IN PARTNERSHIP WITH:
CNRS

Institut national des sciences appliquées de Rennes

Université Rennes 1

Activity Report 2012

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 Services

Table of contents

Project-Team MYRIADS

Keywords: Autonomic Computing, Cloud Computing, Unconventional Paradigms, Adaptive Systems, Service Oriented Architecture

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

1. Members

Research Scientists

Yvon Jégou [Research Associate (CR), INRIA]
Christine Morin [Research Director (DR), INRIA, Team leader, HdR]
Anne-Cécile Orgerie [Research Associate (CR), CNRS, since October 2012]
Thierry Priol [Research Director (DR), INRIA, collaborator, HdR]

Faculty Members

Jean-Pierre Banâtre [Professor Emeritus, UNIVERSITY RENNES 1, since November 2012, HdR]
Nikos Parlavantzas [Assistant Professor, INSA]
Jean-Louis Pazat [Professor, INSA, HdR]
Guillaume Pierre [Professor, UNIVERSITY RENNES 1, since November 2012]
Cédric Tedeschi [Assistant Professor, UNIVERSITY RENNES 1]

Engineers

Amine Belhaj [INRIA Associate Engineer (IA), ADT XtreemOS Easy, until October 2012]
Alexandra Carpen-Amarie [INRIA Engineer, EIT ICT Lab activity, since March 2012]
Roberto-Gioacchino Cascella [INRIA Engineer, Contrail European project]
Ghislain Charrier [INRIA Associate Engineer (IA), ADT Aladdin, until November 2012]
Florian Dudouet [INRIA Engineer, Contrail European project]
Maxence Dunnewind [INRIA Engineer BonFire IP Project]
Rémy Garrigue [INRIA Associate Engineer (IA), ADT XtreemOS Easy, until October 2012]
Filippo Gaudenzi [INRIA Engineer, Contrail European project, since May 2012]
Piyush Harsh [INRIA Engineer, Contrail European project]
Ancuta Iordache [INRIA Engineer, EIT ICT Labs RMAC activity, since February 2012]
Sajith Kalathingal [INRIA Engineer, fixed term contract, Cloud ANR project, until September 2012]
Nicolas Lebreton [INRIA Engineer, fixed term contract, 50% Bonfire, 50% Fed4FIRE, since December 2012]
David Margery [INRIA Research Engineer (IR, ADT Aladdin technical director)]
Pascal Morillon [UNIVERSITY RENNES 1 Engineer (IE, ADT Aladdin)]
Eric Poupart [INRIA Engineer, fixed term contract, BonFIRE European Project, since April 2012]
Yann Radenac [INRIA Engineer, COOP ANR project]
Cyril Rohr [INRIA Engineer BonFire IP Project, until March 2012]
Matthieu Simonin [INRIA Research Engineer (IR, ADT Snooze), since October 2012]

PhD Students

Stefania Costache [CIFRE EDF Grant]
Erwan Daubert [INRIA Grant]
Djawida Dib [INRIA and Brittany Regional Council Grant]
Eugen Feller [INRIA Grant]
Héctor Fernández [INRIA Grant, until July 2012]
Guillaume Gauvrit [MENRT Grant until August 2012]
André Lage [INRIA Grant until April 2012]
Marko Obrovac [INRIA and Brittany Regional Council Grant]
Chen Wang [INRIA Grant]

Administrative Assistants

Maryse Fouché [Team administrative assistant (TR) INRIA]
Katarzyna Kucharczyk [INRIA Contrail Scientific Coordinator Assistant, until June 2012]
Anne Pasteur [INRIA Contrail Scientific Coordinator Assistant, part time 50%, since June 2012]

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. MYRIADS originates from the PARIS project-team (1999-2009), which focused on the programming of parallel and distributed systems for large scale simulation.

The objective of MYRIADS is to design and implement systems and environments for autonomous service and resource management in distributed virtualized infrastructures. The team tackles the challenges of dependable application execution and efficient resource management in the future Internet of Services.

2.2. Context

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

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

Outsourcing. Software is provided as a service over the Internet. Myriads of applications are available online 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 *Linkedln* link millions of users who share various kinds of information within communities. Virtual organizations tightly connected to Grids allow scientists to share computing resources aggregated from different institutions (universities, computing centers...). The EGEE European Grid is an example of production Grid shared by thousands of scientists all over Europe.

[1] According to World Stats, there are 1.59 billion Internet users i.e. nearly a quarter of the total world population in March 2009 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 and 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 and 3 engineers. 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

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

organizations necessitates to provide means to the organization administrators to manage and monitor the activity of their members on the cloud: authorization to access services, resource usage and quotas.

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

2.4.1.3. 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 CDMI...) to deal with heterogeneous IaaS clouds and also to attempt to limit the vendor lock-in that is prevalent today. When deploying a VEP, we need to take into account the SLA terms negotiated between the cloud provider and customer. For instance, resource reservation mechanisms will be studied in order to provide guarantees in terms of resource availability. Moreover, we will develop the monitoring and measurement mechanisms needed to assess relevant SLA terms and detect any SLA violation. We also plan to develop efficient mechanisms to support VEP horizontal and vertical elasticity in the framework of cloud federations.

We envision that in the future Internet, a VEP or part of a VEP may migrate from one IaaS cloud to another one. While VM migration has been extensively studied in the framework of a single data center, providing efficient VM migration mechanisms in a WAN environment is still challenging [67], [64]. 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.

2.4.2. 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 [68]. 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, $\cdots$). 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 loosing (*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

- Support for research and innovation Inria Award received in June 2012 by David Margery, Grid'5000 technical director, hosted in Myriads team;
- ASF (ACM SIGOPS de France) Best PhD thesis in system for Anne-Cécile Orgerie who has joined Myriads team in October 2012 (PhD advised by Laurent Lefèvre and Isabelle Guérin-Lassous at ENS de Lyon and defended in September 2011) ;
- Best Paper finalist at CloudCom 2012 (paper presenting the results of the work of the Master internship of Armel Esnault, co-advised by Eugen Feller and Christine Morin);
- Highly successful second annual review by the European Commission of the Contrail project, coordinated by Christine Morin, assisted by Roberto Cascella, technical manager.

3. Scientific Foundations

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

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

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

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

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

3.3. Future Internet and SOA

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

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

However, SOA and Web services technologies are mainly market-driven and sometimes far from the state-of-the-art of distributed systems. Achieving dependability or being able to guarantee Service Level Agreement (SLA) needs much more agility of software elements. Dynamic adaptability features are necessary at many different levels (business processes, service composition, service discovery and execution) and should be coordinated. When addressing very large scale systems, autonomic 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 XtreemOS 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 [65], [62] 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 will work on the design and implementation of system services 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 [63] 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 [66] 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. Application Domains

The MYRIADS research activities address a broad range of applications 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 simulation.

5. Software

5.1. SAFDIS

Contact: Jean-Louis Pazat, Jean-Louis.Pazat@irisa.fr

URL: http://www.irisa.fr/myriads/software/folder.2011-12-13.8949308917/

Status: Version 1.0

License: TBD

Presentation: SAFDIS (Self Adaptation for Distributed Services) is a generic framework allowing the self-adaptation of distributed service based applications within a highly volatile context. Compared to other adaptation frameworks, the main advantages of SAFDIS are its genericity, its distributed nature and the focus on SOAs. SAFDIS is in its final implementation and testing phase within the Myriads team and is being used with a real life use case for emergency services.

The current implementation of SAFDIS is based on a Java OSGi implementation. SAFDIS is written in Java and organized into OSGi bundles. SAFDIS is not tight to any specific operating system and work within any JAVA 1.6 platform. An OSGi implementation is needed (such as the Apache Felix http://felix.apache.org or Equinox eclipse.org/equinox implementations). In order to benefit from the reactive adaptation tools, the Jess engine is also needed as an OSGi bundle (http://www.jessrules.com).

Active contributors (from Myriads project-team): Erwan Daubert, Guillaume Gauvrit, Jean-Louis Pazat.

5.2. HOCL-tools

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

Status: Version 1.0 to be released

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 team. 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 (refer to Section 6.3), 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): Héctor Fernández, 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.

5.3. XtreemOS

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

URL: http://www.xtreemos.eu, http://gforge.inria.fr/projects/xtreemos

Status: Version 3.0

License: GPL-2/BSD depending on software packages composing the system

Presentation: XTREEMOS is a Grid Operating system based on Linux with native support for virtual organizations. Three flavours of XTREEMOS were developed for individual PCs, clusters and mobile devices (PDA, notebooks and smartphones). XTREEMOS has been developed by the XTREEMOS consortium.

XTREEMOS software is a set of services developed in Java, C++ and C. XTREEMOS cluster version leverages KERRIGHED single system image operating system. A permanent testbed composed of computers provided by several XTREEMOS partners has been public since fall 2010. The third public version of XTREEMOS has been released in February 2012 for the OpenSuse Linux distribution. Ready-to-use XtreemOS virtual machine images have been made available for the community.

Active contributors (from Myriads project-team): Amine Belhaj, Rémy Garrigue, Yvon Jégou, Christine Morin, Yann Radenac.

Impact: XtreemOS software has been used as part of the COOP ANR project. It was also used in the ANR CLOUD project.Some services such as XtreemFS are used in various R&D projects including Contrail European project.

5.4. Contrail Virtual Execution Platform (VEP)

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

URL: http://vep.gforge.inria.fr/index.php?title=Main_Page

Status: Version 1.0

License: BSD

Presentation: Virtual Execution Platform (VEP)[32] is a Contrail service that sits just above IaaS layer at the service provider end of the Contrail cloud federation. The VEP 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 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.

