INRIA

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Project-Team Myriads

Design and Implementation of Autonomous Distributed Systems

Rennes - Bretagne-Atlantique

Theme : Distributed Systems and Services

Table of contents

1. Team

Research Scientists

Françoise André [Professor, on leave from UNIVERSITY RENNES 1, HdR]
Yvon Jégou [Research Associate (CR), INRIA]
Christine Morin [Research Director (DR), INRIA, Team leader, HdR]
Thierry Priol [Research Director (DR), INRIA, part time, 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, since October 2010]
Ghislain Charrier [INRIA Associate Engineer (IA), ADT Aladdin, since December 2010]
Rémy Garrigue [INRIA Associate Engineer (IA), ADT XtreemOS Easy, since November 2010]
Peter Linnell [INRIA Engineer, XTREEMOS IP Project, until November 2010]
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, since April 2010]
Thomas Ropars [INRIA Engineer, XTREEMOS IP Project, until March 2010]
Cyril Rohr [INRIA Associate Engineer (IA) ADT Aladdin until September 2010, Engineer BonFire IP Project since October 2010]

PhD Students

Stefania Costache [CIFRE EDF Grant, since May 2010]
Erwan Daubert [INRIA Grant]
Djawida Dib [INRIA and Brittany Regional Council Grant, since October 2010]
Eugen Feller [INRIA Grant]
Hector Fernandez [INRIA Grant]
Jérôme Gallard [INRIA Grant]
Guillaume Gauvrit [MENRT Grant]
Sylvain Jeuland [INRIA Grant, until September 2010]
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]

Visiting Scientists

Liviu Iftode [Rutgers University, USA, Invited Professor at INSA Rennes, June 2010]
Archana Nottamkandath [VU Amsterdam, The Netherlands, Master Student, summer internship, June-August 2010]
Sajith Kalathingal [VU Amsterdam, The Netherlands, Master Student, summer internship, June-August 2010]
Guillermo Vigueras Gonzales [Valencia University, Spain, PhD Student, internship, september-December 2010]

Administrative Assistants

Maryse Fouché [Team administrative assistant (TR) INRIA]
Sandrine L'Hermitte [XTREEMOS Scientific Coordinator (until September 2010) and Contrail Coordinator (since October 2010) Assistant INRIA]

Other

Jean-Pierre Banâtre [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 examplified 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. 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.

2.4.2. 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, state of resources (networks, processors, memory,...) have 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.

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

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

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 like fault-tolerance, but also raise some societal and environmental issues, like 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. To our knowledge, 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 while abstracting out details of the implementation to the users. 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. We have started a research activity aiming at emulating a DSM based on a Distributed Hash Table, whose load balancing property is of interest in the distribution of work. In this way, we have launched the development of a prototype based on these concepts. 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

In 2010, the team has the following highlights:

- The XTREEMOS European project, coordinated by the project-team (Ch. Morin) announced the third public version of the XTREEMOS Grid Operating System available for clusters, PC, notebooks, PDA and smartphones.
- Christine Morin was made a Knight in France's National Order of Merit.
- Pierre Riteau won the Large Scale Deployment Challenge that took place during the Grid'5000 Spring School 2010. He demonstrated the deployment of a sky computing environment based on nimbus in 31min! This included the deployment of the nodes using kadeploy3 and their configuration to form a sky computing environment located in 3 sites and managing 278 VMMs providing 1628 virtual cpus and 2 124 Gb of memory.

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, Service Oriented Architectures, cloud computing, distributed operating systems and unconventional and nature-inspired programming models.

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 complexity is reaching a level beyond human ability to manage and secure.

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 have to be far less involved in running them. Ideally, we would like computing systems to entirely manage themselves.

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

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

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

Transparent. 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 an with 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 companies applications.

Web services and Service Oriented Architectures 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 garantee Service Level Agreement 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. Cloud Computing

Cloud computing is internet-based computing where Internet providers provision computation power, data, software and services on a pay-as-you-use base to users. Cloud computing users avoid huge investment in hardware and services and allow rapid adaptation to the workload when they pay a provider only for what they use.

The cloud computing model 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 cloud, need to cooperate with resources and services provisioned from the cloud in order to cope with workload variations.

3.5. 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.

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 have a long experience in the distributed operating system, for instance in VIGNE and XTREEMOS projects.

The advent of cloud and green computing introduces new challenges in this domain: 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.

3.6. 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's 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 [49] that observing nature could be of great interest to tackle the problem of modeling and programming complex computing platforms, and overcome the limits of traditional programming models. Innovating, unconventional programming paradigms are requested to provide a high-level view of these interactions, then allowing to clearly separate what is a matter of expression from what is a question of implementation. Towards this, nature is of high inspiration, providing examples of self-organizing, fully decentralized coordination of complex and large scale systems.

As an example, chemical computing [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.

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,
- Business applications,
- Bio-informatics applications,
- Computational science applications,
- Numerical simulation,
- High performance computing.

5. Software

5.1. Kerrighed

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

URL: http://www.kerrighed.org/

Status: Registered at APP, under Reference IDDN.FR.001.480003.006.S.A.2000.000.10600.

