Activity Report 2015

Project-Team MYRIADS

Design and Implementation of Autonomous Distributed Systems

Table of contents

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

Keywords:

Computer Science and Digital Science:

      1.1.13. - Virtualization
      1.1.6. - Cloud
      1.1.9. - Fault tolerant systems
      1.2.5. - Internet of things
      1.3. - Distributed Systems
      1.4. - Ubiquitous Systems
      1.6. - Green Computing
      2.2.3. - Run-time systems
      2.6.1. - Operating systems
      2.6.2. - Middleware
      4.9.1. - Intrusion detection
      6.2. - Scientific Computing, Numerical Analysis & Optimization
      7.1. - Parallel and distributed algorithms
      7.11. - Performance evaluation

Other Research Topics and Application Domains:

      3.1.1. - Resource management
      3.2. - Climate and meteorology
      4.4.1. - Green computing
      6.3.1. - Web
      6.4. - Internet of things
      8.2. - Connected city
      9.6. - Reproducibility

1. Members

Research Scientists

Deborah Agarwal [Lawrence Berkeley National Laboratory, Berkeley, USA, Senior Scientist, Inria International Chair]
Yvon Jégou [Inria, Researcher]
Christine Morin [Team leader, Inria, Senior Researcher, HdR]
Anne-Cécile Orgerie [CNRS, Researcher]
Thierry Priol [Inria, Senior Researcher, part-time, HdR]
Louis Rilling [DGA-MI, Associated team member]

Faculty Members

Jean-Pierre Banâtre [Univ. Rennes I, Professor Emeritus]
Nikolaos Parlavantzas [INSA Rennes, Associate Professor]
Jean-Louis Pazat [INSA Rennes, Professor, HdR]
Guillaume Pierre [Univ. Rennes I, Professor]
Martin Quinson [Normale Sup Rennes, Professor, from September 2015, HdR]

Cédric Tedeschi [Univ. Rennes I, Associate Professor]

Engineers

David Margery [Inria]

Pascal Morillon [Univ. Rennes I]

Matthieu Simonin [Inria]

Anirvan Basu [Inria]

Teodor Crivat [Univ. Rennes I]

Julien Lefeuvre [Inria, granted by FP7 ECO2Clouds project]

Younes Ouchtouban [Inria]

Aboozar Rajabi [Inria, until Mar 2015]

Arnab Sinha [Inria]

Genc Tato [Univ. Rennes I, until Nov 2015]

PhD Students

Ismael Cuadrado Cordero [Inria]

Anna Giannakou [Inria]

Baptiste Goupille–lescar [Inria, from Nov 2015]

David Guyon [Univ. Rennes I, from Sep 2015]

Sabbir Hasan Rochi [INSA de Rennes]

Ancuta Iordache [Univ. Rennes I]

Yunbo Li [Ecole des Mines de Nantes]

Edouard Outin [Inst. de Recherche Technologique B-COM]

Genc Tato [Inria, from Dec 2015]

Amir Teshome Wonjiga [Inria, from Oct 2015]

Post-Doctoral Fellows

Maria Del Mar Callau Zori [Univ. Rennes I, until Oct 2015]

Javier Rojas Balderrama [Inria, until Aug 2015]

Damian Serrano [Univ. Rennes I, until Sep 2015]

Visiting Scientists

Gene Cooperman [Northeastern University, June 2015]

Georgios Ioannidis [Ecole Polytechnique Fédérale de Lausanne, Switzerland, from June 2015 until Dec 2015]

Carlos Ruiz Diaz [University of Guadalajara, Mexico, from Sep 2015]

Anita Sobe [University of Neuchâtel, Switzerland, Apr 2015]

Palakiyem Wallah [University of Kara, Togo, from Sep 2015 until December 2015]

Administrative Assistant

Maryse Fouché [Inria]

Others

Philippe Fabian [Inria, from May 2015 until Aug 2015]

David Guyon [Univ. Rennes I, M2 intern, from Feb 2015 until June 2015]

Timothée Haudebourg [Inria, L3 intern, from May 2015 until Jul 2015]

Vlad Mirel [Univ. Rennes I, M1 intern, from May 2015 until Jul 2015]

Ghada Moualla [Inria, M2 intern, from Feb 2015 until June 2015]

Akshat Puri [Inria, M2 intern, from Feb 2015 until June 2015]

Benjamin Soulas [Inria, M2 intern, from March 2015 until Aug 2015]

Lodewijck Vogelzang [Univ. Rennes I, M1 intern, from May 2015 until Jul 2015]

2. Overall Objectives

2.1. General Objectives

MYRIADS is a joint team with INRIA, CNRS, UNIVERSITY RENNES 1, and INSA 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 and environments for autonomous service and resource management in distributed virtualized infrastructures. The team tackles the challenges of dependable application execution and efficient resource management in the future Internet of Services.

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 *LinkedIn* 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

Dependable application execution in the future Internet raises a number of scientific challenges. The MYRIADS team aims at the design, programming and implementation of autonomous distributed systems and applications.

---

[1]According to World Stats, there are 2,41 billion Internet users i.e. nearly a quarter of the total world population in March 2012 http://www.internetworldstats.com/stats.htm.

The underlying computing infrastructure for the Internet of Services is characterized by its very large scale, dynamic nature and heterogeneity. The system scale is to be measured in terms of number of users, services, computers and geographical wingspan. The Internet of Services infrastructure spans multiple sites in multiple administrative domains. Its dynamic nature results from a number of factors such as Internet node volatility (due to computer or network failures, voluntarily connections and disconnections), services evolution (services appearing, disappearing, being modified), and varying demand depending on human being activities.

In a world in which more and more personal, business, scientific and industrial activities rely on services, it is essential to guarantee the high availability of services despite failures in the underlying continuously evolving (dynamic) execution environment. Multiple actors are involved in service provision. Also, computing infrastructures used for service execution are naturally distributed on multiple geographically distant sites belonging to different institutions. On the one hand, service execution infrastructures are often shared by different service providers (which might be competitors) and on the other hand services are accessed by multiple independent, and sometimes unknown, customers. In such an environment, providing confidence to the involved parties is of utmost importance.

Delivering a service depends on myriads of physical and virtualized resources, ranging from memory and CPU time to virtual machines, virtual clusters and other local or remote resources. Providing Quality of Service guarantees to users requires efficient mechanisms for discovering and allocating resources as well as dynamically adjusting resource allocations to accommodate workload variations. Moreover, efficient resource management is essential for minimizing resource supply costs, such as energy costs.

The Internet of Services is characterized by its uncertainty. It is an incommensurate and unpredictable system. Dependable application execution in such a distributed system can only be achieved through autonomic resource and service management. The MYRIADS project-team's objectives are to design and implement systems and environments for autonomous service and resource management in distributed virtualized infrastructures. We intend to tackle the challenges of dependable application execution and efficient resource management in the future Internet of Services.

Experiment-driven research in such a context is in itself a challenge. Confidence in scientific results for such large-scale systems can be greatly improved when they are verified on large-scale experimental testbeds. The Myriads project-team is therefore deeply involved in the management of the Grid'5000 testbed, by hosting its budget, technical director (David Margery), 1 engineer for Grid'5000 (Pascal Morillon) and some engineers for the European activities based on Grid'5000 know-how (Julien Lefeuvre). Here, the same challenges are faced at a smaller but nevertheless relevant scale for the project, with operational constraints for its experimenters and administrators.

2.4. Research Directions

The Myriads project-team aims at dependable execution of applications, particularly, but not exclusively, those relying on Service Oriented Architectures and at managing resources in virtualized infrastructures in order to guarantee service level agreement (SLA) terms to resource users and efficient resource management (energy efficiency, business efficiency...) to resource suppliers.

Our research activities are organized along three main work directions (structuring the remainder of this section): (i) autonomous management of virtualized infrastructures, (ii) dynamic adaptation of service-based applications and (iii) investigation of an unconventional, chemically-inspired, programming model for autonomous service computing.

2.4.1. *Autonomous Management of Virtualized Infrastructures*

Clouds can be defined as platforms for on-demand resource provisioning over the Internet. These platforms rely on networked computers. Three flavors of cloud platforms have emerged corresponding to different kinds of service delivery:

- IaaS (Infrastructure as a Service) refers to clouds for on-demand provisioning of elastic and customizable execution platforms (from physical to virtualized hardware).

- PaaS (Platform as a Service) refers to clouds providing an integrated environment to develop, build, deploy, host and maintain scalable and adaptable applications.
- SaaS (Software as a Service) refers to clouds providing customers access to ready-to-use applications.

2.4.1.1. Federation of IaaS clouds

With Infrastructure-as-a-Service (IaaS) cloud providers offer plain resources like x86 virtual machines (VM), IP networking and unstructured storage. These virtual machines can be already configured to support typical computation frameworks such as bag of tasks, MapReduce, etc. integrating autonomous elasticity management. By combining a private cloud with external resources from commercial or partner cloud providers, companies will rely on a federation of clouds as their computing infrastructure. A federation of clouds allows them to quickly add temporary resources when needed to handle peak loads. Similarly, it allows scientific institutions to bundle their resources for joint projects. We envision a peer-to-peer model in which a given company or institution will be both a cloud provider during periods when its IT infrastructure is not used at its maximal capacity and a cloud customer in periods of peak activity. Moreover it is likely that, in the future, huge data centers will reach their limits in term of size due to energy consumption considerations leading to a new landscape with a wide diversity of clouds (from small to large clouds, from clouds based on data centers to clouds based on highly dynamic distributed resources). We can thus anticipate the emergence of highly dynamic federations of virtualized infrastructures made up of different clouds. We intend to design and implement system services and mechanisms for autonomous resource management in federations of virtualized infrastructures.

2.4.1.2. SLA-driven PaaS over Cloud Federations

Platform as a Service (PaaS) promises to ease building and deploying applications, shielding developers from the complexity of underlying federated clouds. To fulfill its promise, PaaS should facilitate specifying and enforcing the QoS objectives of applications (e.g., performance objectives). These objectives are typically formalized in Service Level Agreements (SLAs) governing the interactions between the PaaS and hosted applications. The SLAs should be enforced automatically, which is essential for accommodating the dynamism of application requirements and of the capabilities of the underlying environment. Current PaaS offerings, such as Google App Engine and Microsoft Azure, include some form of SLA support, but this support is typically ad-hoc, limited to specific software stacks and to specific QoS properties.

Our main goal is to integrate flexible QoS support in PaaS over cloud federations. Specifically, we will develop an autonomous management solution for ensuring application SLAs while meeting PaaS-provider objectives, notably minimizing costs. The solution will include policies for autonomously providing a wide range of QoS guarantees to applications, focusing mainly on scalability, performance, and dependability guarantees. These policies will handle dynamic variations in workloads, application requirements, resource costs and availabilities by taking advantage of the on-demand elasticity and cloud-bursting capabilities of the federated infrastructure. The solution will enable performing in a uniform and efficient way diverse management activities, such as customizing middleware components and migrating VMs across clouds; these activities will build on the virtualized infrastructure-management mechanisms, described in the following paragraphs.

Several research challenges arise in this context. One challenge is translating from SLAs specifying properties related to applications (e.g., fault-tolerance) to federation-level SLAs specifying properties related to virtualized resources (e.g., number and type of VMs). This translation needs to be configurable and compliant with PaaS objectives. Another challenge is supporting the necessary decision-making techniques. Investigated techniques will range from policy-based techniques to control-theory and utility-based optimization techniques as well as combined approaches. Designing the appropriate management structure presents also a significant challenge. The structure must scale to the size of cloud-based systems and be itself dependable and resilient to failures. Finally, the management solution must support openness in order to accommodate multiple objectives and policies and to allow integration of different sensors, actuators, and external management solutions.

2.4.1.3. Virtual Data Centers

Cloud computing allows organizations and enterprises to rapidly adapt the available computational resources to theirs needs. Small or medium enterprises can avoid the management of their own data center and rent computational as well as storage capacity from cloud providers (outsourcing model). Large organizations already managing their own data centers can adapt their size to the basic load and rent extra capacity from cloud providers to support peak loads (cloud bursting model). In both forms, organization members can expect a uniform working environment provided by their organization: services, storage, ...This environment should be as close as possible to the environment provided by the organization' own data centers in order to provide transparent cloud bursting. A uniform environment is also necessary when applications running on external clouds are migrated back to the organization resources once they become free after a peak load. Supporting organizations necessitates to provide means to the organization administrators to manage and monitor the activity of their members on the cloud: authorization to access services, resource usage and quotas.

To support whole organizations, we will develop the concept of Elastic Virtual Data Center (VDC). A Virtual Data Center is defined by a set of services deployed by the organization on the cloud or on the organization's own resources and connected by a virtual network. The virtual machines supporting user applications deployed on a VDC are connected to the VDC virtual network and provide access to the organization's services. VDCs are elastic as the virtual compute resources are created when the users start new applications and released when these applications terminate. The concept of Virtual Data Center necessitates some form of Virtual Organization (VO) framework in order to manage user credentials and roles, to manage access control to services and resources. The concept of SLA must be adapted to the VDC context: SLA are negotiated by the organization administrators with resource providers and then exploited by the organization members (the organization receives the bill for resource usage). An organization may wish to restrict the capability to exploit some form of cloud resources to a limited group of members. It should be possible to define such policies through access rights on SLAs based on the user credential in a VO.

