

informatics mathematics
Irria

IN PARTNERSHIP WITH:
CNRS

**Institut national des sciences
appliquées de Rennes**

Université Rennes 1

Activity Report 2014

Project-Team MYRIADS

Design and Implementation of Autonomous
Distributed Systems

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

RESEARCH CENTER
Rennes - Bretagne-Atlantique

THEME
Distributed Systems and middleware

Table of contents

1. Members	1
2. Overall Objectives	2
2.1. General Objectives	2
2.2. Context	2
2.3. Challenges	3
2.4. Research Directions	4
2.4.1. Autonomous Management of Virtualized Infrastructures	4
2.4.1.1. Federation of IaaS clouds	4
2.4.1.2. SLA-driven PaaS over Cloud Federations	5
2.4.1.3. Virtual Data Centers	5
2.4.1.4. Virtualized Infrastructure Management	6
2.4.1.5. Heterogeneous Cloud Infrastructure Management	6
2.4.2. Multilevel Dynamic Adaptation of Service-based Applications	6
2.4.3. Exploration of unconventional programming models for the Internet of services	7
3. Research Program	8
3.1. Introduction	8
3.2. Autonomic Computing	8
3.3. Future Internet and SOA	8
3.4. Distributed Operating Systems	9
3.5. Unconventional/Nature-inspired Programming	9
4. New Software and Platforms	10
4.1. ConPaaS	10
4.2. HOCL-tools	10
4.3. Merkat	11
4.4. Meryn	11
4.5. Resilin	11
4.6. Snooze	12
4.7. Virtual Execution Platform (VEP)	13
5. New Results	13
5.1. Highlights of the Year	13
5.2. Dependable Cloud Computing	14
5.2.1. Deployment of distributed applications in a multi-provider environment	14
5.2.2. Checkpointing for multi-cloud environments	14
5.2.3. Towards a distributed cloud inside the backbone	14
5.2.4. A multi-objective adaptation system for the management of a Distributed Cloud	14
5.2.5. Multi-cloud application deployment in ConPaaS	15
5.2.6. Application Performance Modeling in Heterogeneous Cloud Environments	15
5.2.7. Dynamic reconfiguration for multi-cloud applications	15
5.2.8. Self-adaptable Monitoring for Security in the Cloud	15
5.2.9. Fog Computing	16
5.3. Heterogeneous Resource Management	16
5.3.1. Cross-resource scheduling in heterogeneous cloud environments	16
5.3.2. Maximizing private cloud provider profit in cloud bursting scenarios	17
5.3.3. Data life-cycle management in clouds	17
5.4. Energy-efficient Resource Infrastructures	17
5.4.1. Energy-efficient IaaS clouds	17
5.4.2. Energy-aware IaaS-PaaS co-design	17
5.4.3. Energy-efficient and network-aware resource allocation in Cloud infrastructures	18
5.4.4. Simulating Energy Consumption of Wired Networks	18

5.4.5.	Resource allocation in a Cloud partially powered by renewable energy sources	19
5.4.6.	SLA driven Cloud Auto-scaling for optimizing energy footprint	19
5.4.7.	Simulating the impact of DVFS within SimGrid	19
5.5.	Decentralised and Adaptive workflows	20
5.5.1.	Template workflows	20
5.5.2.	Adaptive Workflows with Chemical Computing	20
5.5.3.	Best-effort decentralised workflow execution	20
5.6.	Experimental Platforms	20
5.6.1.	Energy measurement	20
5.6.2.	BonFIRE	20
5.6.3.	Fed4FIRE	21
6.	Partnerships and Cooperations	21
6.1.	Regional Initiatives	21
6.1.1.	CominLabs EPOC project (2013-2016)	21
6.1.2.	EcoPaaS, Brittany region SAD project(2014-2015)	21
6.1.3.	Monitoring for Cloud Security, collaboration with DGA-MI (2014-2017)	22
6.1.4.	IRT B-Com	22
6.2.	National Initiatives	22
6.2.1.	Inria ADT GinFlow (2014-2016)	22
6.2.2.	Inria ADT Snooze (2012-2014)	22
6.2.3.	HEMERA Inria AEN (2010-2014)	22
6.2.4.	Inria IPL CityLab (under submission) (2014-2018)	22
6.2.5.	MIHMES ANR Investissements d’Avenir (2012 - 2018)	23
6.2.6.	CNRS GDS EcoInfo	23
6.3.	European Initiatives	23
6.3.1.	FP7 & H2020 Projects	23
6.3.1.1.	CONTRAIL	23
6.3.1.2.	ECO ₂ Clouds	24
6.3.1.3.	Fed4FIRE	24
6.3.1.4.	HARNESS	25
6.3.1.5.	PaaSage	25
6.3.2.	Collaborations in European Programs, except FP7 & H2020	26
6.3.2.1.	NESUS	26
6.3.2.2.	MC-DATA	26
6.3.2.3.	VEP-S	27
6.4.	International Initiatives	27
6.4.1.	Inria International Labs	27
6.4.2.	Inria Associate Teams	28
6.4.3.	Inria International Partners	28
6.5.	International Research Visitors	29
6.5.1.	Visits of International Scientists	29
6.5.2.	Visits to International Teams	29
7.	Dissemination	30
7.1.	Promoting Scientific Activities	30
7.1.1.	Program committees	30
7.1.2.	Evaluation committees, consulting	30
7.2.	Teaching - Supervision - Juries	31
7.2.1.	Teaching	31
7.2.2.	Supervision	32
7.2.3.	Juries	32
7.3.	Popularization	32

7.4. Miscellaneous	33
8. Bibliography	34

Project-Team MYRIADS

Keywords: Autonomic Computing, Cloud Computing, Energy Consumption, Unconventional Paradigms, Service Oriented Architecture

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

1. Members

Research Scientists

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]

Faculty Members

Jean-Pierre Banâtre [Université de Rennes 1, Professor Emeritus]
Nikolaos Parlavantzas [INSA Rennes, Associate Professor]
Jean-Louis Pazat [INSA Rennes, Professor, HdR]
Guillaume Pierre [Univ. Rennes I, Professor]
Cédric Tedeschi [Univ. Rennes I, Associate Professor]

Engineers

Alexandra Carpen-Amarie [Inria, until Feb 2014, granted by FP7 CONTRAIL project]
Roberto-Gioacchino Cascella [Inria, granted by FP7 Contrail project until Mar 2014 and by EIT ICT LABS VEP-S project from Apr 2014]
Stéphane Chevalier [Univ. Rennes I, until Aug 2014, granted by EIT ICT LABS MC-DATA project]
Stefania Costache [Inria, until Jan 2014, granted by FP7 CONTRAIL project]
Teodor Crivat [Univ. Rennes I, from Sep 2014, granted by EIT ICT LABS MC-DATA project]
Florian Dudouet [Inria, until Jan 2014, granted by FP7 CONTRAIL project]
Maxence Dunnewind [Inria, until Jan 2014, granted by FP7 FED4FIRE project]
Filippo Gaudenzi [Inria, until Feb 2014, granted by FP7 CONTRAIL project]
Nicolas Lebreton [Inria, until Nov 2014, granted by FP7 ECO2Clouds project]
Julien Lefeuvre [Inria, granted by FP7 ECO2Clouds and FED4FIRE projects]
Usman Malik [Univ. Rennes I, until Jun 2014]
David Margery [Inria]
Pascal Morillon [Univ. Rennes I]
Morteza Neishaboori [Univ. Rennes I, until Jun 2014]
Younes Ouchtouban [Inria, supported by D2T for Grid'5000]
Eric Poupart [Inria, until Mar 2014, granted by FP7 ECO2Clouds project]
Aboozar Rajabi [Inria, from Oct 2014, granted by FP7 PAASAGE project]
Matthieu Simonin [Inria, ADT Snooze (until Sep 2014, ADT GinFlow since Oct 2014)]
Arnab Sinha [Inria, from July 2014, granted by EIT ICT LABS VEP-S project]
Dzenan Softic [Univ. Rennes I, until Jun 2014]
Genc Tato [Univ. Rennes I, from Feb 2014, granted by FP7 HARNESS project]

PhD Students

Ismael Cuadrado Cordero [Inria]
Djawida Dib [Inria, until Jun 2014]
Anna Giannakou [Inria, from Mar 2014]
Sabbir Hasan Rochi [INSA de Rennes]
Ancuta Iordache [Univ. Rennes I]
Yunbo Li [EMN]

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

Post-Doctoral Fellows

Eliya Buyukkaya [Univ. Rennes I, until Aug 2014]

Maria Del Mar Callau Zori [Univ. Rennes I, from May 2014]

Bogdan Florin Cornea [Inria, until Nov 2014]

Eugen Feller [Inria, until Apr 2014]

Javier Rojas Balderrama [Inria, from Jun 2014]

Visiting Scientists

Jiajun Cao [Inria, from May 2014 until Aug 2014]

Gene Cooperman [Univ. Rennes I, from Mar 2014 until Apr 2014]

Georgios Ioannidis [Ecole Polytechnique Fédérale de Lausanne, from Oct 2014]

Palakiyem Wallah [INSA de Rennes, from Oct 2014 to Dec 2014]

Qian Zhang [PhD student, Australian National University, Oct 2014]

Administrative Assistants

Maryse Fouché [Inria]

Anne Pasteur [Inria, part-time, FP7 Contrail Project assistant until Feb 2014]

Others

Guillaume Chapel [Inria, intern, from Apr 2014 until Aug 2014]

Vishrut Mehta [Inria, intern, from May 2014 until Jul 2014]

Céline Merlet [Inria, intern, from May 2014 until Aug 2014]

Louis Rilling [DGA-MI, external collaborator from Mar 2014]

Lavinia Samoila [Univ. Rennes I, intern, from May 2014 until Aug 2014]

Benjamin Soulas [Inria, intern, from May 2014 until Aug 2014]

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.