License: GNU General Public License (GPL) version 2. KERRIGHED is a registered trademark.

Presentation: KERRIGHED is a *Single System Image* (SSI) operating system for high-performance computing on clusters. It provides the user with the illusion that a cluster is a virtual SMP machine. KERRIGHED is based on Linux which is slightly patched and extended with a kernel module. It is Posix compliant. Legacy sequential or parallel applications running on Linux can be executed without modification on top of KERRIGHED. Professional support is provided by Kerlabs http://www.kerlabs.com, a spin-off from PARIS project-team created in 2006. KERRIGHED is used in the cluster flavour of XTREEMOS Grid operating system. In 2010, we have further automatized the installation and configuration process of KERRIGHED when used in the framework of XTREEMOS system.

Active contributors (from Myriads team): Marko Obrovac, Eugen Feller.

5.2. Dynaco

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

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

Status: Version 0.2 is available.

License: GNU Lesser General Public License (LGPL) version 2.1.

Presentation: DYNACO (*Dynamic Adaptation for Components*) is a framework that helps in designing and implementing dynamically adaptable components. This framework is developed by the PARIS Project-Team. The implementation of DYNACO is based on the *Fractal Component Model* and its formalism.

Active contributors (from Myriads team): Françoise André, Jean-Louis Pazat

5.3. Vigne

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

URL: http://www.irisa.fr/paris/software/vigne/index_html#software

Status: Prototype

License: GNU General Public License (GPL).

Presentation: VIGNE is a prototype of a grid-aware operating system for grids, whose goal is to ease the use of computing resources in a grid for executing distributed applications. VIGNE is made up of a set of operating system services based on a peer-to-peer infrastructure. This infrastructure currently implements a structured overlay network inspired from *Pastry* and an unstructured overlay network inspired from *Scamp* for join operations. On top of the structured overlay network, a transparent data-sharing service based on the sequential consistency model has been implemented. It is able to handle an arbitrary number of simultaneous reconfigurations. An application execution management service has also been implemented including resource discovery, resource allocation, and application monitoring services. In 2010, we completed the development of the SEMIAS framework that provides statefull services with high availability and self healing based on active replication on top of a structured overlay. The VIGNE application execution management and application monitoring services have been made highly available and self-healing using the SEMIAS framework. The VIGNE prototype has been developed in C and includes around 30,000 lines of code. This prototype is coupled with a discrete-event simulator.

Active contributors (from Myriads team): Stefania Costache, Thomas Ropars

5.4. 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 [52]. 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 [51]. 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. In 2010, Saline has been integrated with Entropy developed by the Ascola project-team in order to consolidate virtual machines in a Grid environment.

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

5.5. HOCL

Contact: Thierry Priol, `Thierry.Priol@inria.fr`

Status: V2.3

License: GPL-2

Presentation: HOCL (Higher Order Chemical Language) is a chemical programming language based on the chemical metaphor presented before (see Section 3.6). 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 following features:

- Possibility of communication between two multisets
- Support for the communication between two different chemical programs

- Support for passing chemical rules between two multisets
- Support for the dynamic creation and addition of chemical rules at runtime

These aspects will allow to have a software support to experiment the chemical implementation of service coordination.

Active contributors (from Myriads team): Thierry Priol, Cédric Tedeschi, Chen Wang

5.6. Xtreemos

Contact: Christine Morin, `Christine.Morin@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 demonstrated at EuroPar 2010, Ischia, Italy, in August 2010 (Y. Jégou), at the XTREEMOS summer school, Günsburg, Germany in July 2010 (Y. Jégou), SC'10, New Orleans, USA (Y. Jégou, P. Riteau, Ch. Morin) in November 2010. 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 made public in fall 2010.

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

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

6.2.1. Systems for Clusters

Participants: Eugen Feller, Christine Morin, Thomas Ropars.

6.2.1.1. Energy management in clusters

In 2010, we worked on the design of energy-aware workload placement algorithms and a scalable middleware for workload management of software heterogeneous (i.e. virtualization and SSI technology) clusters. The energy-aware workload placement problem was modeled as one-dimensional bin-packing problem. According to this model, a heuristic (i.e. Better-fit) was selected as one possible approach to find a solution in reasonable time. Finally, a novel generic, scalable and fault-tolerant middleware called Snooze has been designed to manage software heterogeneous clusters. This work was published in GreenCom 2010 [18]. The implementation of Snooze is currently in progress and will be based on the cloud management software OpenNebula.

6.2.1.2. Checkpointing Parallel Applications

To improve message logging protocols scalability, we proposed distributed event logging. We implemented an efficient version of our distributed event logger in the Open MPI library and ran extensive tests on Grid'5000 to show that it can improve the performance of different message logging protocols [27].

6.2.2. Grid operating systems

Participants: Amine Belhaj, Stefania Costache, Jérôme Gallard, Rémy Garrigue, Yvon Jégou, Sylvain Jeuland, Peter Linnell, Christine Morin, Marko Obrovac, Yann Radenac, Pierre Riteau, Thomas Ropars.