2.4.1.4. Virtualized Infrastructure Management

In the future, service-based and computational applications will be most likely executed on top of distributed virtualized computing infrastructures built over physical resources provided by one or several data centers operated by different cloud providers. We are interested in designing and implementing system mechanisms and services for multi-cloud environments (e.g. cloud federations).

At the IaaS level, one of the challenges is to efficiently manage physical resources from the cloud provider view point while enforcing SLA terms negotiated with cloud customers. We will propose efficient resource management algorithms and mechanisms. In particular, energy conservation in data centers is an important aspect to take into account in resource management.

In the context of virtualized infrastructures, we call a virtual execution platform (VEP) a collection of VMs executing a given distributed application. We plan to develop mechanisms for managing the whole life-cycle of VEPs from their deployment to their termination in a multi-cloud context. One of the key issues is ensuring interoperability. Different IaaS clouds may provide different interfaces and run heterogeneous hypervisors (Xen, VMware, KVM or even Linux containers). We will develop generic system level mechanisms conforming to cloud standards (e.g. DMTF OVF & CIMI, OGF OCCI, SNIA CDMI...) to deal with heterogeneous IaaS clouds and also to attempt to limit the vendor lock-in that is prevalent today. When deploying a VEP, we need to take into account the SLA terms negotiated between the cloud provider and customer. For instance, resource reservation mechanisms will be studied in order to provide guarantees in terms of resource availability. Moreover, we will develop the monitoring and measurement mechanisms needed to assess relevant SLA terms and detect any SLA violation. We also plan to develop efficient mechanisms to support VEP horizontal and vertical elasticity in the framework of cloud federations.

We envision that in the future Internet, a VEP or part of a VEP may migrate from one IaaS cloud to another one. While VM migration has been extensively studied in the framework of a single data center, providing efficient VM migration mechanisms in a WAN environment is still challenging [48], [42]. In a multi-cloud context, it is essential to provide mechanisms allowing secure and efficient communication between VMs belonging to the same VEP and between these VMs and their user even in the presence of VM migration.

2.4.1.5. Heterogeneous Cloud Infrastructure Management

Today's cloud platforms are missing out on the revolution in new hardware and network technologies for realizing vastly richer computational, communication, and storage resources. Technologies such as Field Programmable Gate Arrays (FPGA), General-Purpose Graphics Processing Units (GPGPU), programmable network routers, and solid-state disks promise increased performance, reduced energy consumption, and lower cost profiles. However, their heterogeneity and complexity makes integrating them into the standard Platform as a Service (PaaS) framework a fundamental challenge.

Our main challenge in this context is to automate the choice of resources which should be given to each application. To execute an application a cloud user submits an SLO document specifying non-functional requirements for this execution, such as the maximum execution latency or the maximum monetary cost. The goal of the platform developed in the HARNESS European project is to deploy applications over well-chosen sets of resources such that the SLO is respected. This is realized as follows: (i) building a performance model of each application; (ii) choosing the implementation and the set of cloud resources that best satisfy the SLO; (iii) deploying the application over these resources; (iv) scheduling access to these resources.

2.4.2. Multilevel Dynamic Adaptation of Service-based Applications

In the Future Internet, most of the applications will be built by composing independent software elements, the services. A Service Oriented Architecture (SOA) should be able to work in large scale and open environments where services are not always available and may even show up and disappear at any time.

Applications which are built as a composition of services need to ensure some Quality of Service (QoS) despite the volatility of services, to make a clever use of new services and to satisfy changes of needs from end-users.

So there is a need for dynamic adaptation of applications and services in order to modify their structure and behavior.

The task of making software adaptable is very difficult at many different levels:

- At business level, processes may need to be reorganized when some services cannot meet their Service Level Agreement (SLA).

- At service composition level, applications may have to change dynamically their configuration in order to take into account new needs from the business level or new constraints from the services and the infrastructure level. At this level, most of the applications are distributed and there is a strong need for *coordinated adaptation*.

- At the infrastructure level, the state of resources (networks, processors, memory,...) has to be taken into account by service execution engines in order to make a clever use of these resources such as taking into account available resources and energy consumption. At this level there is a strong requirement for *cooperation* with the underlying operating system.

Moreover, the adaptations at these different levels need to be coordinated. In the Myriads project-team we address mainly the infrastructure and service composition layers.

So our main challenge is to build generic and concrete frameworks for self-adaptation of services and service based applications at run-time. The basic steps of an adaptation framework are Monitoring, Analysis/decision, Planning and Execution, following the MAPE model proposed in [53]. We intend to improve this basic framework by using models at runtime to validate the adaptation strategies and establishing a close cooperation with the underlying Operating System.

We will pay special attention to each step of the MAPE model. For instance concerning the Monitoring, we will design high-level composite events; for the Decision phase, we work on different means to support decision policies such as rule-based engine, utility function based engine. We will also work on the use of an autonomic control loop for learning algorithms; for Planning, we investigate the use of on-the-fly planning of adaptation actions allowing the parallelization and distribution of actions. Finally, for the Execution step our research activities aim to design and implement dynamic adaptation mechanisms to allow a service to self-adapt according to the required QoS and the underlying resource-management system.

Then we intend to extend this model to take into account proactive adaptation, to ensure some properties during adaptation and to monitor and adapt the adaptation itself.

An important research direction is the coordination of adaptation at different levels. We will mainly consider the cooperation between the application level and the underlying operating system in order to ensure efficient and consistent adaptation decisions. This work is closely related to the activity on autonomous management of virtualized infrastructures.

We are also investigating the Chemical approach as an alternative way to frameworks for providing autonomic properties to applications.

2.4.3. *Exploration of unconventional programming models for the Internet of services*

Facing the complexity of the emerging ICT landscape in which highly heterogeneous digital services evolve and interact in numerous different ways in an autonomous fashion, there is a strong need for rethinking programming models. The question is "*what programming paradigm can efficiently and naturally express this great number of interactions arising concurrently on the platform?*".

It has been suggested [41] that observing nature could be of great interest to tackle the problem of modeling and programming complex computing platforms, and overcome the limits of traditional programming models. Innovating unconventional programming paradigms are requested to provide a high-level view of these interactions, then allowing to clearly separate what is a matter of expression from what is a question of implementation. Towards this, nature is of high inspiration, providing examples of self-organizing, fully decentralized coordination of complex and large scale systems.

As an example, chemical computing [44] has been proposed more than twenty years ago as a natural way to program parallelism. Even after significant spread of this approach, it appears today that chemical computing exposes a lot of good properties (implicit autonomy, decentralization, and parallelism) to be leveraged for programming service infrastructures.

3. Research Program

3.1. Introduction

The research activity within the MYRIADS team encompasses several areas: distributed systems, middleware and programming models. We have chosen to provide a brief presentation of some of the scientific foundations associated with them: autonomic computing, future internet and SOA, distributed operating systems, and unconventional/nature-inspired programming.

3.2. Autonomic Computing

During the past years the development of raw computing power coupled with the proliferation of computer devices has grown at exponential rates. This phenomenal growth along with the advent of the Internet have led to a new age of accessibility — to other people, other applications and others systems. It is not just a matter of numbers. This boom has also led to unprecedented levels of complexity for the design and the implementation of these applications and systems, and of the way they work together. The increasing system scale is reaching a level beyond human ability to master its complexity.

This points towards an inevitable need to automate many of the functions associated with computing today. Indeed we want to interact with applications and systems intuitively, and we want to be far less involved in running them. Ideally, we would like computing systems to entirely manage themselves.

IBM [53] has named its vision for the future of computing "autonomic computing." According to IBM this new computer paradigm means the design and implementation of computer systems, software, storage and support that must exhibit the following basic fundamentals:

- **Flexibility.** An autonomic computing system must configure and reconfigure itself under varying, even unpredictable, conditions.

- **Accessibility.** The nature of the autonomic system is that it is always on.

- **Transparency.** The system will perform its tasks and adapt to a user's needs without dragging the user into the intricacies of its workings.

In the Myriads team we will act to satisfy these fundamentals.

3.3. Future Internet and SOA

Traditional information systems were built by integrating applications into a communication framework, such as CORBA or with an Enterprise Application Integration system (EAI). Today, companies need to be able to reconfigure themselves; they need to be able to include other companies' business, split or externalize some of their works very quickly. In order to do this, the information systems should react and adapt very efficiently. EAIs approaches did not provide the necessary agility because they were too tightly coupled and a large part of business processes were "hard wired" into company applications.

Web services and Service Oriented Architectures (SOA) partly provide agility because in SOA business processes are completely separated from applications which can only be viewed as providing services through an interface. With SOA technologies it is easily possible to modify business processes, change, add or remove services.

However, SOA and Web services technologies are mainly market-driven and sometimes far from the state-of-the-art of distributed systems. Achieving dependability or being able to guarantee Service Level Agreement (SLA) needs much more agility of software elements. Dynamic adaptability features are necessary at many different levels (business processes, service composition, service discovery and execution) and should be coordinated. When addressing very large scale systems, autonomic behavior of services and other parts of service oriented architectures is necessary.

SOAs will be part of the "Future Internet". The "Future Internet" will encompass traditional Web servers and browsers to support company and people interactions (Internet of services), media interactions, search systems, etc. It will include many appliances (Internet of things). The key research domains in this area are network research, cloud computing, Internet of services and advanced software engineering.

The Myriads team will address adaptability and autonomy of SOAs in the context of Grids, Clouds and at large scale.

3.4. Distributed Operating Systems

An operating system provides abstractions such as files, processes, sockets to applications so that programmers can design their applications independently of the computer hardware. At execution time, the operating system is in charge of finding and managing the hardware resources necessary to implement these abstractions in a secure way. It also manages hardware and abstract resource sharing between different users and programs.

A distributed operating system makes a network of computers appear as a single machine. The structure of the network and the heterogeneity of the computation nodes are hidden to users. Members of the Myriads team members have a long experience in the design and implementation of distributed operating systems, for instance in Kerrighed, Vigne, and XtreemOS projects.

The cloud computing model [43], [40] introduces new challenges in the organization of the information infrastructure: security, identity management, adaptation to the environment (costs). The organization of large IT infrastructures is also impacted as their internal data-centers, sometimes called private clouds, need to cooperate with resources and services provisioned from the cloud in order to cope with workload variations. The advent of cloud and green computing introduces new challenges in the domain of distributed operating systems: resources can be provisioned and released dynamically, the distribution of the computations on the resources must be reevaluated periodically in order to reduce power consumption and resource usage costs. Distributed cloud operating system must adapt to these new challenges in order to reduce cost and energy, for instance, through the redistribution of the applications and services on a smaller set of resources.

The Myriads team works on the design and implementation of system services at IaaS and PaaS levels to autonomously manage cloud and cloud federations resources and support collaboration between cloud users.

3.5. Unconventional/Nature-inspired Programming

Leveraging the computing services available on the Internet requires revisiting programming models, with the idea of expressing decentralized and autonomous behaviors (in particular self-repairing, self-adaptation). More concretely, composing services within large scale platforms calls for mechanisms to adequately discover and select services at run time, upon failure, or unexpected results.

Nature metaphors have been shown to provide adequate abstractions to build autonomic systems. Firstly, we want to explore nature metaphors, such as the chemical programming model as alternative programming models for expressing the interactions and coordination of services at large scale to build applications dynamically.

Within the *chemical* paradigm, a program is seen as a solution in which molecules (data) float and react together to produce new data according to rules (programs). Such a paradigm, implicitly parallel and distributed, appears to be a good candidate to express high level interactions of software components. The language naturally focus on the coordination of distributed autonomous entities. Thus, our first objective is to extend the semantics of chemical programs, in order to model not only a distributed execution of a service coordination, but also, the interactions between the different *molecules* within the Internet of Services (users, companies, services, advertisements, requests, ...). At present, a distributed implementation of the chemical paradigm does not exist. Our second objective is to develop the concepts and techniques required for such an implementation. While the paradigm exhibit several limitations regarding its run-time complexity, revisiting the model and studying its implementation over distributed platforms, and then showing its relevance in concrete settings (such as service coordination) may constitute an innovative research area.

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,
- Numerical simulations,
- Energy and sustainable development,
- Smart cities.

5. Highlights of the Year

5.1. Highlights of the Year

- Christine Morin has been made Knight of the French legion of Honour by decret of the President of French Republic (December 31, 2014) for her contribution to Higher Education and Research. Antoine Petit, President of Inria, presented her with the insignia on February 24th, 2015.

- The HARNESS European project was successfully completed in September 2015. Although the final evaluation report is still pending, the verbal comments by project reviewers were very positive. The HARNESS project has developed a new generation cloud computing platform that integrates heterogeneous hardware (FPGAs, GPGPUs, programmable routers, etc.) and networking resources in order to provide vastly increased performance for a broader array of applications. With HARNESS, cloud providers can profitably manage specialized hardware and network technologies much as they do today's commodity resources, and software engineers can seamlessly integrate them into the design of their cloud-hosted applications.