¹According 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>.

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 have made computing and storage become 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.

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 incommensurable 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 (Eric Poupart, Nicols Lebreton and 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 flavours 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 centres 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 centres 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 their 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's 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: SLAs 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 [50], [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 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.

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 realised 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 behaviour.

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 [58]. 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-organising, 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 [58] 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. EAI 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 behaviour 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 XtremOS 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

Levering the computing services available on the Internet requires to revisit programming models, with the idea of expressing decentralised and autonomous behaviours (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. New Software and Platforms

4.1. ConPaaS

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

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

Status: Version 1.4.2

License: BSD

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

Active contributors (from the Myriads team): Eliya Buyukkaya, Ancuta Iordache, Morteza Neishaboori, Guillaume Pierre, Dzenan Softic, Genc Tato, Teodor Crivat.

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 ICT Labs project.

4.2. HOCL-tools

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

Status: Version 1.0 to be released in open source

License: TBD

Presentation: HOCL (Higher Order Chemical Language) is a chemical programming language based on the chemical metaphor presented before (see Section 3.5). It was developed for several years within the PARIS and Myriads teams. Within HOCL, following the chemical metaphor, computations can be regarded as chemical reactions, and data can be seen as molecules which participate in these reactions. If a certain condition is held, the reaction will be triggered, thus continuing until it gets inert: no more data can satisfy any computing conditions. To realize this program paradigm, a multiset is implemented to act as a chemical tank, containing necessary data and rules. An HOCL program is then composed of two parts: *chemical rule definitions* (reaction rules) and *multiset definition* (data). More specifically, HOCL provides the higher order: reaction rules are molecules that can be manipulated like any other molecules. In other words, HOCL programs can manipulate other HOCL programs.

An HOCL compiler was developed using Java to execute some chemical programs expressed with HOCL. This compiler is based on the translation of HOCL programs to Java code. As a support for service coordination and service adaptation, we recently extended the HOCL compiler so as to support decentralized workflow execution. Works around the implementation of a distributed multiset gave birth to an underlying layer for this compiler, making it able to deploy HOCL programs transparently over large scale platforms. This last part is currently considered to be interfaced with the current HOCL compiler. All these features are planned to be released under the common name of *HOCL-tools*.

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

Impact: The compiler is used as a tool within the team to develop HOCL programs. The decentralized workflow execution support has been extensively used to produce results published and presented at several conferences. It is also used in the framework of the DALHIS ² associated team, as a workflow template executor, integrated with the TIGRES workflow manager developed at the Lawrence Berkeley National Lab. It is supported by the GinFlow ADT funded by Inria.

4.3. Merkat

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

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

Status: Version 1.0

License: TBD

Presentation: Merkat is a market-based private PaaS (Platform-as-a-Service) system, supporting dynamic, fine-grained resource allocation and automatic application management [49], [48] [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. Merkat is currently being evaluated by EDF R&D using EDF high-performance applications. The development was initiated in the framework of Stefania Costache PhD's thesis.

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

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.

4.4. Meryn

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

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

Status: Version 1.0

License: TBD

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 [51]. 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. Meryn was developed in the framework of Djawida Dib's PhD thesis [10].

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

Impact: Meryn is not yet distributed as open source.

4.5. Resilin

²<http://project.inria.fr/dalhis>

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

URL: <http://resilin.inria.fr>

Status: Version 1.0

License: GNU Affero GPL

Presentation: Resilin [6] is an open-source system for creating and managing MapReduce execution platforms over clouds. Resilin is compatible with the Amazon Elastic MapReduce (EMR) API, but it goes beyond Amazon's proprietary EMR solution in allowing users (e.g. companies, scientists) to leverage resources from one or more public and/or private clouds. This enables performing MapReduce computations over a large number of geographically-distributed and diverse resources. Resilin can be deployed across most of the open-source and commercial IaaS cloud management systems (e.g., OpenStack, OpenNebula, Amazon EC2). Once deployed, Resilin takes care of provisioning Hadoop clusters and submitting MapReduce jobs, allowing users to focus on writing their MapReduce applications rather than managing cloud resources. Resilin is implemented in the Python language and uses the Apache Libcloud library to interact with IaaS clouds. Resilin has been evaluated on multiple clusters of the Grid'5000 experimentation testbed. The results show that Resilin enables the use of geographically distributed resources with a limited impact on MapReduce job execution time.

Active contributors (from the Myriads project-team): Ancuta Iordache, Céline Merlet, Christine Morin, Nikolaos Parlavantzas, Matthieu Simonin.

Impact: Resilin is being used in the MOAIS project-team at Inria Grenoble - Rhône Alpes.

4.6. Snooze

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

URL: <http://snooze.inria.fr>

Status: Version 2.1.5

License: GPLv2

Presentation: Snooze [53], [52], [54] [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 [59]), and power management to transition idle servers into a power saving mode. Snooze is a highly modular piece of software. It has been extensively evaluated on the Grid'5000 testbed using realistic applications.

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*, *jeloud* ...). Snooze also provides a web interface to control the system.

Snooze was used as a building box for two internships projects during the summer of 2014. The EC2 interface was used to execute Hadoop jobs configured by Resilin software. As a result we show that (1) the EC2 interface was expressive enough to work with a higher level tool and (2) the control over Snooze allow a better placement of data chunks for Hadoop jobs which leads to a better reliability

of the execution of the different jobs. The second internship topic took part in a collaboration with the Northeastern University of Boston. The goal was to build a *Checkpoint as a Service* system. 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 case of failures. This work is described in [31].

Active contributors (from Myriads team): Jiajun Cao, Gene Cooperman, Eugen Feller, Yvon Jégou, David Margery, Christine Morin, Matthieu Simonin.

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 CCGrid2013. Finally, we know that it was experimented 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.

4.7. Virtual Execution Platform (VEP)

Contact: Yvon Jégou, Yvon.Jegou@inria.fr

URL: <http://project.inria.fr/vep/>

Status: Version 2.2

License: BSD

Presentation: Virtual Execution Platform

(VEP) [57] is a Contrail (<http://contrail-project.eu>) service that sits just above IaaS layer at the service provider end of the Contrail cloud federation. The VEP service provides a uniform interface for managing the whole lifecycle of elastic applications on the cloud and hides the details of the IaaS layer to the user. VEP applications are described in OVF (Open Virtualization Format) standard format. Resource usage is controlled by CEE (Constrained Execution Environment) rules which can be derived from SLAs (Service Level Agreement). The VEP service integrates a monitoring system where the major events about the application, mainly resource usage, are made available to the user.

The VEP service provides a RESTful interface and can be exploited directly by users on top of the provider IaaS. OpenNebula and OpenStack IaaS frameworks were initially supported. During the VEP-S EIT ICT Labs activity in 2014, VEP was extended with a new OCCI IaaS driver which allows to control any IaaS framework providing a standard OCCI API. Support for the new OCCI SLA proposition from OGF has also been added and allows to represent the VEP CEEs in a standard format. Finally, during this activity, the Zabbix open source distributed monitoring system was integrated to VEP.

Active contributors (from Myriads project-team): Roberto-Gioacchino Cascella, Florian Dudouet, Filippo Gaudenzi, Yvon Jégou, Christine Morin, Arnab Sinha.

Impact: VEP is part of Contrail software stack. External users can experiment with it using the open testbed operated by Myriads team. Technology transfer of VEP technology is ongoing in the context of the VEP-S EIT ICT Labs activity.

5. New Results

5.1. Highlights of the Year

- The Contrail project coordinated by Christine Morin received the "Excellent" grade at its final review held on March 14th, 2014 in Brussels.
- Anne-Cécile Orgerie has been awarded the Young Researcher prize of the Lyon city in November 2014.
- Christine Morin has been awarded one of the 12 "Etoile de l'Europe 2014" prizes in December 2014 for the coordination of the Contrail European project.

