clouds*, in "Journal of Parallel and Distributed Computing", February 2017, vol. 100, pp. 85-102 [*DOI :* 10.1016/J.JPDC.2016.10.003], https://hal.archives-ouvertes.fr/hal-01378536

[13] M. DORIER, O. YILDIZ, S. IBRAHIM, A.-C. ORGERIE, G. ANTONIU. *On the energy footprint of I/O management in Exascale HPC systems*, in "Future Generation Computer Systems", March 2016, vol. 62, pp. 17–28 [*DOI :* 10.1016/J.FUTURE.2016.03.002], https://hal.inria.fr/hal-01330735

[14] H. A. DURAN-LIMON, J. FLORES-CONTRERAS, N. PARLAVANTZAS, M. ZHAO, A. MEULENERT-PEÑA. *Efficient execution of the WRF model and other HPC applications in the cloud*, in "Earth Science Informatics", March 2016 [*DOI :* 10.1007/S12145-016-0253-7], https://hal.archives-ouvertes.fr/hal-01288718

[15] H. FERNANDEZ, C. TEDESCHI, T. PRIOL. *A Chemistry-Inspired Workflow Management System for Decentralizing Workflow Execution*, in "IEEE Transactions on Services Computing", 2016, vol. 9, n° 2 [*DOI :* 10.1109/TSC.2013.27], https://hal.inria.fr/hal-01326878

[16] A. LAGE FREITAS, N. PARLAVANTZAS, J.-L. PAZAT. *Cloud resource management driven by profit augmentation*, in "Concurrency and Computation: Practice and Experience", June 2016 [*DOI :* 10.1002/CPE.3899], https://hal.archives-ouvertes.fr/hal-01337996

Invited Conferences

[17] A.-C. ORGERIE. *Green Computing and Sustainability*, in "Energie et radiosciences - Journées scientifiques URSI France", Rennes, France, March 2016, https://hal.inria.fr/hal-01356921

International Conferences with Proceedings

[18] N. CHERIERE, P. DONAT-BOUILLUD, S. IBRAHIM, M. SIMONIN. *On the Usability of Shortest Remaining Time First Policy in Shared Hadoop Clusters*, in "SAC 2016-The 31st ACM/SIGAPP Symposium on Applied Computing", Pisa, Italy, April 2016, https://hal.inria.fr/hal-01239341

[19] I. CUADRADO CORDERO, F. CUADRADO, C. PHILLIPS, A.-C. ORGERIE, C. MORIN. *Microcities: a Platform based on Microclouds for Neighborhood Services*, in "International Conference on Algorithms and Architectures for Parallel Processing (ICA3PP)", Granada, Spain, December 2016, https://hal.inria.fr/hal-01362532

[20] I. CUADRADO-CORDERO, A.-C. ORGERIE, C. MORIN. *Incentives for Mobile Cloud Environments through P2P Auctions*, in "CloudNet: IEEE International Conference on Cloud Networking", Pisa, Italy, October 2016, https://hal.inria.fr/hal-01355685

[21] A. GIANNAKOU, L. RILLING, J.-L. PAZAT, C. MORIN. *AL-SAFE: A Secure Self-Adaptable Application-Level Firewall for IaaS Clouds*, in "CloudCom2016-8th IEEE International Conference on Cloud Computing Technology and Science", luxembourg, Luxembourg, December 2016, https://hal.inria.fr/hal-01363540

[22] D. GUYON, A.-C. ORGERIE, C. MORIN, D. AGARWAL. *How Much Energy can Green HPC Cloud Users Save?*, in "PDP 2017 - 25th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing", Saint Petersbourg, Russia, March 2017, https://hal.inria.fr/hal-01439874

[23] M. S. HASAN, F. ALVARES DE OLIVEIRA, T. LEDOUX, J.-L. PAZAT. *Enabling Green Energy awareness in Interactive Cloud Application*, in "IEEE International Conference on Cloud Computing Technology and Science 2016", Luxembourg, Luxembourg, December 2016, https://hal.inria.fr/hal-01365230

[24] A. IORDACHE, G. PIERRE, P. SANDERS, J. G. D. F. COUTINHO, M. STILLWELL. *High Performance in the Cloud with FPGA Groups*, in "9th IEEE/ACM International Conference on Utility and Cloud Computing (UCC 2016)", Shanghai, China, December 2016, https://hal.inria.fr/hal-01356998

[25] M. KURPICZ, A.-C. ORGERIE, A. SOBE. *How much does a VM cost? Energy-proportional Accounting in VM-based Environments*, in "PDP: Euromicro International Conference on Parallel, Distributed, and Network-Based Processing", Heraklion, Greece, February 2016, 8 p. [*DOI :* 10.1109/PDP.2016.70], https://hal.inria.fr/hal-01276913

[26] Y. LI, A.-C. ORGERIE, J.-M. MENAUD. *Balancing the use of batteries and opportunistic scheduling policies for maximizing renewable energy consumption in a Cloud data center*, in "PDP 2017 - 25th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing", St Petersburg, Russia, March 2017, https://hal.inria.fr/hal-01432752