6. New Software and Platforms

6.1. ConPaaS

Hosting complex applications in the cloud

Contact: Guillaume Pierre, `Guillaume.Pierre@irisa.fr`

URL: http://www.conpaas.eu/

Status: Version 1.5.1

License: BSD

Presentation: ConPaaS [55] 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.

New features developed in 2015 include: a "Generic" service which allows the easy deployment and execution of arbitrary applications in ConPaaS; the "ConPaaS Nutshell edition" which allows users to deploy a fully-functional ConPaaS installation within a single VirtualBox VM; and the "ConPaaS Raspberry Pi edition" which allows users to deploy ConPaaS on a set of Raspberry Pi devices, paving the way toward the development of a fully-featured mobile edge cloud.

Active contributors (from the Myriads team): Ancuta Iordache, Genc Tato, Teodor Crivat, Guillaume Pierre.

Impact: ConPaaS is recognized as one of the major open-source PaaS environments. It is being developed by teams in Rennes, Amsterdam, Berlin and Ljubljana. Technology transfer of ConPaaS technology is ongoing in the context of the MC-DATA EIT Digital project.

6.2. GinFlow

Contact: Cédric Tedeschi, `Cedric.Tedeschi@irisa.fr`

Status: Version 1.0 released in open source, and registered at APP.

License: LGPL-3

Presentation: GinFlow is a decentralized workflow engine. It relies on a set of processes deployed over a cluster. Every task of the workflow is basically encapsulated into a GinFlow worker, workers being able to coordinate with others through read and write operations of a shared space containing the status of the workflow. GinFlow provides the ability to change the workflow logic on-the-fly upon the detection of a non-desired behavior within the execution of some of the tasks of the workflow.

Users typically submit a workflow through its JSON representation (both the by-default workflow and its alternate version if adaptation is needed). It can also use the API to describe its workflow in a more *programmatic* fashion.

Active contributors (from Myriads project-team): Matthieu Simonin, Cédric Tedeschi, Javier Rojas Balderrama.

Impact: GinFlow is a tool meant to provide support for workflow-based applications needing adaptation at run time. It is also targeted at offering a platform for future development and researches around the decentralized execution of workflows. It has been used in the framework of the DALHIS [2] associate team, as a workflow template executor, integrated with the TIGRES workflow manager developed at the Lawrence Berkeley National Lab. It has been supported by the GinFlow ADT funded by Inria since 2014 (see Section 9.2.1).

6.3. Merkat

Contact: Nikolaos Parlavantzas, `Nikolaos.Parlavantzas@irisa.fr`

URL: http://www.irisa.fr/myriads/software/Merkat/

Status: Version 1.0

Presentation: Merkat is a market-based private PaaS (Platform-as-a-Service) system, supporting dynamic, fine-grained resource allocation and automatic application management [47], [46] [3]. Merkat implements a proportional-share auction that ensures maximum resource utilization while providing incentives to applications to regulate their resource usage. Merkat includes generic mechanisms for application deployment and automatic scaling. These mechanisms can be adapted to support diverse performance goals and application types, such as master-worker, MPI, or MapReduce applications. Merkat is implemented in Python and uses OpenNebula for virtual machine management. Experimental results on the Grid'5000 testbed show that using Merkat increases resource utilization and improves application performance. The development was initiated in the framework of Stefania Costache PhD's thesis.

Active contributors (from the Myriads team): Christine Morin, Nikolaos Parlavantzas.

Other contributors: Stefania Costache.

Impact: Merkat has been integrated in EDF R&D portal providing access to internal computing resources and is currently used on a testbed at EDF R&D.

6.4. Meryn

Contact: Nikolaos Parlavantzas, `Nikolaos.Parlavantzas@irisa.fr`

URL: http://www.irisa.fr/myriads/software/Meryn/

Status: Version 1.0

---

[2] http://project.inria.fr/dalhis

Presentation: 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 [49]. The current Meryn prototype was implemented using shell scripts, builds upon the Snooze VM manager software, and supports batch and MapReduce applications using respectively the Oracle Grid Engine OGE 6.2u7 and Hadoop 0.20.2 frameworks.

Active contributors (from the Myriads team): Christine Morin, Nikolaos Parlavantzas.

Other contributors: Djawida Dib.

Impact: Meryn is not yet distributed as an open source software.

6.5. Snooze

Contact: Christine Morin, `Christine.Morin@inria.fr`

URL: http://snooze.inria.fr

Status: Version 2.1.5

License: GPLv2

Presentation: Snooze [51], [50], [52] [4] 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 [54]), 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.

Snooze is fully implemented from scratch in Java and currently comprises approximately 15.000 lines of maintainable abstractions-based code. In order to provide a uniform interface to the underlying hypervisors and support transparent VM monitoring and management, Snooze integrates the *libvirt* virtualization library. Cassandra (since 2.0.0) can be used as base backend, providing reliability and scalability to the database management system. At a higher level Snooze provides its own REST API as well as an EC2 compatible API (since 2.1.0). It can thus be controlled from the command line (using the legacy client or an EC2 compatible tool), or from different language libraries (libcloud, jcloud ...). Snooze also provides a web interface to control the system. In collaboration with the Northeastern University of Bostonwe built the *Checkpoint as a Service* system on top of Snooze. The service allows users to execute their computations in a cloud environment in a reliable way. Periodic checkpoints are saved making it possible to restore the computation from a previous state in the event of failures. This work is described in [16].

Active contributors (from Myriads team): Yvon Jégou, David Margery, Christine Morin, Matthieu Simonin.

Other contributors: Jiajun Cao, Gene Cooperman, Eugen Feller.

Impact: Snooze has been used by students at LIFL, IRIT in France and LBNL in the US in the framework of internships. It has also been deployed and experimented at EDF R&D. Snooze entry won the 2nd prize of the scalability challenge at CCGrid 2013. Finally, we know that it was used by external users from academia and industry as we received feed-back from them. Snooze development was supported by the Snooze ADT funded by Inria from October 2012 to September 2014.

6.6. SimGrid

Scientific Instrument for the study of Large-Scale Distributed Systems.

Contact: Martin Quinson, `Martin.Quinson@ens-rennes.fr`

URL: http://simgrid.gforge.inria.fr/

Status: Version 3.12

License: LGPL 2.1

Presentation: SimGrid is a toolkit providing core features 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.

Active contributors from Myriads project-team: Martin Quinson and Anne-Cécile Orgerie.

Other active contributors: Arnaud Legrand (Mescal project-team, CNRS & Inria Rhone-Alpes), Frédéric Suter (Avalon project-team, CNRS & Inria Rhone-Alpes).

Impact: SimGrid has an active user community of more than one hundred members, mainly composed of researchers and students. In the last decade only, it grounded the experiments of 6 PhDs works, 25 journal papers, and over 100 conference papers. The community gathers every year during the SimGrid User's Days workshop.

SimGrid is also used to run the regression and performance tests and tuning of two large applications: BigDFT (a massively parallel code computing the electronic structure of chemical elements developed by the CEA) and StarPU (a Unified Runtime System for Heterogeneous Multicore Architectures developed by Inria Bordeaux). Both of these programs enjoy large user communities themselves.

7. New Results

7.1. Cloud Resource Management

Participants: Ancuta Iordache, Christine Morin, Ghada Moualla, Guillaume Pierre, Matthieu Simonin, Lodewijck Vogelzang.

7.1.1. *Application Performance Modeling in Heterogeneous Cloud Environments*

Participants: Ancuta Iordache, Lodewijck Vogelzang, Guillaume Pierre.

Heterogeneous cloud platforms offer many possibilities for applications to make fine-grained choices over the types of resources they execute on. This opens for example opportunities for fine-grained control of the tradeoff between expensive resources likely to deliver high levels of performance, and slower resources likely to cost less. We designed a methodology for automatically exploring this performance vs. cost tradeoff when an arbitrary application is submitted to the platform. Thereafter, the system can automatically select the set of resources which is likely to implement the tradeoff specified by the user. We significantly improved the speed at which the system can characterize the performance of an arbitrary application. A first publication on this topic has been published [26], and a second one is in preparation.

7.1.2. *Heterogeneous Resource Management*

Participants: Ancuta Iordache, Guillaume Pierre.

During her internship at Maxeler Technologies, Ancuta Iordache developed an original technique for virtualizing FPGAs such that they can be used as high-performance computing devices in cloud infrastructures. Virtual FPGAs can be accessed remotely by any VM in the system. They can span multiple physical FPGA, they are elastic, and they can also be shared between multiple tenants. A publication on this topic is currently under evaluation.

7.1.3. *Self-adaptatable Hadoop Virtual Clusters*

Participants: Christine Morin, Ghada Moualla, Matthieu Simonin.

In the context of Ghana Moualla's Master internship, we designed the Elastic MapReduce Adaptation (EMRA) system to execute Hadoop MapReduce applications with user-defined deadlines in cloud virtual clusters. EMRA integrates an algorithm to automatically adapt the Hadoop cluster size at runtime in order to meet user-defined deadlines. We proposed an automatic scaling algorithm, which monitors the progress of the Map phase of the application during its execution and estimate if the user-defined deadline can be met. If the current allocated resources are not sufficient to meet the deadline, more resources are provisioned. The adaptation service comprises of three main components: (i) a monitor to check the progress of the running application, (ii) an estimator to predict the time needed to complete the application based on its current progress ; (iii) a controller to adapt the size of the virtual cluster by adding virtual machines as needed. The controller takes into account the start-up overhead of the new virtual machines and the time needed for the VM to fetch their input data from the original nodes over the network in order to start their map tasks. We implemented a prototype of the EMRA system in the context of Sahara, an environment for managing Hadoop virtual clusters on top of OpenStack IaaS clouds. We experimented the EMRA system on Grid'5000 with traditional MapReduce benchmarks. We evaluated the relative error of the estimator, the cost for scaling up or down a virtual cluster and showed that the proposed adaptation algorithm allows user-defined deadlines to be met.

7.2. Distributed Cloud Computing

Participants: Teodor Crivat, Yvon Jégou, Vlad Mirel, Christine Morin, Anne-Cécile Orgerie, Edouard Outin, Nikolaos Parlavantzas, Jean-Louis Pazat, Guillaume Pierre, Aboozar Rajabi, Carlos Ruiz Diaz, Arnab Sinha, Genc Tato, Cédric Tedeschi.

7.2.1. *A multi-objective adaptation system for the management of a Distributed Cloud*

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

In this project, we consider a "Distributed Cloud" made of multiple data/computing centers interconnected by a high speed network and belonging to the same administration domain. Moreover, in the Cloud organization targeted here, the network capabilities can be dynamically configured in order to guarantee QoS for streaming or to negotiate bandwidth for example.

As a first step, we are focusing on a single centralised Cloud.

Due to the dynamic capabilities of the Clouds, often referred to as elasticity, there is a strong need to dynamically adapt both platforms and applications to users needs and environmental constraints such as electrical power consumption.

We address the management of a Cloud in order to consider both optimization for energy consumption and for users' QoS needs. The objectives of this optimization will be negotiated as contracts on Service Level Agreement (SLA). A special emphasis will be put on the distributed aspect of the platform and include both servers and network adaptation capabilities.

The design of the system relies on self-* techniques and on adaptation mechanisms at any level (from IaaS to SaaS). The MAPE-k framework (Monitor-Analysis-Planning-Execution based on knowledge) is used for the implementation of the system. The technical developments are based on the Openstack framework.

We have implemented a system that uses a genetic algorithm to optimize Cloud energy consumption and machine learning techniques to improve the fitness function regarding a real distributed cluster of servers. We have carried out experiments on the OpenStack platform to validate our solution. This experimentation shows that the machine learning produces an accurate energy model, predicting precise values for the simulation.

We are currently refining this model and comparing it to real measurements on the platform.

This work is done in cooperation with the DIVERSE team and in cooperation with Orange under the umbrella of the B-COM Technology Research Center.

7.2.2. *Dynamic reconfiguration for multi-cloud applications*

Participants: Nikolaos Parlavantzas, Aboozar Rajabi, Carlos Ruiz Diaz, Arnab Sinha.

In the context of the PaaSage European project, we are working on model-based, continuous self-optimization of multi-cloud applications. In particular, we are developing a dynamic adaptation system, capable of transforming the currently running application configuration into a target configuration in a cost-effective and safe manner. In 2015, we have improved and extended the Adapter prototype [45]. The system now fully supports dynamic configuration, including detecting changes, generating reconfiguration plans, validating plans based on a cost-benefit calculation, and executing plans in parallel, improving adaptation performance. Moreover, we have performed initial investigations on the use of PaaSage for supporting Internet of Things (IoT) applications [27]. Finally, in the context of Carlos Ruiz's stay, we are defining a model for managing the configuration of cloud applications and environments. This model is based on feature modeling and the derived configurations are mapped to PaaSage models.

7.2.3. *Towards a distributed cloud inside the backbone*

Participants: Christine Morin, Anne-Cécile Orgerie, Genc Tato, Cédric Tedeschi.