BEST PAPER AWARD :

[18] **Robust Performance Control for Web Applications in the Cloud in 4th International Conference on Cloud Computing and Services Science.** H. FERNANDEZ, C. STRATAN, G. PIERRE.

5.2. Dependable Cloud Computing

Participants: Jiajun Cao, Stéphane Chevalier, Gene Cooperman, Teodor Crivat, Roberto-Gioacchino Cascella, Stefania Costache, Florian Dudouet, Filippo Gaudenzi, Anna Giannakou, Yvon Jégou, Ancuta Iordache, Christine Morin, Anne-Cécile Orgerie, Edouard Outin, Nikolaos Parlavantzas, Jean-Louis Pazat, Guillaume Pierre, Aboozar Rajabi, Louis Rilling, Matthieu Simonin, Arnab Sinha, Cédric Tedeschi.

5.2.1. *Deployment of distributed applications in a multi-provider environment*

Participants: Roberto-Gioacchino Cascella, Stefania Costache, Florian Dudouet, Filippo Gaudenzi, Yvon Jégou, Christine Morin, Arnab Sinha.

The move of users and organizations to Cloud computing will become possible when they are able to exploit their own applications, applications and services provided by cloud providers, as well as applications from third party providers in a trustful way on different cloud infrastructures. In the framework of the Contrail European project [2] [46], we have designed and implemented the Virtual Execution Platform (VEP) service in charge of managing the whole life cycle of OVF distributed applications under Service Level Agreement rules on different infrastructure providers [47]. In 2013, we designed the CIMI inspired REST-API for VEP 2.0 with support for Constrained Execution Environment (CEE), advance reservation and scheduling service, and support for SLAs [56], [55] [57]. We integrated support for delegated certificates and developed test scripts to integrate the Virtual Infrastructure Network (VIN) service. VEP 1.1 was slightly modified to integrate the usage control (Policy Enforcement Point (PEP)) solution developed by CNR. The CEE management interface was developed during 2013 and is available through the graphical API as well as through the RESTful API.

5.2.2. *Checkpointing for multi-cloud environments*

Participants: Jiajun Cao, Gene Cooperman, Christine Morin, Matthieu Simonin.

Most cloud platforms currently rely on each application to provide its own fault tolerance. A uniform mechanism within the cloud itself serves two purposes: (a) direct support for long-running jobs, which would otherwise require a custom fault-tolerant mechanism for each application; and (b) the administrative capability to manage an over-subscribed cloud by temporarily swapping out jobs when higher priority jobs arrive.

We propose ([31]) a novel *Checkpointing as a Service* approach, which enables application checkpointing and migration in heterogeneous cloud environments. Our approach is based on a non-invasive mechanism to add fault tolerance to an existing cloud platform *after the fact*, with little or no modification to the cloud platform itself. It achieves its cloud-agnostic property by using an external checkpointing package, independent of the target cloud platform. We implemented a prototype of the service on top of both OpenStack and Snooze IaaS clouds. We conducted a preliminary performance evaluation using the Grid'5000 experimentation platform.

5.2.3. *Towards a distributed cloud inside the backbone*

Participants: Anne-Cécile Orgerie, Cédric Tedeschi.

The DISCOVERY proposal currently in phase of construction and lead by Adrien Lèbre from the ASCOLA team, and currently on leave at Inria aims at designing a distributed cloud, leveraging the resources we can find in the network's backbone.³

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. The design, implementation, and experimentation of the overlay has been described in an article published in 2014 [22].

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

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

³The DISCOVERY website: <http://beyondthecLOUDS.github.io>

In this project, we consider a “Distributed Cloud” made of multiple data/computing centers interconnected by a high speed network. A distributed Cloud is neither a usual Cloud built around a single data center, nor a Cloud Federation interconnecting different data centers owned and run by different administrative entities. Moreover, in the Cloud organization targeted here, the network capabilities can be dynamically configured in order to apply optimizations to guarantee QoS for streaming or negotiated bandwidth for example. 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 the Distributed Cloud in order to consider both optimizations for energy consumption and for users’ QoS needs. The objectives of these optimizations 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 will rely 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) will be used for the implementation of the system. The technical developments are based on the Openstack framework.

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.

5.2.5. Multi-cloud application deployment in ConPaaS

Participants: Stéphane Chevalier, Teodor Crivat, Guillaume Pierre.

We extended ConPaaS to support the deployment of smartphone backend applications in mobile operators’ base stations. The motivation is to reduce the latency compared to a traditional deployment where the backend is located in an external cloud. This requires building a lightweight infrastructure which allows one to easily create containers that can be seamlessly migrated (roaming). A publication on this topic will appear in 2015 [23].

5.2.6. Application Performance Modeling in Heterogeneous Cloud Environments

Participants: Ancuta Iordache, Guillaume Pierre.

Heterogeneous cloud platforms offer many possibilities for applications for make fine-grained choice 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 publication on this topic is currently under review.

5.2.7. Dynamic reconfiguration for multi-cloud applications

Participants: Nikolaos Parlavantzas, Aboozar Rajabi.

In the context of the PaaSage European project, we are working on model-based self-optimisation 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 2014, we have defined the architecture of the adaptation system and produced a first prototype[30].

5.2.8. Self-adaptable Monitoring for Security in the Cloud

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

We aim at designing a self-adaptable system for security monitoring in clouds. The considered system should cope with the dynamic nature of virtual infrastructures in clouds and have a minimal impact on performance. In 2014, we studied the state of the art in cloud security monitoring, which is composed of various approaches for intrusion detection systems (IDS), based on traditional IDS techniques such as signature-based detection and anomaly-based detection.

As a first step towards our goal of making self-adaptable a complete security monitoring architecture for cloud environments, we defined a simple initial monitoring scenario for identifying the impact of the dynamicity of a cloud architecture on the intrusion detection process. In this scenario, the security monitoring infrastructure is composed of two network IDS instances, which are used to monitor the virtual infrastructures network traffic of two cloud clients (one virtual infrastructure per client), and also eventually monitor the physical infrastructure (that is the operator's infrastructure). The virtual network traffic in each host machine is monitored by only one of the IDS instances, so that the IDS instances must be adapted to topology changes (such as migration of VMs) in the cloud environment. The adaptation process includes updates of the rules configured in the instance (deletion or creation).

In 2014, we built our testbed based on OpenStack technology for the underlying IaaS cloud platform and Snort for the network IDS. At this point the testbed consists of only five machines (on the Grid'5000 platform) but we aim to increase the number of host machines and deploy more VMs for having a more realistic representation of a production network. This will allow us to study performance issues and also more complex security monitoring setups. Our goal is also to enable monitoring of other elements, such as resource usage (both per host and per VM) on the cloud provider side.

5.2.9. 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 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 the service distribution.

We first focus on the formal expression of these requirements, by considering first the requirements provided by a Service Operator to the "Fog" Infrastructure (required computing resources, minimal quality of experience (QoE) level, etc.). The resource allocation strategies should also take into account the topology of the "Fog" Infrastructure, the heterogeneous capabilities of the equipments and of the underlying network. Based on this information, strategies and algorithms for resource allocation should be designed that will participate in the process of building an efficient platform for the service distribution. Evaluation of this efficiency will be an important process to justify the relevance of the strategies.

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

5.3. Heterogeneous Resource Management

Participants: Eliya Buyukkaya, Djawida Dib, Eugen Feller, Christine Morin, Nikolaos Parlavantzas, Guillaume Pierre.

5.3.1. Cross-resource scheduling in heterogeneous cloud environments

Participants: Eliya Buyukkaya, Guillaume Pierre.

Allocating resources to applications in a heterogeneous cloud environment is harder than in a homogeneous environment. In a heterogeneous cloud some rare resources are more precious than others, and should be treated carefully to maximize their utilization. Similarly, applications may request groups of resources that exhibit certain inter-resource properties such as the available bandwidth between the assigned resources. We are currently investigating scheduling algorithms for handling such scenarios.

5.3.2. Maximizing private cloud provider profit in cloud bursting scenarios

Participants: Djawida Dib, Christine Morin, Nikolaos Parlavantzas.

Current PaaS offerings either provide no support for SLA guarantees or provide limited support targeting a restricted set of application types. To overcome this limitation, we have developed an open, cloud-bursting PaaS system, called Meryn, designed to be easily extensible to host new application types. The system integrates a decentralized optimization policy that maximises the PaaS provider profit, taking into account the payment of penalties incurred when quality guarantees are unsatisfied. The system was implemented and evaluated on the Grid5000 testbed using batch and MapReduce workloads. The results demonstrated the effectiveness of the policy in increasing provider profit [16] This work was part of Djawida Dib's PhD thesis [10] defended in July 2014.

5.3.3. Data life-cycle management in clouds

Participants: Eugen Feller, Christine Morin.