6.2.2.1. Access Control

XtreemOS aims to provide to grid users an interface similar to their usual Linux desktop interface: application are run on the grid as if they were executed on the local desktop. In order to provide this interface, XTREEMOS must provide means to support single-sign-on (no need to authenticate each time an application is run), delegation (applications running on the grid have the same capabilities as applications running on the desktop). In 2010, we completed the implementation of the XtreemOS single-sign-on and delegation system.

6.2.2.2. Dynamic Virtual Organizations

A key feature of XTREEMOS is its support for Virtual Organizations (VOs). We designed and implemented the KsiVO system, a system managing dynamic virtual organizations in the context of the XtreemOS system. This system allows to establish in an automatic way a secure collaboration environment among several users [42]. It permits users to communicate, exchange data and run jobs in a collaborative workflow.

6.2.2.3. Checkpointing Grid Applications

XTREEMGCP is a service of the XTREEMOS grid system that provides grid applications with fault tolerance. It is able to apply different fault tolerance strategies and to make use of the various kernel checkpointers available on the grid nodes. In 2010, we integrated an independent checkpointing protocol in the XtreemGCP Grid checkpointing service [17].

6.2.2.4. High Availability for Grid Services

Regarding high availability of grid services, we have completed the development of Semias, a framework for active replication on top of structured peer-to-peer overlays. Its self-healing mechanisms ensure the high availability of services despite reconfigurations in the peer-to-peer overlay. We have validated the Semias architecture on two different testbeds, Grid5000 and PlanetLab, and we have carried out performance tests to measure the efficiency of Semias self-healing capabilities in highly dynamic systems [16].

6.2.2.5. XtreemOS Release and Deployment Tools

Yvon Jégou and Peter Linnell, as release manager, coordinated the production and the testing of the third integrated version of XTREEMOS Grid operating system (XtreemOS V3.0 version), publicly released in December 2010 (http://www.xtreemos.eu/software). We have contributed to the XTREEMOS admin and user guides [54], [55]. The public permanent geographically distributed testbed made up of several computers provided by different XTREEMOS partners has been updated with the new XTREEMOS release and used for testing and demonstrating the XTREEMOS prototype (http://www.xtreemos.eu/open-testbed/).

In 2010, we have developed a set of tools and environments to facilitate the deployment of the XTREEMOS system on various infrastructures. We have designed and implemented a configuration tool to facilitate the deployment of XTREEMOS system on a Grid made up of physical or virtual machines. Moreover, pre-configured sets of virtual machines for KVM and VirtualBox hypervisors have been produced, documented and made available for the XTREEMOS consortium and the open source community (http://www.xtreemos.eu/software). Tools to automatically deploy XTREEMOS PC and cluster flavours on Grid'5000 platform have been developed and made available to XTREEMOS consortium.

A guide, [48], detailing all steps necessary to setup a new virtual machine grid from the XtreemOS isos distributed by the consortium has been published in October 2010.

6.2.2.6. Validation of XtreemOS with Grid Applications

In order to validate experimentally XtreemOS features, we ported two applications on top of the system. In the framework of an informal collaboration with INRA, we ported the ParatubSpread_intra application, a stochastic compartmental model which represents the spread of Mycobacterium avium subsp. paratuberculosis (the pathogen of Paratuberulosis) within a dairy cattle herd. This embarrassingly parallel application is based on Scilab. We developed a dedicated runtime in order to efficiently manage the Grid resources

allocated to the application. This runtime will be made generic in the future. We successfully demonstrated ParatubSpread_intra application execution on several Grid'5000 sites where it was deployed on top of Nimbus clouds.

In the context of the COOP ANR project, we ported the Salome numerical simulation environment on top of XtreemOS.

6.2.3. Federated Virtualized Infrastructures

Participants: Stefania Costache, Djawida Dib, Jérôme Gallard, Yvon Jégou, Christine Morin, Nikos Parlavantzas, Thierry Priol, Pierre Riteau.

6.2.3.1. Virtual machine migration

We finalized a first prototype implementation of live virtual machine migration using distributed content addressing to speed up the process on wide area networks. This prototype has been evaluated on the Grid'5000 platform and the results show that it is able to reduce wide area bandwidth usage by 30 to 40% and migration time by 20% [45].

In the context of Djawida Dib's Master internship , we worked on dynamic adaptation of distributed application using network transparent live migration. We implemented a framework based on packet capture to transparently detect communication patterns in distributed application. This system is part of a framework we are building to automatically migrate virtual machines between clouds while taking communication patterns into account.

In collaboration with the ACIS laboratory, we worked on networking support for inter-cloud live virtual machine migration. This work was carried our using ViNe, a virtual network implementation providing all-to-all communication to nodes in grid or clouds environments, even in the presence of firewalls, NAT or private IP addressing. We modified ViNe to detect virtual machine mobility and reconfigure itself so that communications can remain uninterrupted. Our approach is based on standard networking techniques such as ARP proxy and gratuitous ARP messages, and leverages the ViNe infrastructure to establish tunnels between multiple cloud infrastructures. This work has been published at the IEEE MENS 2010 workshop [28].