5.5. Snooze

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

URL: http://snooze.inria.fr

Status: Version 1.0

License: GPLv2

Presentation: Snooze [25], [26], a novel Infrastructure-as-a-Service (IaaS) cloud management system, which is designed to scale across many thousands of servers and virtual machines (VMs) while being easy to configure, highly available, and energy efficient. For scalability, Snooze performs distributed VM management based on a hierarchical architecture. To support ease of configuration and high availability Snooze implements self-configuring and self-healing features. Finally, for energy efficiency, Snooze integrates a holistic energy management approach via VM resource (i.e. CPU, memory, network) utilization monitoring, underload/overload detection and mitigation, VM consolidation (by implementing a modified version of the Sercon algorithm [69]), 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 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. Snooze provides a RESTful command line interface (CLI) to support virtual cluster (VC) definitions and management (i.e., start, shutdown, destroy, suspend, etc.) as well hierarchy visualization and exporting in GraphML format.

Active contributors (from Myriads team): Eugen Feller, Christine Morin.

Impact: Snooze has been used by students at LIFL, IRIT in France and LBNL in the US in the framework of internships during the summer 2012. It has also been deployed and experimented at EDF R&D. Finally, we know that it was experimented by external users from academia and industry as we received feed-back from them.

5.6. Resilin

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

URL: http://resilin.inria.fr

Status: Version 1.0

License: GNU Affero GPL

Presentation: Resilin [51], [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, Nikos Parlavantzas, Christine Morin.

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

5.7. QU4DS

Contact: Jean-Louis Pazat, `Jean-Louis.Pazat@inria.fr`

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

Status: Version 0.1

License: TBD

Presentation: The QU4DS framework provides PaaS (Platform-as-a-Service) support that fills the gap between the higher-level SaaS (Software-as-a-Service) and the underlying IaaS (Infrastructure-as-a-Service). Qu4DS aids service administrators to define high-level objectives that guide execution management in an automatic and transparent fashion. Moreover, Qu4DS supports the full SLA life-cycle while increasing the service provider profit. SLA support includes service negotiation, instantiation and management on the infrastructure. Orthogonally to these features, complementary actions are in charge of increasing the provider profit guided by SLA constraints.

Currently, Qu4DS targets the development of service providers that use the Master/Worker pattern. Qu4DS assists the development of such services by freeing developers from managing workers and by ensuring their proper execution in accordance with time constraints and by reacting to job failures and delays at runtime. At development time, service developers use the Qu4DS library to develop applications and create a Java jar file. The Qu4DS framework uses this jar file to deploy and manage the service instance on the infrastructure according to SLA constraints.

Active contributors (from Myriads team): André Lage Freitas, Nikos Parlavantzas, Jean-Louis Pazat

5.8. Themis

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

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

Status: Version 1.0

License: TBD

Presentation: Themis is a market-based private PaaS (Platform-as-a-Service) system, supporting dynamic, fine-grained resource allocation and automatic application management[19], [20]. Themis implements a proportional-share auction that ensures maximum resource utilization while providing incentives to applications to regulate their resource usage. Themis 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. Themis is implemented in Python and uses OpenNebula for virtual machine management. Experimental results on the Grid'5000 testbed show that using Themis increases resource utilization and improves application performance. Themis is currently installed and being evaluated by EDF R&D using EDF high-performance applications.

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

Impact: Themis is not yet distributed in open source. However, it has been integrated in EDF R&D portal providing access to internal computing resources and is currently experimented on a testbed at EDF R&D.

6. New Results

6.1. Autonomous Management of Virtualized Infrastructures

Participants: Amine Belhaj, Alexandra Carpen-Amarie, Roberto-Gioacchino Cascella, Stefania Costache, Djawida Dib, Florian Dudouet, Eugen Feller, Piyush Harsh, Rémy Garrigue, Filippo Gaudenzi, Ancuta Iordache, Yvon Jégou, Sajith Kalathingal, Christine Morin, Anne-Cécile Orgerie, Nikos Parlavantzas, Yann Radenac.

6.1.1. Application Deployment in Cloud Federations

Participants: Roberto-Gioacchino Cascella, 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 [17], 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 [43]. In 2012, 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 [40], [29], [32]. We integrated support for delegated certificates and provided test scripts to the Virtual Infrastructure Network (VIN) team. VEP 1.1 was slightly modified to integrate the usage control (Policy Enforcement Point (PEP)) solution developed by CNR. Work is in full progress to implement the CEE management interface and a complete web-based platform for all tasks.

6.1.2. Energy Management in IaaS Clouds: A Holistic Approach

Participants: Eugen Feller, Christine Morin.

Energy efficiency has now become one of the major design constraints for current and future cloud data center operators. One way to conserve energy is to transition idle servers into a lower power-state (e.g. suspend). Therefore, virtual machine (VM) placement and dynamic VM scheduling algorithms are proposed to facilitate the creation of idle times. However, these algorithms are rarely integrated in a holistic approach and experimentally evaluated in a realistic environment. We have designed overload and underload detection and mitigation algorithms and implemented them as well as a modified version of the Sercon existing consolidation algorithm [69] and power management algorithms and mechanisms in a novel holistic energy-efficient VM management framework for IaaS clouds called Snooze [25], [39]. In collaboration with David Margery and Cyril Rohr, we have conducted an extensive evaluation of the energy and performance implications of our system on 34 power-metered machines of the Grid'5000 experimentation testbed under dynamic web workloads. The results show that the energy saving mechanisms allow Snooze to dynamically scale data center energy consumption proportionally to the load, thus achieving substantial energy savings with only limited impact on application performance [26], [48]. Snooze has been released as an open source software since May 2012. It will be further developed and maintained as part of the Snooze ADT. This work has been carried out in the framework of Eugen Feller's PhD thesis [24], [8] funded by the ECO-GRAPPE ANR project.

6.1.3. A Case for Fully Decentralized Dynamic VM Consolidation in Clouds

Participants: Eugen Feller, Christine Morin.

One way to conserve energy in cloud data centers is to transition idle servers into a power saving state during periods of low utilization. Dynamic virtual machine (VM) consolidation (VMC) algorithms are proposed to create idle times by periodically repacking VMs on the least number of physical machines (PMs). Existing works mostly apply VMC on top of centralized, hierarchical, or ring-based system topologies, which result in poor scalability and/or packing efficiency with increasing number of PMs and VMs. We have proposed a novel fully decentralized dynamic VMC schema based on an unstructured peer-to-peer (P2P) network of PMs. The proposed schema is validated using three well known VMC algorithms: First-Fit Decreasing (FFD), Sercon, V-MAN, and a novel migration-cost aware ACO-based algorithm we have designed. Extensive experiments performed on the Grid'5000 testbed show that once integrated in our fully decentralized VMC schema, traditional VMC algorithms achieve a global packing efficiency very close to a centralized system. Moreover, the system remains scalable with increasing numbers of PMs and VMs. Finally, the migration-cost aware ACO-based algorithm outperforms FFD and Sercon in the number of released PMs and requires less migrations than FFD and V-MAN [23], [47]. This work has been done in the context of Armel Esnault's Master internship [57].

6.1.4. Market-Based Automatic Resource and Application management in the Cloud

Participants: Stefania Costache, Nikos Parlavantzas, Christine Morin.

Themis is a market-based Platform-as-a-Service system for private clouds. Themis 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. Themis provides users the flexibility to adapt controllers to their application types, and thus it can support diverse application types and performance goals. We have evaluated Themis through simulation and the obtained results demonstrated the effectiveness of the market-based mechanism[19], [20]. We have recently improved Themis in three ways. First, we extended the resource allocation algorithms to support multiple resources (CPU and memory) and to perform load-balancing between physical nodes while considering the migration cost. Second, we improved the management of applications. We added generic support for virtual cluster deployment, configuration and runtime management and also for application monitoring. Finally, we implemented several adaptation policies to scale elastically applications in term of number of provisioned virtual machines and in term of allocated CPU and memory per virtual machine. Themis is implemented in Python and uses OpenNebula for virtual machine operations. We used Themis to scale elastically two resource management frameworks (Torque and Condor) according to their current workload and also MPI scientific codes according to user-given deadlines. Themis has been deployed on Grid'5000 and also on EDF's testbed, HPSLAB. This work is carried out in the fraemwork of Stefania Costache's PhD thesis.

6.1.5. Autonomous PaaS-level resource management

Participants: Djawida Dib, Christine Morin, Nikos Parlavantzas.

PaaS providers host client applications on provider-owned resources or resources leased from public IaaS clouds. The providers have service-level agreements (SLAs) with their clients specifying application quality requirements and prices. A main concern for providers is sharing their private and leased resources among client applications in order to reduce incurred costs. We have proposed a PaaS architecture based on multiple elastic virtual clusters (VCs), each associated with a specific application type (e.g., batch, MapReduce). The VCs dynamically share the private resources using a decentralised allocation scheme and, when necessary, lease remote resources from public clouds. Resource allocation is guided by the SLAs of hosted applications and resource costs. We have implemented a prototype of this architecture that supports batch and MapReduce applications; the application SLAs constrain completion times and prices. The prototype is currently being evaluated on Grid'5000. This work is performed as part of Djawida Dib's thesis.

6.1.6. Elastic MapReduce on Top of Multiple Clouds

Participants: Ancuta Iordache, Yvon Jégou, Christine Morin, Nikos Parlavantzas.