Infrastructure as a Service (IaaS) clouds provide a flexible environment where users can choose and control various aspects of the machines of interest. However, the flexibility of IaaS clouds presents unique challenges for storage and data management in these environments. Users use manual and/or ad-hoc methods to manage storage and data in these environments. FRIEDA is a Flexible Robust Intelligent Elastic Data Management framework that employs a range of data management strategies approaches in elastic environments. This year, our work carried out in the context of the DALHIS associate team⁴, was focused on the extended design and evaluation of the FRIEDA data management system. FRIEDA was tested to work on Amazon EC2 resources. In addition, we layered a commandline utility atop FRIEDA that allows users to plug-in applications to run in FRIEDA. These tools have been adopted by the LBL-ATLAS group to run their experiments on Amazon [29].

5.4. Energy-efficient Resource Infrastructures

Participants: Maria Del Mar Callau Zori, Alexandra Carpen-Amarie, Bogdan Florin Cornea, Ismael Cuadrado Cordero, Djawida Dib, Eugen Feller, Sabbir Hasan Rochi, Yunbo Li, Christine Morin, Anne-Cécile Orgerie, Jean-Louis Pazat, Guillaume Pierre, Lavinia Samoila.

5.4.1. Energy-efficient IaaS clouds

Participants: Alexandra Carpen-Amarie, Christine Morin, Anne-Cécile Orgerie.

Energy consumption has always been a major concern in the design and cost of data centers. The wide adoption of virtualization and cloud computing has added another layer of complexity to enabling an energy-efficient use of computing power in large-scale settings. Among the many aspects that influence the energy consumption of a cloud system, the hardware-component level is one of the most intensively studied. However, higher-level factors such as virtual machine properties, their placement policies or application workloads may play an essential role in defining the power consumption profile of a given cloud system. In this work, we explored the energy consumption patterns of Infrastructure-as-a-Service (IaaS) cloud environments under various synthetic and real application workloads. For each scenario, we investigated the power overhead triggered by different types of virtual machines, the impact of the virtual cluster size on the energy-efficiency of the hosting infrastructure and the tradeoff between performance and energy consumption of MapReduce virtual clusters through typical cloud applications [45].

5.4.2. Energy-aware IaaS-PaaS co-design

Participants: Maria Del Mar Callau Zori, Alexandra Carpen-Amarie, Djawida Dib, Anne-Cécile Orgerie, Guillaume Pierre, Lavinia Samoila.

⁴<http://project.inria.fr/dalhiss>

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 argue that cross-layer cooperation in clouds is a key to achieving an optimized resource management, both performance and energy-wise. To this end, we claim there is a need for a cooperation API between IaaS and PaaS clouds, enabling each layer to share specific information and to trigger correlated decisions. We identified the drawbacks raised by such co-design objectives and discuss opportunities for energy usage optimizations. A position paper has been published on these aspects [15]. Ongoing work is currently conducted in order to quantify the actual possible gains both energy and performance-wise for this IaaS-PaaS co-design approach.

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

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

Cloud computing is increasingly becoming an essential component for Internet service provision, yet at the same time its energy consumption has become a key environmental and economic concern. It becomes urgent to improve the energy efficiency of such infrastructures. Our work aims at designing energy-efficient resource allocation for Cloud infrastructures. Yet, energy is not the only criterion to take into account at risk of losing users. A multi-criteria approach is required in this context to satisfy both users and Cloud providers.

The proposed resource allocation algorithms will take into account not only the computing resources but also the storage and networking resources. Indeed, the ever-growing appetite of new applications for network resources leads to an unprecedented electricity bill for network resources, and for these bandwidth-hungry applications, networks can become a significant bottleneck. This phenomenon is emphasized with the emergence of the big data paradigm. The designed algorithms would thus integrate the data locality dimension to optimize computing resource allocation while taking into account the fluctuating limits of network resources.

In 2014, several experiments were performed to understand and quantify networking energy consumption. These experiments include network protocol energy consumption in the devices, configuration energy consumption in switching/routing devices and associated energy consumption to real cloud computing applications (e.g. Google drive). These experiments have been performed over systems provided by Inria such as Grid'5000 and specific network devices (e.g. level 3 router for a private LAN). Based on this work, we developed an analytic model of networking energy consumption in a cloud computing environment. This analysis will serve as a basis for designing an energy-efficient architecture and related algorithms.

5.4.4. *Simulating Energy Consumption of Wired Networks*

Participants: Bogdan Florin Cornea, Anne-Cécile Orgerie.

Predicting the performance of applications, in terms of completion time and resource usage for instance, is critical to appropriately dimension 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 dimension 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 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. Through this tool, we have studied the energy consumption of data transfers in Clouds [19]. This work has been done in collaboration with the Avalon team from LIP in Lyon.

5.4.5. 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 planned in the future).

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

5.4.6. 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 propose an approach driven by Service Level Agreements (SLAs) for Cloud auto-scaling. A SLA defines a formal contract between a service provider and a service consumer on an expected QoS level. The main idea of this thesis is to exploit the SLA requirements to (i) avoid the interferences between the Cloud autonomic managers by a cross-layer coordination of SLA contracts; (ii) fine-tune the resources needs according to SLA by proposing both dynamic resources provisioning for optimizing the energy footprint and dynamic reconfiguration at the SaaS level to optimize the expected QoS. In particular, we propose to address renewable energy in the SLA contract. The objective is twofold. First, for ecological reasons, it allows Cloud users to express their preferences about the energy provider and the nature of the energy in the data center. Then, for economic reasons, it takes advantage of renewable energy costs (expressed in the SLA) to reconfigure resource allocation and energy usage. The integration of such SLAs in each layer of the Cloud stack and their management by an autonomic manager or by the coordination of autonomic managers still remain open issues.

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

5.4.7. Simulating the impact of DVFS within SimGrid

Participants: Alexandra Carpen-Amarie, Christine Morin, Anne-Cécile Orgerie.

Simulation is 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 experimentations 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.

5.5. Decentralised and Adaptive workflows

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

5.5.1. Template workflows

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

In the framework of the DALHIS associate team ⁵, we started to combine the high-level template workflow language TIGRES ⁶, developed by our partner team from Lawrence Berkeley National Lab (LBL) with the workflow management system developed in the team [5]. The design of this integration and its benefits have been presented in a workshop article [24].

5.5.2. Adaptive Workflows with Chemical Computing

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

We are currently designing a complete programming model for the management of adaptive workflows, based on an extension of the HOCL language, in particular workflows that may evolve at run time in their shape. An article is under preparation.

5.5.3. Best-effort decentralised workflow execution

Participants: Jean-Louis Pazat, Cédric Tedeschi, Palakyiem 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 centralised orchestrator. Such a workflow engine relies on point-to-point cooperation between nodes supporting the execution.

5.6. Experimental Platforms

Participants: Maxence Dunnewind, Nicolas Lebreton, Julien Lefeuvre, David Margery, Eric Poupart.

5.6.1. Energy measurement

Participants: Maxence Dunnewind, Nicolas Lebreton, David Margery, Eric Poupart.

In the context of the ECO₂Clouds project, the BonFIRE infrastructure was updated. At the software layer, the complete monitoring stack was revisited so as to attribute power consumption values to all VMs running on the infrastructure and to expose this information to users. This was used by the project partners to confirm that using an eco-aware scheduler could significantly reduce eco-impact of running a distributed infrastructure.

5.6.2. BonFIRE

Participants: Maxence Dunnewind, Julien Lefeuvre, David Margery, Eric Poupart.

⁵<http://project.inria.fr/dalhis>

⁶<http://tigres.lbl.gov/home>

The project was reviewed in December 2013 during CloudCom 2013 in Bristol and rated Excellent. It has been kept in working state through our commitment to the BonFIRE foundation. The main achievement on this topic was to evolve the cloud reservation system so as to support tracking usage using allocation blocks, as a fragment of the physical machines. Instance types can therefore have a different footprint in number of allocation blocks depending on the hardware they are scheduled on.

5.6.3. Fed4FIRE

Participants: Nicolas Lebreton, 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. Here, we contributed to a proposal to secure usage of OMF and to a design to allow using BonFIRE through SFA. In 2014, the main area of work has been maintenance of the infrastructure and initial prototyping of an SFA API to BonFIRE.

6. Partnerships and Cooperations

6.1. Regional Initiatives

6.1.1. 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.

6.1.2. EcoPaaS, Brittany region SAD project(2014-2015)

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

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 is funding Maria del Mar Callau-Zori's postdoc.

6.1.3. *Monitoring for Cloud Security, collaboration with DGA-MI (2014-2017)*

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

Our study aims at designing a self-adaptable system for security supervision in clouds. The considered system should cope with the dynamic nature of clouds and have a minimal impact on performance. The funding from DGA funds a PhD student, Anna Giannakou, who joined Myriads team in March 2014. Anna Giannakou is co-advised by Christine Morin (Inria), Jean-Louis Pazat (INSA Rennes) and Louis Rilling (DGA-MI). Louis Rilling was formally appointed as external collaborator in Myriads team effective from March 1st 2014.