6.2.3.2. Virtualization techniques for optimal execution of applications under time constraints

In the context of a collaboration with EDF R&D, we are investigating flexible resource management architectures for shared virtualized infrastructures. The objective is to execute applications with strict time constraints while minimizing resource supply costs. We have studied the advantages and limitations of modern virtualization technologies [35], explored current dynamic resource provisioning mechanisms exploiting those technologies, and evaluated the extend to which these mechanisms address our objective.

6.2.3.3. Virtual infrastructure management

The work made this year can be broken down into four parts: (i) extension of an XtreemOS grid with resources provided by a Cloud system (like Nimbus or Amazon EC2), (ii) use of XtreemOS as a system for IaaS Cloud providers, (iii) efficient management of virtual machines at grid level by integrating Saline (developed in the Myriads team) and Entropy (developed in the ASCOLA project-team) and (iv) further investigation of the concept of virtual platform aiming at abstracting the view of the physical resources from the users.

Extension of an XtreemOS Grid with resources provided by Cloud systems. We implemented a first prototype to extend XtreemOS with resources dynamically provided by an external cloud [46]. Currently our prototype relies on a private Nimbus Cloud. In future work, we plan to test our prototype with a public cloud like Amazon EC2. This work was done in collaboration with Eliana-Dina Tîrşa (a PhD student from the "Politehnica" University of Bucharest). Finally, we published our results on deploying XtreemOS over cloud resources as an INRIA technical report [46].

XtreemOS as a system for managing IaaS clouds. We integrated Saline virtual machine management system into XtreemOS Grid operating system, in order to deploy VMs and manage them (by moving them from one site to another) at grid level, transparently for the applications running inside the VMs. This work was done in collaboration with Sajith Kalathingal (master student from VU Amasterdam). A preliminary experimental evaluation has been performed using Grid'5000.

Efficient management of virtual machines at Grid level. We worked on the efficient management of VMs at grid level by integrating the Saline VM manager tool [20], [13] with the Entropy VM efficient manager. A preliminary experimental evaluation has been performed using Grid'5000. This work was done in collaboration with Archana Nottamkandath (master student from VU Amasterdam).

Virtual platform concept. In the framework of the SER-OS associated team, we have worked on the design and the implementation of a novel management tool for managing virtualized infrastructures comprising of high performance computing clusters, massively parallel processing (MPP) systems, and grids [41]. This novel management tool integrates three main concepts: (i) Virtual System Environments (VSEs) describing the application requirements in terms of software configuration (ii) Virtual Organizations (VOs) defining sets of resources shared among users communities, and (iii) Virtual Platforms (VPs) describing the application requirements in terms of hardware platform. In 2010, we have implemented a proof-of-concept prototype and we published a vision paper [21].

6.3. Dynamic Adaptation of Service-based Applications

6.3.1. Adaptation for data management

Participants: Françoise André, Mohamed Zouari.

The usage of context-aware data management in mobile environments has been investigated by Françoise André in collaboration with Mayté Segarra and Jean-Marie Gilliot from Telecom Bretagne Brest (previously known as ENST Bretagne). A context-aware data replication and consistency system that adapts dynamically to changes in the environment has been proposed, based on the use of the DYNACO framework. This work has been supported by a contract (*ReCoDEM*) between ENST Bretagne and Orange Labs (previously known as France-TélécomR&D).

In the *ReCoDEM* project, the distributed aspects of the adaptation system have not been thoroughly investigated. Therefore, a new subject is launched since October 2007 (with M. Zouari as PhD student) to propose a generic distributed adaptation framework. This work focuses on data management in grid and mobile environments; an ambient assisted living application illustrates the approach. Two architectural models have been proposed, one for distributed data management and the other for distributed dynamic adaptation. The models have been implemented as two component-based frameworks and evaluated for the ambient assisted living application. This work is realized in the context of the *ALORAD* project (Architecture LOgicielle pour la Réplication Adaptative de Données), financed by the Brittany council. Mayté Segarra from Telecom Bretagne Brest is co-adviser for the PhD thesis of M. Zouari.

6.3.2. 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 necessity of dynamic self-adaptation of services to satisfy the different users needs on constantly changing environments is evident.

The following works take place in the context of the S-CUBE European Network of Excellence

6.3.2.1. Services adaptation in distributed and heterogeneous systems

Based on our experience on components adaptation, Françoise André, Guillame Gauvrit and Erwan Daubert are now studying the problems of services adaptation in distributed and heterogeneous systems. This work takes place in the context of the S-CUBE European Network of Excellence and 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 build as a set of services, providing functionalities useful to build an adaptation system. The analysis phase takes decisions with multiple temporal scopes. This gives the ability to either react fast or to take proactive 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.

6.3.2.2. Atomic Service Adaptation

Jean-Louis Pazat, Nikos Parlavantzas, and André Lage are studying how to apply dynamic adaptation principles at the level of individual, atomic services running on large-scale distributed infrastructures, such as grids or clouds. This work has resulted in a generic framework that assists service providers in enforcing quality properties through automatically adapting service behaviour and resource usage [23]. A prototype of the framework has been developed, building on the XTREEMOS grid interface (XOSAGA). The initial evaluation demonstrates the effectiveness of the prototype in minimizing violations of service-level agreement terms.