We have worked on the design and implementation of Resilin. To the best of our knowledge Resilin is the first system which is capable of leveraging resources distributed across multiple potentially geographically distinct locations. Unlike the Amazon s proprietary Elastic Map Reduce (EMR) system, Resilin allows users to perform MapReduce computations across a wide range of resources from private, community, and public clouds such as Amazon EC2. Indeed, Resilin can be deployed on top of most of the open-source and commercial IaaS cloud management systems. Once deployed, Resilin takes care of provisioning Hadoop clusters and submitting MapReduce jobs thus allowing the users to focus on writing their MapReduce applications rather than managing cloud resources. In 2012 we designed and implemented a new version of Resilin based on a service-based architecture, which enables system recovery from errors and can be easily extended and maintained. Important functionalities were added to the system: scaling down the platform, deployment of data analysis systems (Apache Hive, Apache Pig). We have also started to work on the design of policies and mechanisms for the autonomous scaling of the virtual Hadoop clusters managed by Resilin. We perfomed an extensive experimental evaluation of Resilin on top of Nimbus and OpenNebula clouds deployed on multiple clusters of the Grid 5000 experimentation testbed. Our results show that Resilin enables the execution of MapReduce jobs across geographically distributed resources with only a limited impact on the jobs execution time, which is the result of intercloud network latencies [51], [31]. Resilin has been released as an open source software since September 2012. This work was carried out in the framework of the RMAC EIT ICT Labs activity.

6.1.7. Adaptation of the CooRM architecture into XtreemOS

Participants: Amine Belhaj, Rémy Garrigue, Yvon Jégou, Christine Morin, Yann Radenac.

In the framework of the COOP ANR project, we have mainly worked on the adaptation and on the implementation of the CooRM architecture (resulting from the work of the Avalon team at Inria Grenoble - Rhône Alpes in the context of the COOP project) into XtreemOS. The main results include a first version of the design of a decentralized version of CooRM, the modification of XtreemOS to support distributed applications (tested with OpenMPI and MPICH2), and the implementation of a launcher of moldable MPI applications using the modified XtreemOS API. A demonstration was presented to the COOP consortium in December 2012.

To get an operational prototype for evaluation purposes, we also had to fix many bugs in XtreemOS, revise its build chain, help clean the distribution package dependencies in collaboration with Rémy Garrigue (engineer from the ADT XtreemOS Easy), rewrite the code generator, help fix issues related to configuration commands in collaboration with Amine Belhaj (engineer from ADT XtreemOS Easy).

6.1.8. Extending a Grid with Virtual Resources Provisioned from IaaS Clouds

Participants: Amine Belhaj, Alexandra Carpen-Amarie, Rémy Garrigue, Sajith Kalathingal, Yvon Jégou, Christine Morin, Yann Radenac.

XtreemOS is a Grid operating system designed to facilitate the execution of grid applications by aggregating resources on multiple sites. XtreemOS provides virtual organization support and enables Grid users to run applications on the resources made available by their virtual organization. As the number of scientific applications that need access to Grid platforms increases, as well as their requirements in terms of processing power, the limited amount of resources that XtreemOS gathers from its virtual organizations may become a bottleneck. To address this limitation, we extended XtreemOS with the capability to acquire virtual resources from cloud service providers. To this end, we enable XtreemOS to provision and configure cloud resources both on behalf of a user and of a virtual organization. This can be done either on-demand, when a user specifically requires cloud resources, or in a dynamic fashion, when the local grid resources cannot comply with the application needs. Furthermore, we devised a selection mechanism for the cloud service providers, allowing users to rent resources from the providers that best match the requirements of their applications. We implemented our approach as a set of extension modules for XtreemOS and we evaluated the prototype in Grid'5000, using cloud resources provisioned from a private OpenNebula cloud. For this evaluation, we made a extensive use of tools developed jointly by Ascola and Myriads project-teams to easily manage large number of VMs on top of IaaS cloud management software (e.g. OpenNebula, Nimbus, OpenStack) deployed on the Grid'5000 platform. This work was carried out as part of the ANR Cloud project [60], [58] and an EIT ICT Labs activity.

6.1.9. Data Management Frameworks for Scientific Applications in Cloud Environments

Participants: Eugen Feller, Christine Morin.

During Eugen Feller's internship at LBNL, we have worked with Lavanya Ramakrishnan from the Advanced Computing for Science department on the evaluation of Hadoop MapReduce jobs in a virtualized environment. We have investigated the performance and power consumption of scientific MapReduce jobs executed in an environment with separated Hadoop compute and data nodes. This enables data sharing across multiple users and is key to support elastic MapReduce. Snooze cloud management stack was used to manage the VMs. Preliminary experimental results on top of Snooze demonstrate the feasibility of our approach.

6.1.10. Energy Consumption Models and Predictions for Large-scale Systems

Participant: Christine Morin.

We have collaborated with Taghrid Samak from the Advanced Computing for Science department at LBNL on the initial investigation of energy consumption models for Grid'5000 sites using Pig and Hadoop, and data from 6 months logs on 135 resources in the Lyon site. The initial results investigate time-series summarization for the entire dataset. For each resource the average power consumption is evaluated and compared with statistically estimated thresholds. A paper is under preparation.

6.1.11. Management of Large Data Sets

Participant: Christine Morin.

Moderate Resolution Imaging Spectroradiometer (MODIS) aboard NASAs satellites continuously generates data important to many scientific analyses. A dataprocessing pipeline that downloads the MODIS products, reprojects them on HPC systems or clouds and make them available to users through a web portal has been developed. In collaboration with Valerie Hendrix and Lavanya Ramakrishnan from the Advanced Computing for Science department at LBNL we have worked on providing community access to MODIS Satellite Reprojection and Reduction Pipeline and Data Sets. In a future version of the system, users will be able to reproject data on demand and/or run algorithms on the reprojected MODIS data such as an evapotranspiration calculation [30].

6.2. Dynamic Adaptation of Service-based Applications

Participants: Djawida Dib, Erwan Daubert, Guillaume Gauvrit, André Lage, Christine Morin, Nikos Parlavantzas, Jean-Louis Pazat, Chen Wang.

6.2.1. Adaptation for Service-Oriented Architectures

Participants: Erwan Daubert, Guillaume Gauvrit, André Lage, Nikos Parlavantzas, Jean-Louis Pazat, Chen Wang.

Service-Oriented Computing is a paradigm that is rapidly spreading in all application domains and all environments - grids, clusters of computers, mobile and pervasive platforms. The following works take place in the context of the S-CUBE European Network of Excellence.

6.2.1.1. Services adaptation in distributed and heterogeneous systems

Participants: Erwan Daubert, Guillaume Gauvrit, Jean-Louis Pazat.

We are still studying service adaptation in distributed and heterogeneous systems. This work covers different aspects such as structural, behavioral and environmental adaptation, distributed decision and planification of adaptation actions, adaptive allocation of resources for services. A framework called SAFDIS for "Self Adaptation For Distributed Services" has been defined and implemented. It is built as a set of services, providing functionalities useful to build an adaptation system. The analysis phase can take reactive as well as proactive decisions. This gives the ability to either react fast or to take decisions for the long term. This implies the ability to analyze the context with a variable depth of reasoning. Our implementation of the SAFDIS analysis phase also distributes and decentralizes its analysis process to spread the computational load and make the analysis process scalable. The planning phase seeks the set of actions (the plan) needed to adapt the system according to the strategy chosen by the analysis phase. It also schedules the selected actions to ensure a coherent and efficient execution of the adaptation. The planning topic is a well known subject in AI research works and many algorithms already exist in that field to produce efficient schedules. With our SAFDIS framework, the planning phase is able to reuse these algorithms. The resulting plan of actions can have actions that can be executed in parallel.

6.2.1.2. Quality Assurance for Distributed Services

Participants: André Lage, Nikos Parlavantzas, Jean-Louis Pazat.

In the context of the service-centric paradigm, we have designed and developed the Qu4DS (Quality Assurance for Distributed Services) system. Qu4DS is a cloud PaaS solution which fills the gap between SaaS service providers and IaaS infrastructures. Qu4DS provides automatic support for service execution management, aiming at increasing service providers' profits by reducing resource costs as well as fines owning to SLA violations. More specifically, Qu4DS dynamically acquires resources according to the customer demand, deploys service instances and implements QoS assurance mechanisms in order to prevent SLA violations. Qu4DS has been evaluated on Grid'5000 and shown to be effective in reducing service provider's costs [33]. This work has been done in the context of André Lage-Freitas' PhD thesis [10].

6.2.1.3. Self-configuration for Cloud Platforms

Participants: Jean-Louis Pazat, Chen Wang.

By definition, cloud computing offers an abstraction to manage various needs and concepts such as distributed software design, the deployment of such software on dynamic resources and the management of this kind of resources. Thus it is possible to reconfigure (adapt) according to some needs the software as well as the use of the resources. However these reconfigurations that are used on different layers may also have impacts on the others. Moreover these layers are independent and so are able to adapt themselves independently of the others. In our work, we propose to use some adaptation capabilities offered for example by the infrastructure (IaaS) that manages the resources to adapt the software (SaaS). We also propose to use planning algorithms to coordinate the adaptations between them to avoid inconsistency or inefficiency due to concurrent adaptations.

6.2.1.4. Dynamic Adaptation of Chemical services

Participants: Jean-Louis Pazat, Chen Wang.

We have proposed a QoS-aware middleware for dynamic service execution. In the context of dynamic execution, a workflow is defined by composing a set of abstract activities as place holders. Each activity is bound to a suitable partner service, which is selected at run-time from a set of functional equivalent candidates with different non-functional properties such as quality of service (QoS). The service selection process is modeled as a series of chemical reactions. This year, we have studied and implemented fragment replacement in workflows within this environment.

6.2.2. Multi-level Adaptation for Distributed Operating Systems

Participants: Djawida Dib, Christine Morin, Nikos Parlavantzas.

This work focused on enhancing distributed operating systems with the ability to continually adapt to their changing environments. Two challenges arise in this context: how to design the distributed operating system (OS) in order to facilitate dynamic adaptation, and how to ensure that OS-level adaptation does not conflict with application-level adaptation. This work proposed to address these challenges by (1) building the distributed OS as an assembly of adaptable services following the service-oriented architecture; and (2) using a common multi-level adaptation framework to adapt both the OS and the application layers in a coordinated way. To demonstrate the usefulness of the proposed architecture, the work focused on distributed shared memory services and provided examples and experimental results using the Kerrighed distributed OS. The work was performed as part of Djawida Dib's thesis [22].

