IN PARTNERSHIP WITH:
Institut national des sciences appliquées de Rennes

Université Rennes 1

Activity Report 2011

Team MYRIADS

Design and Implementation of Autonomous Distributed Systems

RESEARCH CENTER
Rennes - Bretagne-Atlantique

THEME
Distributed Systems and Services

Table of contents

1. Members **1**
2. Overall Objectives **2**
2.1. General Objectives 2
2.2. Context 2
2.3. Challenges 3
2.4. Research Directions 3
2.4.1. Autonomous Management of Virtualized Infrastructures 3
2.4.1.1. SLA-driven PaaS over Cloud Federations 4
2.4.1.2. Virtual Data Centers 4
2.4.1.3. Virtualized Infrastructure Management 5
2.4.2. Multilevel Dynamic Adaptation of Service-based Applications 5
2.4.3. A Chemical Approach for Autonomous Service Computing 6
2.4.3.1. Chemical Expression of Interactions. 6
2.4.3.2. Distributed Implementation of the Chemical Paradigm. 7
2.5. Highlights 7
3. Scientific Foundations **7**
3.1. Introduction 7
3.2. Autonomic Computing 7
3.3. Future Internet and SOA 8
3.4. Distributed Operating Systems 8
3.5. Unconventional/Nature-inspired Programming 9
4. Application Domains **9**
5. Software **10**
5.1. SAFDIS 10
5.2. HOCL-tools 10
5.3. XtreemOS 10
5.4. Contrail Virtual Execution Platform (VEP) 11
5.5. Snooze 11
5.6. Resilin 12
5.7. Saline 12
6. New Results **13**
6.1. Introduction 13
6.2. Autonomous Management of Virtualized Infrastructures 13
6.2.1. Cloud Federations 13
6.2.1.1. Virtual Execution Platforms in Cloud Federations 13
6.2.1.2. Efficient virtual cluster migration 13
6.2.1.3. Elastic Map/Reduce over Cloud Federations 13
6.2.1.4. Sky Computing Experiments 14
6.2.2. Infrastructure as a Service Clouds 14
6.2.2.1. Large scale Energy aware self-healing IaaS 14
6.2.2.2. Resource Management in Private Clouds 14
6.2.2.3. Resilience 14
6.2.3. XtreemOS Grid Distributed Operating System 14
6.2.3.1. Facilitating Experiments with XtreemOS Grid System 14
6.2.3.2. Resource Management for Dynamic Applications 15
6.3. Dynamic Adaptation of Service-based Applications 15
6.3.1. Dynamic Adaptation in a Distributed Operating System 15
6.3.2. Adaptation for Data Management 15
6.3.3. Adaptation for Service-Oriented Architectures 15

6.3.3.1. Services adaptation in distributed and heterogeneous systems 16
6.3.3.2. Quality Assurance for Distributed Services 16
6.3.3.3. Self-configuration for Cloud Platforms 16
6.3.3.4. Dynamic Adaptation of Chemical services 16
6.3.3.5. Prediction of SLA violations and dynamic adaptation in workflows 16
6.4. A Chemical Approach for Autonomous Service Computing 17
6.4.1. Chemistry-Inspired Workflow Management System for e-Science Applications 17
6.4.2. Solving Workflow Patterns Through Molecular Composition 17
6.4.3. Scalable Atomic Capture of Molecules 17
7. Contracts and Grants with Industry **17**
8. Partnerships and Cooperations **18**
8.1. Regional Initiatives 18
8.1.1. ASYST (2010-2013) 18
8.1.2. HOCL4WS (2010-2012) 18
8.2. National Initiatives 18
8.2.1. AUTOCHEM ANR White Project (2007-2011) 18
8.2.2. ECO-GRAPPE ANR ARPEGE Project (2009-2012) 18
8.2.3. COOP ANR COSINUS Project (2009-2012) 18
8.2.4. CLOUD ANR project (October 2011 - September 2012) 19
8.2.5. HEMERA INRIA AEN (2010-2013) 19
8.2.6. INRIA ADT Aladdin (2008-2012) 19
8.2.7. INRIA ADT XtreemOS Easy (2010-2012) 19
8.2.8. INRIA ADT DAUM (2011-2012) 19
8.3. European Initiatives 20
8.3.1. FP7 Projects 20
8.3.1.1. S-CUBE (2008-2012) 20
8.3.1.2. CONTRAIL (2010-2013) 20
8.3.1.3. SCALUS Marie Curie Initial Training Networks (MCITN) (2009-2013) 21
8.3.1.4. BonFire IP Project (2010-2013) 21
8.3.1.5. IC0804 - Energy Efficiency in Large Scale Distributed Systems 22
8.3.2. Collaborations in European Programs, except FP7 22
8.4. International Initiatives 22
8.4.1. INRIA Associate Teams 22
8.4.1.1. SEROS 22
8.4.1.2. DataCloud@Work 23
8.4.2. INRIA International Partners 23
8.4.3. Visits of International Scientists 23
8.4.4. Participation In International Programs 24
9. Dissemination **24**
9.1. Animation of the scientific community 24
9.1.1. Leadership, Steering Committees and community service 24
9.1.2. Program Committees 24
9.1.3. Evaluation committees, consulting 25
9.1.4. Conferences, seminars, invitations 25
9.1.5. Administrative Responsibilities 26
9.1.6. Misc. 27
9.2. Teaching 27
10. Bibliography **28**

Team MYRIADS

Keywords: Autonomic Computing, Cloud Computing, Unconventional Paradigms, Self-Adaptive Systems, Virtualization, Grid'5000

1. Members

Research Scientists

Françoise André [Professor, on leave from UNIVERSITY RENNES 1, passed away in July 2011, HdR]
Yvon Jégou [Research Associate (CR), INRIA]
Christine Morin [Research Director (DR), INRIA, Team leader, HdR]

Faculty Members

Nikos Parlavantzas [Assistant Professor, INSA]
Jean-Louis Pazat [Professor, INSA, HdR]
Cédric Tedeschi [Assistant Professor, UNIVERSITY RENNES 1]

Technical Staff

Amine Belhaj [INRIA Associate Engineer (IA), ADT XtreemOS Easy]
Roberto-Gioacchino Cascella [INRIA Engineer, Contrail European project, since February 2011]
Ghislain Charrier [INRIA Associate Engineer (IA), ADT Aladdin]
Florian Dudouet [INRIA Engineer, Contrail European project, since October 2011]
Maxence Dunnewind [INRIA Engineer BonFire IP Project]
Rémy Garrigue [INRIA Associate Engineer (IA), ADT XtreemOS Easy]
Piyush Harsh [INRIA Engineer, Contrail European project]
David Margery [INRIA Research Engineer (IR, ADT Aladdin technical director)]
Pascal Morillon [UNIVERSITY RENNES 1 Engineer (IE, ADT Aladdin)]
Yann Radenac [INRIA Engineer, COOP ANR project]
Cyril Rohr [INRIA Engineer BonFire IP Project]

PhD Students

Stefania Costache [CIFRE EDF Grant]
Erwan Daubert [INRIA Grant]
Djawida Dib [INRIA and Brittany Regional Council Grant]
Eugen Feller [INRIA Grant]
Héctor Fernandez [INRIA Grant]
Jérôme Gallard [INRIA Grant until May 2011]
Guillaume Gauvrit [MENRT Grant]
André Lage [INRIA Grant]
Marko Obrovac [INRIA and Brittany Regional Council Grant]
Pierre Riteau [MENRT Grant]
Chen Wang [INRIA Grant]
Mohamed Zouari [INRIA and Brittany Regional Council Grant, until June 2011]

Post-Doctoral Fellow

Pierre Lemarinier [ATER until August 2011]

Visiting Scientist

Chokchai Leangsuksun [Louisiana Tech University, USA, Invited Professor at University of Rennes 1, June-July 2011]

Administrative Assistants

Maryse Fouché [Team administrative assistant (TR) INRIA]
Sandrine L'Hermitte [Contrail Scientific Coordinator Assistant (until September 2011), INRIA]
Katarzyna Kucharczyk [Contrail Scientific Coordinator Assistant (since June 2011), INRIA]

Others

Jean-Pierre Banâtre [Collaborator, part time, HdR]
Thierry Priol [Research Director (DR), INRIA, collaborator, part time, HdR]

2. Overall Objectives

2.1. General Objectives

MYRIADS is a joint team with INRIA, 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.