6.1.4. *IRT B-Com*

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

Yvon Jégou and Jean-Louis Pazat are at IRT B-Com ⁷ 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.

6.2. National Initiatives

6.2.1. *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.

6.2.2. *Inria ADT Snooze (2012-2014)*

Participants: Eugen Feller, Yvon Jégou, David Margery, Christine Morin, Anne-Cécile Orgerie, Matthieu Simonin.

The Snooze technological development action funded by INRIA aims at developing an IaaS cloud environment based on the Snooze virtual machine framework developed by the team (<http://snooze.inria.fr>) and to make this new environment available to a wide community. In 2014, we refactored some parts of the code to enable the use of plugins. We also developed the Cloud Agnostic Checkpointing Service (CACS) service on top of Snooze to enable application recovery in the event of the failure of servers hosting virtual machines [31].

6.2.3. *HEMERA Inria AEN (2010-2014)*

Participants: Bogdan Florin Cornea, Yvon Jégou, Anne-Cécile Orgerie.

The Myriads team is involved in the HEMERA large wingspan project funded by INRIA (<http://www.grid5000.fr/mediawiki/index.php/Hemera>). This project aims at demonstrating ambitious up-scaling techniques for large scale distributed computing by carrying out several dimensioning experiments on the Grid'5000 infrastructure, at animating the scientific community around Grid'5000 and at enlarging the Grid'5000 community by helping newcomers to make use of Grid'5000. Yvon Jégou is co-chair of the "Bring Grids Power to Internet-Users thanks to Virtualization Technologies" working group. Anne-Cécile Orgerie is involved in the "Energy" working group which is currently looking at making energy-aware experiments on Grid'5000 easier for the users. This project funded Bogdan Florin Cornea's postdoc supervised by Anne-Cécile Orgerie and Laurent Lefèvre (Inria, LIP, Lyon).

6.2.4. *Inria IPL CityLab (under submission) (2014-2018)*

Participants: Roberto-Gioacchino Cascella, Christine Morin.

⁷<http://b-com.org/wp/>

The Inria Project Lab (IPL) CityLab@Inria (informally started - see <https://citylab.inria.fr>) studies ICT solutions toward smart cities that promote both social and environmental sustainability.

and technology (prnd)3341ensirgndofaI
solu-eancicntal eam pctisirgndofaI
whilensot15(xpanturgnd)14(itons)13(researchnT)]TJ 0 -11.655 Td ajeitocitiestond27954(eam)80(designam)80(anted)80(impleon
sociallutii.a3

Inria contact: Christine Morin

Abstract: The goal of the Conrail project is to design, implement, evaluate and promote an open source system for Cloud Federations. Resources that belong to different operators will be integrated into a single homogeneous federated Cloud that users can access seamlessly. The Conrail project has built a complete Cloud platform which integrates Infrastructure-as-a-Service (IaaS) and Platform-as-a-Service (PaaS) offerings [2].

6.3.1.2. *ECO₂Clouds*

Participants: Maxence Dunnewind, Nicolas Lebreton, Julien Lefevre, David Margery, Eric Poupart.

Type: FP7

Defi: Future internet experimental facility and experimentally-driven research

Instrument: Specific Targeted Research Project

Objectif: Future Internet Research and Experimentation (FIRE)

Duration: October 2012 - September 2014

Coordinator: Atos Spain SA (ATOS, Spain)

Partner: Atos Spain SA (ATOS, Spain) The University of Manchester (UNIMAN, United Kingdom) The University of Edinburgh (UEDIN, United Kingdom) Universitaet Stuttgart (USTUTT, Germany) Politecnico di Milano (POLIMI, Italy)

Inria contact: David Margery

Abstract: In ECO₂Clouds, we add to BonFIRE energy probes to be able to measure power consumption of the infrastructure, combine it with information about energy sources used to produce the power so as to be able to bill CO₂ usage to experimenters running VMs. To allow for scheduling and adaptation of running applications, CO₂ usage is not only billed after the fact but also quoted in advance for a given period for according to resource usage units.

6.3.1.3. *Fed4FIRE*

Participants: Maxence Dunnewind, Julien Lefevre, 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)

Partner: 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.

6.3.1.4. HARNESS

Participants: Eliya Buyukkaya, Georgios Ioannidis, Ancuta Iordache, Guillaume Pierre, Genc Tato.

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 Informationstechnik 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.

6.3.1.5. PaaSage

Participants: Christine Morin, Nikolaos Parlavantzas, Aboozar Rajabi.

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 für Wissenschaftliche Datenverarbeitung MBH Göttingen (Germany), Automotive Simulation Center Stuttgart (Germany), University of Ulm (Germany), Akademia Górniczo-Hutnicza im. Stanisława 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/>

Abstract: PaaSage aims to deliver an open and integrated platform to support both deployment and design of Cloud applications, together with an accompanying methodology that allows model-based application development, configuration, optimisation, and deployment on multiple Cloud infrastructures.

6.3.2. Collaborations in European Programs, except FP7 & H2020

6.3.2.1. 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 than 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 glue 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.

6.3.2.2. MC-DATA

Participants: Stéphane Chevalier, Teodor Crivat, Guillaume Pierre.

Program: EIT ICT Labs

Project acronym: MC-DATA

Project title: Multi-cloud data management

Duration: Jan-Dec 2014

Coordinator: Dr. Peter Pietzuch, Imperial College London

Other partners: SICS, Vodafone

Abstract: In 2014, the continuation of the MC-Data project had two main innovation objectives: (a) to provide and release a novel open-source Platform-as-a-Service (PaaS) cloud computing software stack (MC-ConPaaS) that explicitly targets cloud application deployments across multiple data centre sites; (b) to demonstrate the business value of the MC-ConPaaS platform through a use case of cloud-assisted real-time smartphone applications, thus affecting the future business models of mobile operators. Its expected outcomes are:

- to release the MC-ConPaaS multi-site cloud platform as open source;
- to foster the adoption of the MC-ConPaaS platform by creating tutorials and documentation;
- to transfer the technology of the MC-ConPaaS platform to a mobile operator (VODAFONE), enabling them to offer a cloud infrastructure that supports cloud-assisted real-time applications;
- to develop new business models for mobile operators based on cloud-assisted real-time services running on virtualised mobile base stations.

6.3.2.3. VEP-S

Participants: Roberto-Gioacchino Cascella, Yvon Jégou, Christine Morin, Arnab Sinha.

Program: EIT ICT Labs

Project acronym: VEP-S

Project title: SLA-Aware Heterogeneous Data-Centers Management through Standards

Duration: Jan-Dec 2014

Coordinator: Christine Morin, Inria

Other partners: Intel (Ireland), Reply (Italy)

Abstract: We designed the VEP-S system, which consists of the Virtual Execution Platform (VEP) component with support for OCCI, integration of the OCCI SLA extension, and the monitoring system for deploying and running distributed applications packaged following the Open Virtualization Format (OVF), a DMTF standard, on top of an IaaS cloud. The Virtual Execution Platform (VEP), developed in the framework of Contrail European project, is in charge of provisioning hardware resources from Cloud providers and to deploy and run distributed applications submitted by users under the control of a negotiated Service Level Agreements (SLA). VEP interacts with the underlying IaaS manager to create application networks, register VM images, generate VM templates and manage the lifecycle of virtual machines. The OCCI SLA API extends the OCCI Core Model to implement a SLA management API. This API allows for the creation and management of resources related with the realization of agreements between an OCCI-enabled cloud service provider and potential consumers of the provider's resources. In the context of the VEP-S project, this extension is used to describe SLA terms and map them with the resources and services a cloud provider can offer. The monitoring component will provide three types of services: monitoring the IaaS resource manager to check whether the machine has started or not; monitoring the IaaS to check the usage of the resources; monitoring on the VM (monitoring agent in the VM and activated by the user). The technology used for the monitoring is Zabbix.

6.4. International Initiatives

6.4.1. Inria International Labs

Christine Morin was one of the co-organizers of the BIS 2014 workshop held in Paris in June 2014 in the framework of the Inria@Silicon Valley Inria International Lab. Christine Morin and Deb Agarwal were the co-chairs of the session on computation and communication for the future internet at BIS 2014. Several Myriads team members (Christine Morin, Anne-Cécile Orgerie, Javier Rojas Balderrama, Matthieu Simonin, Cédric Tedeschi) are involved in the DALHIS associate team on data analysis on large-scale heterogeneous infrastructures for science, which is part of the Inria@SiliconValley program.

6.4.2. Inria Associate Teams

6.4.2.1. DALHIS

Participants: Christine Morin, Anne-Cécile Orgerie, Javier Rojas Balderrama, Matthieu Simonin, 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 2014 on scientific workflows, energy efficiency and data management is described respectively in Sections 5.5.1, 5.4.1 and 5.3.3. Christine Morin, Anne-Cécile Orgerie, Javier Rojas Balderrama, Cédric Tedeschi and Deb Agarwal participated in the BIS 2014 workshop held in Paris in June 2014. Christine Morin and Deb Agarwal were the co-chairs of the session on Computation and communication for the Future Internet at BIS 2014. Cédric Tedeschi presented the DALHIS activities during this session focusing on our results on scientific workflows. Deb Agarwal has been awarded an Inria International Chair for the 2015-2019 period enabling long visits in the Myriads team.

