



IN PARTNERSHIP WITH:
CNRS

**Institut national des sciences
appliquées de Rennes**

Université Rennes 1

Activity Report 2013

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	3
2.4.1. Autonomous Management of Virtualized Infrastructures	4
2.4.1.1. SLA-driven PaaS over Cloud Federations	4
2.4.1.2. Virtual Data Centers	4
2.4.1.3. Virtualized Infrastructure Management	5
2.4.1.4. Heterogeneous Cloud Infrastructure Management	5
2.4.2. Multilevel Dynamic Adaptation of Service-based Applications	6
2.4.3. A Chemical Approach for Autonomous Service Computing	7
2.4.3.1. Chemical Expression of Interactions	7
2.4.3.2. Distributed Implementation of the Chemical Paradigm	7
2.5. Highlights of the Year	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	10
4. Application Domains	10
5. Software and Platforms	10
5.1. HOCL-tools	10
5.2. Contrail Virtual Execution Platform (VEP)	11
5.3. Snooze	11
5.4. Resilin	12
5.5. Merkat	12
5.6. ConPaaS	13
5.7. Meryn	13
6. New Results	14
6.1. Dependable Cloud Computing	14
6.1.1. Multi-data Center and Multi-cloud	14
6.1.1.1. Deployment of distributed applications in a multi-provider environment	14
6.1.1.2. Towards a distributed cloud inside the backbone	14
6.1.1.3. Multi-cloud application deployment in ConPaaS	14
6.1.2. Scalability of Snooze Self-healing Cloud Management System	14
6.1.3. Application Performance Modeling in Heterogeneous Cloud Environments	15
6.1.4. Flexible SLA & SLO Management	15
6.2. Heterogeneous Resource Management	15
6.2.1. Cross-resource scheduling in heterogeneous cloud environments	15
6.2.2. Maximizing private cloud provider profit in cloud bursting scenarios	16
6.2.3. Data life-cycle management in clouds	16
6.3. Energy-efficient Resource Infrastructures	16
6.3.1. Energy-efficient IaaS clouds	16
6.3.2. Energy-aware IaaS-PaaS co-design	16
6.3.3. Performance and energy-efficiency evaluation of Hadoop deployment models	17
6.3.4. Energy consumption models and predictions for large-scale systems	17

6.3.5.	Simulating Energy Consumption of Wired Networks	17
6.3.6.	Simulating the impact of DVFS within SimGrid	18
6.4.	Unconventional Models for Large Computations and Platforms	18
6.4.1.	Chemical computing at large scale	18
6.4.2.	Template workflows	18
6.5.	Experimental Platforms	18
6.5.1.	Energy measurement	19
6.5.2.	Deployment of IaaS management system	19
6.5.3.	BonFIRE	19
6.5.4.	Fed4FIRE	19
7.	Bilateral Contracts and Grants with Industry	19
8.	Partnerships and Cooperations	20
8.1.	Regional Initiatives	20
8.2.	National Initiatives	20
8.2.1.	COOP ANR COSINUS Project (2009-2013)	20
8.2.2.	MIHMES ANR Investissements d'Avenir (2012 - 2018)	20
8.2.3.	HEMERA Inria AEN (2010-2013)	20
8.2.4.	Inria ADT Aladdin (2008-2013)	21
8.2.5.	Inria ADT Snooze (2012-2014)	21
8.2.6.	CNRS GDS EcoInfo	21
8.2.7.	CominLabs EPOC project (2013-2016)	21
8.2.8.	IRT B-Com	22
8.3.	European Initiatives	22
8.3.1.	FP7 Projects	22
8.3.1.1.	CONTRAIL	22
8.3.1.2.	BonFIRE	23
8.3.1.3.	PaaSage	23
8.3.1.4.	Fed4FIRE	24
8.3.1.5.	ECO2Clouds	24
8.3.1.6.	HARNESS	24
8.3.2.	Collaborations in European Programs, except FP7	25
8.3.2.1.	ICT COST	25
8.3.2.2.	MC-DATA	26
8.4.	International Initiatives	26
8.4.1.	Inria Associate Teams	26
8.4.2.	Inria International Partners	27
8.4.3.	Inria International Labs	27
8.5.	International Research Visitors	27
9.	Dissemination	27
9.1.	Scientific Animation	27
9.1.1.	Program committees	27
9.1.2.	Evaluation committees, consulting	28
9.2.	Teaching - Supervision - Juries	28
9.2.1.	Teaching	28
9.2.2.	Supervision	29
9.2.3.	Juries	30
9.3.	Popularization and scientific animation	30
9.4.	Miscellaneous	32
10.	Bibliography	32

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

Christine Morin [Team leader, Inria, Senior Researcher, HdR]
Thierry Priol [Inria, Senior Researcher, collaborator, HdR]
Yvon Jégou [Inria, Researcher]
Anne-Cécile Orgerie [CNRS, Researcher]

Faculty Members

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

Engineers

David Margery [Inria, Research Engineer]
Pascal Morillon [Univ. Rennes I, Research Engineer]
Matthieu Simonin [Inria, Research Engineer]
Roberto-Gioacchino Cascella [Inria, granted by FP7 CONTRAIL project]
Stefania Costache [Inria, granted by FP7 CONTRAIL project, from Sep 2013]
Florian Dudouet [Inria, granted by FP7 CONTRAIL project]
Maxence Dunnewind [Inria, granted by FP7 BonFIRE project]
Filippo Gaudenzi [Inria, granted by FP7 CONTRAIL project]
Piyush Harsh [Inria, granted by FP7 CONTRAIL project, until Apr 2013]
Nicolas Lebreton [Inria, granted by FP7 ECO2Clouds project]
Julien Lefeuvre [Inria, granted by FP7 ECO2Clouds project, from Feb 2013]
Usman Malik [Univ. Rennes I, from Sep 2013]
Morteza Neishaboori [Univ. Rennes I, from Sep 2013]
Eric Poupart [Inria, granted by FP7 BonFIRE project]
Yann Radenac [Univ. Rennes I, granted by EIT ICT Labs MC-DATA project]
Dzenan Softic [Univ. Rennes I, from Sep 2013]

PhD Students

Stefania Costache [EDF R&D, CIFRE grant, until Jul. 2013]
Ismael Cuadrado Cordero [Inria, from Oct 2013]
Djawida Dib [Inria, granted by Brittany council and FP7 CONTRAIL project]
Ancuta Iordache [Univ. Rennes I, granted by Harness FP7 project, from Feb 2013.]
Yunbo Li [EMN, from Oct 2013]
Sabbir Hasan [INSA Rennes, from Dec 2013]
Édouard Outin [B-COM, INSA Rennes, from Oct 2013]
Marko Obrovac [Inria, granted by FP7 CONTRAIL project, until Mar 2013]
Chen Wang [Inria, granted by EDF, until May 2013]

Post-Doctoral Fellows

Alexandra Carpen-Amarie [Inria]

Eliya Buyukkaya [Univ. Rennes I, granted by Harness FP7 project, from Nov 2013]
Bogdan Florin Cornea [Inria, from Dec 2013]
Eugen Feller [Inria]
Tran Ngoc Minh [Univ. Rennes I, granted by FP7 CONTRAIL project, from Jan 2013 until Aug 2013]

Administrative Assistants

Maryse Fouché [Inria]
Anne Pasteur [Inria, Contrail Coordinator Assistant, part time 50%]

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 users¹ all over the world. Over time, these applications are becoming more and more sophisticated, a given application being a composition of services likely to be executed on various sites located in different geographical locations. The Internet of Services is spreading all domains: home, administration, business, industry and science. Everyone is involved in the Internet of Services: citizens, enterprises, scientists are application, service and resource consumers and/or providers over the Internet.