Scale, dynamic change and heterogeneity. 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.

Dependability. 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.

Efficient resource management. 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.

2.4. Research Directions

Our research activities are organized along three main work directions: autonomous management of virtualized infrastructures, dynamic adaptation of service-based applications and investigation of a chemical approach for autonomous service computing.

2.4.1. Autonomous Management of Virtualized Infrastructures

With virtualized infrastructures (clouds) computing and storage become a utility. Just like water or electricity, they are available in virtually infinite amounts and the consumption can be adapted within seconds like opening or closing a water tap. 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, map-reduce, 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 infrastructures). 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 cloudbursting 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.

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 as 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. 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 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 Internet, 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, 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 both 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 [54], [48]. 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 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 [60]. 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

- Pierre Riteau won one of the three Best Poster Awards in the PhD Forum of the IPDPS 2011 conference.

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 [60] 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 [49], [55] 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 [47] 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 [50] 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).

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 6.4), 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*.

- Version: 1
- ACM: Frameworks

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 are developed for individual PCs, clusters and mobile devices (PDA, notebooks and smartphones). XTREEMOS has been developed by the XTREEMOS consortium. The third public version of XTREEMOS has been released in December 2010.

XTREEMOS has been presented at the Contrail summer school, Presqu'île de Giens, France in July 2011 (Y. Jégou), SC'11, Seattle, USA (Y. Jégou, Ch. Morin) in November 2011. 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. In 2011, XtreemOS has been packaged for the OpenSuse Linux distribution. Ready-to-use XtreemOS virtual machine images have been made available for the community.

Active contributors (from Myriads team): Amine Belhaj, Jérôme Gallard, Rémy Garrigue, Yvon Jégou, Christine Morin, Yann Radenac, Pierre Riteau.

5.4. Contrail Virtual Execution Platform (VEP)

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

URL: http://www.contrail-project.eu

Status: Version 1.0

License: BSD

Presentation: Virtual Execution Platform (VEP) is a Contrail service that sits just above IaaS layer at the service provider end of the Contrail cloud federation. The VEP 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 team): Roberto-Gioacchino Cascella, Florian Dudouet, Piyush Harsh, Yvon Jégou, Christine Morin.

5.5. Snooze

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

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

Status: Version 1.0

License: TBD

Presentation: Snooze [56], is a *novel VM management framework for private clouds* which is designed to scale across thousands of nodes. Unlike the existing cloud management frameworks, Snooze utilizes a *self-organizing hierarchical architecture* and performs *distributed VM management*. Particularly, VM management tasks are performed by multiple managers, with each manager having only a *partial* view of the system. Moreover, fault-tolerance is provided at all levels of the hierarchy by

replication and integrated *leader election algorithm*. Consequently, the system is able to self-heal and continue its operation despite system component failures. In addition, VM monitoring is integrated into the framework and a generic scheduling engine exists to support advanced (e.g., consolidation) scheduling policies. Last but not least, a RESTful command line interface (CLI) exists to support virtual cluster (VC) definitions and management (i.e., start, shutdown, destroy, suspend, etc.) as well hierarchy visualization and exporting in GraphML format.

Snooze is fully implemented from scratch in Java and currently comprises approximately 15.000 lines of maintainable abstractions-based code. The leader election algorithm is built on top of the Apache ZooKeeper [59] highly available and reliable coordination system. In order to provide a uniform interface to the underlying hypervisors and support transparent VM monitoring and management, Snooze integrates the *libvirt* virtualization library.

The first Snooze prototype [41] has been developed and evaluated on Grid5000 experimental testbed.

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

5.6. Resilin

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

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

Status: Version 0.1

License: TBD

Presentation: Resilin [45] is a system to easily create execution platforms over distributed cloud resources for executing MapReduce computations. Resilin implements the Amazon Elastic MapReduce web service API with resources from other clouds than Amazon EC2, such as private and community clouds. Resilin allows users to perform MapReduce computations on other infrastructures than Amazon EC2, and offers more flexibility: users are free to select different types of virtual machines, different operating systems or newer Hadoop versions. Users only have to submit computations to the service through a web service API, and Resilin takes care of provisioning, configuring and managing cloud-based Hadoop execution platforms, potentially using multiple clouds.

Resilin is implemented in the Python language. It uses the `boto` library in order to interact with EC2-compatible clouds. Resilin has been evaluated on the Grid'5000 experimental testbed. Our comparison with the Amazon Elastic MapReduce service shows similar performance.

Active contributors (from Myriads team): Pierre Riteau, Christine Morin.

5.7. Saline

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

URL: https://www.grid5000.fr/mediawiki/index.php/VMdeploy

Status: Version V1.0 (experimental)

License: BSD

Presentation: `Saline` (former called `VMdeploy`) is a generic framework to deploy and manage encapsulated user jobs in virtual machines (VMs) at grid level by moving them from one site to another transparently for the encapsulated jobs [58]. Moreover, `Saline` is non-intrusive and can be used with any non-modified Grid resource management systems (RMSs).

`Saline` deploys and configures a set of VMs according to the user needs. Then, periodically, `Saline` takes snapshots of the running VMs and saves them on a dedicated node in an efficient way [57]. In addition, `Saline` checks the status of the running VMs. If something wrong happens *i.e.* one or more VMs have failed due for instance to a node failure or to the arrival of a higher priority job, `Saline` redeploys the set of VMs from the latest snapshot taken on new available resources provided by the

Grid RMS. Thanks to `Saline`, the redeployment of the snapshot is done in a transparent way from the encapsulated job point of view.

In its current implementation, `Saline` is programmed in `bash` and `C`. It uses `Libvirt` in order to create, to snapshot, and to restart the VMs. It means that `Saline` can deploy and manage KVM and XEN VMs or any other VMs usable with `Libvirt`. In addition, the architecture of `Saline` is very modular in order to have a clear and easily extensible code.

Active contributors (from Myriads team): Jérôme Gallard, Christine Morin

6. New Results

6.1. Introduction

Research results are presented according to the research directions of the MYRIADS team.

6.2. Autonomous Management of Virtualized Infrastructures

Participants: Amine Belhaj, Roberto-Gioacchino Cascella, Stefania Costache, Florian Dudouet, Eugen Feller, Jérôme Gallard, Rémy Garrigue, Piyush Harsh, Yvon Jégou, Chokchai Leangsuksun, Pierre Lemarinier, Christine Morin, Nikos Parlavantzas, Thierry Priol, Yann Radenac, Pierre Riteau.

6.2.1. Cloud Federations

Participants: Roberto-Gioacchino Cascella, Florian Dudouet, Piyush Harsh, Yvon Jégou, Christine Morin, Thierry Priol, Pierre Riteau.