6.4.3. Inria International Partners

6.4.3.1. Informal 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". Gen Cooperman was hosted in Myriads team for a 1.5-month visit in March-April 2014. His PhD student, Jiajun Cao did a 3-month internship in Myriads team from May to August 2014.

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 visited Myriads team in July 2014 to investigate synergetic work directions on cloud security.

Argonne/ Chicago University We collaborate on cloud computing with Kate Keahey from Argonne National Laboratory, USA. She hosted Ismael Cuadrado Cordero in her team for a 12-week summer internship (June-September 2014) on using extended on-availability leases to increase utilization in scientific IaaS clouds.

University of Guadalajara Nikolaos Parlavantzas is collaborating with the team of Prof. Héctor Duran-Limon of the University of Guadalajara, Mexico, preparing a joint ANR-CONACYT project submission.

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 led to two joint publications this year, and another paper has been accepted in 2015.

6.5. International Research Visitors

6.5.1. Visits of International Scientists

Jiajun Cao, PhD student at the Northeastern University (Boston, USA), made a 3-month visit in Myriads team (May-August 2014). He contributed to the design, implementation and evaluation of a cloud agnostic checkpointing service exploiting the DMTCP process-level checkpointing technology developed in Gene Cooperman's team at the Northeastern University. This service was experimented on top of Snooze IaaS cloud management system developed in Myriads team. A paper on this work will be presented at CC-Grid 2015.

Gene Cooperman, Professor at the Northeastern University (Boston, USA), made a 1.5 month sabbatical visit in Myriads team (March-April 2014). His visit was partially funded by the University of Rennes 1.

Georgios Ioannidis (PhD student at EPFL, Switzerland) made a 3-months visit in the Myriads team (Oct-Dec 2014). The goal was to reinforce the collaboration between the two teams in the context of the HARNES FP7 project.

Palakiyem Wallah, assistant professor at the University of Kara (Togo) visited Myriads team from October to December 2014 in the framework of his PhD thesis, which is co-advised by Jean-Louis Pazat and Cédric Tedeschi.

Qian Zhang (PhD student at the Australian National University) spent 3 weeks in Myriads team in October 2014 to learn more about our research activities on SLA management. Her visit was supported by a grant from the Australian-French Association for Science and Technology (AFAS).

6.5.1.1. Internships

Vishrut Mehta Vishrut

Date: May 2014 - Jul 2014

Institution: IITH (India)

6.5.2. Visits to International Teams

6.5.2.1. Research stays abroad

Ancuta Iordache visited Maxeler Technologies (London, U.K.) from May 1st 2014 to July 31st 2014. This visit reinforced the collaboration between the two teams in the context of the HARNES E.U. project, and was funded by the EIT ICT Labs Doctoral Training Center. We plan another 3-months visit in 2015.

Ismael Cuadrado Cordero, who is a student of the EIT ICT Labs Doctoral School, visited the Argonne National Laboratory (USA) for a research internship from June to September 2014. He was hosted in Kate Keahey's team working on resource management in scientific clouds.

7. Dissemination

7.1. Promoting Scientific Activities

7.1.1. Program committees

Roberto-Gioacchino Cascella was program committee member of the ORMaCloud 2014 workshop co-located with ACM HPDC, VTC 2015-Spring and VTC 2015-Fall (track "wireless Networks and Security") and in IEEE TrustCom 2015 (track "privacy"). He co-organized the Contrail Business Day held in Roma, Italy on January 23, 2014. He was General co-chair of 2nd edition of DIHC workshop, co-located with Euro-Par 2014, held in Porto, Portugal, on August 26th, 2014.

Eugen Feller was program committee member of PDP (2014, 2015) and IEEE Cluster 2014.

Yvon Jégou was program committee member of IEEE TrustCom 2014 conference and ORMaCloud workshop co-located with ACM HPDC 2014.

Christine Morin was program committee member of the *International Symposium on Parallel and Distributed Computing (ISPD)* 2014, *IEEE International Parallel & Distributed Processing Symposium (IPDPS)* 2014, *IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CC-GRID)* in 2014 and 2015, *CLOSER* 2014, *Resilience (workshop on Resiliency in High Performance Computing in Clusters, Clouds, and Grids)* co-located with EuroPar 2014, *VTDC (International Workshop on Virtualization Technologies in Distributed Computing)* co-located with ACM HPDC in 2015, *ScienceCloud* co-located with ACM HPDC in 2014 and 2015, *OrmaCloud* co-located with ACM HPDC 2014, *CrossCloud (workshop on cloud interoperability and federation)* co-located with INFOCOM in 2014, *Distributed Cloud Computing (DCC) workshop* co-located with SIGCOMM in 2014, *VHPC 2014 and 2015 (workshop on Virtualization in High-Performance Cloud Computing)* co-located with EuroPar, *Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP 2015)*. She is also a member of the SC'15 technical poster committee.

She co-organized the Contrail Business Day held in Roma, Italy on January 23, 2014. She was a member of the scientific committee for the Inria Scientific Days held in Lille in June 2014. She was General co-chair of 2nd edition of DIHC workshop, co-located with Euro-Par 2014, held in Porto, Portugal, on August 26th, 2014. She was *Tutorial Chair* and a mentor in the PhD student mentoring track of the IEEE Cluster conference held in Madrid, Spain, in September 2014.

Anne-Cécile Orgerie was program committee member of FedICI workshop co-located with EuroPar 2014, of VTC2014-Fall, of CCGrid 2015, and of the PhD student mentoring track of Cluster 2014. She was also general co-chair for ExtremeGreen workshop co-located with CCGrid 2014, general co-chair and organizer for GreenDays@Rennes, and co-organizer of the Inria booth at SuperComputing 2014.

Nikolaos Parlavantzas was program committee member of DAIS 2014 and CSB 2014.

Guillaume Pierre was program committee member of Middleware 2014, IC2E 2014, MobileCloud 2014, ComPAS 2014, CloudDP 2014, DADS track of SAC2015, IC2E 2015, AINA 2015 and CSCS20-2015. He was also workshop chair of EuroSys 2014.

Cédric Tedeschi was a program committee member of ICWS 2014, ICCS 2014, and CloSer 2014.

7.1.2. Evaluation committees, consulting

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.

C. Morin acted as an expert to review ERC starting grant proposals for the European Commission. She is 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 jury for the 2014 Inria prizes. She was a consultant for SAFRAN.

Anne-Cécile Orgerie is expert at the HCERES for evaluating the doctoral schools.

Nikolaos Parlavantzas acted as an expert to review proposals for Alpes Grenoble Innovation Recherche (AGIR) 2014.

7.2. Teaching - Supervision - Juries

7.2.1. 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.

Christine Morin:

- Master : Internet of Services: Programming Models & Things and Clouds (ISI), 12 hours ETD, M2RI, University of Rennes 1, France.
- Master : Energy management in computing infrastructures (as part of the Eco-STIC module), 4.5 hours ETD, M1, Supelec, France.

Anne-Cécile Orgerie (at ENS de Rennes):

- License 3: ARCSYS2 - architecture et systèmes 2 (60 hours ETD)

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, Networking and SOA.

Jean-Louis Pazat (at INSA Rennes for 2012-2013):

- 4th year: Advanced Operating Systems (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 ETD)
- Master 1: Service Technologies (24 ETD)
- Master 2: Approche algorithmique des applications et systèmes répartis (32 ETD)

Cédric Tedeschi (at the University of Rennes 1 for 2013-2014):

- Master 1: Multitask Operating Systems (65 hours ETD)
- Master 1: Concurrency in Systems and Networks (70 hours ETD)
- Master 1: Parallel Programming (36h ETD)
- Master 1: Service Technologies (30h ETD)
- Master 2 (research): Internet of Services (6 hours ETD)

E-learning

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

7.2.2. Supervision

PhD : Djawida Dib, Optimizing PaaS Provider Profit under SLA constraints, defended on July 7th 2014, Christine Morin, Nikolaos Parlavantzas.

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 Pizat.