Outsourcing. Software is provided as a service over the Internet. Myriads of applications are available on-line to billions of users as, for instance, *GoogleApps* (Gmail). After decades in which companies used to host their entire IT infrastructures in-house, a major shift is occurring where these infrastructures are outsourced to external operators such as Data Centers and Computing Clouds. In the Internet of Services, not only software but also infrastructure are delivered as a service. Clouds 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.

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>.

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 (that 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 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 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

With virtualized infrastructures (clouds) computing and storage become a utility. 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.

2.4.1.1. 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.

2.4.1.2. 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.

2.4.1.3. 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 CDML...) 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 [52], [47]. 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.

2.4.1.4. 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 (see Section 8.3.1.6) 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 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. A Chemical Approach for Autonomous Service Computing

While the very nature of Internet is the result of a decentralized vision of the numeric world, the Internet of Services tends today to be supported by highly centralized platforms and software (data centers, application infrastructures like Google or Amazon, *etc.*). These architectures suffer from technical problems such as lack of fault-tolerance, but also raise some societal and environmental issues, such as privacy or energy consumption. Our key challenge is to promote a decentralized vision of service infrastructures, clearly separating *expression* (description, specification) of the platform from its *implementation*.

2.4.3.1. Chemical Expression of Interactions

As *programming* service infrastructures (in the user's point of view) mainly means expressing the coordination of services, we need an expressive and high level language, abstracting out *low level* implementation details to the user, while being able to model in a simple way the nature of service infrastructures.

Existing standardized languages do not provide this level of abstraction (mixing expression of the service coordination and implementation details). 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 behaviors. 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, ...). Finally, expressing the quality of services in a chemical context is investigated.

2.4.3.2. Distributed Implementation of the Chemical Paradigm

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. Molecules will be distributed among the underlying platform and need to meet to react. To achieve this, we will consider several research tracks. A first track will be algorithmic solutions for information dissemination and retrieval over decentralized (peer-to-peer) networks, allowing nodes to exchange some molecules according to some probabilistic rules. A second track is the development of a shared virtual space gathering the molecules, similar to the series of works conducted around the *Distributed Shared Memory* (DSM) approach, which simulates a global virtual shared memory on top of a distributed memory platform. In both tracks, we will finally consider fault-tolerance, as we cannot afford losing (*too many*) molecules pertained by some reactions of the program, when nodes storing them are unreliable. For example, one of the techniques envisioned for fault-tolerance is replication. Replication must be manipulated with care, as replicating molecules should ensure reactions fulfillment while avoiding to trigger too many reactions (several replicas of the same molecules could trigger a reaction, generating more reactions than specified by the program).

2.5. Highlights of the Year

- Eugen Feller has been awarded the second PhD prize of the MATISSE doctoral school by the Fondation Rennes 1 in March 2013 for his thesis entitled *Autonomic and Energy-Efficient Management of Large-Scale Virtualized Data Centers* [53] defended in December 2012 under the supervision of Christine Morin.

- Matthieu Simonin, Eugen Feller, Yvon Jégou, David Margery, Christine Morin and Anne-Cécile Orgerie have been awarded the second prize at the Scale Challenge organized with the ACM/IEEE CC-Grid 2013 conference held in Delft, the Netherlands in May 2013. They demonstrated the scalability and resilience of Snooze IaaS management system [26], developed as part of Eugen Feller's PhD thesis and that has been supported by an Inria technology development action since October 2012.
- The paper entitled *Resilin: Elastic MapReduce over Multiple Clouds* presented by Ancuta Iordache was amongst the three best paper finalists at the CCGRID'2013 conference [31].

3. Research Program

3.1. Introduction

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's 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 companies 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 computer 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.

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.

The cloud computing model [48], [45] introduces new challenges in the organization of the information infrastructure: security, identity management, adaptation to the environment (costs). The organization of large organization 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

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 [46] 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 [49] has been proposed more than twenty years ago for 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.

The Myriads team will investigate nature-inspired programming such as chemical computing for autonomous service computing.

4. Application Domains

4.1. Overview

The Myriads research activities address a broad range of application domains. We validate our research results with selected use cases from the following application domains:

- Web services, Service oriented applications,
- Business applications,
- Bio-informatics applications,
- Computational science applications,
- Numerical simulations.

5. Software and Platforms

5.1. HOCL-tools

Contact: Cédric Tedeschi, Cedric.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 high 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 with the support of 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): Marko Obrovac, Cédric Tedeschi.

Impact: The compiler is used as a tool within the team to develop HOCL programs. The decentralized workflow execution support has been used extensively to produce results published and presented at several conferences. It is also used in the framework of the DALHIS² associated team.

5.2. Contrail Virtual Execution Platform (VEP)

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

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

Status: Version 2.1

License: BSD

Presentation: Virtual Execution Platform (VEP) [56] is a Contrail service that sits just above IaaS layer at the service provider end of the Contrail cloud federation. The VEP service provides a uniform interface for managing the whole lifecycle of elastic applications on the cloud and hides the details of the IaaS layer to the user. VEP applications are described in OVF (Open Virtualization Format) standard format. Resource usage is controlled by CEE (Constrained Execution Environment) rules which can be derived from SLAs (Service Level Agreement). The VEP service integrates a monitoring system where the major events about the application, mainly resource usage, are made available to the user.

The VEP service provides a RESTful interface and can be exploited directly by users on top of the provider IaaS. OpenNebula and OCCI-based IaaS interfaces are currently supported.

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

Impact: VEP is part of Contrail software stack. Several Contrail partners experiment use cases on top of VEP. External users can experiment with it using the open testbed operated by Myriads team.

5.3. Snooze

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

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

Status: Version 2.1.1

License: GPLv2

Presentation: Snooze [26] [53] [5], [4], 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

²<http://project.inria.fr/dalhis>

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 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 of 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 langage libraries (libcloud, jcloud ...). Snooze also provides a web interface to control the system.

Active contributors (from Myriads team): Eugen Feller, Yvon Jégou, David Margery, Christine Morin, Anne-Cécile Orgerie, 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.

5.4. Resilin

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

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

Status: Version 1.0

License: GNU Affero GPL

Presentation: Resilin [31] 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, Christine Morin, Nikos Parlavantzas.

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

5.5. Merkat

Contact: Nikos Parlavantzas, Nikos.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[23], [22], [11]. 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.

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

5.6. ConPaaS

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

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

Status: Version 1.3.1

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, Yann Radenac, Dzenan Softic.

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.

5.7. Meryn

Contact: Nikos Parlavantzas, Nikos.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 [24]. The current Meryn prototype is implemented in shell script, 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 is developed in the framework of Djawida Dib's PhD thesis.

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

Impact: Meryn is not yet distributed as open source.

6. New Results

6.1. Dependable Cloud Computing

Participants: Roberto-Gioacchino Cascella, Stefania Costache, Florian Dudouet, Eugen Feller, Filippo Gaudenzi, Yvon Jégou, Ancuta Iordache, David Margery, Christine Morin, Anne-Cécile Orgerie, Guillaume Pierre, Nikos Parlavantzas, Yann Radenac, Matthieu Simonin, Cédric Tedeschi.

6.1.1. Multi-data Center and Multi-cloud

6.1.1.1. Deployment of distributed applications in a multi-provider environment

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

The move of users and organizations to Cloud computing will become possible when they will be 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 [39] [50], 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 [51]. 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 [55], [54] [56]. 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.

6.1.1.2. 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 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 basis for this overlay is described as part of a recent research report [44].

6.1.1.3. Multi-cloud application deployment in ConPaaS

Participants: Guillaume Pierre, Yann Radenac.