6.2.1.1. Virtual Execution Platforms in Cloud Federations

In the context of the Contrail European project, we have defined the overall architecture of the Contrail software stack for cloud computing on top of cloud federations [51]. We have focused on the design and the implementation of a first basic prototype of the Virtual Excution Plaform (VEP) component [52]. It is in charge of provisioning hardware resources from Cloud providers and to deploy and run distributed applications submitted by users under the control of a negotiated Service Level Agreement (SLA) [16]. Within VEP software, REST interface, OVF parsing, SSL security, Authorization modules are under active development and at various levels of integration. A first demo version of VEP running on top of OpenNebula IaaS cloud has been successfully demonstrated at the first annual project review.

6.2.1.2. Efficient virtual cluster migration

We continued our work on Shrinker, a system providing efficient live migration of virtual clusters on wide area networks. The design has been improved to coordinate the deduplication on the source site of the migration. Deduplication is now performed only within an individual virtual cluster, in order to reduce security issues and avoid performance impact on virtual machines of other users. We performed a comprehensive performance evaluation of Shrinker. An article presenting the design, implementation, and performance of Shrinker was published in [28].

6.2.1.3. Elastic Map/Reduce over Cloud Federations

We worked on the development of Resilin, a system to easily create execution platforms over distributed cloud resources for executing MapReduce computations. Resilin implements the Amazon Elastic MapReduce web service API and uses resources from private and community clouds. Resilin takes care of provisioning, configuring and managing cloud-based Hadoop execution platforms, potentially using multiple clouds. An initial implementation of Resilin was presented at the CCA '11 workshop [36]. Further development was performed in the context of Ancuta Iordache's master internship. The results of this work were published as a research report [45].

6.2.1.4. Sky Computing Experiments

We continued our collaboration with the University of Florida on sky computing experiments, which led to the publication of a book chapter [38].

6.2.2. Infrastructure as a Service Clouds

Participants: Stefania Costache, Eugen Feller, Yvon Jégou, Christine Morin, Nikos Parlavantzas, Pierre Riteau.

6.2.2.1. Large scale Energy aware self-healing IaaS

The research done in 2011 was two fold. A prototype of the previously proposed scalable, fault-tolerant and energy-aware virtual machine (VM) management framework called Snooze was implemented and evaluated on the Grid5000 testbed [41]. In 2011, we have focused on the implementation of the self-healing mechanisms and protocols, and on integrating in Snooze the system-level mechanisms (e.g. for automatically switching on/off cluster nodes) to support energy aware resource management algorithms. Our experimental results show that the fault-tolerance features of the framework do not impact application performance. Moreover, negligible cost is involved in performing distributed VM management and the system remains highly scalable with increasing amounts of resources. A nature-inspired VM placement algorithm [24] based on the Ant Colony Optimization (ACO) meta-heuristic was developed and evaluated by means of simulations.

6.2.2.2. Resource Management in Private Clouds

We focused on the design of a resource management system for private clouds that provides support for different application SLAs while maximizing the resource utilization of the infrastructure. As we also considered the need of providing users the incentives to truthfully express their valuation for the performance of their application we investigated existing economic models of allocating resources. As a result, we proposed a novel resource management architecture based on a virtual economy. In this system, independent agents monitor the application's performance and dynamically provision virtual machines from the infrastructure under user's budget constraints. To provision virtual machines, a proportional share auction is used, allowing a fine-grain resource sharing at a low complexity cost. This work was done as part of a collaboration with EDF R&D and was published at the VHPC 2011 workshop [22]. We have also implemented a first prototype of this proposal. In collaboration with Vydia Rajagopalan (Master student at VU Amsterdam) we have implemented the proportional-share auction scheduler and integrated it with the OpenNebula Virtual Infrastructure Manager. Then, we have extended this work with the design of agents that execute scientific applications (MPI and Bag-of-Task applications) under deadline and budget constraints. Experimental evaluations are currently performed on Grid5000.

6.2.2.3. Resilience

We initiated a collaboration with Box Leangsuksun's group on high availability of cloud infrastructures. We carried out a preliminary study on how the HA-OSCAR environment developed at the Louisiana Tech University could be used to ensure the high availability of critical services in Nimbus IaaS clouds [30].

6.2.3. XtreemOS Grid Distributed Operating System

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

6.2.3.1. Facilitating Experiments with XtreemOS Grid System

XtreemOS Grid system that has been developed in the framework of the XtreemOS European project is now evolving as an open source software in a community driven by INRIA in the framework of the XtreemOS Easy ADT. We have provided first level support to users and maintained XtreemOS website, wiki and mailing-lists. We have updated XtreemOS documentation to reflect the evolution of XtreemOS system. We facilitated the access to XtreemOS in three different ways: maintaining an open public testbed runnning XtreemOS, providing ready-to-use XtreemOS virtual machines and developing tools to automatically deploy XtreemOS Grid system on the Grid'5000 large-scale experimentation platform. In 2011, we have finalized a new 3.0 version of XtreemOS system and ported it on top of the OpenSuse 11.4 Linux distribution. We performed a number of tests to validate the installation, configuration and execution of the new XtreemOS version based on

openSuSE Linux distribution. An incremental integration process has been set up to facilitate the integration of patches and bug fixes. We have run a number of experiments with XtreemOS 3.0 based on Mandriva Linux distribution: MPI programs, Salomé numerical simulation platform, bio-informatics applications. Yann Radenac, in the framework of the COOP project funded by ANR contributed to XtreemOS's code by fixing bugs, cleaning the source code to improve maintainability, and adding a few minor features.

6.2.3.2. Resource Management for Dynamic Applications

In the framework of the COOP project funded by ANR, we compared the features offered by the CooRM resource manager for dynamic applications developed by Christian Perez and Cristian Klein at ENS Lyon with those provided by the XtreemOS Grid system. A plan has been set to adapt CooRM to XtreemOS system and to extend XtreemOS's API to include a CooRM-like interface [53]. We worked on the definition of a variant of CooRM that can work in the context of XtreemOS Grid operating system.

6.3. Dynamic Adaptation of Service-based Applications

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

6.3.1. Dynamic Adaptation in a Distributed Operating System

Participants: Françoise André, Djawida Dib, Christine Morin, Nikos Parlavantzas.

We have studied the feasibility to dynamic adapt the features of a distributed operating system using a framework for self-adaptation of service oriented distributed applications [46]. We have focused on the consistency protocols for replicated data in distributed shared memory systems. We have considered two strict consistency protocols, one based on invalidation and one based on broadcast on write operations. The adaptation framework selects one of these two algorithms based on the inter-node data transfer delay. We have implemented a prototype based on Kerrighed single system image operating system for clusters and the SAFDIS adaptation framework. We have integrated a broadcast based consistency protocol in Kerrighed that already implements a write invalidation consistency protocol. We have implemented the adaptation policy in the SAFDIS framework and the needed adaptation mechanisms in Kerrighed as well as a component for monitoring data transmission delays. An experimental evaluation is being carried out.

6.3.2. Adaptation for Data Management

Participants: Françoise André, Mohamed Zouari.

The usage of context-aware data management in mobile environments has been previously investigated by Françoise André in collaboration with Mayté Segarra and Jean-Marie Gilliot from Telecom Bretagne Brest (previously known as ENST Bretagne). This work focuses on data management in grid and mobile environments; an ambient assisted living application illustrates the approach. This work was realized in the context of the *ALORAD* project (Architecture LOgicie lle pour la Réplication Adaptative de Données), financed by the Brittany council. Mayté Segarra from Telecom Bretagne Brest was co-adviser for the PhD thesis of M. Zouari [12].