The DISCOVERY proposal officially started at the end of 2015. It is an Inria Project Lab (IPL) led by Adrien Lebre from the ASCOLA team, and currently on leave at Inria. It aims at designing a distributed cloud, leveraging the resources we can find in the network backbone. [3] In practice, this work is intended to get integrated within the OpenStack software https://www.openstack.org/ so as to decentralize its whole architecture.

In this context, and in collaboration with ASCOLA and ASAP teams, we started the design of an overlay network whose purpose is to be able, with a limited cost, to locate geographically-close nodes from any point of the network. In this framework, the PhD thesis of Genc Tato started in December 2015. It aims at developing locality mechanisms at the data management layer.

We have also started an energy/cost-benefit analysis of a decentralized Cloud infrastructure like the one proposed within Discovery. This work is conducted by Anthony Simonet, a post-doctoral researcher on an Inria contract for the Discovery IPL and co-supervised by Adrien Lebre from the ASCOLA team and Anne-Cécile Orgerie from Myriads team.

7.2.4. *Mobile edge cloud computing with ConPaaS*

Participants: Teodor Crivat, Vlad Mirel, Guillaume Pierre.

Interactive multi-user applications usually rely on intermediate cloud servers to mediate the inter-user interaction. However, current mobile networks exhibit network latencies in the order of 50-150 ms between the device and any cloud. Such latencies make it impossible to create smooth interactions with the end user. To enable an "instantaneous" feeling, augmented reality applications require that end-to-end latencies should remain below 20 ms.

To address these issues, we extended ConPaaS to support the deployment of cloud applications in a distributed set of Raspberry Pi machines. The motivation is to reduce the latency compared to a traditional deployment where the backend is located in an external cloud: 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.

---

[3]The DISCOVERY website: http://beyondtheclouds.github.io

7.2.5. *Fog Computing*
Participant: Jean-Louis Pazat.

The concept of "Fog Computing" is currently developed on the idea of hosting instances of services not on centralized datacenters (i.e. the "Cloud"), but on a highly distributed infrastructure: the Internet Edge (i.e. the "Fog"). This infrastructure consists in geographically distributed computing resources with relatively small capabilities. Compared with datacenters, a "Fog" infrastructure is able to offer to Service Providers a shorter distance from the service to the user but with the same flexibility of software deployment and management.

This work focus on the problem of resource allocation in such infrastructure when considering services in the area of Internet of Things, Social Networks or Online Gaming. For such use-cases, service-to-user latency is a critical parameter for the quality of experience. Optimizing such a parameter is an objective for the platform built on top of the Fog Infrastructure that will be dedicated to the deployment of the considered service. In order to achieve such a goal, the platform needs to select some strategies for the allocation of network and computing resources, based on the initial requirements for service distribution.

We are designing a prototype based on micro services and we are considering low overhead virtualization systems using containers. This prototype is intended to run inside an Internet Box or inside a LAN disk server at user's home. The whole system will be intended to be used very small or medium size user communities willing to share devices and data. The main characteristics of the system will be reliable distributed storage and distributed execution of services.

This work is part of Bruno Stevant's PhD thesis, which began in December 2014. It is done in cooperation with the REOP team, Institut Mines telecom/IRISA.

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

7.3.1. *Security Monitoring of Clouds*
Participants: Anna Giannakou, Christine Morin, Jean-Louis Pazat, Louis Rilling, Amir Teshome Wonjiga.

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 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.

In 2015 we started to study the state of the art about SLA for security monitoring in clouds, as well as about evaluating security monitoring setups in clouds.

In 2015 we also studied the self-adaptation issues of security monitoring with two kinds of security monitoring components: a network intrusion detection system (NIDS), and a secured application-level firewall. Moreover a new approach to secure an application-level firewall has been proposed.

To experiment with both kinds of components, a prototype called SAIDS has been implemented in the OpenStack-based IaaS cloud testbed that was setup in 2014. The NIDS software used is Snort. The application-level firewall is based on Linux nftables and Open vSwitch. In order to study more complex security monitoring setups, SAIDS will be extended in 2016.

A preliminary evaluation of SAIDS has been published in the doctoral symposium of CCGrid 2015. A more complete evaluation of SAIDS as well as the evaluation of the application-level firewall will be done in 2016.

7.4. Greening Clouds

Participants: Maria Del Mar Callau Zori, Ismael Cuadrado Cordero, David Guyon, Sabbir Hasan Rochi, Yunbo Li, Christine Morin, Anne-Cécile Orgerie, Jean-Louis Pazat, Guillaume Pierre.

7.4.1. *Energy-aware IaaS-PaaS co-design*

Participants: Maria Del Mar Callau Zori, Anne-Cécile Orgerie, Guillaume Pierre.

The wide adoption of the cloud computing paradigm plays a crucial role in the ever-increasing demand for energy-efficient data centers. Driven by this requirement, cloud providers resort to a variety of techniques to improve energy usage at each level of the cloud computing stack. However, prior studies mostly consider resource-level energy optimizations in IaaS clouds, overlooking the workload-related information locked at higher levels, such as PaaS clouds. We conducted an extensive experimental evaluation of the effect of a range of Cloud infrastructure operations (start, stop, migrate VMs) on their computing throughput and energy consumption, and derived a model to help drive cloud reconfiguration operations according to performance/energy requirements. A publication on this topic is in preparation.

7.4.2. *Energy-efficient cloud elasticity for data-driven applications*

Participants: David Guyon, Anne-Cécile Orgerie, Christine Morin.

Distributed and parallel systems offer to users tremendous computing capacities. They rely on distributed computing resources linked by networks. They require algorithms and protocols to manage these resources in a transparent way for users. Recently, the maturity of virtualization techniques has allowed for the emergence of virtualized infrastructures (Clouds). These infrastructures provide resources to users dynamically, and adapted to their needs. By benefiting from economies of scale, Clouds can efficiently manage and offer virtually unlimited numbers of resources, reducing the costs for users.

However, the rapid growth for Cloud demands leads to a preoccupying and uncontrolled increase of their electric consumption. In this context, we will focus on data driven applications which require to process large amounts of data. These applications have elastic needs in terms of computing resources as their workload varies over time. While reducing energy consumption and improving performance are orthogonal goals, this internship aims at studying possible trade-offs for energy-efficient data processing without performance impact. As elasticity comes at a cost of reconfigurations, these trade-offs will consider the time and energy required by the infrastructure to dynamically adapt the resources to the application needs.

The master internship work of David Guyon on this topic has been presented at IEEE GreenCom 2015 [39]. This work will be continued during David's PhD thesis.

7.4.3. *Energy-efficient and network-aware resource allocation in Cloud infrastructures*

Participants: Ismael Cuadrado Cordero, Christine Morin, Anne-Cécile Orgerie.

Energy consumption is cloud computing has become a key environmental and economic concern. Our work aims at designing energy-efficient resource allocation for 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. New algorithms have to be designed integrating the data locality dimension to optimize computing resource allocation while taking into account the fluctuating limits of network resources. Towards this end, we proposed GRaNADA, a semi-decentralized Platform-as-a-service (PaaS) architecture for real-time multiple-users applications. Our architecture 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. GRaNADA implements the concept of micro-cloud, a fully autonomous energy-efficient subnetwork of clients of the same service, designed to keep the greenest path between its nodes. Along with GRaNADA, we proposed DEEPACC, a cloud-aware routing protocol which distributes the connection between the nodes. Our system GRaNADA targets services where the geographical distribution of clients working on the same data is limited - for example, a shared on-line document

- or those services where, even if the geographical distribution of clients is high, the upload data communication to the cloud is small - for instance a light social network like Twitter. We compared our approach with two main existing solutions - replication of data in the edge and traditional centralized cloud computing. Our approach based on micro-clouds exhibits interesting properties in terms of QoS and especially latency. Simulations show that, using the proposed PaaS, one can save up to 75% of the spent network energy compared to traditional centralized cloud computing approaches. Our approach is also more energy-efficient than the most popular semi-decentralized solutions, like nano data centers. This work has been presented at IEEE GreenCom 2015 [18].

We also evaluated the suitability of using micro-clouds in the context of smart cities. We investigated the idea to build a local cloud on top of networking resources spread across a defined area and including the mobile devices of the users. This local cloud is managed by lightweight mechanisms in order to handle users who can appear/disappear and move. We used a scenario considering a platform for neighborhood services and showed that micro-clouds make better use of the network, reducing the amount of unnecessary data traveling through external networks. This work is currently under review for a conference.

7.4.4. *Resource allocation in a Cloud partially powered by renewable energy sources*
Participants: Yunbo Li, Anne-Cécile Orgerie.

We propose here to design a disruptive approach to Cloud resource management which takes advantage of renewable energy availability to perform opportunistic tasks. To begin with, the considered Cloud is mono-site (i.e. all resources are in the same physical location) and performs tasks (like web hosting or MapReduce tasks) running in virtual machines. 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 needs to integrate a prediction model to be able to forecast these extra-power periods of time in order to schedule more work during these periods. Batteries will be used to guarantee that enough energy is available when switching on a new server working exclusively on renewable energy. Given a reliable prediction model, it is possible to design a scheduling algorithm that aims at optimizing resource utilization and energy usage, problem known to be NP-hard. The proposed heuristics will thus schedule tasks spatially (on the appropriate servers) and temporally (over time, with tasks that can be planed in the future).

This work is done in collaboration with Ascola team from LINA in Nantes. Two publications have been accepted this year on this topic for: SmartGreens 2015 [15] and IEEE GreenCom 2015 [21].

7.4.5. *SLA driven Cloud Auto-scaling for optimizing energy footprint*
Participants: Sabbir Hasan Rochi, Jean-Louis Pazat.

As a direct consequence of the increasing popularity of Internet and Cloud Computing services, data centers are amazingly growing and hence have to urgently face energy consumption issues. At the Infrastructure-as-a-Service (IaaS) layer, Cloud Computing allows to dynamically adjust the provision of physical resources according to Platform-as-a-Service (PaaS) needs while optimizing energy efficiency of the data center.

The management of elastic resources in Clouds according to fluctuating workloads in the Software-as-a-Service (SaaS) applications and different Quality-of-Service (QoS) end-user's expectations is a complex issue and cannot be done dynamically by a human intervention. We advocate the adoption of Autonomic Computing (AC) at each XaaS layer for responsiveness and autonomy in front of environment changes. At the SaaS layer, AC enables applications to react to a highly variable workload by dynamically adjusting the amount of resources in order to keep the QoS for the end users. Similarly, at the IaaS layer, AC enables the infrastructure to react to context changes by optimizing the allocation of resources and thereby reduce the costs related to energy consumption. However, problems may occur since those self-managed systems are related in some way (e.g. applications depend on services provided by a cloud infrastructure): decisions taken in isolation at given layer may interfere with other layers, leading whole system to undesired states.

We have defined a scheme for green energy management in the presence of explicit and implicit integration of renewable energy in datacenter [13]. More specifically we propose three contributions: i) we introduce the concept of virtualization of green energy to address the uncertainty of green energy availability, ii) we extend the Cloud Service Level Agreement (CSLA) language to support Green SLA introducing two new threshold parameters and iii) we introduce greenSLA algorithm which leverages the concept of virtualization of green energy to provide per interval specific Green SLA. Experiments were conducted with real workload profile from PlanetLab and server power model from SPECpower to demonstrate that, Green SLA can be successfully established and satisfied without incurring higher cost.

This work is done in collaboration with Ascola team from LINA in Nantes.

7.5. Energy-efficient Computing Infrastructures

Participants: Christine Morin, Anne-Cécile Orgerie, Martin Quinson.

7.5.1. *Simulating the impact of DVFS within SimGrid*

Participants: Christine Morin, Anne-Cécile Orgerie, Martin Quinson.

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. Furthermore, studying the impact of application configuration choices on energy is a difficult task, as not many platforms are equipped with the proper power measurement tools. The goal of this work is to enable energy-aware experimentation within the SimGrid simulation toolkit, by introducing a model of application energy consumption and enabling the use of Dynamic Voltage and Frequency Scaling (DVFS) techniques for the simulated platforms. We provide the methodology used to obtain accurate energy estimations, highlighting the simulator calibration phase. The proposed energy model is validated by means of a large set of experiments featuring several benchmarks and scientific applications. This work is available in the latest SimGrid release. This work is done in collaboration with the Mescal team from LIG in Grenoble. A paper is currently under preparation on this work.

7.5.2. *Simulating Energy Consumption of Wired Networks*

Participants: Timothée Haudebourg, Anne-Cécile Orgerie.

Predicting the performance of applications, in terms of completion time and resource usage for instance, is critical to appropriately dimensioning resources that will be allocated to these applications. Current applications, such as web servers and Cloud services, require lots of computing and networking resources. Yet, these resource demands are highly fluctuating over time. Thus, adequately and dynamically dimensioning these resources is challenging and crucial to guarantee performance and cost-effectiveness. In the same manner, estimating the energy consumption of applications deployed over heterogeneous cloud resources is important in order to provision power resources and make use of renewable energies. Concerning the consumption of entire infrastructures, some studies show that computing resources represent the biggest part in the Cloud's consumption, while others show that, depending on the studied scenario, the energy cost of the network infrastructure that links the user to the computing resources can be bigger than the energy cost of the servers.

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.