We extended ConPaaS to support application deployment over multiple clouds. There are two main reasons for this: first, it is a necessary mechanism to allow application migration from one cloud to another, without any service interruption. Second, for some applications it may be useful to execute over multiple clouds on a permanent basis, for reliability reasons for example. The main challenges to address were ensuring full network connectivity between resources acquired in multiple clouds. We addressed these issues by integrating the IPOP virtual network in ConPaaS. Second, we designed protocols to ensure application and data migration without any service interruption during the migration.

6.1.2. Scalability of Snooze Self-healing Cloud Management System

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

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

We evaluated the scalability and resilience of Snooze IaaS management system [26]. Unlike existing systems, for scalability, ease of configuration, and high availability, Snooze is based on a self-organizing and self-healing hierarchical architecture of system services [36], [27], [27]. In Snooze hierarchy, each compute server is managed by a local controller that interacts with one of the group managers to which it is dynamically assigned and the set of group managers is coordinated by a group leader elected among them. We performed an extensive scalability study of Snooze across over 500 servers of the Grid'5000 experimentation testbed. We evaluated the Snooze self-organizing and self-healing hierarchy with thousands of system services. The results show that the resource consumption of the Snooze system services is bounded both during the hierarchy construction and system operation. We also show that Snooze prototype implementation is robust enough to manage thousands of servers and hundreds of VMs. Moreover, its autonomic behavior allows to achieve high availability in the presence of a large number of simultaneous system services failures. Indeed, as long as at least two group managers remain operational the system remains alive. We also demonstrated the application deployment scalability across hundreds VMs on the example of a Hadoop MapReduce application. We participated in the Scale Challenge organized in the framework of the ACM/IEEE CC-Grid 2013 conference [26] and won the second prize.

6.1.3. *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-grain 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. A publication on this topic is currently in preparation.

6.1.4. *Flexible SLA & SLO Management*

Participants: Stefania Costache, Christine Morin, Nikos Parlavantzas.

Merkat is a market-based, SLO-driven, PaaS system for private clouds. Merkat dynamically shares resources between competing applications to ensure a fair resource utilization in terms of application priority and actual resource needs. Resources are allocated through a proportional-share auction while autonomous controllers apply elasticity rules to scale application demand according to resource availability and user priority. Merkat provides users the flexibility to adapt controllers to their application types, and it can support diverse application types and performance goals. Merkat is implemented in Python and uses OpenNebula for virtual machine operations.

We evaluated Merkat in simulation and we analyzed the behavior of the system for multiple user types [23]. Furthermore, we deployed Merkat on Grid'5000 and EDF's tested and tested it with applications representative to EDF [22]. Results showed that: (i) the system provides flexible support for different application types (static and malleable) and different SLOs (deadline and performance); (ii) the system provides good user satisfaction achieving acceptable performance degradation, compared to existing centralized solutions. Furthermore, we extended Merkat to manage different clusters and run MPI applications on them. We also submitted a survey on evolution of resource management systems for shared virtualized computing infrastructures to an international journal. This work was carried out in the framework of Stefania Costache's PhD thesis [11].

6.2. **Heterogeneous Resource Management**

Participants: Eliya Buyukkaya, Djawida Dib, Eugen Feller, Tran Ngoc Minh, Christine Morin, Nikos Parlavantzas, Guillaume Pierre.

6.2.1. *Cross-resource scheduling in heterogeneous cloud environments*

Participants: Eliya Buyukkaya, Tran Ngoc Minh, 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.

6.2.2. *Maximizing private cloud provider profit in cloud bursting scenarios*

Participants: Christine Morin, Djawida Dib, Nikos 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 are developing an open, SLA-driven PaaS system, called Meryn, that aims at providing SLA guarantees to diverse application types while maximizing the PaaS provider profit. Meryn supports cloud bursting and applies a decentralized protocol for selecting cloud resources, trying to minimize the cost of running applications without affecting their agreed quality properties. We have performed a preliminary evaluation of Meryn [24] and worked on optimising the system and performing further experiments on the Grid5000 testbed. This work is part of Djawida Dib's PhD thesis.

6.2.3. *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. In the context of the DALHIS associate team ⁴, we evaluated the importance of this framework on multiple cloud testbeds. Our evaluation showed that storage planning needs to be performed in coordination with compute planning and the specific configuration of virtual machine had a strong impact on the application (e.g., some applications performed better on small instances than large instances) [40].

6.3. Energy-efficient Resource Infrastructures

Participants: Alexandra Carpen-Amarie, Bogdan Florin Cornea, Ismael Cuadrado Cordero, Djawida Dib, Eugen Feller, Yunbo Li, Christine Morin, Anne-Cécile Orgerie, Guillaume Pierre.

6.3.1. *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 [21].

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

Participants: Alexandra Carpen-Amarie, Djawida Dib, Guillaume Pierre, Anne-Cécile Orgerie.

⁴<http://project.inria.fr/dalhis>

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, and plan to start the research to address these issues in 2014.

6.3.3. Performance and energy-efficiency evaluation of Hadoop deployment models

Participants: Eugen Feller, Christine Morin.

The exponential growth of scientific and business data has resulted in the evolution of the cloud computing and the MapReduce parallel programming model. Cloud computing emphasizes increased utilization and power savings through consolidation while MapReduce enables large scale data analysis. The Hadoop framework is the most popular open source software implementing the MapReduce model. In our work, we evaluated Hadoop performance in two modes – the traditional model of collocated data and compute services and separated mode where the services are deployed on separate services. The separation of data and compute services provides more flexibility in environments where data locality might not have a considerable impact such as virtualized environments and clusters with advanced networks. In this work, we also conducted an energy efficiency evaluation of Hadoop on physical and virtual clusters in different configurations. The experiments were performed on the Grid’5000 experimentation testbed. To enable virtual machine management, the Snooze cloud stack developed by the Myriads project-team was used. Our extensive evaluation shows that: (1) performance on physical clusters is significantly better than on virtual clusters; (2) performance degradation due to separation of the services depends on the data to compute ratio; (3) application completion progress correlates with the power consumption and power consumption is heavily application specific [28].

6.3.4. Energy consumption models and predictions for large-scale systems

Participant: Christine Morin.

Responsible, efficient and well-planned power consumption is becoming a necessity for monetary returns and scalability of computing infrastructures. While there is a variety of sources from which power data can be obtained, analyzing this data is an intrinsically hard task. In our work, we described a generic approach to analyze large power consumption datasets collected from computing infrastructures. As a first step, we proposed a data analysis pipeline that can handle the large-scale collection of energy consumption logs, apply sophisticated modeling to enable accurate prediction, and evaluate the efficiency of the analysis approach. We presented the analysis of a power consumption data set collected over a 6-month period from two clusters of the Grid’5000 experimentation platform used in production. We used Hadoop with Pig to handle the large volume of data. Our data processing generated a summary of the data that provides basic statistical aggregations, over different time scales. The aggregate data was then analyzed as a time series using sophisticated modeling methods with R statistical software. We exploited time series to detect outliers and derive hourly and daily power consumption predictive models. We demonstrated the accuracy of the predictive models and the efficiency of the data processing performed on a 55-node cluster at NERSC [34]. Energy models from such large dataset can help in understanding the evolution of consumption patterns, predicting future energy trends, and providing basis for generalizing the energy models to similar large-scale systems.

6.3.5. Simulating Energy Consumption of Wired Networks

Participant: 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.

6.3.6. *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 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.

6.4. Unconventional Models for Large Computations and Platforms

Participants: Marko Obrovac, Christine Morin, Cédric Tedeschi.

6.4.1. *Chemical computing at large scale*

Participants: Marko Obrovac, Cédric Tedeschi.