6.3.3. Adaptation for Service-Oriented Architectures

Participants: Françoise André, 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.

6.3.3.1. Services adaptation in distributed and heterogeneous systems

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.

6.3.3.2. Quality Assurance for Distributed Services

In the context of the service-centric paradigm, we have designed and developed the Qu4DS (Quality Assurance for Distributed Services) research prototype. Qu4DS is a cloud PaaS solution which fills the gap between the conception of higher-level SaaS service providers over the resource-level PaaS layer. Qu4DS provides an automatic support for service execution management by aiming at increasing the service provider's profit. More specific, Qu4DS dynamically acquires resources according to the customer demand, deploys service instances and implements QoS assurance mechanisms in order to prevent SLA violations. Moreover, Qu4DS has been evaluated on Grid5000 and showed to be effective on reducing service provider's costs [27].

6.3.3.3. Self-configuration for Cloud Platforms

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.

6.3.3.4. Dynamic Adaptation of Chemical services

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.

6.3.3.5. Prediction of SLA violations and dynamic adaptation in workflows

During execution, run-time QoS is determined by the dynamic execution environment and thus the expected QoS is not always ensured. In addition, infrastructure failures can make a service undeliverable. The adaptive execution reflects the capability to recompose a (part of a) workflow on-the-fly in case that global SLA violation is predicted. Most techniques for predicting global SLA violation require past experiences on executions of a business process. All historical execution instances have the same structure as well as the same bindings. These solutions do not adapt to the case of dynamic execution, where for each execution, partner services are selected and bound at run time.

In order to predict global SLA violation in the context of dynamic service execution, we proposed a 2-phase prediction technique, which is fit for generic workflow composition. The prediction method works with a high accuracy for simple workflows, but when the workflow composes complicated structures (such as loops and exclusive branches), the performance degrades. The reason is that the estimation of global SLA is based on the critical path, which is not definitely executed. To solve this problem, we propose to use data mining technique to predict workflow branches and the number of loop execution. Based on predicted branches, the prediction of global SLA violation is much more accurate. The numerical evaluation will be carried out in the near future.

6.4. A Chemical Approach for Autonomous Service Computing

Participants: Héctor Fernandez, Marko Obrovac, Thierry Priol, Cédric Tedeschi.

6.4.1. Chemistry-Inspired Workflow Management System for e-Science Applications

Participants: Héctor Fernandez, Cédric Tedeschi, Thierry Priol.

In the research track that aims at leveraging the properties of the chemical Programming models for autonomic computing, we have built a software based on the HOCL compiler (part of the HOCL-tools) that was actually deployed and experimented over the Grid'5000 platform. The experiments have shown, that envisioning the execution workflow as an autonomic chemical process is actually viable in practice. Experimented with different well known workflow-based e-Science applications, the software showed a performance level comparable to current top-rated scientific workflow management systems [25].

6.4.2. Solving Workflow Patterns Through Molecular Composition

Participants: Héctor Fernandez, Cédric Tedeschi, Thierry Priol.

In the same area, but on a more conceptual point of view, we have shown how the expressive power of the chemical model can be leveraged to solve complex workflow patterns. This aspect was also integrated into the HOCL-tools and experimented over the Grid'5000 platform, following two architectures with a different level of decentralization, showing the advantages and drawbacks of decentralizing the workflow execution using a chemical workflow engine [26].

6.4.3. 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. We proposed a protocol solving this issue on top of a distributed hash table (DHT). The DHT secures the scalability of the communications and provides a scalable discovery of reactants. Our protocol is triggered once reactants are found. It is made of two sub-protocols being used at different stages of the computation, according to the density of possible reactions. The protocol is validated through its proof of liveness and simulations showing that it is able to switch from one sub-protocol to the other efficiently, according to the execution's conditions [18].

7. Contracts and Grants with Industry

7.1. Contracts with Industry

7.1.1. EDF R&D

Participants: Christine Morin, Nikos Parlavantzas.

In the context of the INRIA - EDF R&D general agreement, we have a contract (2010 - 2013) with the EDF R&D SINETICS department to investigate resource management in virtualized distributed platforms in order to efficiently execute distributed applications with stringent time constraints for providing their results. The goal is to federate computing resources and use them simultaneously for research and exploitation purposes, the latter class of applications being executed with a higher priority than the former. This project funds a CIFRE PhD grant (Stefania Costache).

8. Partnerships and Cooperations

8.1. Regional Initiatives

8.1.1. ASYST (2010-2013)

Participants: Françoise André, 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).

8.1.2. HOCL4WS (2010-2012)

Participants: Thierry Priol, Cédric Tedeschi.

The objective of the HOCL project funded by the Brittany council is to develop a prototype of a middleware system for the distributed execution of chemical programs (targeted for large scale platforms). This project funds 50% of a PhD grant (Marko Obrovac).

8.2. National Initiatives

8.2.1. AUTOCHEM ANR White Project (2007-2011)

Participants: Thierry Priol, Cédric Tedeschi.

The goal of the AUTOCHEM project funded under the ANR white program is the programming of desktop Grids using the chemical programming model (http://www.irisa.fr/myriads/collaborate/national/anr/autochem/). This project funds a PhD grant (Héctor Fernandez).

8.2.2. ECO-GRAPPE ANR ARPEGE Project (2009-2012)

Participant: Christine Morin.

The goal of the ECO-GRAPPE project (http://www.irisa.fr/myriads/collaborate/national/anr/autochem/) 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.

8.2.3. COOP ANR COSINUS Project (2009-2012)

Participants: Christine Morin, Yvon Jégou.

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).

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

Participant: 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 described in Section 8.3.2. It funds a research engineer.

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.

8.2.6. INRIA ADT Aladdin (2008-2012)

Participant: Yvon Jégou.

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: Christine Morin, Yvon Jégou.

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.

8.2.8. INRIA ADT DAUM (2011-2012)

Participants: Erwan Daubert, Jean-Louis Pazat.

We participate to the ADT DAUM which is coordinated by the Triskell 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.

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

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.

8.3. European Initiatives

8.3.1. FP7 Projects

8.3.1.1. S-CUBE (2008-2012)

Françoise André, Jean-Louis Pazat

Title: S-Cube: Software Services and Systems Network

Type: COOPERATION (ICT)

Defi: Service & software architectures, infrastructures and engineering

Instrument: Network of Excellence (NoE)

Duration: March 2008 - February 2012

Coordinator: Universität Duisburg-Essen (Germany)

Others partners: University of Duisburg-Essen, Paluno - Ruhr Institute for Software Technology (Germany); Tilburg University (Netherlands); City University London (U.K.); Consiglio Nazionale delle Ricerche (Italy); Center for Scientific and Technological Research, FBK (Italy); Lero - The Irish Software Engineering Research Centre (Ireland); Politecnico di Milano (Italy); MTA SZTAKI – Computer and Automation Research Institute (Hungary); Vienna University of Technology (Austria); Université Claude Bernard Lyon (France); University of Crete (Greece); Universidad Politécnica de Madrid (Spain); University of Stuttgart (Germany); University of Hamburg (Germany); Vrije University Amsterdam (Netherlands)

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

Abstract:S-Cube is the European network of excellence in software services and systems federating the software engineering and distributed system research communities to shape the Internet of Services (http://www.s-cube-network.eu/). The MYRIADS team is involved in service discovery, coordination and adaptation. Three PhD thesis grants are funded by the S-Cube project.