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

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

7.2.3. Juries

- Christine Morin is a reviewer of the HdR committee of Brice Goglin, University of Bordeaux (April 15th, 2014).
- Christine Morin is a member (chair) of the PhD defense committee of Arnab Sinha, Université de Rennes 1 (May 28th 2014).
- Christine Morin is a reviewer in the PhD defense committee of Jonathan Lejeune, University of Paris 6 (September 19th, 2014).
- Christine Morin is a member (chair) of the PhD defense committee of Alessandro Ferreira Leite, University of Paris 11 (December 2nd, 2014).
- Christine Morin is a reviewer of the HdR committee of Jacques Jorda, University of Toulouse 3 - Paul Sabatier (December 11th, 2014).
- Guillaume Pierre is a member of the PhD defense committee of Wubin Li, University of Umea (April 25th 2014).
- Guillaume Pierre is a member of the PhD defense committee of Cassidy Clark, TU Delft (May 14th 2014).
- Guillaume Pierre is a member of the PhD defense committee of Pierpaolo Cincilla, Université Paris 6 (September 15th 2014).
- Anne-Cécile Orgerie is a member of the PhD defense committee of Guillaume Le Louët, Ecole des Mines de Nantes (May 12th 2014)
- Nikolaos Parlavantzas is a member of the PhD defense committee of Chrysostomos Zeginis, University of Crete (October 20th, 2014)
- Jean-Louis Pizat is a reviewer for the PhD thesis of the PhD thesis of Alessandro Leite (December 2, 2014)
- Jean-Louis Pizat is the president of the Habilitation defense of Olivier Barais, Université de Rennes 1 (December 8th 2014)

7.3. Popularization

- Roberto-Gioacchino Cascella presented a demonstration of the VEP-S system at the EIT ICT Labs result day, Helsinki, Finland, December 2014.
- Christine Morin gave a talk entitled "Contrail Project Overview" at the Contrail Business Days, Roma, Italy, January 2014.
- Christine Morin was invited to give a talk on Myriads activities on cloud computing at the EIT ICT Labs meeting on future clouds, Rennes, France, March 2014.
- Christine Morin gave a talk entitled "Contrail: Interoperability and Dependability in a Cloud Federation" at the first workshop of the Joint Laboratory for Extreme Scale Computing, Sophia-Antipolis, France, June 2014.
- Christine Morin was invited to give a talk entitled "Contrail: Toward a Trusted Cloud Market Place" at the EIT ICT Labs Symposium on Future Cloud Computing, Rennes, France, June 2014.
- Christine Morin was invited to give a talk entitled "Contrail: Toward a Trusted Cloud Market Place" at the Inria Scientific Days, Lille, France, June 2014.
- Christine Morin gave an invited talk entitled "Snooze: A Scalable Energy-Efficient IaaS Cloud Management System" at the Next-GWiN 2014 Workshop, Rennes, France, October 2014.
- Christine Morin gave a talk entitled "Checkpointing as a Service in Heterogeneous Clouds" at the second workshop of the Joint Laboratory for Extreme Scale Computing, Chicago, USA, November 2014.
- Guillaume Pierre organized the 2nd ConPaaS workshop (Rennes, May 26th 2014) and gave the opening presentation
- Guillaume Pierre organized the EIT ICT Labs Future Cloud Symposium (Rennes, June 23-24th 2014).
- Anne-Cécile Orgerie gave a talk entitled "Green IT : maîtrise de la consommation énergétiques des équipements informatiques", at Rencontres de la recherche, Fondation Rennes 1, Rennes, September 19, 2014.
- Anne-Cécile Orgerie gave a talk entitled "Toward Energy-Efficient Cloud Computing" at the Inria-Technicolor Workshop, Rennes, December 4, 2014.
- Anne-Cécile Orgerie gave a tutorial entitled "Green Big Data Processing using Hadoop: An Introductory" with Shadi Ibrahim from KerData team, at the ACM/IFIP/USENIX Middleware conference, Bordeaux, December 2014.
- Arnab Sinha presented a demonstration of the VEP-S system at the EIT ICT Labs result day, Helsinki, Finland, December 2014.
- Cédric Tedeschi gave a talk about workflow activities within the DALHIS associate team at the Berkeley Inria Stanford workshop, Paris, June 2014.

7.4. Miscellaneous

Roberto-Gioacchino Cascella is a member of the Inria "Commission locale formation".

Christine Morin is a member of the Project-Team Committee of INRIA RENNES – BRETAGNE ATLANTIQUE (Comité des projets), *Référent Chercheur* for INRIA RENNES – BRETAGNE ATLANTIQUE. She is a member of the Irisa/Inria "Commission Personnel" being in charge of post-docs and "*délé-gations*". She is correspondant for North America relationships in the Inria Direction of International Relationships. She was the coordinator of Contrail European project. She participated in a researchers-journalists exchange organized by the AJSPI (*Association des Journalistes Scientifiques de la Presse d'Information*). In this context, she hosted Sophy Caulier, free-lance journalist in Myriads team, January 13-17, 2014 and she was hosted March 24-28, 2014 as an intern by Guillaume Mézières, editor in chief of *Le laboratoire des Savoirs* hosted at Radio Prun in Nantes. She contributed to a radio programme (<http://labodessavoirs.fr/emissions-du-labo/un-monde-plurilingue/>).

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 ICT Labs master school in Rennes. He is 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 co-leader of the Operating Systems Activity in the CNRS ASR research Group (*Groupement de Recherche*). 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 Partnership department.

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

8. 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, n^o 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, n^o 5, pp. 88–92, http://www.globule.org/publi/CPHECA_ic2012.html

Publications of the year

Doctoral Dissertations and Habilitation Theses

- [10] D. DIB. *Optimizing PaaS Provider Profit under Service Level Agreement Constraints*, Université Rennes 1, July 2014, <https://hal.inria.fr/tel-01091660>

Articles in International Peer-Reviewed Journals

- [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", 2015, 24 p. , <https://hal.inria.fr/hal-01102203>
- [12] J. LIU, S. AHMAD, E. BUYUKKAYA, R. HAMZAOU, G. SIMON. *Resource allocation in underprovisioned multioverlay peer-to-peer live video sharing services*, in "Peer-to-Peer Networking and Applications", April 2014, pp. 1-15 [DOI : 10.1007/s12083-014-0260-8], <https://hal.archives-ouvertes.fr/hal-01073055>
- [13] A.-C. ORGERIE, M. DIAS DE ASSUNCAO, L. LEFÈVRE. *A Survey on Techniques for Improving the Energy Efficiency of Large Scale Distributed Systems*, in "ACM Computing Surveys", December 2014, vol. 46, n^o 4, 31 p. , <https://hal.inria.fr/hal-00767582>

International Conferences with Proceedings

- [14] J. CAO, M. SIMONIN, G. COOPERMAN, C. MORIN. *Checkpointing as a Service in Heterogeneous Cloud Environments*, in "CC-GRID - 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing", Shenzhen, Guangdong, China, Proceedings of the 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, May 2015, <https://hal.inria.fr/hal-01102094>
- [15] A. CARPEN-AMARIE, D. DIB, A.-C. ORGERIE, G. PIERRE. *Towards Energy-Aware IaaS-PaaS Co-design*, in "SMARTGREENS: International Conference on Smart Grids and Green IT Systems", Barcelone, Spain, April 2014, <https://hal.inria.fr/hal-01088946>
- [16] D. DIB, N. PARLAVANTZAS, C. MORIN. *SLA-based Profit Optimization in Cloud Bursting PaaS*, in "14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing", Chicago, United States, May 2014, <https://hal.inria.fr/hal-00940299>
- [17] H. FERNANDEZ, G. PIERRE, T. KIELMANN. *Autoscaling Web Applications in Heterogeneous Cloud Infrastructures*, in "IEEE International Conference on Cloud Engineering", Boston, MA, United States, IEEE, October 2014, <https://hal.inria.fr/hal-00937944>

- [18] *Best Paper*
H. FERNANDEZ, C. STRATAN, G. PIERRE. *Robust Performance Control for Web Applications in the Cloud*, in "4th International Conference on Cloud Computing and Services Science", Barcelona, Spain, April 2014, Best paper award, <https://hal.inria.fr/hal-01006607>.
- [19] B. FLORIN CORNEA, A.-C. ORGERIE, L. LEFÈVRE. *Studying the energy consumption of data transfers in Clouds: the Ecofen approach*, in "CloudNet: IEEE International Conference on Cloud Networking", Luxembourg, October 2014, <https://hal.inria.fr/hal-01088944>
- [20] M. S. HASAN, Y. KOUKI, T. LEDOUX, J.-L. PAZAT. *Cloud Energy Broker: Towards SLA-driven Green Energy Planning for IaaS Providers*, in "HPCC - IEEE International Conference on High Performance Computing and Communications", France, August 2014, pp. 1-8, <https://hal.archives-ouvertes.fr/hal-01015811>
- [21] D. MARGERY, E. MOREL, L. NUSSBAUM, O. RICHARD, C. ROHR. *Resources Description, Selection, Reservation and Verification on a Large-scale Testbed*, in "TRIDENTCOM - 9th International Conference on Testbeds and Research Infrastructures for the Development of Networks & Communities", Guangzhou, China, May 2014, <https://hal.inria.fr/hal-00965708>
- [22] J. PASTOR, M. BERTIER, F. DESPREZ, A. LÈBRE, F. QUESNEL, C. TEDESCHI. *Locality-aware Cooperation for VM Scheduling in Distributed Clouds*, in "Euro-Par 2014", Porto, Portugal, August 2014, <https://hal.inria.fr/hal-00991530>
- [23] 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>
- [24] M. SIMONIN, C. MORIN, H. VALERIE, L. RAMAKRISHNAN, A. DEBORAH, C. TEDESCHI, J. ROJAS BALDERRAMA. *Combining Workflow Templates with a Shared Space-based Execution Model*, in "WORKS '14 : Proceedings of the 9th Workshop on Workflows in Support of Large-Scale Science", New Orleans, United States, November 2014, pp. 50-58, <https://hal.inria.fr/hal-01090670>
- [25] T. WAUTERS, B. VERMEULEN, W. VANDENBERGHE, P. DEMEESTER, S. TAYLOR, L. BARON, M. SMIRNOV, Y. AL-HAZMI, A. WILLNER, M. SAWYER, D. MARGERY, T. RAKOTOARIVELO, F. L. VILELA, D. STAVROPOULOS, C. PAPAGIANNI, F. FRANCOIS, C. BERMUDO, A. GAVRAS, D. DAVIES, J. LANZA, S.-Y. PARK. *Federation of Internet experimentation facilities: architecture and implementation*, in "European Conference on Networks and Communications (EuCNC 2014)", Bologna, Italy, June 2014, <https://hal.inria.fr/hal-01087607>