In 2015, this simulator has been extended to integrate two more green levers: low power idle (IEEE 802.3az) and adaptive link rate. This work has been done during the internship of Timothée Haudebourg (L3 ENS Rennes) and a publication is currently under preparation.

7.5.3. *Multicriteria scheduling for large-scale HPC environments*

Participant: Anne-Cécile Orgerie.

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. Energy efficiency must be taken into account and must be combined with other criteria like performance, resilience, Quality of Service.

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 will explore and propose a multicriteria scheduling model and framework for large scale HPC systems. Based on real energy measurements and calibrations, we will propose some performance and energy models and will build a multi criteria scheduler. Simulation on selected scenario will be explored and a prototype will be designed for ensuring experimental validation.

This work is done in collaboration with ROMA and Avalon teams from LIP in Lyon.

7.6. Decentralized and Adaptive workflows

Participants: Jean-Louis Pazat, Javier Rojas Balderrama, Matthieu Simonin, Cédric Tedeschi, Palakiyem Wallah.

7.6.1. *Adaptive Workflows with Chemical Computing*

Participants: Javier Rojas Balderrama, Matthieu Simonin, Cédric Tedeschi.

We have designed a high-level programming model based on the HOCL rule-based language to express workflow adaptation. It was specifically designed to support changes in the workflow logic at run time. This mechanism was implemented within the GinFlow software and experimented over the Grid'5000 platform. An article was just accepted for publication at the IPDPS 2016 conference.

7.6.2. *Best-effort decentralized workflow execution*

Participants: Jean-Louis Pazat, Cédric Tedeschi, Palakiyem Wallah.

We are currently proposing a simple workflow model for workflow execution in platforms with limited computing resources and services. The key idea is to devise a best-effort workflow engine that does not require a strong centralized orchestrator. Such a workflow engine relies on point-to-point cooperation between nodes supporting the execution. A minimalistic demonstrator of these concepts has been devised and implemented. Early experiments have been conducted on a single machine.

7.7. Experimental Platforms

Participants: Julien Lefeuvre, David Margery.

7.7.1. *Contribution to Fed4FIRE testbed*

Participants: Julien Lefeuvre, David Margery.

In Fed4FIRE, two key technologies have been adopted as common protocols to enable experimenters to interact with testbeds: Slice Federation Architecture (SFA), to provision resources, and Control and Management Framework for Networking Testbeds (OMF) to control them. In 2015, the main area of work has been the implementation of an SFA API to BonFIRE, still on-going. In the process, we wrote the reference documentation to write a new delegate for geni-tools, the reference implementation of SFA maintained by the GENI project office. This codebase has now been made public on github, in part because of our interactions with the code and suggested changes to ease writing new delegates. We have also contributed to the design of a service layer proxy mechanisms so that testbeds with http based APIs can be queried by any Fed4FIRE user using a standard authentications mechanism. The BonFIRE API has been made available through that mechanism, based on XML documents signed using the XML Signature specification.

8. Bilateral Contracts and Grants with Industry

8.1. Bilateral Contracts with Industry

8.1.1. *Thales Research and Technology*
Participants: Baptiste Goupille–lescar, Christine Morin, Nikolaos 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. From November 2015 to December 2015, we performed a state of the art study on resource management in virtualized computing infrastructures to cope with cyber-physical system constraints.

9. Partnerships and Cooperations

9.1. Regional Initiatives

9.1.1. *Security Monitoring of Clouds, collaboration with DGA-MI in the framework of the Cybersecurity Pole of Excellence (2014-2018)*
Participants: Anna Giannakou, Christine Morin, Jean-Louis Pazat, Louis Rilling, Amir Teshome Wonjiga.

Our study 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.

9.1.2. *CominLabs EPOC project (2013-2016)*
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. EcoPaaS, Brittany region SAD project(2014-2015)

Participants: Maria Del Mar Callau Zori, Anne-Cécile Orgerie, Guillaume Pierre.

Many research efforts have been dedicated to reducing cloud energy consumption, in particular by optimizing the Infrastructure-as-a-Service layer of the Cloud. Infrastructure-as-a-Service (IaaS) is the layer in charge of the virtualization of physical resources, and therefore has direct control over energy-related elements. However, the IaaS layer has no knowledge about the nature of applications which run over these resources, which limits the scope of decisions it can take.

The EcoPaaS project therefore aim at making the IaaS layer (in charge of resources) and the PaaS layer (in charge of applications) collaborate to further reduce the Cloud energy consumption. The idea is to define standard interfaces that allow both layers to exchange relevant information and to coordinate their actions. Exchanging information will for example allow the PaaS layer to estimate the energy consumption of each application it is running. Coordinating actions will in turn allow the system to avoid situations where both layers simultaneously take mutually-damaging actions. This project has been funding Maria del Mar Callau-Zori's postdoc.

9.1.4. IRT B-Com

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

Yvon Jégou and Jean-Louis Pazat are at IRT B-Com [4] 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.

9.2. National Initiatives

9.2.1. Inria ADT GinFlow (2014-2016)

Participants: Christine Morin, Javier Rojas Balderrama, 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 Héctor 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).

9.2.2. Inria IPL Discovery (2015-2019)

Participants: Anne-Cécile Orgerie, 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.

9.2.3. Inria IPL CityLab (2015-2018)

Participant: Christine Morin.

---

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

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 2015, Christine Morin contributed to the preparation of the MOOC entitled *Villes Intelligentes : défis technologiques et sociétaux* (Smart cities : technological and social challenges) to be run on the FUN platform starting in January 2016. She prepared eight sequences on urban data management in clouds.

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

Participants: Yvon Jégou, Christine Morin.

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 will provide software tools to efficiently manage and ease the use of a distributed computing infrastructure for the execution of different simulation applications.

In 2015, we developed a distributed framework which allows do 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. A first version is available since summer 2015. The framework is currently being repackaged to facilitate its reuse for new workloads.

9.2.5. 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.

9.2.6. 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. Fed4FIRE

Participants: Julien Lefeuvre, David Margery.

Type: FP7

Defi: Future internet experimental facility and experimentally-driven research

Instrument: Integrated Project

Objectif: ICT-2011.1.6 Future Internet Research and Experimentation (FIRE) with a specific focus on b) FIRE Federation

Duration: October 2012 - September 2016

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

Partners: Interdisciplinary institute for broadband technology (iMinds, Belgium), University of Southampton (It Innovation, United Kingdom) Universite Pierre et Marie Curie - paris 6 (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: David Margery

Abstract: 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. The baseline architectural choice for this project is to use the key concepts of the Slice Federation Architecture (SFA) to provision resources on experimental platforms, a Control and Management Framework for Networking Testbeds named OMF for experiment control and OML, the OMF Measurement library for data collection. We investigate whether these can be used to run experiments on BonFIRE and how they need to be extended to support the operating model of BonFIRE or Grid'5000.

9.3.1.2. HARNESS

Participants: Ancuta Iordache, Guillaume Pierre, Damian Serrano, Genc Tato, Georgios Ioannidis, Guillaume Pierre.

Type: COOPERATION

Defi: Pervasive and Trusted Network and Service Infrastructures

Instrument: Small or medium-scale focused research project

Objectif: ICT-2011.1.2 Cloud Computing, Internet of Services and Advanced Software Engineering

Duration: October 2012 - September 2015

Coordinator: Imperial College London (IMP, United Kingdom)

Partner: Ecole polytechnique fédérale de Lausanne (EPFL, Switzerland), Université de Rennes 1 (UR1, France), Konrad-Zuse-Zentrum für Informationtechniek Berlin (ZIB, Germany), Maxeler Technologies (MAX, United Kingdom), SAP AG (SAP, Germany)

UR1 contact: Guillaume Pierre

Abstract: The HARNESS FP7 project aims to incorporate innovative hardware and network technologies seamlessly into data centres that provide platform-as-a-service cloud infrastructures.

The dominant approach in offering cloud services today is based on homogeneous commodity resources: large numbers of inexpensive machines, interconnected by off-the-shelf networking equipment, supported by stock disk drives. However, cloud service providers are unable to use this platform to satisfy the requirements of many important and high-value classes of applications.

Today's cloud platforms are missing out on the revolution in new hardware and network technologies for realising vastly richer computational, communication, and storage resources. Technologies such as Field Programmable Gate Arrays (FPGA), General-Purpose Graphics Processing Units (GPGPU), programmable network routers, and solid-state disks promise increased performance, reduced energy consumption, and lower cost profiles. However, their heterogeneity and complexity makes integrating them into the standard Platform as a Service (PaaS) framework a fundamental challenge.

The HARNESS project brings innovative and heterogeneous resources into cloud platforms through a rich programme of research, validated by commercial and open source case studies.

9.3.1.3. PaaSage

Participants: Christine Morin, Nikolaos Parlavantzas, Aboozar Rajabi, Arnab Sinha.

Type: COOPERATION

Objectif: ICT-2011.1.2 Cloud Computing, Internet of Services and Advanced Software Engineering

Instrument: Collaborative Project

Duration: October 2012 - September 2016

Coordinator: GEIE ERCIM (France)

Partner: SINTEF (Norway), Science and Technology Facilities Council (UK), University of Stuttgart (Germany), Inria (France), Centre d'Excellence en Technologies de l'Information et de la Communication (Belgium), Foundation for Research and Technology Hellas (Greece), BE.Wan SPRL (Belgium), EVRY AS (Norway), SysFera SAS (France), Flexiant Limited (UK), Lufthansa Systems AG (Germany), Gesellschaft fur Wissenschaftliche Datenverarbeitung MBH Gottingen (Germany), Automotive Simulation Center Stuttgart (Germany), University of Ulm (Germany), Akademia Górniczo-Hutnicza im. Stanislawa Staszica (Poland), University of Cyprus (Cyprus), IBSAC-Intelligent Business Solutions ltd (Cyprus), University of Oslo (Norway)

Inria contact: Nikolaos Parlavantzas

See also: http://www.paasage.eu/

Currently there exist several open source and commercial services at the Infrastructure as a Service (IaaS) level. Software developers targeting the cloud would ideally want to develop their software once and be able to deploy it on any of the available services, reaping the benefits of a cloud market without losing on performance, availability, or any other service properties. The impediment to this objective is that IaaS platforms are heterogeneous, and the services and APIs that they provide are not standardized. Porting an existing application to one of these platforms or switching between platforms is thus a challenging task and involves a high risk that the results do not meet the expected requirements.

PaaSage delivers a development and deployment platform, with an accompanying methodology, with which developers of enterprise systems can access services of cloud platforms in a technology neutral approach while guiding developers to configure their applications for best performance. The consortium brings together ERCIM for management and STFC as scientific coordinator together with experts in different aspects of clouds ranging from software and services (SINTEF), High Performance Computing (HLRS) and systems development environments (Inria) to a group of SMEs working on cloud systems and end-user organisations with requirements in the cloud domain.

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: Professor 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.

9.3.2.2. MC-DATA

Participants: Teodor Crivat, Guillaume Pierre.

Program: EIT Digital

Project acronym: MC-DATA

Project title: Multi-cloud data management

Duration: Jan-Dec 2015

Coordinator: Dr. Peter Pietzuch, Imperial College London

Other partners: SICS, U-Hopper, VTT, Proxible.

Abstract: The goal of this activity is to increase developer uptake and commercial exploitation of the previously-developed MC-ConPaaS mobile edge cloud platform through new products in the area of location-based advertising services. The activity will (a) integrate MC-ConPaaS with the Android mobile platform to encourage mobile developer adoption; (b) pilot a location-based interactive advertising service with augmented reality and 3D tracking; (c) commercialise the pilot, demonstrating the business value of a mobile edge cloud model; and (d) model and shape the ecosystem of mobile edge cloud services, enabling new revenue streams for mobile operators.

9.4. International Initiatives

9.4.1. Inria International Labs

Christine Morin contributed to the edition of the 2011-2014 activity report of the Inria@SiliconValley Inria International Lab (https://project.inria.fr/siliconvalley/files/2015/06/Inria@SV_Activity_Report_2011_2014.pdf).

Christine Morin was one of the co-organizers of the BIS 2015 workshop held in Berkeley in May 2015 in the framework of the Inria@Silicon Valley Inria International Lab. Deb Agarwal co-chaired the panel on Big Data Science. Christine Morin chaired one of the two keynotes sessions.

9.4.2. *Inria Associate Teams*

9.4.2.1. *DALHIS*

Participants: Christine Morin, Anne-Cécile Orgerie, Javier Rojas Balderrama, Matthieu Simonin, Arnab Sinha, Cédric Tedeschi.