8.3.1.2. CONTRAIL (2010-2013)

Yvon Jégou, Christine Morin

Title: Contrail, Open Computing Infrastructures for Elastic Services

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://www.contrail-project.eu

Abstract: The goal of the Contrail project is to design, implement, evaluate and promote an open source system for Cloud Federations. To allow open access to shared computing resources, the vision of the Contrail project is that any organization should be able to be both a Cloud provider when its IT infrastructure is not used at its maximal capacity, and a Cloud customer in periods of peak activity. Resources that belong to different operators will be integrated into a single homogeneous Federated Cloud that users can access seamlessly.

Contrail will vertically integrate an open-source distributed operating system for autonomous resource management in Infrastructure-as-a-Service environments, and high level services and runtime environments as foundations for Platform-as-a-Service. Contrail will address key technological challenges in existing commercial and academic Clouds: the lack of standardized rich and stable interfaces; limited trust from customers; and relatively poor Quality of Service (QoS) guarantees regarding the performance and availability of Cloud resources.

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

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.

8.3.1.4. BonFire IP Project (2010-2013)

David Margery

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.

8.3.1.5. IC0804 - Energy Efficiency in Large Scale Distributed Systems

Françoise André, Jean-Louis Pazat

This COST Action will propose realistic energy-efficient alternate solutions to share IT distributed resources (http://www.irit.fr/cost804/). 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 will characterize the energy consumption and energy efficiencies of distributed applications. In his PhD thesis work, Eugen Feller is contributing to this COST action. He is one of the two co-proposers of a working group on energy and QoS-aware workload management in clouds.

8.3.2. Collaborations in European Programs, except FP7

Program: EIT ICT Labs

Project acronym: Activity 10239

Project title: Boosting Innovative Software in Cloud and Service Computing

Duration: January - December 2011

Coordinator: Christine Morin

Other partners: Institut Telecom Sud Paris (France), VTT (Sweden), SAP (Germany)

Abstract: This activity aims at building testbeds integrating XtreemOS and Contrail open source software and at experimenting these testbeds on large scale experimentation platforms such as Grid'5000 and BonFire. We organized an EIT ICT Labs summer school (jointly with Contrail European project) held June 27-July 1, 2011 in Presqu'île de Giens, France.

8.4. International Initiatives

8.4.1. INRIA Associate Teams

8.4.1.1. SEROS

Title: Scalable, Efficient, and Resilient Operating Systems

INRIA principal investigator: Christine Morin

International Partner: Stephen Scott

Institution: Oak Ridge National Laboratory (United States)

Laboratory: System Research Team

Duration: 2009 - 2011

See also: http://www.irisa.fr/myriads/ser-os/

The main objectives of the collaboration in the area of operating systems and system tools for HPC are: operating systems for HPC (focusing on system-level virtualization), system management tools for HPC platforms, and resilience for HPC systems.

8.4.1.2. DataCloud@Work

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. INRIA International Partners

Nikos Parlavantzas is collaborating with the team of Prof. Héctor Duran-Limon of the University of Guadalajara, Mexico on resource management in virtualised environments. We have produced a joint publication and are preparing an ANR project proposition[23].

Héctor Fernandez and Cédric Tedeschi collaborate with the team of Dr. Patricia Lago from Vrije University Amsterdam. Héctor spent 3 weeks there, working on the use of the Chemical Programming Model to implement a proof of concept of a method for Global Software Development developed in VUA's team, financed by the S-Cube project.

Christine Morin has visited Professor Vladimir Getov, Westminster University, UK in February 2011. A PhD thesis topic in the area of cloud computing has been jointly proposed. Professor Vladimir Getov visited Myriads team in May 2011.

8.4.3. Visits of International Scientists

Chokchai Leangsuksun, Professor at Louisina Tech University (USA) spent one month and a half in the Myriads team in June and July 2011, as an invited professor, funded by the University of Rennes 1. We investigated potential joint research directions related to resilience in cloud computing [30].

8.4.3.1. Internships

Neha Jatav

Subject: Planning software deployment on heterogeneous and distributed infrastructures

Institution: IIT Bombay (India)

Vidya Lakshmi Rajagopalan

Subject: Dynamic vertical scaling of parallel applications in clouds

Institution: Vrije Universiteit, Amsterdam (Netherlands)

Ancuta Iordache

Subject: Elastic MapReduce in cloud federations

Institution: West Timisoara University (Romania)

8.4.4. Participation In International Programs

Christine Morin was involved in the Master school education action line to participate in the definition of the Distributed Systems and Services (DSS) major of the EIT ICT Labs Master in ICT innovation.

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). During her sabbatical visit at the Lawrence Berkeley National Laboratory, Christine Morin is in charge of the INRIA@SiliconValley program jointly with the INRIA international affairs department.

9. Dissemination

9.1. Animation of the scientific community

9.1.1. Leadership, Steering Committees and community service

Ch. Morin served as the scientific chair of the first Contrail summer school held in Presqu'île de Giens, France from June 27 to July 1st, 2011 (http://contrail-project.eu/summerschool-2011).

N. Parlavantzas was a member of the Organizing Committee of RenPar'20 conference.

J.-L. Pazat served as the Chair of the Organizing Committee of the RenPar, CFSE and Sympa federated conference series. He is the chairman of the Steering Committee of RenPar http://www.renpar.org/).

Th. 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 is program vice-chair for ICPP 2012, in charge of the topic "Programming Models, Languages and Environments". He is is the Chairman of the ServiceWave Steering Committee.

C. Tedeschi coordinated the organization of the RenPar'20/SympA'14/CFSE 8 joint French-speaking conferences held in Saint-Malo, France, in May 2011.

9.1.2. Program Committees

Ch. Morin served in the Program Committees of the following conferences: IPDPS 2011, Cluster 2011, SC 2011 (technical papers and posters), VTDC 2011, HPCVirt 2011, Cloud for High Performance Computing Workshop at ICCSA 2011, ISPDC 2011, ServiceWave 2011, CFSE 2011, ISPA 2012, CC-GRID 2012, IPDPS 2012, EuroPar 2012, VTDC 2012.

N. Parlavantzas served in the Program Committees of RTSOAA 2011 and INFOCOMP 2011 international conferences.

Th. Priol served in the Program Committees of the following conferences: Cloud 2011, CloudCom 2011, ICWS 2011, ICPP 2012, CCGRID 2012, VECPAR 2012.

C. Tedeschi is a member of the Program Committees of ICCS 2011, ICCS 2012, AINA 2012 and Closer 2012 international conferences, and was the PC chair of RenPar'20 french-speaking conference.

9.1.3. Evaluation committees, consulting

J.-P. Banâtre is the INRIA representative in the Contrail General Assembly.

Ch. Morin has served as a member for a selection committee of Lille 1 University and of Ensimag, Grenoble to appoint an assistant Professor in Computer Science. She acted as a referee for the HDR Committee of Jean-Marc Menaud (Ecole des Mines de Nantes), Eric Renault (University of Paris 6) and was a member of the HDR Committee of Laurent Réveillère (University of Bordeaux 1) and as a member for the PhD Committee of Anne-Cécile Orgerie (ENS Lyon). She was a member of the selection committee for the PES (scientific excellence award) for junior and senior researchers at INRIA. She was a member of the IRISA selection committee for the PhD grants. She acted as an expert to review proposals for the French Research Agency (ANR). She acted as expert to review a proposal for a CIFRE PhD grant for the ANRT. She is a member of the cloud computing expert group in the Software & Service Architectures and Infrastructures unit of the European Commission.