Scientific Books (or Scientific Book chapters)

- [26] L. LOPEZ, J. ZILINSKAS, A. COSTAN, R. G. CASCELLA, G. KECSKEMETI, E. JEANNOT, M. CANNATARO, L. RICCI, S. BENKNER, S. PETIT, V. SCARANO, J. GRACIA, S. HUNOLD, S. L. SCOTT, S. LANKES, C. LENGAUER, J. CARRETERO, J. BREITBART, M. ALEXANDER. *Euro-Par 2014: Parallel Processing Workshops, Part I*, Lecture Note In Computer Science, Springer, December 2014, vol. 8805, <https://hal.inria.fr/hal-01110069>

- [27] L. LOPEZ, J. ZILINSKAS, A. COSTAN, R. G. CASCELLA, G. KECSKEMETI, E. JEANNOT, M. CANNATARO, L. RICCI, S. BENKNER, S. PETIT, V. SCARANO, J. GRACIA, S. HUNOLD, S. L. SCOTT, S. LANKES, C. LENGAUER, J. CARRETERO, J. BREITBART, M. ALEXANDER. *Euro-Par 2014: Parallel Processing Workshops, Part II*, Lecture Note In Computer Science, Springer, December 2014, vol. 8806, <https://hal.inria.fr/hal-01110071>
- [28] A. LÈBRE, J. PASTOR, M. BERTIER, F. DESPREZ, J. ROUZAUD-CORNABAS, C. TEDESCHI, A.-C. ORGERIE, F. QUESNEL, G. FEDAK. *Beyond The Clouds, How Should Next Generation Utility Computing Infrastructures Be Designed?*, in "Cloud Computing: Challenges, Limitations and R&D Solutions", M. ZAIGHAM (editor), Springer, November 2014, <https://hal.inria.fr/hal-01067888>
- [29] L. RAMAKRISHNAN, D. GHOSHAL, V. HENDRIX, E. FELLER, P. MANTHA, C. MORIN. *Storage and Data Lifecycle Management in Cloud Environments with FRIEDA*, in "Cloud Computing for Data Intensive Applications", Springer, January 2014, <https://hal.inria.fr/hal-00932366>

Research Reports

- [30] 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>
- [31] J. CAO, M. SIMONIN, G. COOPERMAN, C. MORIN. *Checkpointing as a Service in Heterogeneous Cloud Environments*, Inria-IRISA Rennes Bretagne Atlantique ; Northeastern University (Boston, Mass), November 2014, n^o RR-8633, 10 p. , <https://hal.inria.fr/hal-01086834>
- [32] R. G. CASCELLA, S. COSTACHE, F. DUDOUE, F. GAUDENZI, Y. JÉGOU, C. MORIN. *Multi-Cloud Portable Application Deployment with VEP*, Inria, February 2014, <https://hal.inria.fr/hal-01076925>
- [33] R. G. CASCELLA, C. MORIN, J.-P. BANÂTRE, T. PRIOL. *Private-by-Design: Towards Personal Local Clouds*, Inria Rennes, November 2014, n^o RR-8634, <https://hal.inria.fr/hal-01087558>
- [34] S. COSTACHE, S. KORTAS, C. MORIN, N. PARLAVANTZAS. *Market-based Autonomous Application and Resource Management in the Cloud*, Inria , December 2014, n^o RR-8648, 35 p. , <https://hal.inria.fr/hal-01091280>
- [35] E. FELLER, L. RAMAKRISHNAN, C. MORIN. *Performance and Energy Efficiency of Big Data Applications in Cloud Environments: A Hadoop Case Study*, Lawrence Berkeley National Laboratory, California, USA, December 2014, n^o LBNL-6861E, <https://hal.inria.fr/hal-01102198>
- [36] E. FELLER, M. SIMONIN, Y. JÉGOU, A.-C. ORGERIE, D. MARGERY, C. MORIN. *Snooze: A Scalable and Autonomic Cloud Management System*, Inria Rennes, December 2014, n^o RR-8649, 31 p. , <https://hal.inria.fr/hal-01091755>
- [37] D. GHOSHAL, H. VALERIE, E. FELLER, C. MORIN, P. BETH, L. RAMAKRISHNAN. *Data Management Strategies for Scientific Applications in Cloud Environments*, Lawrence Berkeley National Laboratory, California, USA, December 2014, n^o LBNL-6860E, <https://hal.inria.fr/hal-01102191>

Other Publications

- [38] S. ARNOUX, Y. JÉGOU, G. BEAUNÉE, P. EZANNO. *A generic framework to model infection dynamics in a metapopulation of cattle herds*, September 2014, ECCB'14, the 13th European Conference on Computational Biology, <https://hal.inria.fr/hal-01097441>
- [39] R. CARVAJAL-GÓMEZ, D. FREY, M. SIMONIN, A.-M. KERMARREC. *WebGC: Browser-Based Gossiping [Live Demo/Poster]*, December 2014, Middleware 2014 [DOI : 10.1145/2678508.2678515], <https://hal.inria.fr/hal-01080032>

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, n^o 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. CARPEN-AMARIE, A.-C. ORGERIE, C. MORIN. *Experimental Study on the Energy Consumption in IaaS Cloud Environments*, in "6th IEEE/ACM International Conference on Utility and Cloud Computing", Dresden, Germany, December 2013, <https://hal.inria.fr/hal-00905118>
- [46] R. G. CASCELLA, C. MORIN, P. HARSH, Y. JÉGOU. *Contrail: a reliable and trustworthy cloud platform*, in "1st European Workshop on Dependable Cloud Computing", Sibiu, Romania, ACM, May 2012 [DOI : 10.1145/2365316.2365322], <http://hal.inria.fr/hal-00762967>
- [47] CONTRAIL CONSORTIUM. *Deliverable D5.2 Revised specification of computational resource management functionalities for virtual platforms*, Contrail European Integrated Project, July 2012
- [48] S. COSTACHE. *Market-based Autonomous Application and Resource Management in the Cloud*, Université Rennes 1, July 2013, <https://tel.archives-ouvertes.fr/tel-00925352>
- [49] 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>
- [50] 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>

- [51] 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>
- [52] 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>
- [53] 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>
- [54] 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>
- [55] P. HARSH, F. DUDOUET, R. G. CASCELLA, Y. JÉGOU, C. MORIN. *Using Open Standards for Interoperability - Issues, Solutions, and Challenges facing Cloud Computing*, in "6th International DMTF Academic Alliance Workshop on Systems and Virtualization Management: Standards and the Cloud", Las Vegas, United States, October 2012, <http://hal.inria.fr/hal-00720636>
- [56] Y. JÉGOU, P. HARSH, R. G. CASCELLA, C. MORIN. *Open Computing Infrastructures for Elastic Services: Contrail Federation*, in "Invitation to a Journey in the ERA of Cloud Computing", D. PETCU, J. L. VÁZQUEZ-POLETTI (editors), Cambridge & Scholars Publishing, January 2012, <http://hal.inria.fr/hal-00717422>
- [57] Y. JÉGOU, P. HARSH, R. G. GIOACCHINO CASCELLA, F. DUDOUET, C. MORIN. *Managing OVF applications under SLA constraints on Contrail Virtual Execution Platform*, in "6th International DMTF Academic Alliance Workshop on Systems and Virtualization Management: Standards and the Cloud", Las Vegas, United States, October 2012, <http://hal.inria.fr/hal-00721786>
- [58] J. O. KEPHART, D. M. CHESS. *The Vision of Autonomic Computing*, in "Computer", 2003, vol. 36, n^o 1, pp. 41-50
- [59] 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_Computing
- [60] 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