> Title: Data Analysis on Large Heterogeneous Infrastructures for Science
>
> International Partner (Institution - Laboratory - Researcher):
>
> > Lawrence Berkeley National Laboratory, Berkeley, USA
> >
> > Data Science and Technology department
> >
> > French PI: Christine Morin
> >
> > American PI: Deb Agarwal, head of the Data Science and Technology department
>
> Duration: 2013 - 2015
>
> See also: https://project.inria.fr/dalhis/

The worldwide scientific community is generating large datasets at increasing rates causing data analysis to emerge as one of the primary modes of science. Existing data analysis methods, tools and infrastructure are often difficult to use and unable to handle the "data deluge". A scientific data analysis environment needs to address three key challenges: a) programmability: easily composable and reusable programming environments for analysis algorithms and pipeline execution, b) agility: software that can adapt quickly to changing demands and resources, and, c) scalability: take advantage of all available resource environments including desktops, clusters, grids, clouds and HPC environments. The goal of the DALHIS associated team is to coordinate research and create together a software ecosystem to facilitate data analysis seamlessly across desktops, HPC and cloud environments. Specifically, our end goal is to build a dynamic environment that is user-friendly, scalable, energy-efficient and fault tolerant through coordination of existing projects. We plan to design a programming environment for scientific data analysis workflows that will allow users to easily compose their workflows in a programming environment such as Python and execute them on diverse high-performance computing (HPC) and cloud resources. We will develop an orchestration layer for coordinating resource and application characteristics. The adaptation model will use real-time data mining to support elasticity, fault-tolerance, energy efficiency and provenance. We investigate how to provide execution environments that allow users to seamlessly execute their dynamic data analysis workflows in various research environments.

The work done in 2015 on scientific workflows and energy efficiency is described respectively in 7.6.1 and 7.4.2.

The recent results of the DALHIS associate team were presented by Lavanya Ramakrishnan from LBNL during the working session on Scientific and Large Scale Computing. Christine Morin, Anne-Cécile Orgerie and Deb Agarwal participated in the BIS 2015 workshop held in Berkeley in May 2015.

Deb Agarwal has been awarded an Inria International Chair for the 2015-2019 period enabling long visits in the Myriads team. She was hosted in Myriads team during 2.5 months from May 1st to July 10th 2015. During this visit, we initiated the work on the design of a mobile application for reliable field data collection for FluxNet. Critical to the interpretation of global Fluxnet carbon flux dataset is the ancillary information and measurements taken at the measurement tower sites (e.g. vegetation species, leaf area index, instrument calibrations, etc). The submission and update of this data using excel sheets is difficult and error prone. In 2015, the team developed some initial sketches of the User Interface design for a mobile application for the reliable collection of FluxNet data and Arnab Sinha, Deb Agarwal, and Christine Morin performed an initial usability feedback interview with Chris Flechard (INRA Rennes), a CarboEurope participant who collects carbon flux data at several sites in Brittany. M. Sandesh (LBL) simultaneously performed a couple of usability interviews at Berkeley. We updated the design based on the combined feedback. Currently, the mobile application

prototype development is in progress. The design was presented by Dario Papale at the ICOS meeting in September 2015. The expectation is that the design will be adopted by ICOS (European flux towers) and AmeriFlux (flux towers in the Americas). A first basic working demonstration prototype has been developed.

9.4.3. Inria International Partners

Northeastern University We started a collaboration with Professor Gene Cooperman, Northeastern University, Boston, USA on the design of a cloud agnostic checkpointing service on top of IaaS clouds for reliable application execution, inter-cloud application migration and easing application "cloudification". Gene Cooperman was hosted in Myriads team during a week in June 2015.

ORNL/TTU We collaborate on cloud computing with Stephen Scott, Professor at Tennessee Tech University (TTU) and researcher at Oak Ridge National Laboratory (ORNL) in the USA. He participated in Anna Giannakou's mid-PhD thesis defense in October 2015.

University of Guadalajara Nikolaos Parlavantzas is collaborating with the team of Prof. Héctor Duran-Limon of the University of Guadalajara, Mexico, on adaptive resource management in cloud environments.

VU University Amsterdam We collaborate with Thilo Kielmann's research group at VU University Amsterdam on research and development around the ConPaaS system. This collaboration has lead to a joint publication this year [24].

EPFL 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. This collaboration has been supported thanks to the extended visit of Georgios Ioannidis in Rennes in the context of the HARNESS project, and it is expected to continue after the end of HARNESS. At least one joint publication on this topic is currently in preparation.

9.5. International Research Visitors

9.5.1. Visits of International Scientists

Gene Cooperman, Professor at the Northeastern University (Boston, USA), made a one-week visit in Myriads team in June 2015.

Georgios Ioannidis (PhD student at EPFL, Switzerland) made a 7-months visit in the Myriads team (Jun-Dec 2015). The goal was to reinforce the collaboration between the two teams in the context of the HARNESS FP7 project.

Carlos Ruiz Diaz (PhD student at the University of Guadalajara, Mexico) is visiting Myriads for 6 months (Sep 2015-Feb 2016) in the context of his PhD thesis, directed by Héctor Duran-Limon and co-advised by Nikolaos Parlavantzas. The visit is supported by a grant from Rennes Metropole.

Palakiyem Wallah, assistant professor at the University of Kara (Togo) visited Myriads team from September to December 2015 in the framework of his PhD thesis, which is co-advised by Cédric Tedeschi and Jean-Louis Pazat.

Anita Sobe, post-doctoral researcher at the University of Neuchâtel (Switzerland) visited Myriads team for two weeks in April 2015 in the context of the Nesus COST Action. During her stay, she has worked with Anne-Cécile Orgerie and their work has been accepted in the PDP 2016 conference.

9.5.1.1. Internships

Philippe Fabian (M1 Université Rennes 1) did his internship of first year of Master under the supervision of Marin Bertier and Cédric Tedeschi. Philippe devised and tested in a simulator heuristics for speeding-up the execution of chemical programs on top of an unstructured decentralized platform.

Timothée Haudebourg (L3 ENS Rennes) has done a two-month internship (June - July 2015) under the supervision of Anne-Cécile Orgerie. He has worked on quantifying the energy-efficiency of green leverages in wired networks.

David Guyon (M2 Université Rennes 1) has done his master internship under the supervision of Anne-Cécile Orgerie and Christine Morin. He has worked on energy-efficient cloud elasticity for data-driven applications. He has presented this work at the IEEE GreenCom conference (December 2015).

Ghada Moualla (M2 Université Rennes 1) has done his master internship under the supervision of Christine Morin and Matthieu Simonin. She has worked on reliable and efficient data processing in a cloud environment.

Akshat Puri (M2, EIT ICT Labs, Université Rennes 1) has done his internship under the supervision of Nikolaos Parlavantzas and Guillaume Pierre. He worked on elasticity of cloud applications and approaches for application migration across different cloud vendors.

Benjamin Soulas (M2 Université Rennes 1) did its internship of second year of Master under the supervision of Matthieu Simonin and Cédric Tedeschi. Benjamin developed the *Storm Watcher* prototype, a tool exposing monitoring information about the execution of programs running within the Storm framework.

9.5.1.2. Research stays abroad

Following a first visit which took place in 2014, Ancuta Iordache visited Maxeler Technologies (London, U.K.) again from February 2015 to April 2015. This visit reinforced the collaboration between the two teams in the context of the HARNESS E.U. project, and was funded by the EIT Digital Doctoral Training Center. A concrete outcome of this visit is a joint research paper which is currently under evaluation.

Ismael Cuadrado Cordero, who is a student of the EIT ICT Labs Doctoral School, visited the Queen Mary University of London (UK) for a research internship from June to August 2015. He was hosted in Chris Phillip's team working on micro-clouds architectures for neighborhood services. A joint research paper is currently under review on this topic.

10. Dissemination

10.1. Promoting Scientific Activities

10.1.1. Scientific events organisation

10.1.1.1. General chair, scientific chair

- Christine Morin and Deb Agarwal were co-chairs of the Science Data Ecosystem workshop, held in the framework of the Data Science Symposium organized by IRISA for the 40th anniversary celebration (November 26th, 2015, Rennes, France)
- Louis Rilling was also co-organizer of the SEC2 2015 workshop, co-located with Compas 2015.
- Anne-Cécile Orgerie was co-chair of the ExtremeGreen workshop, in conjunction with CCGrid 2015 (Shenzhen, China, May 2015).

10.1.1.2. Member of the organizing committees

- Martin Quinson was the main organizer of the 2015 SimGrid's User Days (Lyon, France, 8-12 June 2015), gathering the community of users.
- Christine Morin was a member of the scientific committee for the Inria Scientific Days held in Nancy in June 2015.

10.1.2. Scientific events selection

10.1.2.1. Chair of conference program committees

- Guillaume Pierre was program committee co-chair of the EIT Digital Symposium on the Future of Cloud Computing (Rennes, France, 19-20 October 2015).

10.1.2.2. Member of the conference program committees

- Christine Morin was a member of the program committee of the ACM/IEEE CCGrid 2015 conference.

- Christine Morin was a member of the program committee of the first IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN 2015) conference.

- Christine Morin was a member of the program committee of the 23rd Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP 2015).

- Christine Morin was a member of the technical posters committee of the 2015 ACM/IEEE International Conference on High Performance Computing, Networking, Storage and Analysis (SC15).

- Christine Morin was a member of the program committee of the Compas 2015 conference.

- Christine Morin was a member of the program committee of the 6th International workshop on Scientific Cloud Computing (ScienceCloud 2015) co-located with the ACM HPDC 2015 conference.

- Christine Morin was a member of the program committee of the 8th International workshop on Virtualization Technologies in Distributed Computing (VTDC-2015) co-located with the ACM HPDC 2015 conference.

- Christine Morin was a member of the program committee of the 10th international workshop on Virtualization in High-Performance Cloud Computing (VHPC 2015) co-located with the EuroPar 2015 conference.

- Christine Morin was a member of the program committee of the Resilience 2015 workshop co-located with the EuroPar 2015 conference.

- Anne-Cécile Orgerie was a member of the program committee of the CCGrid 2015 conference.

- Anne-Cécile Orgerie was a member of the program committee of the CloudNet 2015 conference.

- Nikolaos Parlavantzas was a member of the program committee of the CLOSER 2015 conference.

- Guillaume Pierre was a member of the program committee of the DADS track of the SAC 2015 conference.

- Guillaume Pierre was a member of the program committee of the IC2E 2015 conference.

- Guillaume Pierre was a member of the program committee of the AINA 2015 conference.

- Guillaume Pierre was a member of the program committee of the CSCS20-2015 conference.

- Guillaume Pierre was a member of the program committee of the CloudDP 2015 workshop.

- Guillaume Pierre was a member of the program committee of the VTDC 2015 workshop.

- Guillaume Pierre was a member of the program committee of the HotWeb 2015 workshop.

- Martin Quinson was a member of the program committee of the ACM SIGSIM Principles of Advanced Discrete Simulation (PADS 2015) conference.

- Martin Quinson was a member of the program committee of the High Performance Computing & Simulation (HPCS 2015) conference.

- Martin Quinson was a member of the program committee of the SimulTech 2015 conference.

- Martin Quinson was a member of the program committee of the Parallel and Distributed Computing Education for Undergraduate Students (Euro-EduPar 2015) workshop.

- Louis Rilling was program committee member of the DIHC 2015 workshop co-located with EuroPar 2015.

- Cédric Tedeschi was a member of the program committee of the ICWS 2015 conference.

- Cédric Tedeschi was a member of the program committee of the ICCS 2015 conference.

- Cédric Tedeschi was a member of the program committee of the CloSer 2015 conference.

10.1.3. Journal

10.1.3.1. Member of the editorial boards

Christine Morin has been appointed associate editor in the IEEE's Transactions on Parallel and Distributed Systems' Editorial Board since June 2015.

- Martin Quinson is a member since 2013 of the editorial board of the Interstices journal (edited by Inria in collaboration), aiming at increasing the scientific outreach of informatics.

10.1.4. Invited talks

- *Breaking the cloud : Energy-Efficient & Network-aware Distributed Cloud Infrastructure.* Ismael Cuadrado-Cordero, talk at GreenDays. Toulouse, France, 16-17 March 2015.

- *Resource allocation in a Cloud partially powered by renewable energy sources.* Yunbo Li, talk at GreenDays. Toulouse, France, 16-17 March 2015.

- *Breaking the cloud: Energy-efficient cloud PaaS.* Ismael Cuadrado-Cordero, poster presentation at the EIT Digital Symposium on the Future of Cloud Computing. Rennes, France, 19-20 October 2015.

- *Breaking the cloud: Energy-efficient cloud PaaS.* Ismael Cuadrado-Cordero, poster presentation at Green ICT day at IRISA. Rennes, France, 30 November 2015.

- *Self Adaptation for Security Monitoring in IaaS Clouds.* Anna Giannakou, 3 minute talk at the D1 Department Day. Rennes, France, 26 May 2015.

- *Self Adaptation in Security Monitoring for IaaS clouds.* Anna Giannakou, 5 minute talk at the SEC2-Premier atelier sur la Securite dans les Clouds. Lille, France, 30 June 2015.

- *Self Adaptable Security Monitoring in IaaS Clouds.* Anna Giannakou, poster presentation at the EIT Digital Symposium on the Future of Cloud Computing. Rennes, France, 19-20 October 2015

- *Self Adaptable Security Monitoring in IaaS Clouds.* Anna Giannakou, poster presentation at Club recherche du Pôle d'Excellence Cyber. Rennes, France 8 December 2015.

- *Private-by-design: towards Personal Local Clouds.* Christine Morin, invited presentation at the Dagstuhl seminar on the Distributed Cloud Computing (DCC), Dagstuhl, Germany, February 10, 2015.