6.3. A Chemical Approach for Autonomous Service Computing

Participants: Héctor Fernández, Marko Obrovac, Cédric Tedeschi.

6.3.1. Chemical Computing for the Simulation of Agile-Based Software Engineering

Participants: Héctor Fernández, Cédric Tedeschi.

In the framework of Héctor Fernández' internship at Vrije University, we applied the chemical programming model to simulate the behavior of a team developing software with Agile methods. Although an unexpected application, it has been the occasion to widen the range of applications and users of the software prototype developed during Héctor's thesis.

6.3.2. Scalable Atomic Capture of Molecules

Participants: Marko Obrovac, Cédric Tedeschi.

Capturing the reactants involved in a reaction constitutes one of the main challenges in the execution of chemical programs. Doing it at large scale is one of the essential barriers hindering the actual execution of chemical programs at large scale. While the problem resembles the classic resource allocation problem, it differs from it by different aspects. One of the main difference stands in the fact that the probability of a conflict varies during the course of execution. When the number of possible reactions is high, then there is no need for a complex conflict resolution scheme, as it would lead to a useless additional cost. In contrary, when this number drops, the probability of a conflict increases, and a systematic conflict resolution is mandatory to ensure at least one reaction will take place.

An adaptive protocol has been proposed, based on the dynamic combination of several strategies. Based on simulations, we have shown that, by dynamically switching from one strategy to another one, even by locally deciding which protocol to use, it is possible to combine the good properties of the strategies without suffering from their drawbacks [18].

The work was recently extended to take several rules into account. Rules have been defined to be able, not only to choose a strategy, but also to choose the rule to be executed, with the constant objective of maximizing the number of reactions executed in a given time.

6.3.3. DHT-based Runtime for the Chemical Programming Model

Participants: Marko Obrovac, Cédric Tedeschi.

The development of a distributed chemical machine entered its experimental phase with the development of a software prototype containing the following building blocks:

- A distributed hash table structures the network and allows any node to communicate with any other node in a logarithmic number of hops in this logical overlay.
- On top of the distributed hash table, a set of discovery mechanisms allows to find molecules needed in reactions, whatever their location is. These mechanisms are based on complex distribution and retrieval scheme borrowed from the P2P literature.
- The atomic capture protocol described before has been fully integrated in this framework.
- The discovery of molecules has been extended in order to detect the termination of the program and to be able to send the results of the computation back to the requester.

This software prototype has been deployed over the Grid'5000 platform [36].

7. Bilateral Contracts and Grants with Industry

7.1. Bilateral Grants with Industry

7.1.1. EDF R&D (2010-2013)

Participants: Stefania Costache, Christine Morin, Nikos Parlavantzas.

In our collaboration with EDF R&D we investigate resource management in virtualized computing platforms in order to efficiently execute distributed applications with stringent time constraints. Our goal 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 is funded through a CIFRE grant with EDF R&D.

In 2012, we have completed the implementation of Themis prototype and evaluated it with realistic applications provided by EDF R&D and task farming and batch scheduling environments such as Condor and Torque [20], [19].

8. Partnerships and Cooperations

8.1. Regional Initiatives

8.1.1. HOCL4WS (2010-2012)

Participants: Thierry Priol, Cédric Tedeschi, Marko Obrovac.

The objective of the HOCL4WS project is to develop a prototype of a middleware system for the distributed execution of chemical programs (targeted for large scale platforms). It partially funds Marko Obrovac's PhD grant.

8.1.2. 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 provide the view of an Operating System as an "Infrastructure as a Service" (IaaS) and even more as a set of adaptable services. The main functionalities of an Operating System such as memory allocation or job scheduling have to be dynamically adapted to cope with the ever changing environment. This project funds 50% of a PhD grant (Djawida Dib).

In 2012, we have worked on the design and implementation of a PaaS framework for scaling up and down virtual clusters under SLA constraints (price and completion time).

8.2. National Initiatives

8.2.1. ECO-GRAPPE ANR ARPEGE Project (2008-2012)

Participants: Eugen Feller, Christine Morin.

The goal of the ECO-GRAPPE project (http://ecograppe.inria.fr/) funded under the ANR ARPEGE program is to design, implement and validate energy saving policies in clusters.This project funds a PhD grant (Eugen Feller). Partners involved in the ECO-GRAPPE project are EDF R&D and Kerlabs.

In 2012, we completed the implementation of the energy saving algorithms and mechanisms in Snooze and evaluated them experimentally with an elastic web service [39], [26], [48], [24], [8]. We also studied a fully decentralize approach to VM consolidation [47], [23]

8.2.2. 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 will be conducted in particular with the SALOME platform and TLSE as examples of programming environments and Marcel, DIET and XtreemOS as examples of RMS. Partners involved in the COOP project are the GRAAL and RUNTIME INRIA EPI, IRIT and EDF R&D. This project funds a research engineer (Yann Radenac).

In 2012, we completed the design and implementation of the modifications needed in XtreemOS Grid distributed operating system in order to integrate the CooRM architecture defined by the Avalon Inria team to support dynamic applications.

8.2.3. CLOUD ANR project (October 2011 - September 2012)

Participants: Sajith Kalathingal, Christine Morin.

The CLOUD project aims at extending an XtreemOS Grid with resources dynamically provisioned from IaaS clouds. An algorithm to select resources in a multi-cloud environment will be defined. A prototype based on XtreemOS Grid and OpenNebula and Nimbus clouds will be built. This project is related to the EIT ICT labs activity 10239 on cloud computing. It funds a research engineer.

In 2012, we augmented XtreemOS Grid distributed system with the capability to acquire virtual resources from cloud service providers. To this end, we enable XtreemOS to provision and configure cloud resources both on behalf of a user and of a virtual organization. We implemented our approach as a set of extension modules for XtreemOS and we evaluated the prototype on Grid'5000 experimentation platform using cloud resources provisioned from a private OpenNebula cloud [58], [60].

8.2.4. MIHMES ANR Investissements d'Avenir (January 2012 - December 2018)

Participants: Christine Morin, Yvon Jégou.

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 2012, we collected the requirements from the bio-informatics applications and defined a workplan to experiment them on top of the cloud technologies developed by Myriads project-team.

8.2.5. HEMERA Inria AEN (2010-2013)

Participants: Christine Morin, Yvon Jégou.

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.

In 2012, several members of Myriads project-team performed large scale experiments to evaluate the systems and services they proposed. The results obtained are described in our publications.

8.2.6. Inria ADT Aladdin (2008-2012)

Participants: Ghislain Charrier, 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). It structures INRIA's leadership role as the institute is present in 8 of the 9 Grid'5000 sites distributed across France.

An executive committee, where each of the 10 project-teams supporting Grid'5000 in the 8 research centers is represented, meets every month. It gives recommendations to the directors on scientific animation, access policy to the instrument as well as for the hardware and software development according to the resources devoted to this ADT. Yvon Jégou represents INRIA Rennes in this executive committee.

The technical team is now composed of 12 engineers, of which 3 are hosted in the Myriads team (David Margery, technical director, (SED [2] member), Pascal Morillon (SED member), Ghislain Charrier).This technical team is structured in a sysadmin team, managing the instrument, and a development team building the tools to build, execute and analyze experiments.

8.2.7. Inria ADT XtreemOS Easy (2010-2012)

Participants: Amine Belhaj, Rémy Garrigue, Yvon Jégou, David Margery, Christine Morin.

The XtreemOS EASY technological development action funded by INRIA aims at developing a set of tools and environments to ease the installation, configuration, deployment, experimentation and use of the XtreemOS Grid operating system and at providing support to the XtreemOS open source community. Two associate engineers are involved in this project: Amine Belhaj and Rémy Garrigue. David Margery (SED) is tutoring them in software development.

In 2012, we completed a major release of XtreemOS system for the OpenSuse Linux distribution. We operated the open testbed and built ready-to-use virtual machine images to ease the use of the system. We also provided support to the user community.

8.2.8. Inria ADT DAUM (2011-2012)

Participants: Erwan Daubert, Jean-Louis Pazat.

[2] The SED is the INRIA Experimentation and Development Service.

We participate to the ADT DAUM which is coordinated by the Triskell project-team. DAUM is a a Technology Development Action (ADT) by INRIA aiming at providing an integrated platform for distributed dynamically adaptable component based applications. DAUM unites and integrates results and software from the Triskell EPI and the Myriads team. More precisely, DAUM extends the Kevoree component framework designed by Triskell with adaptation mechanisms from the SAFDIS framework designed by Myriads.

DAUM will evaluate this integration by designing a full scale system for a tactical assistant for firefighter officers, in collaboration with the firefighters organization of Ille et Vilaine department (2800 firefighters).

Project duration: October 2011 - September 2012

Triskell budget share: One associated engineer shared with the Triskell EPI

Project Coordinator: Noël Plouzeau, Triskell INRIA Project.

Participants: Myriads, Triskell.

8.2.9. Inria ADT Snooze (2012-2014)

Participants: 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 wider community.

In 2012, we validated Snooze on top of Xen hypervisor. We also started re-implementing Snooze based on the Akka library providing asynchronous data communication. We studied how to re-use in Snooze some OpenStack components such as the image repository storage. We deployed Snooze on multiple sites of the Grid'5000 platform. We implemented the libcloud driver for Snooze.

8.2.10. CNRS GDS EcoInfo

Participant: Anne-Cécile Orgerie.

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

8.2.11. Competitivity Clusters

The COOP ANR project is recognized by the Images & Réseaux cluster.

8.3. European Initiatives

8.3.1. FP7 Projects

8.3.1.1. S-CUBE

Participants: Erwan Daubert, Guillaume Gauvrit, André Lage, Jean-Louis Pazat, Chen Wang.