Th. Priol acts as an expert to review FP7 proposals and projects for the European Commission. He is member of the Scientific Advisory Board for the Contrail and VENUS-C FP7 projects. He was also member of the AERES panel of experts to evaluate Institut Télécom.

9.1.4. Conferences, seminars, invitations

R. Cascella gave a keynote presentation entitled "CONTRAIL: Bringing trust in Clouds" at the 1st International Workshop on Cloud Computing Platforms (CloudCP) co-located with Eurosys 2011, Salzburg, Austria, 10 Apr. 2011. He gave a talk entitled "Enabling a collaborative environment via the Contrail project", at the Aristote Seminar, Paris, 11 Oct. 2011. He also gave talks on Contrail at the Internet of Service Collaboration Meeting, Brussels, 28 Sep. 2011, at the OW2con Session on Collaborative Projects, Paris Issy les Moulineaux, 24 Nov. 2011, and at the Open World Forum, Session Open Cloud in Europe: When Interoperability Matters, Paris, 22 Sep. 2011.

S. Costache gave a talk entitled "Virtualization Techniques for Optimal Execution of Applications under Time Constraints", at the PHD Student Day at EDF R&D, Paris, France, November 2011. She gave a seminar entitled "An Economic Approach for Application QoS Management in Clouds" at ORNL and Tennessee Technical University in the framework of the SER-OS associated team, December 2011.

D. Dib gave a talk entitled "Adaptation dynamique des fonctionnalités d'un système d'exploitation large échelle" at the Adapt workshop, Saint-Malo, May 2010.

E. Feller gave a talk entitled "Focus group: Energy and QoS-aware Workload Management in Clouds" at a COST IC0804 workshop (http://www.cost804.org/), Budapest, Hungary, May 2011. He gave a talk entitled "Energy-Aware Ant Colony Based Workload Placement" at a COST IC0804 workshop (http://www.cost804.org/), Thessaloniki, Greece, November 2011. He gave a talk entitled "Towards Agent-Based Energy-Aware Workload Consolidation in Clouds" at the GreenDays meeting (http://perso.ens-lyon.fr/laurent.lefevre/greendaysparis/), Paris, France, June 2011. He gave an overview of his research directions at the CominLabs meeting (http://www.cominlabs.ueb.eu/), Rennes, France, October 2011.

H. Fernandez gave a talk on "Chemically-inspired Workflow Management Systems", at the Vrije University Amsterdam (VUA) on November 17, 2011 during his visit at VUA in the context of the S-Cube Project.

Y. Jégou gave a talk entitled “Open Computing Infrastructures for Elastic Services: Contrail Approach” at the 2nd WoSS Workshop in Timisoara, Romania in June 2011. He gave a lecture on XTREEMOS at the first Contrail summer school on Grid and Cloud computing held in Presqu’île de Giens, France in June 2011. He presented the exploitation of the OVF application format in Contrail at the Open World Forum meeting, Paris in September 2011. He gave a talk entitled “Contrail Cloud Federations” at the OCCI working-group meeting, Lyon in September 2011. He presented XTREEMOS system and Contrail project at SC’11, Seattle, USA (November 2011).

A. Lage presented a talk entitled "A QoS Assurance Framework for Distributed Infrastructures", REDIS research group, Federal University of Pernambuco. Recife, Brazil. May, 2011.

Ch. Morin gave a talk entitled "From Grids to Federated Clouds: Virtualization of Large-scale Distributed Computing Infrastructures" at a workshop organized for the third anniversary of the INRIA - Alcatel Lucent Bell Labs joint laboratory, Rocquencourt, France, in January 2011. She presented a talk entitled "*Le centre INRIA Rennes : un acteur européen dans le domaine des grilles et des clouds*" at the AERES evaluation of the INRIA Rennes Bretagne Atlantique research center in February 2011. She was invited to give a seminar entitled "Building Large Scale Dynamic Computing Infrastructures over Clouds" at UC Berkeley and at the Lawrence Berkeley National Laboratory, Berkeley, USA in May 2011. She gave an invited talk entitled "Open Computing Infrastructures for Elastic Services: Contrail Approach" at the 5th international workshop on Virtualization Technologies in Distributed Computing (VTDC), co-located with HPDC 2011, San Jose, USA, June 2011. She gave a talk entitled "Introduction to Contrail school & Contrail FP7 project" at the first Contrail summer school on Grid and Cloud computing held in Presqu’île de Giens, France in June 2011. She gave a talk entitled " Open Computing Infrastructures for Elastic Services: The Contrail European Project Approach" at the Advanced Computing for Science department seminar at Lawrence Berkeley National Laboratory, Berkeley, USA, in November 2011. She presented Contrail project at SC’11, Seattle, USA (November 2011).

Yann Radenac gave a talk entitled "XtreemOS and COOP: concepts, use cases and limits" at the COOP bi-annual project meeting, Rennes, France, April 2011.

P. Riteau gave a talk entitled "Sky Computing on FutureGrid and Grid’5000 with Nimbus" at Journée Cloud France Grilles, ENS Lyon (Lyon, France), December 13, 2010. He gave a talk entitled "Large Scale Sky Computing Applications with Nimbus" at the workshop on High Performance Applications of Cloud and Grid Tools (HPACG) at Argonne National Laboratory (Illinois, USA), April 14, 2011. He gave a talk entitled "Deploying and Using IaaS Clouds on Grid’5000" (with Alexandra Carpen-Amarie and Houssem-Eddine Chihoub), in the Practical Session during the Grid’5000 Spring School 2011, Reims, France April 2011. He gave a talk entitled "Creating Customized Grid’5000 Environments with Chef", in the Practical Session during the Grid’5000 Spring School 2011, Reims, France April 2011. He gave a presentation entitled "*Anatomie d’une expérience dimensionnante de Sky Computing*" at the Journée tutoriel Grid’5000 (RenPar 2011) , Saint-Malo, France, May 2011. He gave a talk entitled "An Overview of Virtualization Technologies", at the Contrail Summer school 2011, "Presqu’île de Giens", Hyères-les-Palmiers (France), June 2011. He gave a presentation entitled "Virtual Machine Migration", at the Atelier France Grilles - Opérations, IPNL (Lyon, France), October 2011.

9.1.5. Administrative Responsibilities

C. Morin is a member of the Project-Team Committee of INRIA - Rennes Bretagne Atlantique (Comité des projets).

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.

T. Priol is Scientific Deputy Director at INRIA, responsible for monitoring all research, development, transfer and partnership activities in the thematic areas of Networks and Telecommunication, Distributed and High Performance Computing, Distributed Systems and Services. He was the Research Director of EIT ICT Labs till October 2011.

P. Riteau was elected member of the INRIA Rennes - Bretagne Atlantique Center Council (Comité de Centre).

C. Tedeschi was elected as a member of the administration council of the EECS departement of the University of Rennes 1.

9.1.6. Misc.

Piyush Harsh animates the *Interoperability Working Group* created on the OpenNebula ecosystem by INRIA and UCM.

Christine Morin participated in the "*A la découverte de la recherche*" action presenting her researcher profession to a class at lycée Jean Macé (high school) in Rennes, France (March 2011).

9.2. Teaching

Christine Morin is responsible for a graduate teaching Module *Distributed Systems: from networks to Grids* of the Master Program in Computer Science, UNIVERSITY RENNES 1.