One of the commonly cited problem when dealing with chemistry-inspired computing is its lack of experimental validation. The DHT-based runtime developed recently, in the framework of Marko Obrovac's PhD thesis [13], has been deployed over the Grid'5000 platform with promising results. This runtime is now mature enough for being considered as a viable candidate to underlie a distributed workflow engine [32].

6.4.2. *Template workflows*

Participants: Christine Morin, Cédric Tedeschi.

In the framework of the DALHIS associate team ⁵, we plan 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 [17]. This work started with the development of a parser of TIGRES.

6.5. Experimental Platforms

Participants: Alexandra Carpen-Amarie, Maxence Dunnewind, Nicolas Lebreton, Julien Lefeuvre, David Margery, Eric Poupart.

⁵<http://project.inria.fr/dalhis>

⁶<http://tigres.lbl.gov/home>

6.5.1. Energy measurement

Participants: David Margery, Maxence Dunnewind, Nicolas Lebreton.

In the context of the ECO₂Clouds project, the BonFIRE infrastructure was updated. At the hardware level power distribution units that report electricity usage for each outlet were installed. At the software layer, a probe reporting energy sources used was configured. This probe gets its information from RTE, the French Electricity transport network, and allows publication of CO₂ metrics for each machine in the testbed. Moreover, access to these metrics was abstracted through the general API to access BonFIRE.

6.5.2. Deployment of IaaS management system

Participant: Alexandra Carpen-Amarie.

The Grid'5000 platform has become one of the most complete testbeds for designing or evaluating large-scale distributed systems, playing an essential role in enabling experimental research at all levels of the Cloud Computing stack and providing configurable cloud platforms similar to commercially available clouds.

However, the complexity of managing the deployment and tuning of large-scale private clouds emerged as a major drawback. Typically, users study specific cloud components or carry out experiments involving applications running in cloud environments. A key requirement in this context is seamless access to ready-to-use cloud platforms, as well as full control of the deployment settings.

To address these needs, we developed a set of deployment tools for open-source IaaS environments, capable of installing and tuning fully-functional clouds on the Grid'5000 testbed [20]. The deployment tools support four widely-used IaaS clouds, namely OpenNebula, CloudStack, Nimbus and OpenStack.

They rely on the concept of extensible engines for defining experiments. Such engines implement all the stages of an experiment: physical node reservations in Grid'5000, environment deployment, configuration and experiment execution. We designed generic engines for nodes reservation and deployment according to a set of requirements specified in a cloud configuration file. Thus, these engines do not require any prior knowledge of lower-level Grid'5000 tools, allowing the user to easily achieve multi-site Grid'5000 deployments based on multiple environments.

6.5.3. BonFIRE

Participants: Maxence Dunnewind, David Margery, Eric Poupart.

The project was reviewed in December 2013 during CloudCom 2013 in Bristol and rated Excellent. The main achievement this year is the introduction of a reservation system for resources on the BonFIRE platform.

6.5.4. Fed4FIRE

Participants: Julien Lefeuvre, Nicolas Lebreton, David Margery.

In Fed4FIRE, two key technologies have been adopted as common protocols to enable experimenter to interact with testbeds. SFA, to provision resources, and 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.

7. Bilateral Contracts and Grants with Industry

7.1. Bilateral Contracts with Industry

7.1.1. EDF R&D (2010-2013)

Participants: Stefania Costache, Christine Morin, Nikos Parlavantzas.

The objective of our collaboration with EDF R&D is to design a resource management system for private clouds that provides support for different application SLAs while maximizing the resource utilization of the infrastructure. Stefania Costache's PhD work [11] is funded through a CIFRE grant with EDF R&D. In 2013, we have completed the implementation of the Merkat prototype and evaluated it with realistic applications provided by EDF R&D and with task farming and batch scheduling environments such as Condor and Torque [23], [22].

8. Partnerships and Cooperations

8.1. Regional Initiatives

8.1.1. ASYST (2010-2013)

Participants: Djawida Dib, Christine Morin, Nikos Parlavantzas.

The objective of the ASYST project (*Adaptation dynamique des fonctionnalités d'un SYSTème d'exploitation large échelle*) funded by the Brittany council is to propose building distributed operating systems as sets of adaptable services. This project funds 50% of a PhD grant (Djawida Dib). In 2013, we have worked on the design and implementation of Meryn [24], a flexible PaaS system that supports dynamically resizing virtual clusters to satisfy SLAs involving completion time and prices.

8.2. National Initiatives

8.2.1. COOP ANR COSINUS Project (2009-2013)

Participants: Yvon Jégou, Christine Morin, Yann Radenac.

The COOP project (<http://coop.gforge.inria.fr/>) funded under the ANR COSINUS program relates to multi level cooperative resource management. The two main goals of this project are to set up a cooperation as general as possible with respect to programming models and resource management systems (RMS) and to develop algorithms for efficient resource selection. Experimentations were conducted in particular with the SALOME platform and TLSE as examples of programming environments and Marcel, DIET and XtreamOS as examples of RMS. Partners involved in the COOP project are the AVALON and RUNTIME INRIA EPI, IRIT and EDF R&D. This project funds a research engineer (Yann Radenac). In 2013, we completed the design and implementation of the modifications needed in XtreamOS Grid distributed operating system in order to integrate the CoORM architecture defined by the Avalon Inria team to support dynamic applications.

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

Participants: Yvon Jégou, Christine Morin.

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

In 2013, we interacted with the INRA/BioEpAR research team in order to improve the initial software prototype and to make it ready for parallelisation. The code has been re-written in C++. In 2014, Inria is in charge of developing a parallel version of the code.

8.2.3. HEMERA Inria AEN (2010-2013)

Participants: Bogdan Florin Cornea, Yvon Jégou, Christine Morin, 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 funds Bodgan Florin Cornea's postdoc.

8.2.4. *Inria ADT Aladdin (2008-2013)*

Participants: Yvon Jégou, David Margery, Pascal Morillon.

The Aladdin technological development action funded by INRIA aims at the construction of a scientific instrument for experiments on large-scale parallel and distributed systems, building on the Grid'5000 platform (<http://www.grid5000.fr>).

As governing body of Grid'5000, it was superseded by a national GIS (Scientific interest group) that was signed in 2012.

As the host of engineers contributed to Grid'5000's technical team by INRIA, it finished operating in 2013. Two engineers of this technical team who are SED⁷ members are still hosted in the Myriads team: David Margery, technical director and Pascal Morillon. Pascal Morillon is the chief engineer for all tasks automating Grid'5000 administration.

8.2.5. *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 2013, we validated Snooze at large scale on the Grid'5000 testbed. A poster was presented at CCGRID 2013 [35] and the results of the study were awarded the second prize at CCGRID2013 scale challenge [26]. We introduced the Apache Cassandra system as database backend in Snooze. We have also started to refactor some parts of the code to enable the use of plugins. We implemented an EC2 interface and a web GUI. Puppet recipes were also released as well as a capistrano based deployment script for Grid'5000.

8.2.6. *CNRS GDS EcoInfo*

Participant: Anne-Cécile Orgerie.

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

8.2.7. *CominLabs EPOC project (2013-2016)*

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

⁷The SED is the INRIA Experimentation and Development Service.

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.

8.2.8. IRT B-Com

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

Yvon Jégou and Jean-Louis Pazat are at IRT B-Com ⁸ one day per week, Yvon Jégou since October 1st and Jean-Louis Pazat since September 1st. 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.