- *Autonomic and Energy-Efficient Management of Large-Scale Virtualized Data Centers* and *Towards High Performance Computing in Clouds.* Christine Morin, invited tutorials at the *Seconda Università di Napoli*, Napoli, Italy, April 19-20 2015.

- *ConPaaS Generic Service for Automated Application Deployment in Distributed Clouds.* Christine Morin, presentation at the Inria Industry Meeting on Smart City & Mobility Innovations: Towards environmental and social sustainability, San Francisco, USA, May 11, 2015.

- *Security of Information Systems hosted in Clouds: SLA Definition and Enforcement in a Dynamic Environment.* Christine Morin, invited presentation at the SEC 2 workshop on security in clouds, co-located with Compas 2015, Lille, France, June 30, 2015.

- *Debate on Ecological Transitions.* Christine Morin, panelist, Nantes, France, September 24, 2015.

- *Energy Efficiency & Energy Saving in Clouds.* Christine Morin, presentation at the IRISA Green IT meeting, Rennes, France, November 30, 2015.

- *Self-adaptation for Security Monitoring in IaaS clouds.* Anna Giannakou, Christine Morin, and Louis Rilling, invited poster at the Cyber Pole of Excellence's research club meeting, Rennes, France, November 30, 2015.

- Anne-Cécile Orgerie gave a keynote entitled "Energy efficiency in Cloud computing infrastructures", at ExtremeGreen workshop in conjunction with CCGrid 2015, Shenzhen, China, May 4, 2015.

- Anne-Cécile Orgerie gave a talk entitled "Le Green IT à l'IRISA", Matinale Rennes Atalante on Smart Grids, Rennes, July 8, 2015.

- Anne-Cécile Orgerie gave a talk entitled "Green computing" and participated to a debate with Françoise Berthoud (CNRS), at the Forum du CNRS, Paris, December 12, 2015.

- *HARNESS: bringing real hardware heterogeneity to the cloud.* Guillaume Pierre, invited presentation at the 5th International Workshop on Cloud Data and Platforms. Co-located with EuroSys 2015 in Bordeaux, France. April 21, 2015.

- *HARNESS: bringing real hardware heterogeneity to the cloud.* Guillaume Pierre, invited presentation at the 8th International Workshop on Virtualization Technologies in Distributed Computing. Co-located with HPDC 2015, Portland, OR, USA, June 15th 2015.
- *Accelerating clouds with FPGA virtualization.* Ancuta Iordache, poster presentation at the EIT Digital Symposium on the Future of Cloud Computing. Rennes, France, 19-20 October 2015.
- *Accelerating clouds with FPGA virtualization.* Ancuta Iordache, presentation at the 5th Workshop on Storage and Data Analytics. Cesson-Sévigné, France, 19 November 2015.
- *Cloud application performance modeling in ConPaaS.* Guillaume Pierre, invited presentation at the Paris Open Source Summit. Paris, France, 18-19 November 2015.
- *An Edge Cloud Architecture for Latency-Critical Interactive Applications.* Teodor Crivat, poster presentation + demo at the EIT Digital Symposium on the future of Cloud computing. Rennes, France, 19-20 October 2015.
- *An Edge Cloud Architecture for Latency-Critical Interactive Applications.* Teodor Crivat, poster presentation + demo at the EIT Digital results day. Helsinki, Finland, December 1st 2015.

10.1.5. Scientific expertise

- Christine Morin acted as a reviewer for the University College Cork (UCC) Research Quality Review (April 2015). She was one of the scientific experts to examine a promotion as Senior University Lecturer in Computing Science and Engineering at the Faculty of Science and Technology of Umeå University (August 2015). She was a member of the ModaClouds European project Advisory Board. She is a member of the scientific council of ENS Cachan. She was a member of the admission jury for the 2015 Inria senior researcher competition.

  She was a member of the selection committees for a Professor position in Computer Science at ESIR engineering school at the University of Rennes 1 and for a Professor position in Computer Science at ENS Rennes.
- Anne-Cécile Orgerie participated to the evaluation of three Doctoral schools in January 2015 as a young researcher expert for HCERES.
- Guillaume Pierre participated in the evaluation panels of ongoing national (ANR) projects LISP-LAB, MOEBUS, PROFIL, ABCD, and DataTweet. Paris, France, 27-27 October 2015.
- Nikolaos Parlavantzas acted as an expert reviewer for ANRT CIFRE, PHC Orchid, and FFCR projects.

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 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 et systèmes 2 (60 hours ETD at ENS Rennes)
- Master 2: Green ICT - 3 hours of invited lecture at Telecom SudParis

Nikolaos Parlavantzas (at INSA Rennes):

- 4th year: Operating Systems (40 hours ETD)
- 4th year: Big Data and Applications (25 hours ETD)
- 4th year: Networking and SOA (12 hours ETD)
- 4th year: Advanced Operating Systems (12 hours ETD)
- 4th year: Parallel programming (12 hours ETD)
- 4th year: Software Development Project (30 hours ETD)
- 5th year: Component-based Software Engineering (16 hours ETD)

Jean-Louis Pazat is responsible for the following graduate teaching Modules: Advanced operating Systems, Parallel Computing, Clouds, Networking and SOA.

Jean-Louis Pazat (at INSA Rennes):

- 4th year: Advanced Operating Systems and Clouds (32 hours ETD)
- 4th year: Parallel Programming (48 hours ETD)
- 4th year: Networking and SOA (48 hours ETD)
- 4th year: Software development project (60 hours ETD)

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 ETC)
- 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)

Martin Quinson:

- Licence 3: Programming and Software Engineering (30 hours ETD at ENS Rennes)
- Master 1: Concurrency in Systems and Networks (20 hours ETD at University of Rennes 1)
- Master 2: Pedagogy and Scientific Mediation for Computer Science (30 hours EDT at ENS Rennes)

Cédric Tedeschi (at the University of Rennes 1):

- Master 1: Concurrency in Systems and Networks (70 hours ETD)
- Master 1: Parallel Programming (36h hours ETD)
- Master 1: Service Techonologies (30h hours ETD)
- Master 2 (research): Internet of Services (6 hours ETD)

    E-learning

    Christine Morin contributed to a MOOC entitled "*Villes intelligentes : défis technologiques et sociaux*" that will be available on the FUN platform in January 2016.

    Cédric Tedeschi is responsible for the Operating Systems class within the eMiage online teaching program.

    Martin Quinson extended the Programmer's Learning Machine (PLM) with another lesson on recursive programming. The 12 new exercises revisit the old problem of Hanoï towers with several less common variants.

10.2.2. Supervision

PhD in progress : Ancuta Iordache, Multi-resource optimization for application hosting in heterogeneous clouds, started in February 2013, 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: 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: 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: Sabbir Hasan, SLA Driven Cloud autoscaling for optimizing energy footprint, started in December 2013, Thomas Ledoux (Ascola), Jean-Louis Pazat.

PhD in progress: Anna Giannakou, A Self-adaptable Security Supervision System for Clouds, started in March 2014, Christine Morin, Jean-Louis Pazat, Louis Rilling.

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

PhD in progress: Amir Teshome Wonjiga, Definition and Enforcement of Service Level Agreements for Cloud Security Supervision, started in October 2015, Christine Morin, Louis Rilling.

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: Palakiyem Wallah, Best-Effort Execution of Service Workflows, started in October 2013, Ayité Ajavon (University of Lomé, Togo), Jean-Louis Pazat, Cédric Tedeschi.

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: 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).

10.2.3. Juries

- Christine Morin was a member (chair) of the PhD defense committee of Sylvain Gault, ENS Lyon (March 23rd 2015).
- Christine Morin was a reviewer of the HdR committee of Gilles Fedak, ENS Lyon (May 28th, 2015).
- Christine Morin was a reviewer of the HdR committee of Vania Marangozova-Martin, Grenoble University (June 12th, 2015).
- Christine Morin was a reviewer of the PhD defense committee of Tatiana Martsinkevich, Paris-Sud University (September 22nd 2015).
- Christine Morin was a member (chair) of the PhD defense committee of Arnaud Lefray, ENS Lyon (October 3rd 2015).
- Christine Morin was a reviewer of the HdR committee of Georges DaCosta, Toulouse University (November 12th, 2015).
- Christine Morin was a reviewer of the PhD defense committee of Aline Bousquet, Orléans University (December 2nd 2015).
- Christine Morin was a reviewer of the PhD defense committee of Damien Riquet, Lille 1 University (December 3rd 2015).
- Guillaume Pierre was a reviewer in the PhD committee of Mihai Dobrescu, EPFL, 3 July 2015.

- Guillaume Pierre was a reviewer in the PhD committee of Kaveh Razavi, Vrije Universiteit Amsterdam, 5 November 2015.
- Guillaume Pierre was a reviewer in the PhD committee of Diego Didona, Lisbon University, 20 November 2015.
- Martin Quinson was a member in the PhD committee of Luka Stanisic, Grenoble University, 30 October 2015.
- Louis Rilling was a member of the PhD defense committee of Pierrick Buret, Université de Limoges (December 1st 2015).
- Louis Rilling was a member of the PhD defense committee of Anis Bkakria, Télécom Bretagne (December 17th 2015).
- Cédric Tedeschi was a member in the PhD committee of Maurice Faye, ENS de Lyon, 10 November 2015.

10.3. Popularization

- Martin Quinson organized a 2-days workshop for secondary maths teachers on how algorithms could be used to reinforce the pupils abilities to verbalize, set out arguments and conduct rigorous demonstrations (Nancy, March 2-3 2015). This workshop was co-organized by the IREM (Institut de Recherche en Enseignement des Mathématiques) of Nancy.
- Martin Quinson gave a talk entitled "Computational Science: the real societal impact of computers", during the workshop presented in previous item (Nancy, March 3th 2015).
- Martin Quinson co-organized a one day workshop for about 100 secondary computer science teachers. This event spreads the best practices to teach our topic (Nancy, March 12th 2015).
- Martin Quinson was a scientific expert in an experiment in which we explored how Scratch can be used to teach Computer Science in after school activities every week for the whole week.
- Martin Quinson is a scientific expert in a teaching manual of Computer Science for primary schools, currently authored within the «La main à la pâte» research network. This manual will be released in 2016.
- Martin Quinson organized a collaborative workshop on Wikipedia gathering students and researchers to improve the French speaking pages on Computer Science in the free encyclopedia (Rennes, December 10, 2015).

10.4. Miscellaneous

Christine Morin is a member of the Project-Team Committee of INRIA RENNES – BRETAGNE ATLANTIQUE (Comité des projets), a member of the board of the Project-Team Committee of INRIA RENNES – BRETAGNE ATLANTIQUE (since November 2015), *Référent Chercheur* for INRIA RENNES – BRETAGNE ATLANTIQUE. She was a member of the Irisa/Inria "Commission Personnel" being in charge of post-docs and "*délégations*" (until October 2015). She was corresponding for North America relationships in the Inria Direction of European and International Relationships (until November 2015).

Anne-Cécile Orgerie is in charge (chargée de mission) of the "Green IT" transversal axis of IRISA.

Nikolaos Parlavantzas is the local coordinator for the international exchange of students at the computer science department of Insa.

Guillaume Pierre is the local coordinator of the EIT Digital master school in Rennes, and the local coordinator of the Erasmus exchange program between Université de Rennes 1 and Politehnica University Bucharest.

Jean-Louis Pazat is the leader of the "Large Scale Systems" department of IRISA. He is the leader of the computer Science Lab at INSA (IRISA-INSA). He is the coordinator for reviews of international bilateral cooperation projects at the Ministry of Research and Higher Education in the STIC domain.

Thierry Priol is the director of the Inria European and International Partnership department.

Cédric Tedeschi is a member of the administration council of the EECS departement of the University of Rennes 1.

Yvon Jégou is a member of the Grid 5000 executive committee. He is a member of the *Comité de Sélection et de Validation (CSV)* of the *Images & Réseaux* cluster.