6.3.2.3. Dynamic Adaptation of Chemical services

Jean-Louis Pazat and Chen Wang are studying how to express and implement dynamic adaptation of services within the context of chemical programming. Workflow enactement have been studied withing the S-Cube network and dynamic service binding of workflows. A prototype environnement connected to real web services is being implemented. Some extensions of the HoCl compiler have been done in order to remove some limitations and add new dynamic features such as communication.

6.4. A Chemical Approach for Autonomous Service Computing

Participants: Jean-Pierre Banâtre, Hector Fernandez, Jean-Louis Pazat, Thierry Priol, Cédric Tedeschi, Chen Wang.

In the context of web services, we have started to study how to take advantage of the good properties of the chemical computing paradigm (implicit parallelism, autonomic behavior) to express service coordination. For this purpose, we are defining the S-HOCL language, since the basic HOCL language does not support service coordination, and we are extending the HOCL compiler in several ways (refer to Section 5.5).

In collaboration with STFC (UK) under the framework of the CoreGRID Network of Excellence, we are studying the security engineering of distributed systems, such as Grids, when following the chemical-programming paradigm. We have analysed how to model secure systems using HOCL. Since our emphasis is on modularity, we advocate the use of aspect-oriented techniques, where security is seen as a cross-cutting concern impacting the whole system. We showed how HOCL can be used to model Virtual Organisations (VOs), exemplified by a VO system for the generation of digital products. We also develop security patterns for HOCL, including patterns for security properties such as authorisation, integrity and secure logs [3].

As mentioned earlier (refer to Section 6.3.2.3), we are also studying how to express and implement service adaptation through a chemical approach.

Finally, we started to study the development of the prototype of a middleware system for the distributed execution of chemical programs (targeted for large scale platforms).

7. Contracts and Grants with Industry

7.1. Contracts with Industry

7.1.1. Collaboration with EDF R&D

Participants: Christine Morin, Nikos Parlavantzas.

In the context of the INRIA - EDF R&D general agreement, we have a contract 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. Other Grants and Activities

8.1. Regional Initiatives

8.1.1. ASYST (2010-2013)

Participant: Françoise André.

The objective of the ASYST project (Adaptation dynamique des fonctionnalités d'un SYSTème d'eexploitation 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 (Hector Fernandez).

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

Participant: Christine Morin.

The goal of the ECO-GRAPPE project funded under the ANR ARPEGE program is to design, implement and validate energy saving policies in clusters (http://ecograppe.inria.fr/). Experimentations will be conducted using Kerrighed single system image cluster operating system. 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: Yvon Jégou, Christine Morin.

The COOP project funded under the ANR COSINUS program relates to multi level cooperative resource management (http://coop.gforge.inria.fr/). 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. HEMERA AEN

Participants: Yvon Jégou, Christine Morin.

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.

8.2.5. ADT Aladdin

Participants: Yvon Jégou, David Margery, Pascal Morillon.

The Aladdin technological development action funded by INRIA aims at the construction of a scientific instrument for experiments on large-scale parallel and distributed systems, building on the Grid'5000 platform. It structures INRIA's leadership role as the institute is present in 8 if the 9 Grid'5000 sites distributed across France.

Frédéric Desprez is the scientific director of this ADT and David Margery the technical director. 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 – BRETAGNE ATLANTIQUE in this executive committee.

The technical team is now composed of 12 engineers, of which 3 are from the MYRIADS team (David Margery, Pascal Morillon, Cyril Rohr). 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.

30 project-teams not including the 10 mentioned above are using or have shown interest in using such an instrument.

8.2.6. ADT XtreemOS Easy

Participants: Yvon Jégou, David Margery, Pascal Morillon, Christine Morin.

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

8.3. European Initiatives

8.3.1. XtreemOS IP Project (2006-2010)

Participants: Yvon Jégou, Christine Morin, Thierry Priol.

Christine Morin is the *Scientific Coordinator* (SCO) of the XTREEMOS Integrated Project (IP) http://www.xtreemos.eu in the IST 2006 Work Programme. XTREEMOS involves 19 academic and industrial partners from 7 European countries plus China. The XTREEMOS project aims at the design, implementation, evaluation and distribution of an open source Grid operating system with native support for virtual organizations and capable of running on a wide range of underlying platforms, from clusters to mobiles. The approach we propose in this project is to investigate the construction of a new Grid OS, XTREEMOS, based on the existing general-purpose OS Linux [12]. *Yvon Jégou* leads the WP4.3 Work-Package, aiming at setting up XTREEMOS testbeds. The GRID'5000 experimental grid platform is used as a large-scale testbed by XTREEMOS partners. Christine Morin leads WP1.1, Project management, WP2.2 Federation management and WP5.3, Collaboration with other IST Grid-related projects. Thierry Priol is a member of the *Scientific Advisory Committee*. Peter Linnell is the release manager and Sandrine L'Hermitte the project assistant. This project funds two PhD grants (Jérôme Gallard and Sylvain Jeuland). In 2010, we organized the second XtreemOS summit co-located with Europar

2010 and the first XtreemOS challenge in August 2010. We were also involved in the organization committee for the second XtreemOS summer school held in Günzburg, Germany in July 2010. The final review meeting was held December 1-2, 2010 in Barcelona, Spain.