Title: S-CUBE

Type: COOPERATION (ICT)

Defi: Service & SW architectures, infrastructures and engineering

Instrument: Network of Excellence (NoE)

Duration: March 2008 - February 2012

Coordinator: Universität Duisburg-Essen (Germany)

Others partners: Tilburg University (The Netherlands), City University London (UK), Consiglio Nazionale delle Ricerche (Italy), Center for Scientific and Technological Research, The French National Institute for Research in Computer Science and Control, Lero - The Irish Software Engineering Research Centre (Ireland), Politecnico di Milano (Italy), MTA SZTAKI - Computer and Automation Research Institute, Vienna University of Technology (Austria), Université Claude Bernard Lyon (France), University of Crete,Universidad Politécnica de Madrid (Spain), University of Stuttgart(Germany)

See also: http://www.s-cube-network.eu/

Abstract: S-Cube, the Software Services and Systems Network, will establish an integrated, multidisciplinary, vibrant research community which will enable Europe to lead the software-services revolution, thereby helping shape the software-service based Internet which is the backbone of our future interactive society. An integration of research expertise and an intense collaboration of researchers in the field of software services and systems are needed to address the following key problems:

Research fragmentation: Current research activities are fragmented and each research community (e.g., grid computing or software engineering) concentrates mostly on its own specific techniques, mechanisms and methodologies. As a result the proposed solutions are not aligned with or influenced by activities in r

8.3.1.2. CONTRAIL

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

Title: Contrail

Type: COOPERATION (ICT)

Defi: Internet of Services, Software & Virtualisation

Instrument: Integrated Project (IP)

Duration: October 2010 - September 2013

Coordinator: INRIA (France)

Others partners: 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; EBM WebSourcing, France.

See also: http://contrail-project.eu/

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 [44], [55].

In 2012, we led the revision of Contrail overall architecture [42]. We also progressed on the design and implementation of VEP advanced features such as the reservation manager and scheduler [43]. We defined a revised version of the API. 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. She leads WP 10. Yvon Jégou leads WP 5 on VEP and WP 13 on testbeds.

8.3.1.3. HARNESS

Participant: Guillaume Pierre.

Title: Hardware- and Network-Enhanced Software Systems for Cloud Computing

Type: COOPERATION (ICT)

Defi: Pervasive and Trusted Network and Service Infrastructures

Instrument: STREP

Duration: October 2012 - September 2015

Coordinator: Imperial College (United Kingdom)

Others partners: École Polytechnique Fédérale de Lausanne, Konrad-Zuse-Zentrum für Informationstechnik Berlin, Maxeler Technologies and SAP AG.

See also: http://www.harness-project.eu/

Abstract: Cloud computing systems are currently composed of large numbers of relatively inexpensive computers, interconnected by standard IP routers and supported by stock disk drives. However, many demanding applications have now reached a fundamental limit in their ability to scale out using traditional machines. Future performance improvements will derive from the use of high-end specialized equipment in addition to standard hardware: GPUs of course, but also FPGAs, programmable routers, and advanced storage technologies. In this context the European project HARNESS investigates: (i) how cloud providers may offer such extremely heterogeneous hardware to its users; and (ii) how cloud customers may make use of these heterogeneous resources to run their applications such that they exhibit the best possible price-performance tradeoff.

8.3.1.4. PaaSage

Participants: Christine Morin, Nikos Parlavantzas.

Title: PaaSage - Model-Based Cloud Platform Upperware

Type: ICT

Instrument: Large Scale Integrated Project

Duration: October 2012 - September 2016

Coordinator: ERCIM

Other partners: SINTEF, STFC, University of Stuttgart, CETIC, FORTH, BE.Wan, EVRY Solutions, SysFera, Flexiant, Lufthansa Systems AG, GWDG, Automotive Simulation Center Stuttgart

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

Abstract: Software developers targeting the Cloud want an easy way to develop their software in a fashion that exploits the full potential of the clouds, and still is able to run on any of the available offerings. Current platforms are heterogeneous and tend to impose a specific architecture on deployed applications. Accordingly, there is a significant dependency between client applications and the services provided by the platform. It is generally up to the developer to specify and exploit platform services to her best knowledge. However, the typical developer will neither know how to use these characteristics, nor how they impact on the overall behaviour and, what is more, how they relate to a given Cloud infrastructure. To address this complexity, PaaSage will deliver an open and integrated platform to support model-based lifecycle management of Cloud applications. The platform and the accompanying methodology will allow model-based development, configuration, optimisation, and deployment of existing and new applications independently of the existing Cloud infrastructures.

In 2012, we studied the state of the art and worked on the system requirements and specifications.

8.3.1.5. Eco2Clouds

Participants: David Margery, Christine Morin, Anne-Cécile Orgerie, Nicolas Lebreton.

Title: Experimental Awareness of CO2 in Federated Cloud Sourcing

Type: ICT

Instrument: STREP

Duration: October 2012 - September 2014

Coordinator: ATOS

Other partners: The University of Manchester, EPCC, HLRS, Politecnico Di Milano, Inria.

See also: http://eco2clouds.eu

Abstract: The ECO2CLOUD project tackles CO_2 emission awarness in virtualized infrastructures, applying its results to the BonFIRE facility. We specifically tackle the question of predicatable costs for the user despite the varying load on the infrastructure and tractable cost models and APIs to enable application deployment optimization and adaptation.

8.3.1.6. BonFire

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

Title: BonFIRE, Building service testbeds on FIRE

Type: COOPERATION (ICT)

Defi: Future Internet experimental facility and experimentally-driven research

Instrument: Integrated Project (IP)

Duration: June 2010 - November 2013

Coordinator: ATOS SPAIN SA (Spain)

Others partners: 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.) Techniche 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 Technologico De Supercomputacion De Galicia (Spain); Centre d'Excellence en technologies de l'Information et de la communication (Belgium); University of Manchester (U.K.);

See also: http://www.bonfire-project.eu/

Abstract:he BonFIRE (Building service testbeds for Future Internet Research and Experimentation) project will design, build and operate 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.

8.3.1.7. FED4FIRE

Participants: Nicolas Lebreton, David Margery.

Title: Federation for Future Internet Research and Experimentation

Type: ICT

Instrument: Integrated Project

Duration: October 2012 - September 2016

Coordinator: iMinds

Other partners: IT Innovation, UPMC, Fraunhofer, TUB, UEDIN, Inria, NICTA, ATOS, UTH, NTUA, UNIVBRIS, i2CAT, EUR, DANTE Limited, UC, NIA.

See also: http://www.fed4fire.eu

Abstract: The key outcome of Fed4FIRE will be an open federation solution supporting all stakeholders of FIRE. Fed4FIRE is bringing together key players in Europe in the field of experimentation facilities and tool development who play a major role in the European testbeds of the FIRE initiative projects.

8.3.1.8. SCALUS Marie Curie Initial Training Networks (MCITN) (2009-2013)

Participant: Christine Morin.

Title: SCALUS - SCALing by means of Ubiquitous Storage

Type: PEOPLE (ICT)

Defi: elevating education, research, and development inside the area of storage architectures with a focus on cluster, grid, and cloud storage

Instrument: Marie Curie Initial Training Networks (MCITN)

Duration: 4 years

Coordinator: Padeborn University, Germany

Others partners: Paderborn Center for Parallel Computing (PC2), Germany ; BSC, Spain ; Durham University, UK ; Goethe Universität Frankfurt, Germany ; FORTH-ICS, Greece ; Universidad Politecnica De Madrid, Spain ; Ecole des Mines de Nantes, France ; XLAB, Slovenia ; Universität Hamburg, Germany ; Xyratex, UK ; Fujitsu Technology Solutions GmbH, Germany (associated partner) ; CERN, Switzerland (associated partner) ; Microsoft Research, UK (associated partner) ; NEC, Germany (associated partner) ; ORACLE, Germany (associated partner).

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

Abstract: The consortium of this Marie Curie Initial Training Network (MCITN) SCALing by means of Ubiquitous Storage (SCALUS) aims at elevating education, research, and development inside the area of storage architectures with a focus on cluster, grid, and cloud storage. The vision of the SCALUS MCITN is to deliver the foundation for ubiquitous storage systems, which can be scaled in arbitrary directions (capacity, performance, distance, security, . . .). The consortium involves 8 full academic partners, 2 full industrial partners and 5 additional associated industrial partners. Christine Morin participates in this project by co-advising with Professor Ludwig from the University of Hamburg a PhD student (Amandine Pignier) working on Load Balancing and Scheduling in Parallel and Cluster File Systems.

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 March 2012, Eugen Feller organized a meeting for the participants in the focus group on "Energy and QoS-Aware Workload Management in Clouds" in Rennes.

8.3.2.2. RMAC

Participants: Ancuta Iordache, Yvon Jégou, Christine Morin, Nikos Parlavantzas.

Program: EIT ICT Labs

Project acronym: RMAC

Project title: Resource Management Across Clouds

Duration: January-December 2012

Coordinator: Dick Epema, TU Delft and TU Eindhoven

Other partners: Institut Telecom, KTH, TU Delft and TU Eindhoven

See also: http://www.pds.ewi.tudelft.nl/ghit/projects/rmac/

Abstract: The main goal of this activity is to provide solutions for effective, efficient, elastic resource management across multiple clouds at the IaaS level for a wide range of application types (e.g., applications that fit the MapReduce paradigm and data-intensive applications) in federated public and private cloud infrastructures as extensions of the current systems of the partners.

In 2012, we implemented a new version of Resilin, a software which provides the Amazon Elastic MapReduce API and allows users to leverage resources from one or multiple public and/or private clouds. Resilin is now implemented as a distributed and loosely-coupled system whose business logic is separated into distinct services that can be distributed over the network, combined and reused. We also performed an extensive experimental evaluation conducted on multiple clusters of the Grid'5000 experimentation testbed [59], [31].