Christine Morin:

Master : Grid Computing (as part of the Distributed Systems: from networks to Grids teachning unit), 6.5 ETD, M2RI, University of Rennes 1, France

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

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

Master : XtreemOS in the Clouds, 2 ETD, M1, University of Westminster, UK.

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)

Cédric Tedeschi (at the University of Rennes 1):

- Licence 3: Organization of Operating System (77 hours ETD)
- Licence 3: Algorithmic methods (22 hours ETD)
- Master 1: Multitask Operating Systems (65 hours ETD)
- Master 2: Muti-users systems (50 hours ETD)
- Master 2: Internet of Services (6 hours ETD)

PhD & HdR :

PhD : Jérôme Gallard, Flexible Management of Distributed Computing Infrastructures, University of Rennes 1, May 6, 2011, Christine Morin

PhD : Mohamed Zouari, Software architecture for dynamic adaptation; application to data replication, University of Rennes 1, June 28, 2011, Françoise André, Mayte Segarra-Montesinos

PhD : Pierre Riteau, Dynamic Execution Platforms over Federated Clouds, University of Rennes 1, December 2, 2011, Thierry Priol, Christine Morin

PhD in progress : Eugen Feller, Energy aware resource management in IaaS clouds, December 2009, Christine Morin

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

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

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

PhD in progress : André Lage, Dynamic Adaptation of Services on Distributed Infrastructures, October 2008, Nikos Parlavantzas, 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 : Héctor Fernández, Chemistry-inspired Workflow Management for Autonomic Service Composition, April 2009, Thierry Priol, Cédric Tedeschi

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

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] A. ARENAS, J.-P. BANÂTRE, T. PRIOL. *Developing Autonomic and Secure Virtual Organisations with Chemical Programming*, in "The 11th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS 2009)", France Lyon, A. K. DATTA, R. GUERRAOUI, F. PETIT (editors), Springer, 2009, vol. 5873, http://hal.inria.fr/inria-00409502/en/.

[3] J.-P. BANÂTRE, D. LE MÉTAYER. *Programming by Multiset Transformation*, in "Communications of the ACM", January 1993, vol. 36, n^{o} 1, p. 98–111.

[4] 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.

[5] 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.

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

[7] 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.

[8] 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.

[9] XTREEMOS CONSORTIUM. *Annex 1 - Description of Work*, April 2006, Contract funded by the European Commission, XtreemOS Integrated Project, IST-033576.

Publications of the year

Doctoral Dissertations and Habilitation Theses

[10] J. GALLARD. *Flexibilité dans la gestion des infrastructures informatiques distribuées*, Université Rennes 1, May 2011, http://hal.inria.fr/tel-00625278/en.

[11] 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.

[12] M. ZOUARI. *Architecture logicielle pour l'adaptation distribuée : Application à la réplication de données*, Université Rennes 1, June 2011, http://hal.inria.fr/tel-00652046/en.

Articles in International Peer-Reviewed Journal

[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/en.

[14] J. GALLARD, A. LÈBRE, C. MORIN, T. NAUGHTON, S. SCOTT, G. VALLÉE. *Architecture for the Next Generation System Management Tools*, in "Future Generation Computer Systems", 2011 [*DOI :* 10.1016/J.FUTURE.2011.06.003], http://hal.inria.fr/inria-00605143/en.

[15] J.-L. PAZAT, T. PRIOL, C. TEDESCHI. *Towards a Chemistry-Inspired Middleware to Program the Internet of Services*, in "ERCIM News", April 2011, vol. 2011, n^{o} 85, p. 34-35, http://hal.inria.fr/inria-00600706/en.

Invited Conferences

[16] P. HARSH, J. YVON, R. CASCELLA, C. MORIN. *Contrail Virtual Execution Platform - Challenges in Being Part of a Cloud Federation*, in "4th European Conference Servicewave 2011", Poznan, Poland, October 2011, http://hal.inria.fr/inria-00626639/en.

International Conferences with Proceedings

[17] F. ANDRÉ, E. DAUBERT, G. GAUVRIT. *Distribution and Self-Adaptation of a Framework for Dynamic Adaptation of Services*, in "The Sixth International Conference on Internet and Web Applications and Services (ICIW)", St. Maarten, Netherlands Antilles, IARIA, March 2011, p. 16-21, http://hal.inria.fr/inria-00553539/en.

[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, January 2012, http://hal.inria.fr/hal-00644262/en.

[19] M. BERTIER, M. OBROVAC, C. TEDESCHI. *A Protocol for the Atomic Capture of Multiple Molecules on Large Scale Platforms*, in "13th International Conference ON Distributed Computing AND Networking (ICDCN 2012)", Hong Kong, China, January 2012.

[20] L. BHOOBALAN, P. HARSH. *An Experimental Study and Analysis of Crowds based Anonymity*, in "The 2011 International Conference on Internet Computing", Las Vegas, United States, July 2011, http://hal.inria.fr/inria-00626653/en.

[21] M. CAEIRO RODRIGUEZ, Z. NÉMETH, T. PRIOL. *A chemical model for dynamic workflow coordination*, in "The 19th Euromicro International Conference on Parallel, Distributed and Network-Based Computing", Ayia Napa, Cyprus, Y. COTRONIS, G. A. PAPADOPOULOS, M. DANELUTTO (editors), Conference Publishing Services, February 2011, http://hal.inria.fr/inria-00534561/en.

[22] S. COSTACHE, N. PARLAVANTZAS, C. MORIN, S. KORTAS. *An Economic Approach for Application QoS Management in Clouds*, in "6th Workshop on Virtualization in High-Performance Cloud Computing (VHPC)", Bordeaux, France, August 2011, http://hal.inria.fr/inria-00625983/en.

[23] H. DURAN-LIMON, N. PARLAVANTZAS, M. ZHAO, L. SILVA-BAÑUELOS, V. TELLEZ-VALDEZ. *Using Lightweight Virtual Machines to Run High Performance Computing Applications: The Case of the Weather Research and Forecasting Model*, in "4th IEEE/ACM International Conference on Utility and Cloud Computing (UCC 2011)", Melbourne, Australia, December 2011, http://hal.inria.fr/hal-00625910/en.

[24] E. FELLER, L. RILLING, C. MORIN. *Energy-Aware Ant Colony Based Workload Placement in Clouds*, in "The 12th IEEE/ACM International Conference on Grid Computing (GRID-2011)", Lyon, France, September 2011, http://hal.inria.fr/inria-00626042/en.

[25] H. FERNANDEZ, C. TEDESCHI, T. PRIOL. *A Chemistry-Inspired Workflow Management System for Scientific Applications in Clouds*, in "7th IEEE International Conference on e-Science", Stockholm, Sweden, December 2011, http://hal.inria.fr/inria-00625240/en.

[26] H. FERNANDEZ, C. TEDESCHI, T. PRIOL. *Decentralized Workflow Coordination Through Molecular Composition*, in "7th International Workshop on Engineering Service-Oriented Applications", Paphos, Cyprus, December 2011, http://hal.inria.fr/hal-00644261/en.

[27] A. LAGE FREITAS, N. PARLAVANTZAS, J.-L. PAZAT. *Cost Reduction Through SLA-driven Self-Management*, in "European Conference on Web Services (ECOWS)", Lugano, Switzerland, September 2011 [*DOI :* 10.1109/ECOWS.2011.23], http://hal.inria.fr/inria-00600289/en.