8.3.2. Contrail European Integrated Project (2010-2013)

Participants: Yvon Jégou, Christine Morin.

The goal of the Contrail project is to design, implement and validate an open source cloud computing software stack (http://www.contrail-project.eu/). This project whose coordinator is Christine Morin is carried out by a consortium 10 academic and industrial partners. The Myriads team is mainly involved in research related to virtual cluster platform management.

8.3.3. S-Cube European Network of Excellence (2008-2012)

Participants: Françoise André, Jean-Louis Pazat.

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.

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

Participant: Christine Morin.

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 (http://www.scalus.eu/). 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, etc). The consortium involves 8 full academic partners, 2 full industrial partners and 4 additional associated industrial partners. The Myriads team participates in this project by co-advising with Professor Thomas Ludwig from Hamburg University (Germany) a PhD student working on load balancing and scheduling in parallel and cluster File Systems.

8.3.5. BonFire IP project (2010-2013)

The 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/).

8.3.6. IC0804 - Energy efficiency in large scale distributed systems

Participants: Françoise André, Jean-Louis Pazat.

This COST Action will propose realistic energy-efficient alternate solutions to share IT distributed resources. As large scale distributed systems gather and share more and more computing nodes and Storage resources, their energy consumption is exponentially increasing. While much effort is nowadays put into hardware specific solutions to lower energy consumptions, the need for a complementary approach is necessary at the distributed system level, i.e. middleware, network and applications. The Action will characterize the energy consumption and energy efficiencies of distributed applications.

8.3.7. CoreGrid ERCIM Working Group

Participant: Thierry Priol.

The Myriads team is involved in the CoreGrid ERCIM working group, a follow-up of the CoreGrid NoE Project that was led by Thierry Priol (http://www.coregrid.net/mambo/content/view/747/418/). The research activities of interest in the CoreGrid working group are evolving from Grid computing to cloud computing.

8.4. International Initiatives

8.4.1. DataCloud@work Associated Team with University Polytechnical Bucharest (2009-2011)

We are involved in the DataCloud@work INRIA associated team involving the team of Prof. Valentin Cristea from University Polytechnical Bucharest (UPB), the INRIA KERDATA and Myriads teams (http://www.irisa.fr/kerdata/doku.php?id=cloud_at_work:start). DataCloud@work is led by Gabriel Antoniu (KERDATA). 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. SER-OS Associated Team with ORNL (2009-2011)

Participant: Christine Morin.

The MYRIADS team is involved in a collaboration in the area of operating systems and system tool for HPC with the SRT team lead by Stephen Scott at ORNL. The main objectives of the SER-OS associated team (http://www.irisa.fr/myriads/ser-os/index.html) are the design of system services for the management of virtualized platforms and of resilience services for high performance computational applications. In 2010, we worked on the design of a novel management tool for HPC systems. We implemented a proof-of-concept prototype by integrating Saline and OSCAR (cluster management tool developed at ORNL). Jérôme Gallard visited ORNL during 3 months from October 2009 to January 2010 and two weeks in May 2010.

8.4.3. Argonne National Lab and the University of Florida

Participants: Yvon Jégou, Christine Morin, Thierry Priol.

We initiated a collaboration with Kate Keahey from Argonne National Laboratories and José Fortes from the ACIS laboratory at the University of Florida. In this context, we experimented with creating and using Sky Computing infrastructures, which allow to build virtual cluster on top of multiple federated clouds. This work was performed using the FutureGrid and Grid'5000 testbeds, and was presented as a demo at the OGF'29 meeting and as a poster at the Teragrid'10 conference. Pierre Riteau spent three months at the University of Florida (June-August 2010) and visited Argonne National Laboratory in June 2010. His internship has been supported by a grant from Rennes Metropole.

8.5. Exterior research visitors

Liviu Iftode, Professor at Rutgers University (USA) spent one month in the Myriads team in June 2010, as an invited professor, funded by INSA Rennes. We investigated potential joint research directions in the area of social clouds.

9. Dissemination

9.1. Animation of the scientific community

9.1.1. Leadership, Steering Committees and community service

European XTREEMOS IST-FP6 Integrated Project. Ch. Morin is the *Scientific Coordinator* of the XTREEMOS Integrated Project (http://www.xtreemos.eu/). Y. Jégou is a member of XTREEMOS Executive Committee. Th. Priol is a member of XTREEMOS Scientific Advisory Committee. J.-P. Banâtre is the INRIA representative in the XTREEMOS Governing Board.

European Contrail ICT-FP7 Integrated Project. Ch. Morin is the *Coordinator* of the Contrail Integrated Project (http://www.contrail-project.eu/). Y. Jégou is a member of Contrail Executive Committee. Th. Priol is a member of Contrail Scientific Advisory Board. J.-P. Banâtre is the INRIA representative in the Contrail General Assembly.