[27] A. Lèbre, A. Simonet, A.-C. Orgerie. *Deploying Distributed Cloud Infrastructures: Who and at What Cost?*, in "Proceedings of the fifth IEEE International Workshop on Cloud Computing Interclouds, Multiclouds, Federations, and Interoperability", Berlin, Germany, Intercloud 2016, April 2016, 6 p. [*DOI :* 10.1109/IC2EW.2016.48], https://hal.inria.fr/hal-01404594

[28] E. Outin, J.-E. Dartois, O. Barais, J.-L. Pazat. *Seeking for the Optimal Energy Modelisation Accuracy to Allow Efficient Datacenter Optimizations*, in "16th International Symposium on Cluster, Cloud and Grid Computing", Cartagena, Italy, IEEE/ACM 16th International Symposium on Cluster, Cloud and Grid Computing, May 2016 [*DOI :* 10.1109/CCGrid.2016.67], https://hal.archives-ouvertes.fr/hal-01356099

[29] I. Rais, L. Lefèvre, A. Benoit, A.-C. Orgerie. *An Analysis of the Feasibility of Energy Harvesting with Thermoelectric Generators on Petascale and Exascale Systems*, in "Workshop Optimization of Energy Efficient HPC & Distributed Systems (OPTIM 2016) - The 2016 International Conference on High Performance Computing & Simulation (HPCS 2016)", Innsbruck, Austria, Proceedings of the 2016 International Conference on High Performance Computing & Simulation (HPCS 2016), July 2016, https://hal.inria.fr/hal-01348554

[30] I. Rais, A.-C. Orgerie, M. Quinson. *Impact of Shutdown Techniques for Energy-Efficient Cloud Data Centers*, in "ICA3PP: 16th International Conference on Algorithms and Architectures for Parallel Processing", Granada, Spain, December 2016, pp. 203-210, https://hal.inria.fr/hal-01362530

[31] J. Rojas Balderrama, M. Simonin, C. Tedeschi. *Chemistry-Inspired Adaptive Stream Processing*, in "16th International Conference on Membrane Computing", Valencia, Spain, August 2016, https://hal.inria.fr/hal-01326884

[32] J. Rojas Balderrama, M. Simonin, C. Tedeschi. *GinFlow: A Decentralised Adaptive Workflow Execution Manager*, in "30th IEEE International Parallel & Distributed Processing Symposium", Chicago, United States, Proceedings of the 30th IEEE International Parallel & Distributed Processing Symposium, May 2016, https://hal.archives-ouvertes.fr/hal-01292135

[33] C. Ruiz, H. A. Duran-Limon, N. Parlavantzas. *Towards a software product line-based approach to adapt IaaS cloud configurations*, in "9th International Conference on Utility and Cloud Computing (UCC '16)", Shanghai, China, December 2016 [*DOI :* 10.1145/2996890.3007893], https://hal.archives-ouvertes.fr/hal-01433820

[34] A. Teshome, L. Rilling, C. Morin. *Including Security Monitoring in Cloud Service Level Agreements*, in "Symposium on Reliable Distributed Systems (SRDS), PhD forum", Budapest, Hungary, September 2016, https://hal.inria.fr/hal-01354975

Conferences without Proceedings

[35] A. Giannakou, L. Rilling, C. Morin, J.-L. Pazat. *AL-SAFE: A Secure Self-Adaptable Application-Level Firewall for IaaS Clouds*, in "SEC2 2016 - Second workshop on Security in Clouds", Lorient, France, July 2016, https://hal.inria.fr/hal-01340494

[36] A. Giannakou, L. Rilling, C. Morin, J.-L. Pazat. *How to Secure Application-Level Firewalls in IaaS Clouds*, in "CloudDP, 6th International Workshop on Cloud Data and Platforms", London, United Kingdom, Guillaume Pierre, Rui Oliveira , April 2016, https://hal.inria.fr/hal-01340473

[37] I. RAIS, A. BENOIT, L. LEFÈVRE, A.-C. ORGERIE. *An Analysis of the Feasibility of Energy Harvesting with Thermoelectric Generators on Petascale and Exascale Systems*, in "Conférence d'informatique en Parallélisme, Architecture et Système (COMPAS 2016)", Lorient, France, Actes de COMPAS, la Conférence d'informatique en Parallélisme, Architecture et Système, July 2016, https://hal.inria.fr/hal-01348555

[38] L. RILLING, C. MORIN. *INDIC : Supervision de sécurité à la carte pour les utilisateurs de cloud IaaS*, in "Rendez-vous de la Recherche et de l'Enseignement de la Sécurité des systèmes d'information (RESSI 2016)", Toulouse, France, LAAS-CNRS, INSA Toulouse, May 2016, https://hal.inria.fr/hal-01403774

[39] A. TESHOME, L. RILLING, C. MORIN. *Including Security Monitoring in Cloud SLA*, in "SEC2 2016 - Second workshop on Security in Clouds", Lorient, France, July 2016, https://hal.inria.fr/hal-01354974