8.3. European Initiatives

8.3.1. FP7 Projects

8.3.1.1. CONTRAIL

Type: COOPERATION

Defi: Internet of Services, Software & Virtualisation

Instrument: Integrated Project

Objectif: Internet of Services, Software and Virtualisation

Duration: October 2010 - September 2013

Coordinator: Inria

Partner: XLAB Razvoj Programske Opreme In Svetovanje d.o.o., Slovenia; Italian National Research Council, ISTI-CNR & IIT-CNR, Italy; Vrije Universiteit Amsterdam, The Netherlands; Science and Technology Facilities Council, STFC, UK; Genias Benelux bv, The Netherlands; Tiscali Italia SpA, Italy; Konrad-Zuse-Zentrum für Informationstechnik Berlin, ZIB, Germany; Hewlett Packard Italiana S.r.l - Italy Innovation Center, Italy; Country Constellation Technologies Ltd, UK; Linagora, France.

Inria contact: Christine Morin

Abstract: The goal of the Contrail 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 Contrail project will provide a complete Cloud platform which integrates Infrastructure-as-a-Service (IaaS) and Platform-as-a-Service (PaaS) offerings [39].

In 2013, we led the evaluation of Contrail software stack [57]. We also completed the design and implementation of VEP [56], [51] advanced features such as the reservation manager and scheduler. We defined a revised version of the API and implemented the CIMI interface. We ported VEP on top of the OpenStack IaaS management system. We worked on the integration of VEP with the other Contrail components. We set up an open permanent testbed for VEP and a testbed running Contrail software stack for internal use by consortium members to allow extensive tests with applications. Christine Morin is the coordinator of Contrail project and Roberto Cascella is the technical manager. Christine Morin leads WP 10 on Contrail global architecture. Yvon Jégou leads WP 5 on VEP and WP 13 on testbeds.

⁸<http://b-com.org/wp/>

8.3.1.2. BonFIRE

Type: COOPERATION

Defi: Future internet experimental facility and experimentally-driven research

Instrument: Integrated Project

Objectif: ICT-2009.1.6

Duration: June 2010 - December 2013

Coordinator: Atos Spain SA (Spain)

Partner: The university of Edinburgh (U.K.); SAP AG (Germany); Universitaet Stuttgart (Germany); Fraunhofer-Gesellschaft zur Foaerung der Angewandten Forshung E.V (Germany); Interdisciplinary Institute for Broadband Technology (Belgium); Universidad Complutense De Madrid (Spain) ; Fundacio Privada I2CAT, Internet I Innovacio Digital A Catalunya (Spain); Hewlett-Packard Limited (U.K.); The 451 Group Limited (U.K.) Technische Universitat Berlin (Germany); University of Southampton (U.K.); Inria (France); Instytut Chemii Bioorganicznej Pan (Poland); Nextworks (Italy); Redzinc Services Limited (Ireland); Cloudium systems Limited (Ireland); Fundacio Centro Tecnologico De Supercomputacion De Galicia (Spain); Centre d'Excellence en technologies de l'Information et de la communication (Belgium); University of Manchester (U.K.);

Inria contact: David Margery

Abstract: The BonFIRE (Building service testbeds for Future Internet Research and Experimentation) project has designed, built and operated a multi-site cloud facility to support applications, services and systems research targeting the Internet of Services community within the Future Internet (<http://www.bonfire-project.eu>). The Myriads team is involved in this project as it hosts the Aladdin ADT.

In the context of BonFIRE, we operate one of the five cloud sites integrated into the BonFIRE cloud federation. This cloud site is based on OpenNebula and can be extended on-request to all the machines of the local Grid'5000 site. We have also contributed to the cloud federation layer and host the integration infrastructure for the project, generated from configuration management tools using puppet.

8.3.1.3. PaaSage

Type: COOPERATION

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

Instrument: Collaborative Project

Duration: October 2012 - September 2016

Coordinator: GEIE ERCIM (France)

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

Inria contact: Nikos 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.

8.3.1.4. Fed4FIRE

Type: COOPERATION

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: June 2010 - December 2013

Coordinator: ATOS SPAIN SA (Spain)

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 to operating model of BonFIRE.

8.3.1.5. ECO2Clouds

Type: COOPERATION

Defi: Future internet experimental facility and experimentally-driven research

Instrument: Specific Targeted Research Project

Objectif: ICT-2011.1.6 – Target outcome c) FIRE Experimentation

Duration: October 2012 - September 2014

Coordinator: Atos Spain SA (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.

8.3.1.6. HARNESS

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.

8.3.2. Collaborations in European Programs, except FP7

8.3.2.1. ICT COST

Participants: Eugen Feller, Christine Morin, Anne-Cécile Orgerie.

Program: ICT COST

Project acronym: IC0804

Project title: Energy efficiency in large scale distributed systems

Duration: 23/01/2009 - 04/05/2013

Coordinator: Professor Jean-Marc PIERSON, IRIT, France, <http://www.irit.fr/cost804/>

Other partners: 22 COST countries and 7 non-COST institutions

Abstract: The COST Action IC0804 proposes realistic energy-efficient alternate solutions to share IT distributed resources. As large scale distributed systems gather and share more and more computing nodes and storage resources, their energy consumption is exponentially increasing. While much effort is nowadays put into hardware specific solutions to lower energy consumptions, the need for a complementary approach is necessary at the distributed system level, i.e. middleware, network and applications. The Action characterizes the energy consumption and energy efficiencies of distributed applications. Then based on the current hardware adaptation possibilities and innovative algorithms it proposes adaptive and alternative approaches taking into account the energy saving dimension of the problem. The Action characterizes the trade-off between energy savings and functional and non-functional parameters, including the economic dimension.

In April 2013, Anne-Cécile Orgerie presented a demonstration of Snooze system at the final COST workshop [36].

8.3.2.2. MC-DATA

Program: EIT ICT Labs

Project acronym: MC-DATA

Project title: Multi-Cloud Data Management

Duration: Jan 2013 - Dec 2014

Coordinator: Imperial College London (IMP, United Kingdom)

Other partners: Université de Rennes 1 (UR1, France), Konrad-Zuse-Zentrum für Informationstechnik Berlin (ZIB, Germany), Swedish Institute of Computer Science (SICS, Sweden), Vodafone (Germany)

Abstract: the MC-DATA project has 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 centers; (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.

8.4. International Initiatives

8.4.1. Inria Associate Teams

8.4.1.1. DALHIS

Title: Data Analysis on Large Heterogeneous Infrastructures for Science

Inria principal investigator: Christine Morin

International Partner:

Lawrence Berkeley National Laboratory (United States) - Advanced Computing for Science department led by Deb Agarwal

Duration: 2013 - 2015

See also: <http://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 user 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 will investigate how to provide execution environments that allow users to seamlessly execute their dynamic data analysis workflows in various research environments.

8.4.2. Inria International Partners

8.4.2.1. Informal International Partners

We collaborate on cloud computing with Stephen Scott, Professor at Tennessee Tech University (TTU) and researcher at Oak Ridge National Laboratory (ORNL). He visited Myriads team in September 2013 to investigate research directions for future joint work on cloud computing for scientific applications. We also collaborate on cloud computing with Kate Keahey from Argonne National Laboratory. She chairs the Contrail European project scientific advisory board. Nikos Parlavantzas is involved in an informal collaboration with Héctor Duran Limon, Professor at the University of Guadalajara, Mexico, who came for a 1 week visit in February 2013.

8.4.3. Inria International Labs

Christine Morin was the Inria@Silicon Valley scientific manager until August 2013. She co-organized with Eric Darve, professor at Stanford University and the Inria international relations department the Berkeley-Inria-Stanford workshop (BIS 2013) held at Stanford University in May 2013. Several Myriads team members (Eugen Feller, Christine Morin, Anne-Cécile Orgerie, 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. She was also involved in an informal collaboration with the CITRIS Social Apps Lab, led by James Holston and Greg Niemeyer from UC Berkeley. Collaboration opportunities between Inria and the Social Apps Lab on smart cities and social sustainability were investigated.