European S-CUBE IST-FP7 Network of Excellence J.-L. Pazat is the *Leader* of the Infrastructure JRA 2.3 workpackage of this project. J.-P. Banâtre is the INRIA representative in the S-CUBE Governing Board.

CNRS, GDR ASR. J.-L. Pazat is member of the steering committee of the CNRS Research Co-operative Federation (*Groupement de recherche*, GDR) ASR (*Architectures, Systems and Networks*). F. André serves as the coordinator of the ADAPT action (*Dynamic Adaptation*) of the GSP working group.

ALADDIN-G5K D. Margery is the technical director of ALADDIN-G5K. Y. Jégou is a member of the Grid'5000 executive committee.

9.1.2. Editorial boards, direction of program committees

Ch. Morin served as the General Chair of the ACM EuroSys 2010 conference held April 13-16 in Paris (http://eurosys2010.sigops-france.fr/), as the Program Co-Chair of the 2nd International ICST Conference on Cloud Computing (CloudComp 2010) held October 25-26 in Barcelona, Spain (http://www.cloudcomp.eu/). She has been involved in the Organization Committee of the second XTREEMOS summer school held in Günsburg, Germany in July 2010 (http://www.xtreemos.eu/project/xtreemos-events/summer-school-2010). She organized the first INRIA - Microsoft Research workshop on extreme operating systems held in Paris, France on November 3-4, 2010.

J.-L. Pazat was co-chair of the IEEE CIT2010 program committee, he serves 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 *Future Internet Journal*, Scalable Computing and Communications Journal,*International Journal of Intelligent Engineering Informatics* and of the *International Journal of Web Services Research*.

9.1.3. Program Committees

Ch. Morin served in the Program Committees of the following conferences: CIT 2010, LADIS 2010, SNAPI 2010, Resilience 2010, VTDC 2010, ServiceWave 2010, IPDPS 2011, SC 2011, VTDC 2011, Cloud for High Performance Computing Workshop at ICCSA 2011.

J.-L. Pazat served in the Program Committees of the following conferences: CIT 2009, GPC 2009, RenPar 19.

Th. Priol served in the Program Committees of the following conferences: CBHPC10, CLOUD10, ICWS10, IWHGA10, VECPAR10.

Th. Priol served as the General Chair of ServiceWave 2010. He is also the Chairman of the ServiceWave Steering Committee.

C. Tedeschi served in the program Committee of ICCS 2010, and was P2P track chair for SSS 2010.

9.1.4. Evaluation committees, consulting

F. André acts as an AERES (Agence d'évaluation de la recherche et de l'enseignement supérieur) expert to evaluate a research laboratory.

Ch. Morin acted as a referee for the HDR Committee of Sara Bouchenak (University of Grenoble I) and as a member for the PhD Committee of Simon Duquennoy (LIFL).

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 XtreemOS, ADMIRE and DRIHMS FP7 projects. He is also an evaluator for the French-Isreaeli Multicomputing initiative.

J.-L. Pazat is the responsible of the Executive Board of Computer Science Selections Committees at INSA (*Bureau de section*. He is member of the Administrative Board at INSA.

9.2. Academic teaching

Only the teaching contributions of project-team members on non-teaching positions are mentioned below.

Ch. Morin is responsible for a graduate teaching Module *Distributed Systems: from networks to Grids* of the Master Program in Computer Science, UNIVERSITY RENNES 1. Within this module, she gave lectures on Grid computing. She gave a lecture on cluster single system image operating systems within the *Parallelism* Module of the 3rd-year students of Telecom Sud Paris.

9.3. Conferences, seminars, and invitations

Only the events not listed elsewhere are listed below.

Y. Jégou gave a lecture on *Support for interactive applications* at the tutorial on Easing Application Execution in Grids with XtreemOS Operating System, at OGF 28 held in Munich, Germany, in March 2010. He gave a lecture *Testbed Deployment* at the XTREEMOS Summer School, Günsburg, Germany, July 2010. He presented demonstrations of the XTREEMOS system at the EGEE 5th user meeting, Uppsala, Sweden, in April 2010, at the CCGRid 2010 conference, Melbourne, Australia, in May 2010, at Ter@tec forum on high performance computing, in Paris, France, June 2010, at the second XtreemOS summit, co-located with EuroPar 2010, Ischia, Italy in August 2010 and at SC'10, New Orleans, USA (November 2010).

E. Feller gave the following talks: *Snooze: A Scalable, Fault-Tolerant and Distributed Consolidation Manager for Large-Scale Clusters* at the COST IC0804 workshop, Choimbra, Portugal, November 2010,

EcoGrappe - Energy Consumption Management in Clusters at the DataCloud@Work INRIA associated team workshop, Rennes, France, September 2010

at the SER-OS INRIA associated team workshop, Oak Ridge National Labaratory, USA, May 2010 and

at the COST IC0804 workshop, Passau, Germany, April 2010.

J. Gallard gave a talk entitled *Everything as a Service our Vision of the Flexibility in Computing Systems* at the SER-OS Workshop, Oak Ridge, Tennessee, US, May 2010) http://www.irisa.fr/myriads/ser-os/workshop10/index.html.