11. Bibliography

Major publications by the team in recent years

[1] M. BERTIER, M. OBROVAC, C. TEDESCHI. *Adaptive atomic capture of multiple molecules*, in "Journal of Parallel and Distributed Computing", September 2013, vol. 73, n⁰ 9, pp. 1251-1266, http://hal.inria.fr/hal-00915220

[2] R. G. CASCELLA, B. LORENZO, Y. JÉGOU, M. COPPOLA, C. MORIN. *Contrail: Distributed Application Deployment under SLA in Federated Heterogeneous Clouds*, in "FIA book 2013", A. GALIS, A. GAVRAS (editors), Lecture Notes in Computer Science, Springer, May 2013, vol. 7858, https://hal.inria.fr/hal-00805713

[3] S. COSTACHE, N. PARLAVANTZAS, C. MORIN, S. KORTAS. *Merkat: A Market-based SLO-driven Cloud Platform*, in "5th IEEE International Conference on Cloud Computing Technology and Science (CloudCom 2013)", Bristol, United Kingdom, December 2013, http://hal.inria.fr/hal-00862509

[4] E. FELLER, L. RILLING, C. MORIN. *Snooze: A Scalable and Autonomic Virtual Machine Management Framework for Private Clouds*, in "12th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing (CCGrid 2012)", Ottawa, Canada, May 2012, http://hal.inria.fr/hal-00664621

[5] H. FERNANDEZ, C. TEDESCHI, T. PRIOL. *A Chemistry-Inspired Workflow Management System for a Decentralized Workflow Execution*, in "Transactions on Services Computing", January 2013, http://hal.inria.fr/hal-00915222

[6] A. IORDACHE, C. MORIN, N. PARLAVANTZAS, E. FELLER, P. RITEAU. *Resilin: Elastic MapReduce over Multiple Clouds*, in "13th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing", Delft, Netherlands, ACM, 2013, http://hal.inria.fr/hal-00790406

[7] A. LAGE FREITAS, N. PARLAVANTZAS, J.-L. PAZAT. *An Integrated Approach for Specifying and Enforcing SLAs for Cloud Services*, in "The IEEE 5th International Conference on Cloud Computing (CLOUD 2012)", Honolulu, United States, June 2012, http://hal.inria.fr/hal-00703129

[8] J.-L. PAZAT, C. WANG. *A Chemistry-Inspired Middleware for Self-Adaptive Service Orchestration and Choreography*, in "CCGRID", Delft, Netherlands, P. BALAJI, D. EPEMA, T. FAHRINGER (editors), IEEE, May 2013, pp. 426 - 433 [*DOI :* 10.1109/CCGRID.2013.51], http://hal.inria.fr/hal-00934099

[9] G. PIERRE, C. STRATAN. *ConPaaS: a Platform for Hosting Elastic Cloud Applications*, in "IEEE Internet Computing", September-October 2012, vol. 16, n⁰ 5, pp. 88–92, http://www.globule.org/publi/CPHECA_ic2012.html

Publications of the year

Articles in International Peer-Reviewed Journals

[10] M. BAGEIN, J. BARBOSA, V. BLANCO, I. BRANDIC, S. CREMER, S. FREMAL, H. KARATZA, L. LEFÈVRE, T. MASTELIC, A. OLEKSIAK, A.-C. ORGERIE, G. L. STAVRINIDES, S. VARRETTE. *Energy Efficiency for Ultrascale Systems: Challenges and Trends from Nesus Project*, in "Supercomputing frontiers and innovations ", September 2015, vol. 2, n⁰ 2, pp. 105-131 [*DOI :* 10.14529/JSFI150206], https://hal.inria.fr/hal-01196798

[11] 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", May 2015, vol. 79-80, pp. 80-89, https://hal.inria.fr/hal-01102203

[12] Y. GLOUCHE, A. SINHA, P. COUDERC. *A Smart Waste Management with Self-Describing Complex Objects*, in "International Journal On Advances in Intelligent Systems", June 2015, vol. 8, n⁰ 1 & 2, pp. 1-16, Invited journal, https://hal.inria.fr/hal-01198382

[13] M. S. HASAN, Y. KOUKI, T. LEDOUX, J.-L. PAZAT. *Exploiting Renewable sources: when Green SLA becomes a possible reality in Cloud computing*, in "IEEE Transactions on Cloud Computing", July 2015, vol. PP, n⁰ 99, 1 p. [*DOI :* 10.1109/TCC.2015.2459710], https://hal.archives-ouvertes.fr/hal-01187907

[14] G. L. KLEIN, G. PIERRE, M.-N. BELLON-FONTAINE, M. GRABER. *Inverse Gas Chromatography with Film Cell Unit: An Attractive Alternative Method to Characterize Surface Properties of Thin Films*, in "Journal of Chromatographic Science", February 2015, vol. 53, n⁰ 8 [*DOI :* 10.1093/CHROMSCI/BMV008], https://hal.archives-ouvertes.fr/hal-01247054

International Conferences with Proceedings

[15] N. BELDICEANU, B. DUMAS FERIS, P. GRAVEY, M. 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, https://hal.archives-ouvertes.fr/hal-01131602

[16] J. CAO, M. SIMONIN, G. COOPERMAN, C. MORIN. *Checkpointing as a Service in Heterogeneous Cloud Environments*, in "15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CC-GRID 2015)", Shenzhen, Guangdong, China, May 2015, https://hal.inria.fr/hal-01102094

[17] 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

[18] 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

[19] 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

[20] Y. KOUKI, M. S. HASAN, T. LEDOUX. *Delta Scaling: How Resources Scalability/Termination Can Be Taken Place Economically?*, in "11th World Congress on Services (SERVICES)", New York, United States, IEEE, June 2015, pp. 55-62 [*DOI :* 10.1109/SERVICES.2015.17], https://hal.archives-ouvertes.fr/hal-01162745

[21] 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

[22] A.-C. ORGERIE. *Interconnecting Smart Grids and Clouds to save Energy*, in "International Conference on Smart Cities and Green ICT Systems (SMARTGREENS)", Lisbon, Portugal, May 2015, 6 p. , https://hal.inria.fr/hal-01160234

[23] E. OUTIN, J.-E. DARTOIS, O. BARAIS, J.-L. PAZAT. *Enhancing Cloud Energy Models for Optimizing Datacenters Efficiency*, in "IEEE International Conference on Cloud and Autonomic Computing (ICCAC)", Cambridge, MA, United States, IEEE, September 2015, 8 p. [*DOI :* 10.1109/ICCAC.2015.10], https://hal.inria.fr/hal-01243146

[24] 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

Conferences without Proceedings

[25] A. GIANNAKOU, L. RILLING, J.-L. PAZAT, F. MAJORCZYK, C. MORIN. *Towards Self Adaptable Security Monitoring in IaaS Clouds*, in "15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CC-GRID 2015)", Shenzen, China, May 2015, https://hal.inria.fr/hal-01165134

[26] 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

[27] T. KIRKHAM, A. SINHA, N. PARLAVANTZAS, B. KRYZA, P. FREMANTLE. *Privacy Aware On-Demand Resource Provisioning for IoT Data Processing*, in "IoTaaS 2015 - 2nd EAI International Conference on IoT as a Service", Rome, Italy, October 2015, https://hal.inria.fr/hal-01198381

Scientific Books (or Scientific Book chapters)

[28] A. GAZO CERVERO, M. CHINCOLI, L. DITTMANN, A. FISCHER, A. E. GARCIA, J. GALÁN-JIMÉNEZ, L. LEFÈVRE, H. DE MEER, T. MONTEIL, P. MONTI, A.-C. ORGERIE, L.-F. PAU, C. PHILLIPS, S. RICCIARDI, R. SHARROCK, P. STOLF, T. TRINH, L. VALCARENGHI. *Green Wired Networks*, in "Large-Scale Distributed Systems and Energy Efficiency", Wiley, March 2015, pp. 41-80 [*DOI :* 10.1002/9781118981122.CH3], https://hal.inria.fr/hal-01196819

Research Reports

[29] 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

[30] V. ISSARNY, T. CASTRO, H. KIRCHNER, C. MORIN. *Inria@SiliconValley Activity Report 2011-2014*, Inria, June 2015, https://hal.inria.fr/hal-01199428

[31] A. LEBRE, J. PASTOR, . THE DISCOVERY CONSORTIUM. *The DISCOVERY Initiative - Overcoming Major Limitations of Traditional Server-Centric Clouds by Operating Massively Distributed IaaS Facilities*, Inria, September 2015, n⁰ RR-8779, 14 p. , https://hal.inria.fr/hal-01203648

[32] J. ROJAS BALDERRAMA, M. SIMONIN, C. TEDESCHI. *A Chemistry-inspired Programming Model for Adaptive Decentralised Workflows*, , February 2015, n⁰ RR-8691, https://hal.inria.fr/hal-01120274

[33] J. ROJAS BALDERRAMA, M. SIMONIN, C. TEDESCHI. *GinFlow: A Decentralised Adaptive Workflow Execution Manager*, Inria ; Rennes 1, November 2015, n⁰ RR-8813, 27 p. , https://hal.inria.fr/hal-01232682

Scientific Popularization

[34] F. BERTHOUD, E. DREZET, L. LEFÈVRE, A.-C. ORGERIE. *L'épidémie du smartphone : prolifération et dissémination des composants électroniques*, in "Interstices", June 2015, https://hal.inria.fr/hal-01250139

[35] F. BERTHOUD, E. DREZET, L. LEFÈVRE, A.-C. ORGERIE. *La déferlante des données*, in "Interstices", July 2015, https://hal.inria.fr/hal-01250133

[36] F. BERTHOUD, E. DREZET, L. LEFÈVRE, A.-C. ORGERIE. *Le syndrome de l'obésiciel: des applications énergivores*, in "Interstices", July 2015, https://hal.inria.fr/hal-01250135

[37] F. BERTHOUD, E. DREZET, L. LEFÈVRE, A.-C. ORGERIE. *Sciences du numérique et développement durable : des liens complexes*, in "Interstices", June 2015, https://hal.inria.fr/hal-01250136

Other Publications

[38] R. CARVAJAL-GÓMEZ, D. FREY, M. SIMONIN, A.-M. KERMARREC. *WebGC Gossiping on Browsers without a Server [Live Demo/Poster]*, November 2015, Web Information System Engineering, Poster, https://hal.inria.fr/hal-01251787

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

References in notes

[40] K. JEFFERY, B. NEIDECKER-LUTZ (editors). *The Future of Cloud Computing - Opportunities for European Cloud Computing Beyond 2010*, European Commission, January 2010

[41] G. AGHA. *Computing in pervasive cyberspace*, in "Commun. ACM", 2008, vol. 51, n⁰ 1, pp. 68-70

[42] S. AL-KISWANY, D. SUBHRAVETI, P. SARKAR, M. RIPEANU. *VMFlock: virtual machine co-migration for the cloud*, in "HPDC", A. B. MACCABE, D. THAIN (editors), ACM, 2011, pp. 159-170, http://dblp.uni-trier.de/db/conf/hpdc/hpdc2011.html#Al-KiswanySSR11

[43] M. ARMBRUST, A. FOX, R. GRIFFITH, A. D. JOSEPH, R. KATZ, A. KONWINSKI, G. LEE, D. PATTERSON, A. RABKIN, I. STOICA, M. ZAHARIA. *A view of cloud computing*, in "Commun. ACM", April 2010, vol. 53, pp. 50–58, http://doi.acm.org/10.1145/1721654.1721672

[44] J.-P. BANÂTRE, P. FRADET, D. LE MÉTAYER. *Gamma and the Chemical Reaction Model: Fifteen Years After*, in "WMP", 2000, pp. 17-44

[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] S. COSTACHE. *Market-based Autonomous Application and Resource Management in the Cloud*, Université Rennes 1, July 2013, https://tel.archives-ouvertes.fr/tel-00925352

[47] S. COSTACHE, N. PARLAVANTZAS, C. MORIN, S. KORTAS. *On the Use of a Proportional-Share Market for Application SLO Support in Clouds*, in "Europar -19th International European Conference on Parallel and Distributed Computing - 2013", Aachen, Germany, August 2013, https://hal.inria.fr/hal-00821558

[48] U. DESHPANDE, X. WANG, K. GOPALAN. *Live gang migration of virtual machines*, in "HPDC", A. B. MACCABE, D. THAIN (editors), ACM, 2011, pp. 135-146, http://dblp.uni-trier.de/db/conf/hpdc/hpdc2011.html#DeshpandeWG11

[49] D. DIB, N. PARLAVANTZAS, C. MORIN. *Meryn: Open, SLA-driven, Cloud Bursting PaaS*, in "International Workshop on Optimization techniques for Resources Management in Clouds (ORMaCloud'13)", New York, United States, June 2013, https://hal.inria.fr/hal-00816293

[50] E. FELLER. *Autonomic and Energy-Efficient Management of Large-Scale Virtualized Data Centers*, Université Rennes 1, December 2012, http://tel.archives-ouvertes.fr/tel-00785090

[51] E. FELLER, C. MORIN, M. SIMONIN, A.-C. ORGERIE, Y. JÉGOU. *Scalability of the Snooze Autonomic Cloud Management System*, in "The 6th IEEE International Scalable Computing Challenge (SCALE 2013)", Delft, Netherlands, May 2013, Short paper and demo, https://hal.inria.fr/hal-00925664

[52] E. FELLER, C. ROHR, D. MARGERY, C. MORIN. *Energy Management in IaaS Clouds: A Holistic Approach*, in "5th IEEE International Conference on Cloud Computing (CLOUD)", Honolulu, Hawaii, United States, June 2012, http://hal.inria.fr/hal-00695038

[53] J. O. KEPHART, D. M. CHESS. *The Vision of Autonomic Computing*, in "Computer", 2003, vol. 36, n^o 1, pp. 41-50

[54] A. MURTAZAEV, S. OH. *Sercon: Server Consolidation Algorithm using Live Migration of Virtual Machines for Green Computing*, in "IETE Technical Review", 2011, vol. 28 (3), http://www.researchgate.net/publication/271113653_Sercon_Server_Consolidation_Algorithm_using_Live_Migration_of_Virtual_Machines_for_Green_Com

[55] G. PIERRE, C. STRATAN. *ConPaaS: A Platform for Hosting Elastic Cloud Applications*, in "IEEE Internet Computing", 2012, vol. 16, n^o 5, pp. 88-92