8.5. International Research Visitors

8.5.1. Visits to International Teams

Christine Morin was on sabbatical until August 2013 in the Advanced Computing for Science department at the Lawrence Berkeley National Laboratory (USA). Eugen Feller has been a post-doc in the Advanced Computing for Science department at the Lawrence Berkeley National Laboratory (USA) as part of the Inria@Silicon Valley program since February 2013. He is involved in the DALHIS associate team.

9. Dissemination

9.1. Scientific Animation

9.1.1. Program committees

Roberto Cascella was program committee member for IEEE VTC 2013-Fall and IEEE VTC 2013-Spring.

Eugen Feller was program committee member of EE-LSDS 2013 and PDP (2013, 2014).

Yvon Jégou was program committee member of IEEE TrustCom 2013 conference, GC2013 conference and ORMaCloud workshop co-located with ACM HPDC 2013.

Christine Morin was program committee member of the *International Symposium on Parallel and Distributed Computing* (ISPDC) in 2013 and 2014, *IEEE International Parallel & Distributed Processing Symposium* (IPDPS) in 2014, *IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing* (CC-GRID) in 2013 and 2014, *CLOSER* in 2013 and 2014, *International Conference on Cloud and Green Computing* (CGC) in 2013, *Resilience (Workshop on Resiliency in High Performance Computing in Clusters, Clouds, and Grids)* co-located with EuroPar in 2013, *VTDC (International Workshop on Virtualization Technologies in Distributed Computing)* co-located with ACM HPDC in 2013, *ScienceCloud* co-located with ACM HPDC in 2013, *OrmaCloud* co-located with ACM HPDC in 2013 and 2014, *ESPAS (Workshop on Extreme Scale Parallel Architectures and Systems)* in 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. She was *General Co-Chair* of the IEEE ISPA conference held in Melbourne, Australia, in

July 2013. She was *Program vice-chair* of the IEEE GreenCom conference held in Beijing, China, in August 2013. She was *General co-chair* of the DIHC international workshop co-located with the EuroPar conference held in Aachen, Germany, in August 2013. She is *Tutorial Chair* of the IEEE Cluster conference to be held in Madrid, Spain, in September 2014.

Anne-Cécile Orgerie was program committee member of FedICI workshop co-located with EuroPar 2013, OPTIM workshop co-located with HPCS 2013, IEEE GreenCom 2013, CGC 2013, ITC SSEEGN seminar co-located with ATNAC 2013, TIWDC 2013, HSNCE workshop co-located with COMPSAC 2013, EEHPDC workshop co-located with HPDC 2013, EE-LSDS 2013, and E2HPC2 workshop co-located with EuroMPI 2013. She was also publication co-chair for ICPP 2013, general co-chair for ExtremeGreen workshop co-located with CCGrid 2013, and demo chair for EE-LSDS 2013.

Guillaume Pierre was program committee member of CCGrid 2013, CFSE 2013, IC2E 2013, SAC 2013, DAIS 2013, CSCS19, CloudDP 2013, ISPD 2013, ORMaCloud 2013, BigDataCloud 2013, DIHC 2013, IC2E 2014, DADS track of SAC 2014, and MobileCloud 2014. He was also workshop chair of EuroSys 2014.

Thierry Priol was SC'13 "Grids and Clouds" track co-chair.

Cédric Tedeschi was a program committee member of ICWS 2013, ICCS 2013, CloSer 2013 and ComPas2013.

9.1.2. Evaluation committees, consulting

Roberto Cascella was invited to join the final review of the FUI11 CompatibleOne French project.

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 proposals for the French Research Agency (ANR). She is a member of the ModaClouds European project Advisory Board. Since September 2013, she has been a member of the Irisa/Inria "Commission Personnel" being in charge of post-docs and "*délégations*". She is a member of the scientific council of ENS Cachan. She was a member of a professor position selection committee at *Ecole des Mines de Nantes*.

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

9.2. Teaching - Supervision - Juries

9.2.1. Teaching

Roberto Cascella (at the University of Rennes 1):

- License 3: Inter-networking (30h ETD)

Christine Morin is responsible for the ISI teaching unit.

Christine Morin:

- Master : Internet of Services: Programming Models & Infrastructures (ISI), 12 hours ETD, M2RI, University of Rennes 1, France.
- Master : Cluster single system image operating systems, 4.5 hours ETD, M2, Institut Telecom Sud Paris, 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):

- Master 1: projet logiciel (24 hours ETD)
- License 3: architecture et systèmes 2 (36 hours ETD)

Nikos 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 2012-2013):

- Licence 3: Organization of Operating System (38 hours ETD)
- Licence 3: Algorithmic methods (22 hours ETD)
- Master 1: Multitask Operating Systems (65 hours ETD)
- Master 1: Concurrency in Systems and Networks (56 hours ETD)
- Master 2: Muli-users systems (50 hours ETD)
- Master 2 (research): Internet of Services (6 hours ETD)

9.2.2. Supervision

PhD : Marko Obrovac, Chemical Computing for Distributed Systems: Algorithms and Implementation, Université de Rennes 1, March, 28 2013, Thierry Priol, Cédric Tedeschi.

PhD : Erwan Daubert, Environmental adaptation of services in large scale distributed architectures, INSA de Rennes, May 24 2013, Françoise André, Olivier Barais (Triskell), Jean-Louis Pazat.

PhD : Chen Wang, Using Chemical Metaphor to Express Workflow and Service Orchestration, INSA de Rennes May, 28 2013, Jean-Louis Pazat.

PhD : Stefania Costache, Market-based Autonomous Resource and Application Management in the Cloud, defended on 3 July 2013, Christine Morin, Nikos Parlavantzas.

PhD in progress : Djawida Dib, Dynamic adaptation in distributed systems, October 2010, Christine Morin, Nikos Parlavantzas.

PhD in progress : Ancuta Iordache, Multi-resource optimization for application hosting in heterogeneous clouds, February 2013, Guillaume Pierre.

PhD in progress: Yunbo Li, Resource allocation in a Cloud partially powered by renewable energy sources, 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, October 2013, Christine Morin, Anne-Cécile Orgerie.

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

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

9.2.3. *Juries*

- Christine Morin is a reviewer in the PhD defense committee of Marco Meoni, EPFL (January 8, 2013).
- Christine Morin is a member of the PhD defense committee of Flavien Quesnel, Ecole des Mines de Nantes (February 20th, 2013).
- Christine Morin is a reviewer in the PhD defense committee of Damien Borghetto, Université de Toulouse 1 (June 3rd 2013).
- Christine Morin is a reviewer in the PhD defense committee of Enric Tejedor Saavedra, Université Polytechnique de Catalogne (UPC), Spain (July 15th, 2013).
- Christine Morin is a member of the PhD defense committee of Gylfi Gudmundsson, Université de Rennes 1 (September 12th 2013).
- Christine Morin is a member of the PhD defense committee of Joseph Emeras, Université de Grenoble (October 1st, 2013).
- Christine Morin is the president the HDR committee of Laurent Lefèvre, ENS Lyon (November 14th 2013).
- Christine Morin is a reviewer in the PhD defense committee of Sylvain Lefebvre, CNAM, Paris (December 10th 2013).
- Christine Morin is the president of the PhD defense committee of Minh Tuan Ho, University of Rennes 1 (December 18th, 2013).
- Guillaume Pierre is a member of the PhD defense committee of Jiayi Liu, Telecom Bretagne (November 4th 2013).
- Guillaume Pierre is a member of the promotion committee to the rank of Assistant Professor of George Pallis, university of Cyprus (October 30th 2013).
- Guillaume Pierre is a member of the PhD defense committee of Alexandre van Kampen, Inria (March 8th 2013).
- Guillaume Pierre is a member of the PhD defense committee of Viet-Trung Tran, Inria (January 21st 2013).
- Cédric Tedeschi is a member of the PhD defense committee of Chen Wang, Université Rennes 1 (May 28th 2013)
- Jean-Louis Pizat is a reviewer for the PhD thesis of Frederico Alvares, Ecole des Mines de Nantes (April 9th 2013)
- Jean-Louis Pizat is the president of the PhD defense of Rémy Druilhe, Université de Lille (December 5th 2013)
- Jean-Louis Pizat is the president of the PhD defense of Tan Le Nhan, Université de Rennes 1 (December 10th 2013)