Research Reports

[40] M. CALLAU-ZORI, L. SAMOILA, A.-C. ORGERIE, G. PIERRE. *An experiment-driven energy consumption model for virtual machine management systems*, IRISA ; Université de Rennes 1 ; CNRS, January 2016, n^{o} RR-8844, https://hal.inria.fr/hal-01258766

[41] R.-A. CHERRUEAU, A. LEBRE, D. PERTIN, A. SIMONET, M. SIMONIN. *ENOS: a HolisticFramework forConducting ScientificEvaluations of OpenStack*, Inria Rennes Bretagne Atlantique ; Nantes, November 2016, n^{o} RT-0485, https://hal.inria.fr/hal-01415522

[42] I. CUADRADO-CORDERO, F. CUADRADO, C. PHILLIPS, A.-C. ORGERIE, C. MORIN. *Microcities: a Platform based on Microclouds for Neighborhood Services*, inria, February 2016, n^{o} RR-8885, https://hal.inria.fr/hal-01291536

Other Publications

[43] A. DEGOMME, A. LEGRAND, M. S. MARKOMANOLIS, M. QUINSON, M. S. STILLWELL, F. S. SUTER. *Simulating MPI applications: the SMPI approach*, November 2016, Under review in IEEE TPDS, https://hal.inria.fr/hal-01415484

References in notes

[44] N. BELDICEANU, B. DUMAS FERIS, P. GRAVEY, S. HASAN, C. JARD, T. LEDOUX, Y. LI, D. LIME, G. MADI-WAMBA, J.-M. MENAUD, P. MOREL, M. MORVAN, M.-L. MOULINARD, A.-C. ORGERIE, J.-L. PAZAT, O. ROUX, A. SHARAIHA. *The EPOC project: Energy Proportional and Opportunistic Computing system*, in "International Conference on Smart Cities and Green ICT Systems (SMARTGREENS)", Lisbonne, Portugal, May 2015, 15397, https://hal.archives-ouvertes.fr/hal-01131602

[45] A. BSILA, N. FERRY, G. HORN, T. KIRKHAM, M. MALAWSKI, N. PARLAVANTZAS, C. PÉREZ, J. ROUZAUD-CORNABAS, D. ROMERO, A. ROSSINI, A. SOLBERG, H. SONG. *D3.1.1 - Upperware Prototype*, March 2014, https://hal.inria.fr/hal-00991340

[46] I. CUADRADO CORDERO, A.-C. ORGERIE, C. MORIN. *GRaNADA: A Network-Aware and Energy-Efficient PaaS Cloud Architecture*, in "IEEE International Conference on Green Computing and Communications (GreenCom)", Sydney, Australia, December 2015, https://hal.archives-ouvertes.fr/hal-01205905

[47] D. GUYON. *Energy-efficient Cloud Elasticity for Data-driven Applications*, Université de Rennes 1, June 2015, 46 p. , https://hal.inria.fr/hal-01212524

[48] D. GUYON, A.-C. ORGERIE, C. MORIN. *Energy-efficient User-oriented Cloud Elasticity for Data-driven Applications*, in "IEEE International Conference on Green Computing and Communications (GreenCom)", Sydney, Australia, December 2015, https://hal.archives-ouvertes.fr/hal-01205915

[49] A. IORDACHE, E. BUYUKKAYA, G. PIERRE. *Heterogeneous Resource Selection for Arbitrary HPC Applications in the Cloud*, in "Distributed Applications and Interoperable Systems", Grenoble, France, June 2015 [*DOI :* 10.1007/978-3-319-19129-4_9], https://hal.inria.fr/hal-01159024

[50] A. IORDACHE. *Performance-Cost Trade-Offs in Heterogeneous Clouds*, Université de Rennes 1, France, September 2016, https://hal.inria.fr/tel-01419975

[51] Y. LI, A.-C. ORGERIE, J.-M. MENAUD. *Opportunistic Scheduling in Clouds Partially Powered by Green Energy*, in "IEEE International Conference on Green Computing and Communications (GreenCom)", Sydney, Australia, December 2015, https://hal.archives-ouvertes.fr/hal-01205911

[52] A. LÈBRE, J. PASTOR, M. BERTIER, F. DESPREZ, J. ROUZAUD-CORNABAS, C. TEDESCHI, P. ANEDDA, G. ZANETTI, R. NOU, T. CORTES, E. RIVIÈRE, T. ROPARS. *Beyond The Cloud, How Should Next Generation Utility Computing Infrastructures Be Designed?*, Inria, July 2013, n^o RR-8348, https://hal.inria.fr/hal-00854204

[53] K. RAZAVI, A. ION, G. TATO, K. JEONG, R. FIGUEIREDO, G. PIERRE, T. KIELMANN. *Kangaroo: A Tenant-Centric Software-Defined Cloud Infrastructure*, in "Proceedings of the IEEE International Conference on Cloud Engineering", Tempe, AZ, USA, United States, March 2015, https://hal.inria.fr/hal-01100774