8.4. International Initiatives

8.4.1. Inria Associate Teams

8.4.1.1. DataCloud@Work

Participants: Alexandra Carpen-Amarie, Christine Morin.

Title: DataCloud@Work

INRIA principal investigator: Gabriel Antoniu, Kerdata

International Partner: Valentin Cristea

Institution: University Polytechnical Bucharest (UPB)

Laboratory: Team of Prof. Valentin Cristea

Duration: 2010 - 2012

See also: http://www.irisa.fr/kerdata/doku.php?id=cloud_at_work:start

The goal of the Associated team is to study massive data management in cloud based service infrastructures. In this context, the Myriads team is involved in a study aiming at the integration of the BlogSeer large scale storage system in XtreemOS distributed system in a vision where XtreemOS is used for the management of IaaS clouds.

8.4.2. Participation In International Programs

Since September 2011, Christine Morin has been an affiliate at Lawrence Berkeley National Laboratory working in the Advanced Computing for Science (ACS) department of the Computational Research Division (CRD) headed by Deb Agarwal. She is actively engaged in three research collaborations with ACS personnel including data management frameworks for scientific applications in cloud environments (with Lavanya Ramakrishnan), use of data-mining and machine-learning techniques to improve resource and failure management in large-scale infrastructures (with Taghrid Samak), and providing community access to MODIS Satellite Reprojection and Reduction Pipeline and Data Sets [30](with Valerie Hendrix and Lavanya Ramakrishnan). During her 2-year sabbatical visit at the Lawrence Berkeley National Laboratory, Christine Morin is the scientific manager of the Inria@SiliconValley program [54]. Deb Agarwal visited Myriads team in May 2012. The Dalhis associate team proposal was submitted in September 2012.

8.5. International Research Visitors

8.5.1. Visits to International Teams

Eugen Feller did a 3-month internship at the Lawrence Berkeley National Laboratory from July to Septmeber 2012. This internship was partially funded by a fellowship from Ecole Doctorale Matisse. E. Feller has worked with L. Ramakrishnan and C. Morin on the evaluation of Hadoop MapReduce jobs in a virtualized environment.

Héctor Fernández did a 1-month internship at Vrije University in November 2011. This internship was funded by the S-Cube network of Excellence. H. Fernández worked with P. Lago on the simulation through the use of the chemical programming model of Agile Software engineering.

9. Dissemination

9.1. Scientific Animation

9.1.1. Leadership, Steering Committees and community service

C. Morin is general co-chair for IEEE ISPA 2013, co-chair of track 4 on Storage, Big Data, and Cloud Computing of IEEE GreenCom 2013, tutorial chair for IEEE Cluster 2014 and the organizer of a BOF session on interoperability in cloud federations for scientific applications at SC'12 (November 2012).

A.-C. Orgerie is publication co-chair for ICPP 2013, demonstration chair for EE-LSDS 2013 (http://www.irit.fr/~Georges.Da-Costa/ee-lsds2013/) and general co-chair for ExtremeGreen 2013 (workshop in conjunction with CCGrid 2013 http://www.ens-lyon.fr/LIP/RESO/extremegreen/).

T. Priol is a member of the Editorial Board of the *International Journal of Intelligent Engineering Informatics* and of the *International Journal of Web Services Research*. He chaired the organization committee the *Let's imagine the Future* seminar organized for Jean-Pierre Banâtre's retirement celebration, November 8-9, Rennes, France.

J.-L. Pazat is the chairman of the Steering Committee of RenPar (http://www.renpar.org/).

C. Tedeschi is a member of the steering committee of the ComPAS (RenPar/SympA/CFSE) joint conference series.

9.1.2. Program Committees

R. Cascella is a program committee member of Q2SWinet 2012 (8th ACM International Symposium on QoS and Security for Wireless and Mobile Networks), IEEE VTC 2013 spring track "Wireless Networks, Access Control, and Resource Management", IEEE VTC 2013 fall track "Wireless Networks and Security".

Y. Jégou is a program committee member of IEEE TrustCom 2013 conference, and ORMaCloud workshop co-located with ACM HPDC 2013.

C. Morin is a program committee member of IEEE CC-GRID 2012, EuroPar 2012, IEEE CloudCom 2012, the International Symposium on Parallel and Distributed Computing (ISPDC) 2012 and 2013, the 4th workshop on scientific computing (ScienceCloud) and of the 4th workshop on Virtualization Technologies in Distributed Computing (VTDC) co-located with ACM HPDC in 2013, CFSE 2013, CLOSER 2013, and IEEE CC-GRID 2013.

A.-C. Orgerie is a program committee member of E2HPC2 2013, HSNCE 2013 and EE-LSDS 2013.

J.-L. Pazat is a program committee member of RenPar'21, GPC2012, CPC2012.

G. Pierre is a program committee member of CCGrid 2013, CFSE 2013, IC2E 2013, SAC 2013, DAIS 2013, CSCS19, CloudDP 2013 and ISPDC 2013.

T. Priol is a program committee member of CCGRID12, CLOUD12, CloudCom12, HPCS12, ICPP12, SC12 and VECPAR12. He will act as one of the co-chairs of the "Grids and Clouds" track of SC13.

C. Tedeschi is a program committee member of ICCS 2013, CloSer 2013 and RenPar'21

9.1.3. Evaluation committees, consulting

J.-P. Banâtre was the INRIA representative in the Contrail General Assembly until September 2012. Since October 2012, he has been a member of the Contrail Scientific Advisory Board.

Y. Jégou is a member of the Grid 5000 executive committee;

Y. Jégou 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 cloud computing expert group in the Software & Service Architectures and Infrastructures unit of the European Commission. She was a member of the committee for the Inria Boost Your Code competition. She is a member of the ModaCloud European project Advisory Board.

J.-L. Pazat is the coordinator for the evaluation of submissions to international bilateral programs in the Information technology domain at Ministry of Higher Education and Research (MESR).

T. Priol has been the INRIA representative in the Contrail General Assembly since October 2012. He was a member of the Contrail Scientific Advisory Board until September 2012.

9.1.4. Conferences, seminars, invitations

Roberto Cascella gave presentation *Leveraging Contrail to support reliable IoT based applications* at the 4th EU-Japan Symposium on the "New Generation Network" and the "Future Internet", January 19, 2012, Tokyo Japan;

Roberto Cascella gave talk *Contrail: A reliable and trustworthy cloud platform* at the 1st European Workshop on Dependable Cloud Computing, EWDCC'2012, May 8 2012, Sibiu, Romania;

Roberto Cascella gave talk *The future of the Cloud* at the Contrail Summerschool, July 23, 2012, Almere, The Netherlands;

Roberto Cascella was invited to give talk *The Cloud and the Contrail project: open computing infrastructures for elastic services* at Politecnico di Bari, September 5, 2012, Bari, Italy;

Roberto Cascella and Yvon Jégou presented the Contrail project on the Contrail booth during the *Journée professionnelle Cloud Computing - MEITO*, INRIA-Rennes, december 11;

Roberto Cascella was invited to give talk *Contrail: A trustworthy cloud platform providing interoperability* at Seconda Università di Napoli, December 19, 2012, Aversa, Italy;

Eugen Feller presented a demonstration entitled *Snooze : un système auto-organisant pour la gestion de l'énergie dans les clouds* at the *Sciences Numériques et efficacité énergétique* Inria industry meeting, Grenoble, France, March 2012.

Eugen Feller gave a talk entitled *Scalable, Autonomic, and Energy-Aware Management of Large-Scale Cloud Infrastructures*, at COST IC0804 Energy and QoS-aware Workload Management in Clouds focus group meeting, Rennes, France, March 2012.

Eugen Feller gave a talk entitled *Energy Management in IaaS Clouds: A Holistic Approach* at EDF R&D (SINETICS department), Paris, France, April 2012.

Eugen Feller gave a talk entitled *Autonomic and Energy-Efficient Management of Large-Scale Virtualized Infrastructures* in the Lawrence Berkeley National Laboratory CS seminar, Berkeley California, USA, August 2012.

Piyush Harsh and Yvon Jégou presented the Contrail VEP on the INRIA booth during the "Solutions Linux" event in Paris, June 19-21, 2012;

Ancuta Iordache presented a talk entitle *Elastic MapReduce on multiple clouds* at the EIT ICT Labs RMAC kick-off meeting in Delft, February 20, 2012;

Ancuta Iordache gave a talk entitled *Elastic MapReduce on Cloud Federations* at the EIT Cloud workshop in Paris, November 29, 2012;

Yvon Jégou gave a talk *The benefits of mixing grid and cloud workloads for genomics* at the "Workshop on Science Applications and Infrastructure in Clouds and Grids" held in conjunction with OGF 34, March 15-16, 2012, Oxford, UK;

Yvon Jégou gave a talk about Contrail Virtual Execution Platform at the "EIT ICT Labs - Contrail workshop", March 19, Paris;

Yvon Jégou gave a talk on *Interoperability in Cloud Federations* and participated to the panel discussion during the "Future Internet Assembly" meeting, May 10-11 2012, Aalborg, Denmark;

Yvon Jégou gave a talk entitled *The Need of Interoperability between Clouds* during the "2nd FIA roadmapping workshop", June 26, 2012, Bruxelles;

Yvon Jégou gave talk *Contrail: Open Computing Infrastructures for Elastic Services* during the "IoS Collaboration Days", October 16-18, 2012, Bruxelles;

Yvon Jégou gave a talk *Interoperability support in cloud federations: Contrail approach based on Virtual Execution Platforms* at the "Interoperability in Scientific Cloud Federations" Birds-of-a-Feather Sessions, SC'12, Salt Lake City, November 13, 2012;

Yvon Jégou, Piyush Harsh and Christine Morin presented Contrail software stack on Contrail European project booth at SC'12, Salt Lake City, November 12-15, 2012.