9.3. Popularization and scientific animation

- Roberto Cascella gave a talk entitled “The Contrail approach for interoperable and dependable Clouds” at Create-Net, 12 December 2013 (Trento, Italy)

- Roberto Cascella gave a talk entitled “The Contrail Approach for Interoperable Clouds” at Cloud Interoperability Workshop, 18 September 2013 (Madrid, Spain — <http://www.cloudplugfest.org/cloud-interoperability-week/workshop>)
- Roberto Cascella gave a talk entitled “Open Social Lab: Ad Hoc Clouds for Crowd Social Computing” at Séminaire thématique Inria “Smart Cities”, 17 July 2013 (Paris, France)
- Roberto Cascella was a panelist in the panel “la sécurité dans le monde ouvert d’Internet” at journée CominLabs (Rennes, France)
- Roberto Cascella gave a talk entitled “Contrail: Open Computing Infrastructures for Elastic Services” at pre-FIA workshop “Multi-Cloud Scenarios for the Future Internet”, 7 May 2013 (Dublin, Ireland)
- Djawida Dib gave a talk entitled "Meryn - Open, SLA-driven, Cloud Bursting PaaS " at Lawrence Berkeley National Laboratory, Berkeley, USA, in June 2013.
- Stefania Costache gave a talk entitled "Market-based Autonomous Resource and Application Management in the Cloud" at Queen’s University of Belfast, Belfast, UK, in November 2013.
- Florian Dudouet gave a talk about VEP at the Contrail summer school (Almere, the Netherlands, July 2013)
- Eugen Feller gave a talk entitled “From Autonomic Cloud Management to Storage and Data Management for Scientific Applications” at Ericsson Research Silicon Valley, October 2013 (San Jose, CA, USA).
- Eugen Feller gave a talk entitled “DALHIS – Data Analysis on Large Heterogeneous Infrastructures for Science” at the Berkeley-Inria-Stanford workshop, May 2013 (Stanford University, CA, USA).
- Yvon Jégou gave a talk entitled *Mise en œuvre des SLAs dans la pile logicielle Contrail* at the Workshop “SLA Management in Cloud computing”, colocated with ComPAS’2013.
- Yvon Jégou demonstrated the VEP component of Contrail on the Contrail booth during the *EGI Community Forum 2013*, April 8-12, Manchester UK.
- Yvon Jégou gave talk *Contrail Virtual Execution Platform* at the *OpenNebulaConf* event, September 24-26, Berlin, Germany.
- Christine Morin was invited to give a tutorial on *Energy Management in Large-scale Distributed Systems* at the *Green Wireless Communications school* of the NEWCOM Network of Excellence, Poznan, Poland, in September 2013.
- Christine Morin gave a talk entitled "Dependable Cloud Computing in Contrail" on Inria booth at SC’13, Denver, USA, November 2013.
- 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, November 2013.
- Christine Morin gave a talk on Myriads research activities on cloud computing at CREATE-NET, Trento, Italy in December 2013.
- Guillaume Pierre gave a talk about ConPaaS and organized a hands-on session at Contrail summer school (Almere, the Netherlands, July 2013)
- Guillaume Pierre organized the 1st ConPaaS workshop (Amsterdam, the Netherlands, June 13th 2013) and gave the opening presentation
- Anne-Cécile Orgerie gave a talk entitled “Economies d’énergie dans les systèmes distribués à grande échelle”, at ENS de Rennes, December 2013.
- Anne-Cécile Orgerie gave a talk entitled “Energy-Efficiency in Large-Scale Distributed Systems: an On-Off Approach” at the ACS department at LBNL, Berkeley, USA, May 2013.
- Anne-Cécile Orgerie gave a talk entitled “Snooze: an autonomic and energy-efficient management system for virtualized clusters” at GreenDaysLuxembourg, January 2013.

- Cédric Tedeschi gave a talk *Un protocole pour la capture atomique de molécules* at the “System” track of ComPAS’2013.
- Alexandra Carpen-Amarie gave a talk entitled “Experimental Study of the Energy Consumption in IaaS Cloud Environments ” at GreenDaysLille, November 2013.

9.4. Miscellaneous

Roberto 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, Coordinator of the Inria@Silicon Valley program (in collaboration with Inria DRI) until August 2013, member of the scientific council of ENS Cachan. She is the coordinator of Contrail European project.

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

Jean-Louis Pazat is the leader of the “Large Scale Systems” department of IRISA. He is member of the Steering committee (conseil d’administration) of Insa Rennes. He is a member of the Computer Science Department committee and member of the IRISA-INSA Lab committee. He is the local coordinator for the international exchange of students at the computer science department of Insa.

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.

10. 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] S. V. 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>
- [3] S. COSTACHE, T. ROPARS, C. MORIN. *Semias: Self-Healing Active Replication on Top of a Structured Peer-to-Peer Overlay*, in "29th IEEE International Symposium on Reliable Distributed Systems (SRDS 2010)", New Dehli, India, November 2010, <http://hal.inria.fr/inria-00526101/en>
- [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] E. FELLER, C. ROHR, D. MARGER, 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>

- [6] 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>
- [7] 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>
- [8] 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>
- [9] 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>
- [10] P. RITEAU, C. MORIN, T. PRIOL. *Shrinker: efficient live migration of virtual clusters over wide area networks*, in "Concurrency and Computation: Practice and Experience", June 2012 [DOI : 10.1002/CPE.2861], <http://hal.inria.fr/hal-00714926>

Publications of the year

Doctoral Dissertations and Habilitation Theses

- [11] S. COSTACHE. , *Gestion autonome des ressources et des applications dans un nuage informatique selon une approche fondée sur un marché*, Université Rennes 1, July 2013, <http://hal.inria.fr/tel-00925352>
- [12] E. DAUBERT. , *Adaptation et cloud computing : un besoin d'abstraction pour une gestion transverse*, INSA de Rennes, May 2013, <http://hal.inria.fr/tel-00904364>
- [13] M. OBROVAC. , *Chemical Computing for Distributed Systems: Algorithms and Implementation*, Université Rennes 1, March 2013, <http://hal.inria.fr/tel-00925257>
- [14] C. WANG. , *Chemistry Inspired Middleware for Flexible Service Composition and Application*, INSA de Rennes, May 2013, <http://hal.inria.fr/tel-00932085>

Articles in International Peer-Reviewed Journals

- [15] 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>
- [16] E. CARON, F. CHUFFART, C. TEDESCHI. *When Self-Stabilization Meets Real Platforms: an Experimental Study of a Peer-to-Peer Service Discovery System*, in "Future Generation Computer Systems", June 2013, vol. 29, n^o 6, pp. 1533-1543 [DOI : 10.1016/J.FUTURE.2012.10.003], <http://hal.inria.fr/hal-00762693>
- [17] 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>