[28] P. RITEAU, C. MORIN, T. PRIOL. *Shrinker: Improving Live Migration of Virtual Clusters over WANs with Distributed Data Deduplication and Content-Based Addressing*, in "17th International European Conference on Parallel and Distributed Computing (Euro-Par 2011)", Bordeaux, France, 2011, http://hal.inria.fr/inria-00600707/en.

[29] P. RITEAU. *Building Dynamic Computing Infrastructures over Distributed Clouds*, in "IPDPS 2011 PhD Forum", Anchorage, AK, United States, 2011 [*DOI :* 10.1109/IPDPS.2011.386], http://hal.inria.fr/inria-00596077/en.

[30] T. THANAKORNWORAKIJ, R. SHARMA, B. SCROGGS, C. LEANGSUKSUN, Z. GREENWOOD, P. RITEAU, C. MORIN. *High Availability on Cloud with HA-OSCAR*, in "4th Workshop on Resiliency in High Performance Computing (Resilience) in Clusters, Clouds, and Grids", Bordeaux, France, 2011, http://hal.inria.fr/inria-00625478/en.

[31] M. ZOUARI, M.-T. SEGARRA, F. ANDRÉ, A. THÉPAUT. *An architectural model for building distributed adaptation systems*, in "The 5th International Symposium on Intelligent Distributed Computing", Delft, Netherlands, October 2011, http://hal.inria.fr/inria-00600732/en.

National Conferences with Proceeding

[32] H. FERNANDEZ. *Exécution décentralisée de workflow par composition moléculaire*, in "RenPar 2011", St. Malo, France, May 2011, http://hal.inria.fr/inria-00625273/en.

[33] E. JEANVOINE, N. NICLAUSSE, L. NUSSBAUM, D. MARGERY. *Déploiement et partitionnement dynamique de clusters avec Kadeploy et Kavlan*, in "JRES 2011", Toulouse, France, November 2011, 1, http://hal.inria.fr/hal-00640588/en.

[34] P. MORILLON, L. NUSSBAUM, D. MARGERY. *Gestion d'une infrastructure expérimentale de grande échelle avec Puppet et Git*, in "9èmes Journées Réseaux - JRES 2011", Toulouse, France, November 2011, 1, http://hal.inria.fr/hal-00640589/en.

Conferences without Proceedings

[35] E. DAUBERT, F. ANDRÉ, O. BARAIS. *Adaptation multi-niveaux : l'infrastructure au service des applications*, in "Conférence Française en Systèmes d'Exploitation (CFSE)", St Malo, France, May 2011, http://hal.inria.fr/inria-00603973/en.

[36] P. RITEAU, K. KEAHEY, C. MORIN. *Bringing Elastic MapReduce to Scientific Clouds*, in "3rd Annual Workshop on Cloud Computing and Its Applications: Poster Session", Argonne, Illinois, United States, 2011, http://hal.inria.fr/inria-00624263/en.

Scientific Books (or Scientific Book chapters)

[37] J.-P. BANÂTRE, C. MORIN, T. PRIOL. *Fault Tolerant Autonomic Computing Systems in a Chemical Setting*, in "Dependable and Historic Computing", J. L. JONES (editor), Lecture Notes in Computer Science, Springer, October 2011, vol. 6875, http://hal.inria.fr/inria-00625261/en.

[38] A. MATSUNAGA, P. RITEAU, M. TSUGAWA, J. FORTES. *Crosscloud Computing*, in "High Performance Computing: From Grids and Clouds to Exascale", I. FOSTER, W. GENTZSCH, L. GRANDINETTI, G. R. JOUBERT (editors), Advances in Parallel Computing, IOS Press, 2011, vol. 20, p. 94-108 [*DOI :* 10.3233/978-1-60750-803-8-94], http://hal.inria.fr/inria-00625203/en.

[39] V. VLASSOV, A. AL-SHISHTAWY, P. BRAND, N. PARLAVANTZAS. *Self-Management with Niche, a Platform for Self-Managing Distributed Applications*, in "Formal and Practical Aspects of Autonomic Computing and Networking: Specification, Development and Verification", IGI Global, 2011, http://hal.inria.fr/hal-00519596/en.

Research Reports

[40] E. FELLER, L. RILLING, C. MORIN. *Energy-Aware Ant Colony Based Workload Placement in Clouds*, INRIA, May 2011, n^{o} RR-7622, http://hal.inria.fr/inria-00594992/en.

[41] E. FELLER, L. RILLING, C. MORIN. *Snooze: A Scalable and Autonomic Virtual Machine Management Framework for Private Clouds*, INRIA, December 2011, n^{o} RR-7833, http://hal.inria.fr/hal-00651542/en.

[42] H. FERNANDEZ, C. TEDESCHI, T. PRIOL. *Self-coordination of Workflow Execution Through Molecular Composition*, INRIA, May 2011, n^{o} RR-7610, http://hal.inria.fr/inria-00590357/en.

[43] A. LAGE FREITAS, N. PARLAVANTZAS, J.-L. PAZAT. *Cost Reduction Through SLA-driven Self-Management*, INRIA, May 2011, http://hal.inria.fr/inria-00592037/en.

[44] M. OBROVAC, C. TEDESCHI. *On Distributing the Runtime of the Chemical Programming Model*, INRIA, June 2011, n^{o} RR-7661, http://hal.inria.fr/inria-00604134/en.

[45] P. RITEAU, A. IORDACHE, C. MORIN. *Resilin: Elastic MapReduce for Private and Community Clouds*, INRIA, October 2011, n^{o} RR-7767, http://hal.inria.fr/inria-00632040/en.

Other Publications

[46] D. DIB. *Adaptation dynamique des fonctionnalités d'un système d'exploitation large échelle*, May 2011, Publication de la journée ADAPT 2011, http://hal.inria.fr/inria-00600712/en.

References in notes

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

[48] 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.

[49] 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.

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

[51] CONTRAIL CONSORTIUM. *D10.1, First Specification of the System Architecture*, September 2011, Contrail Deliverable.

[52] CONTRAIL CONSORTIUM. *D5.1, Requirements and Specification of Computational Resource Management Functionalities for Virtual Cluster Platforms*, September 2011, Contrail Deliverable.

[53] COOP CONSORTIUM. *D2.3, Analysis on Elements of Cooperation between Programming Model Frameworks and Resource Management Systems*, July 2011, Coop Deliverable.

[54] 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.

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

[56] E. FELLER, L. RILLING, C. MORIN, R. LOTTIAUX, D. LEPRINCE. *Snooze: A Scalable, Fault-Tolerant and Distributed Consolidation Manager for Large-Scale Clusters*, in "2010 IEEE/ACM International Conference on Green Computing and Communications (GreenCom-2010)", China Hangzhou, October 2010, http://hal.inria.fr/inria-00529702/en.

[57] J. GALLARD. *Kget+: An efficient tool for managing VM image snapshots*, in "Toulouse'2009 (RenPar'19 / SympA'13 / CFSE'7) – Poster Session", Toulouse France, 2009, http://hal.inria.fr/inria-00422651/PDF/main.pdf, http://hal.inria.fr/inria-00422651/en/.

[58] J. GALLARD, A. LÈBRE, C. MORIN. *Saline: Improving Best-Effort Job Management in Grids*, INRIA, 2009, n^{o} RR-7055, http://hal.inria.fr/inria-00422593/PDF/RR-7055.pdf, http://hal.inria.fr/inria-00422593/en/.

[59] P. HUNT, M. KONAR, F. P. JUNQUEIRA, B. REED. *ZooKeeper: wait-free coordination for internet-scale systems*, in "In USENIX Annual Technical Conference", 2010.

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