Christine Morin gave a talk *Contrail: Elastic PaaS Services over Federated IaaS Clouds* by visio-conference for participants in the ATLAS experiment at CERN, February 2, 2012.

Christine Morin gave a talk by visio-conference on Cloud Computing research activities in Rennes at the EIT ICT Labs meeting organized for W. Jonker's visit in Rennes, February 28, 2013.

Christine Morin was invited to give a seminar entitled *Elastic Services on top of Cloud Federations: Contrail Approach* at ORNL and at the Tennessee Technical University in March 2012.

Christine Morin gave a talk entitled *Elastic Services on top of Cloud Federations: Contrail Approach* at Google, Mountain View, April 23, 2012.

Christine Morin presented Myriads project-team research activities at the Inria evaluatio seminar for the "Distributed Systems and Services" theme, October 9, 2012.

Christine Morin gave a talk entitled *Cloud in San Francisco Bay Area* at the CLASS conference, Bled, Slovenia, October 24, 2012.

Christine Morin gave a talk entitled *Work in Progress on Cloud Computing in Myriads Team and Contrail European Project* at the the eighth workshop of the Joint Laboratory for Petascale Computing, Argonne National Laboratory, November 20, 2012.

Christine Morin gave a talk entitled *Destination Inria for researchers!* at the Destination Europe conference, San Francisco, December 7, 2013.

Anne-Cécile Orgerie gave a talk *Energy-efficiency in Cloud infrastructures in Myriads team* at the COST IC0804-IC0805 joint meeting in Cork, Ireland, October 18, 2012;

Anne-Cécile Orgerie gave a talk *ECOFEN: an End-to-end energy Cost mOdel and simulator For Evaluating power consumption in large-scale Networks* at the COST IC0804 meeting in Istanbul, Turkey, November 6;

9.1.5. Administrative Responsibilities

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), 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 committe. 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.

9.2. Teaching - Supervision - Juries

9.2.1. Teaching

Christine Morin is responsible for a graduate teaching Module *Internet of Services: Programming Models & Infrastructures* of the Master Program in Computer Science (research), UNIVERSITY RENNES 1.

Christine Morin:

Master : Internet of Services: Programming Models & Infrastructures, 12 hours ETD, M2RI, University of Rennes 1, France

Master : Cluster single system image operating systems, 4.5 hours ETD, M2, Institut Telcecom Sud Paris, France

Master : Energy management in computing infrastructures (as part of the Eco-STIC module), 4.5 hours ETD, M1, Supelec, France

Nikos Parlavantzas (at INSA Rennes):

- 4th year: Operating Systems (40 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):

- 4th year: Advanced Operating Systems (31 hours ETD)
- 4th year: Parallel Programming (39 hours ETD)
- 4th year: Networking and SOA (39 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'exloitation 2 (67 hours ETC)
- Master 2: Techniques de développement logiciel dans le Cloud (39 ETD)
- Master 2: Approche algorithmique des applications et systèmes répartis (32 ETD)

Cédric Tedeschi (at the University of Rennes 1):

- 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: Muti-users systems (50 hours ETD)
- Master 2 (research): Internet of Services (6 hours ETD)

9.2.2. Supervision

PhD : André Lage, Dynamic Adaptation of Services on Distributed Infrastructures, March 29, 2012, Nikos Parlavantzas, Jean-Louis Pazat

PhD : Héctor Fernández, Flexible Coordination Through the Chemical Programming Model for Service Infrastructures, Université de Rennes 1, June, 20 2012, Thierry Priol, Cédric Tedeschi

PhD : Eugen Feller, Energy aware resource management in IaaS clouds, Université de Rennes 1, December 17, 2012, Christine Morin

PhD in progress : Erwan Daubert, Environmental adaptation of services in large scale distributed architectures, October 2009, Françoise André, Jean-Louis Pazat, Olivier Barais (Triskell)

PhD discontinued : Guillaume Gauvrit, Distributed Dynamic Adaptation of SBAs, October 2008, Françoise André, Jean-Louis Pazat

PhD in progress : Chen Wang, Using Chemical Metaphor to Express Workflow and Service Orchestration October 2009, Jean-Louis Pazat

PhD in progress : Marko Obrovac, Large Scale Runtime of the Chemical Programming Model for Service Orchestration, January 2010, Thierry Priol, Cédric Tedeschi

PhD in progress : Stefania Costache, An economical approach for resource management in private clouds, May 2010, Christine Morin, Nikos Parlavantzas

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

9.2.3. Juries

- Christine Morin is a reviewer for the PhD thesis 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 (defense planned in February, 2013)
- Christine Morin is a reviewer for the PhD thesis of Damien Borghetto, Université de Toulouse 1 (defense planned in Spring 2013)

9.3. Popularization

Contrail summer school (Almere, the Netherlands, July 2012):

- Piyush Harsh gave a talk about the Contrail Virtual Execution Platform and organized a hand-on session on VEP;
- Roberto Cascella gave a presentation entitled *The Future of the Cloud* during the summer school.

Grid'5000 winter school (Nantes, France, December 2012):

- David Margery gave a talk about Grid'5000 and a tutorial about accessing Grid'5000
- Alexandra Carpen-Amarie gave a tutorial on deploying IaaS stacks on Grid'5000, as well as a challenge entry to demonstrate the deployment of CloudStack.
- Marko Obrovac took part in the challenge to demonstrate his experimental approach for his work on chemical computing entitled *Large-scale chemical computing on Grid'5000*

10. Bibliography

Major publications by the team in recent years

[1] M. ALDINUCCI, F. ANDRÉ, J. BUISSON, S. CAMPA, M. COPPOLA, M. DANELUTTO, C. ZOCCOLO. *Parallel program/component adaptivity management*, in "ParCo 2005", Málaga, Spain, 13-16 September 2005.

[2] J. BUISSON, F. ANDRÉ, J.-L. PAZAT. *Supporting Adaptable Applications in Grid Resource Management Systems*, in "8th IEEE/ACM International Conference on Grid Computing", Austin, USA, 19-21 September 2007.

[3] M. COPPOLA, Y. JÉGOU, B. MATTHEWS, C. MORIN, L. P. PRIETO, O. D. SÁNCHEZ, E. YANG, H. YU. *Virtual Organization Support within a Grid-wide Operating System*, in "IEEE Internet Computing", March 2008, vol. 12, n^{o} 2, p. 20-28.

[4] 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.

[5] Y. RADENAC. *Programmation "chimique" d'ordre supérieur*, Université de Rennes 1, April 2007.

[6] L. RILLING. *Système d'exploitation à image unique pour une grille de composition dynamique : conception et mise en oeuvre de services fiables pour exécuter les applications distribuées partageant des données*, Université de Rennes 1, IRISA, Rennes, France, November 2005, In French.

[7] P. RITEAU. *Plates-formes d'exécution dynamiques sur des fédérations de nuages informatiques*, Université Rennes 1, December 2011, http://hal.inria.fr/tel-00651258/en.

Publications of the year

Doctoral Dissertations and Habilitation Theses

[8] 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.

[9] H. FERNÁNDEZ. *Coordination flexible fondée sur la métaphore chimique dans les infrastructures de services*, Université Rennes 1, June 2012, http://hal.inria.fr/tel-00717057.

[10] A. LAGE FREITAS. *Gestion autonome des services dirigée par des accords au niveau des services*, INSA de Rennes, March 2012, http://hal.inria.fr/tel-00715375.

Articles in International Peer-Reviewed Journals

[11] A. BESSANI, R. KAPITZA, D. PETCU, P. ROMANO, S. V. GOGOUVITIS, D. KYRIAZIS, R. CASCELLA. *A look to the old-world sky: EU-funded dependability cloud computing research*, in "Operating Systems Review", July 2012, vol. 46, n^{o} 2, p. 43-56 [*DOI :* 10.1145/2331576.2331584], http://hal.inria.fr/hal-00718716.

[12] 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", October 2012 [*DOI :* 10.1016/J.FUTURE.2012.10.003], http://hal.inria.fr/hal-00762693.

[13] E. FELLER, J. MEHNERT-SPAHN, M. SCHOETTNER, C. MORIN. *Independent Checkpointing in a Heterogeneous Grid Environment*, in "Future Generation Computer Systems", January 2012, vol. 28, n^{o} 1, p. 163-170 [*DOI :* 10.1016/J.FUTURE.2011.03.012], http://hal.inria.fr/inria-00605914.

[14] M. OBROVAC, C. TEDESCHI. *Distributed Chemical Computing : A Feasibility Study*, in "International Journal of Unconventional Computing", July 2012, http://hal.inria.fr/hal-00716124.

[15] A.-C. ORGERIE, M. DIAS DE ASUNCAO, L. LEFEVRE. *On the Road to Energy-Efficient Computing and Network: Where Exactly are We?*, in "ACM Computing Surveys", 2013, To appear, http://hal.inria.fr/hal-00767582.

[16] 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.

Invited Conferences

[17] 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.

International Conferences with Proceedings

[18] M. BERTIER, M. OBROVAC, C. TEDESCHI. *A Protocol for the Atomic Capture of Multiple Molecules at Large Scale*, in "13th International Conference on Distributed Computing and Networking", Hong-Kong, China, Springer, January 2012, http://hal.inria.fr/hal-00644262.

[19] S. COSTACHE, N. PARLAVANTZAS, C. MORIN, S. KORTAS. *Themis: A Spot-Market Based Automatic Resource Scaling Framework*, in "HPDC 2012 Poster Session", Delft, Netherlands, June 2012, http://hal.inria.fr/hal-00721695.

[20] S. COSTACHE, N. PARLAVANTZAS, C. MORIN, S. KORTAS. *Themis: Economy-Based Automatic Resource Scaling for Cloud Systems*, in "14th IEEE International Conference on High Performance Computing and Communications (HPCC 2012)", Liverpool, United Kingdom, 2012, http://hal.inria.fr/hal-00698583.