- [18] M. OBROVAC, C. TEDESCHI. *Distributed Chemical Computing : A Feasibility Study*, in "International Journal of Unconventional Computing", July 2013, vol. 9, n^o 3-4, pp. 203-236, <http://hal.inria.fr/hal-00716124>
- [19] A.-C. ORGERIE, M. DIAS DE ASUNCAO, L. LEFEVRE. *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, To appear, <http://hal.inria.fr/hal-00767582>

International Conferences with Proceedings

- [20] S. BADIA, A. CARPEN-AMARIE, A. LÈBRE, L. NUSSBAUM. *Enabling Large-Scale Testing of IaaS Cloud Platforms on the Grid'5000 Testbed*, in "TTC - 1st International Workshop on Testing The Cloud, co-located with ISSTA 2013", Lugano, Switzerland, ACM, July 2013, pp. 7-12 [DOI : 10.1145/2489295.2489298], <http://hal.inria.fr/hal-00907888>
- [21] 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, <http://hal.inria.fr/hal-00905118>
- [22] S. V. 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>
- [23] S. V. 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, <http://hal.inria.fr/hal-00821558>
- [24] 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, 2013, <http://hal.inria.fr/hal-00816293>
- [25] 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, 2014, <http://hal.inria.fr/hal-00940299>
- [26] 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, <http://hal.inria.fr/hal-00925664>
- [27] E. FELLER, C. MORIN, M. SIMONIN, A.-C. ORGERIE, Y. JÉGOU. *Snooze: A Self-Organizing, Self-Healing and Self-Optimizing Cloud Management System*, in "The 10th USENIX International Conference on Autonomic Computing (ICAC 2013)", San Jose, CA, United States, June 2013, Demo and poster, <http://hal.inria.fr/hal-00925656>
- [28] E. FELLER, L. RAMAKRISHNAN, C. MORIN. *On the Performance and Energy Efficiency of Hadoop Deployment Models*, in "The IEEE International Conference on Big Data 2013 (IEEE BigData 2013)", Santa Clara, United States, October 2013, <http://hal.inria.fr/hal-00856330>

- [29] H. FERNANDEZ, M. OBROVAC, C. TEDESCHI. *Towards Decentralised Workflow Scheduling via a Rule-Driven Shared Space*, in "Distributed Applications and Interoperable Systems - 13th IFIP WG 6.1 International Conference, DAIS 2013", Florence, Italy, June 2013, <http://hal.inria.fr/hal-00915228>
- [30] 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 (editor), March 2014, <http://hal.inria.fr/hal-00937944>
- [31] 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>
- [32] M. OBROVAC, C. TEDESCHI. *Deployment and Evaluation of a Decentralised Runtime for Concurrent Rule-based Programming Models*, in "14th International Conference on Distributed Computing and Networking (ICDCN 2013)", Mumbai, India, LNCS, January 2013, <http://hal.inria.fr/hal-00762659>
- [33] 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>
- [34] T. SAMAK, C. MORIN, D. BAILEY. *Energy consumption models and predictions for large-scale systems*, in "The Ninth Workshop on High-Performance, Power-Aware Computing", Boston, United States, B. R. DE SUPINSKI, D. LI (editors), IEEE, May 2013, <http://hal.inria.fr/hal-00790349>
- [35] M. SIMONIN, E. FELLER, A.-C. ORGERIE, Y. JÉGOU, C. MORIN. *An Autonomic and Scalable Management System for Private Clouds*, in "The 13th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)", Delft, Netherlands, May 2013, Poster, <http://hal.inria.fr/hal-00925726>
- [36] M. SIMONIN, E. FELLER, A.-C. ORGERIE, Y. JÉGOU, C. MORIN. *Snooze: an Autonomic and Energy-Efficient Management System for Private Clouds*, in "Energy Efficiency in Large Scale Distributed Systems", Vienne, Austria, Springer, 2013, <http://hal.inria.fr/hal-00873486>
- [37] M. TRAN NGOC, G. PIERRE. *Failure Analysis and Modeling in Large Multi-Site Infrastructures*, in "13th International IFIP Conference on Distributed Applications and Interoperable Systems", Florence, Italy, IFIP, March 2013, <http://hal.inria.fr/hal-00804747>

Scientific Books (or Scientific Book chapters)

- [38] D. BALOUEK, A. CARPEN-AMARIE, G. CHARRIER, F. DESPREZ, E. JEANNOT, E. JEANVOINE, A. LÈBRE, D. MARGERY, N. NICLAUSSE, L. NUSSBAUM, O. RICHARD, C. PÉREZ, F. QUESNEL, C. ROHR, L. SARZYNIEC. *Adding Virtualization Capabilities to the Grid'5000 Testbed*, in "Cloud Computing and Services Science", I. IVANOV, M. SINDEREN, F. LEYMAN, T. SHAN (editors), Communications in Computer and Information Science, Springer International Publishing, 2013, vol. 367, pp. 3-20 [DOI : 10.1007/978-3-319-04519-1_1], <http://hal.inria.fr/hal-00946971>
- [39] R. 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, <http://hal.inria.fr/hal-00805713>

- [40] 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, <http://hal.inria.fr/hal-00932366>

Research Reports

- [41] M. BERTIER, M. PERRIN, C. TEDESCHI. , *On the Complexity of Concurrent Multiset Rewriting*, Inria, December 2013, n^o RR-8408, 17 p. , <http://hal.inria.fr/hal-00912554>
- [42] S. V. COSTACHE, N. PARLAVANTZAS, C. MORIN, S. KORTAS. , *Merkat: Market-based Autonomous Application and Resource Management in the Cloud*, Inria, August 2013, n^o RR-8343, <http://hal.inria.fr/hal-00850862>
- [43] H. FERNANDEZ, C. TEDESCHI, T. PRIOL. , *A Chemistry-Inspired Workflow Management System for Decentralizing Workflow Execution*, Inria, March 2013, n^o RR-8268, 23 p. , <http://hal.inria.fr/hal-00803406>
- [44] A. LÈBRE, J. PASTOR, M. BERTIER, F. DESPREZ, J. ROUZAUD-CORNABAS, C. TEDESCHI, P. ANEDDA, G. ZANETTI, R. NOU, T. CORTES, E. RIVIÈRE, T. ROPARS. , *Beyond The Cloud, How Should Next Generation Utility Computing Infrastructures Be Designed?*, Inria, July 2013, n^o RR-8348, <http://hal.inria.fr/hal-00854204>

References in notes

- [45] K. JEFFERY, B. NEIDECKER-LUTZ (editors). , *The Future of Cloud Computing - Opportunities for European Cloud Computing Beyond 2010*, European Commission, January 2010
- [46] G. AGHA. *Computing in pervasive cyberspace*, in "Commun. ACM", 2008, vol. 51, n^o 1, pp. 68-70
- [47] 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>
- [48] 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>
- [49] J.-P. BANÂTRE, P. FRADET, D. LE MÉTAYER. *Gamma and the Chemical Reaction Model: Fifteen Years After*, in "WMP", 2000, pp. 17-44
- [50] R. 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>
- [51] CONTRAIL CONSORTIUM. , *Deliverable D5.2 Revised specification of computational resource management functionalities for virtual platforms*, Contrail European Integrated Project, July 2012
- [52] 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>

- [53] 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>
- [54] P. HARSH, F. DUDOUET, R. 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>
- [55] Y. JÉGOU, P. HARSH, R. 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>
- [56] Y. JÉGOU, P. HARSH, R. 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>
- [57] Y. JÉGOU. , *Revised Interface document*, CONTRAIL Project, March 2013, n^o D10.5
- [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://tr.ietejournals.org/text.asp?2011/28/3/212/81230>
- [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