[21] M. DAVIS, B. VILLAIN, J. RIDOUX, A.-C. ORGERIE, D. VEITCH. *An IEEE-1588 compatible RADclock*, in "International IEEE Symposium on Clock Synchronization for Measurement Control and Communication (ISPCS)", San Francisco, United States, September 2012, 6 [*DOI :* 10.1109/ISPCS.2012.6336624], http://hal.inria.fr/hal-00765158.

[22] D. DIB, N. PARLAVANTZAS, C. MORIN. *Towards Multi-Level Adaptation for Distributed Operating Systems and Applications*, in "12th International Conference on Algorithms and Architectures for Parallel Processing (ICA3PP-12)", FUKUOKA, Japan, Y. XIANG, I. STOJMENOVIC, B. O. APDUHAN, G. WANG, K. NAKANO, A. ZOMAYA (editors), 2012, vol. 7440, p. 100–109 [*DOI :* 10.1007/978-3-642-33065-0_11], http://hal.inria.fr/hal-00712309.

[23] E. FELLER, C. MORIN, A. ESNAULT. *A Case for Fully Decentralized Dynamic VM Consolidation in Clouds*, in "4th IEEE International Conference on Cloud Computing Technology and Science (CloudCom)", Taipei, Taiwan, Province Of China, December 2012, http://hal.inria.fr/hal-00734449.

[24] E. FELLER, C. MORIN. *Autonomous and Energy-Aware Management of Large-Scale Cloud Infrastructures*, in "PhD Forum of the 26th IEEE International Parallel & Distributed Processing Symposium (IPDPS PhD Forum)", Shanghai, China, May 2012, http://hal.inria.fr/hal-00676295.

[25] 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.

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

[27] F. FOUQUET, E. DAUBERT, N. PLOUZEAU, O. BARAIS, J. BOURCIER, J.-M. JÉZÉQUEL. *Dissemination of reconfiguration policies on mesh networks*, in "DAIS 2012", Stockholm, Sweden, June 2012, http://hal.inria.fr/hal-00688707.

[28] F. FOUQUET, G. NAIN, B. MORIN, E. DAUBERT, O. BARAIS, N. PLOUZEAU, J.-M. JÉZÉQUEL. *An Eclipse Modelling Framework Alternative to Meet the Models@Runtime Requirements*, in "Models 2012", Innsbruck, Austria, October 2012, http://hal.inria.fr/hal-00714558.

[29] 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.

[30] V. HENDRIX, J. LI, K. JACKSON, L. RAMAKRISHNAN, Y. RYU, K. BEATTIE, C. MORIN, D. SKINNER, C. VAN INGEN, D. AGARWAL. *Community Access to MODIS Satellite Reprojection and Reduction Pipeline and Data Sets*, in "AGU", San Francisco, United States, December 2012, Poster, http://hal.inria.fr/hal-00762851.

[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 (CCGrid 2013)", Delft, The Netherlands, May 2013, to appear.

[32] 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.

[33] 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.

[34] M. OBROVAC, C. TEDESCHI. *On the Feasibility of a Distributed Runtime for the Chemical Programming Model*, in "14th Workshop on Advances in Parallel and Distributed Computational Models", Shanghai, China, July 2012, http://hal.inria.fr/hal-00716114.

[35] M. OBROVAC, C. TEDESCHI. *When Distributed Hash Tables Meet Chemical Programming for Autonomic Computing*, in "15th International Workshop on Nature Inspired Distributed Computing (NIDisC 2012)", Shanghai, China, July 2012, http://hal.inria.fr/hal-00716065.

[36] 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.

[37] C. WANG, J.-L. PAZAT. *A Two-Phase Online Prediction Approach for Accurate and Timely Adaptation Decision*, in "International Conference on Service Computing", honolulu, Hawaii, United States, IEEE, June 2012, http://hal.inria.fr/hal-00705289.

National Conferences with Proceeding

[38] F. FOUQUET, E. DAUBERT, N. PLOUZEAU, O. BARAIS, J. BOURCIER, A. BLOUIN. *Kevoree : une approche model@runtime pour les systèmes ubiquitaires*, in "UbiMob2012", Anglet, France, June 2012, http://hal.inria.fr/hal-00714557.

Scientific Books (or Scientific Book chapters)

[39] E. FELLER, L. RILLING, C. MORIN. *Towards Energy-Efficient, Scalable and Resilient IaaS Clouds*, in "Achieving Federated and Self-Manageable Cloud Infrastructures: Theory and Practice", M. VILLARI, I. BRANDIC, F. TUSA (editors), IGI Global, May 2012, to appear, http://hal.inria.fr/hal-00651740.

[40] 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.

Research Reports

[41] 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 Grid'5000*, Inria, July 2012, n° RR-8026, 18, http://hal.inria.fr/hal-00720910.

[42] CONTRAIL CONSORTIUM. *Deliverable D10.4 Revised specification of the system architecture*, Contrail European Integrated Project, November 2012.

[43] CONTRAIL CONSORTIUM. *Deliverable D5.2 Revised specification of computational resource management functionalities for virtual platforms*, Contrail European Integrated Project, July 2012.

[44] CONTRAIL CONSORTIUM. *Periodic Management Report – P2*, Contrail European Integrated Project, November 2012.

[45] F. DESPREZ, G. FOX, E. JEANNOT, K. KEAHEY, M. KOZUCH, D. MARGERY, P. NEYRON, L. NUSSBAUM, C. PÉREZ, O. RICHARD, W. SMITH, G. VON LASZEWSKI, J. VÖCKLER. *Supporting Experimental Computer Science*, Argonne National Laboratory, March 2012, n° Technical Memo 326, http://hal.inria.fr/hal-00720815.

[46] F. DESPREZ, G. FOX, E. JEANNOT, K. KEAHEY, M. KOZUCH, D. MARGERY, P. NEYRON, L. NUSSBAUM, C. PÉREZ, O. RICHARD, W. SMITH, G. VON LASZEWSKI, J. VÖCKLER. *Supporting Experimental Computer Science*, Inria, July 2012, n° RR-8035, 29, http://hal.inria.fr/hal-00722605.

[47] E. FELLER, C. MORIN, A. ESNAULT. *A Case for Fully Decentralized Dynamic VM Consolidation in Clouds*, Inria, August 2012, n° RR-8032, http://hal.inria.fr/hal-00722245.

[48] E. FELLER, C. ROHR, D. MARGERY, C. MORIN. *Energy Management in IaaS Clouds: A Holistic Approach*, Inria, April 2012, n° RR-7946, http://hal.inria.fr/hal-00692236.

[49] H. FERNÁNDEZ, M. OBROVAC, C. TEDESCHI. *Decentralised Multiple Workflow Scheduling via a Chemically-coordinated Shared Space*, Inria, April 2012, n° RR-7925, 14, http://hal.inria.fr/hal-00690357.

[50] H. FERNÁNDEZ, C. TEDESCHI, T. PRIOL. *A Chemistry-Inspired Workflow Management System for a Decentralized Composite Service Execution*, Inria, April 2012, n° RR-7924, 23, http://hal.inria.fr/hal-00689785.

[51] A. IORDACHE, C. MORIN, N. PARLAVANTZAS, P. RITEAU. *Resilin: Elastic MapReduce over Multiple Clouds*, Inria, October 2012, n° RR-8081, http://hal.inria.fr/hal-00737030.

[52] M. OBROVAC, C. TEDESCHI. *Deployment and Evaluation of a Decentralised Runtime for Concurrent Rule-based Programming Models*, Inria, November 2012, n° RR-8145, 20, http://hal.inria.fr/hal-00755997.

[53] L. RIDEAU, B. SERPETTE, C. TEDESCHI. *Formalization and Concretization of Ordered Networks*, Inria, December 2012, n° RR-8172, 22, http://hal.inria.fr/hal-00762627.

Other Publications

[54] *Inria@silicon Valley*, July 2012, Emergence newsletter, number 22, Interview of Christine Morin, http://emergences.inria.fr/2012/newsletter-n22/L22-SILICON.

[55] *Des grilles vers le cloud*, December 2012, Emergences: Lettre d'information 24, http://emergences.inria.fr/2012/newsletter_24/L24-CONTRAIL.

[56] *Myriads au coeur de la recherche sur le cloud*, October 2012, Meito mag.

[57] A. ESNAULT. *Energy-Aware Distributed Ant Colony Based Virtual Machine Consolidation in IaaS Clouds*, University of Rennes 1, June 2012, http://dumas.ccsd.cnrs.fr/dumas-00725215.

[58] S. KALATHINGAL, A. CARPEN-AMARIE, C. MORIN. *Dynamic resource provisioning in XtreemOS*, September 2012, Deliverable of the ANR Cloud project.

[59] C. MORIN, A. IORDACHE, N. PARLAVANTZAS. *Elastic Data-Intensive Cloud Computing*, December 2012, Deliverable D6 of the EIT ICT Labs RMAC activity.

[60] C. MORIN. *Final report for project ANR-11-EITS-201*, September 2012, Deliverable of the ANR Cloud project.

[61] M. PERRIN. *Algorithmes efficaces de recherche pour une machine chimique pair-à-pair*, University of Rennes 1, June 2012, http://dumas.ccsd.cnrs.fr/dumas-00725210.

References in notes

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

[63] G. AGHA. *Computing in pervasive cyberspace*, in "Commun. ACM", 2008, vol. 51, n^{o} 1, p. 68-70.

[64] 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, p. 159-170, http://dblp.uni-trier.de/db/conf/hpdc/hpdc2011.html#Al-KiswanySSR11.

[65] 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, p. 50–58, http://doi.acm.org/10.1145/1721654.1721672.

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

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

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

[69] 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.