P. Linnell presented XtreemOS at the *Rencontres Mondiales du Logiciel Libre 2010* held in Bordeaux, France, July 2010.

Ch. Morin gave a talk *Service Oriented Infrastructures: from Grids to Clouds & XtreemOS Perspective* at the *KIC-Ouest Infrastructures et Services* seminar on March 4, at Bruz, France. She gave a lecture *XtreemOS Overview* at the tutorial on Easing Application Execution in Grids with XtreemOS Operating System, at OGF 28 held in Munich, Germany, in March 2010. She was invited to give a talk *XtreemOS : des grilles aux nuages informatiques* at the InTech seminar on virtualization and cloud computing at INRIA Grenoble - Rhône Alpes on March 25 2010. She participated in the *A la découverte de la recherche* action giving a talk at Cesson-Sévigné high school (*lycée*) on April 27, 2010. She presented Myriads team research activities on the design and implementation of autonomous distributed systems in the framework of the SER-OS associated team workshop at ORNL, Oak Ridge, USA, in May 2010. She was invited to give a talk entitled *Projet européen XtreemOS Retour d'expérience* at the INRIA Rennes financial service seminar held in Saint Malo, France on June 3, 2010. She was invited to give a talk *XtreemOS: an Open Source Distributed Operating System Approaching Cloud Computing* at the workshop on X-HPC: Experimenting with XtreemOS to provide HPC services organized by CNR, Pisa, Italy June 15, 2010. She gave a talk on *XtreemOS: from Grids to Clouds* at the third workshop of the Joint INRIA-NCSA Laboratory for Petascale Computing, Bordeaux, France, June 23, 2010. She gave a lecture on *XtreemOS: Evolving from a Grid to a Cloud Computing System* at XTREEMOS Summer School, Günsburg, Germany,

July 2010. She gave lectures on XtreemOS and on XtreemGCP Grid checkpointing service at the second XtreemOS summit, co-located with EuroPar 2010, Ischia, Italy in August 2010. She gave an invited talk *XtreemOS European Project: Achievements & Perspectives* at the CCGSC 2010 workshop on Clusters, Clouds and Grids for Computational Science, held in Flat Rock, USA, September 2010. She gave a talk entitled *From XtreemOS Grids to Contrail Clouds* at the INRIA - MSR workshop, Paris, France on November 3, 2010. She presented XtreemOS system on XtreemOS booth at Sc'2010, New Orleans, USA, November 2010. She gave a presentation on XtreemOS at the IRISA AERES evaluation meeting, Rennes, France, November 23, 2010. he gave lectures on XtreemOS Grid system and on the Single System Image technology and presented Contrail project at the training event organized by the SPRERS European project (http://sprers.eu/events/training-on-software-services), Timisoara, Romania, on December 8-9, 2010. She was invited to give a talk *From XtreemOS Grids to Contrail Clouds* at the day on *Des grilles aux clouds, nouveaux problèmes, nouvelles solutions* organized by France Grilles at ENS Lyon, December 13 2010.

P. Riteau gave a seminar at Argonne National Laboratory (June 29, 2010) on *Sky Computing – Challenges, Large-Scale Experiments and Directions*. He presented a demonstration at the OGF'29 meeting (June 21, 2010 in Chicago, IL) on *Sky Computing on FutureGrid and Grid'5000*. He gave a talk at the SER-OS 2010 workshop (May 5, 2010 at Oak Ridge National Laboratory) on *Dynamic Adaptation of Virtualized Distributed Applications* and a talk at the Grid'5000 2010 Spring School (April 8, 2010 in Lille, France) on *Shrinker: Efficient Wide-Area Live Virtual Machine Migration using Distributed Content-Addressing*. He presented demonstrations of the XTREEMOS system at SC'10, New Orleans, USA (November 2010).

S. Costache gave a talk entitled *Dynamic resource provisioning for ensuring application-level QoS*, Paris, Workshop EDF R&D - INRIA, October 2010.

9.4. Administrative responsibilities

T. Priol has been 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.

J. Gallard was elected member of the INRIA Rennes - Bretagne Atlantique Center Council (*Comité de Centre*).

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

9.5. Miscellaneous

F. André has served as a member for a selection committee of IFSIC (Computer Science department of University of Rennes1), to appoint an assistant Professor in Computer Science.

- Jérôme Gallard was volunteer student at Eurosys'10 International Conference, April 2010, Paris, France.
- Sylvain Jeuland was volunteer student at Eurosys'10 International Conference, April 2010, Paris, France.

Ch. Morin served as a member of the INRIA Rennes local committee for attributing INRIA/CORDIS Post-doctoral grants in 2010. She is a member of the Project-Team Committee of INRIA - Rennes Bretagne Atlantique (*Comité des projets*).

J.-L. Pazat is member of the Steering committee (*c*onseil d'administration) of INSA Rennes. He is a member of the Computer Science Department committee. He is the local coordinator for the international exchange of students at the computer science department of INSA. He is member of the technological development committe (CDT) of IRISA.

- Pierre Riteau was volunteer student at Eurosys'10 International Conference, April 2010, Paris, France.

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 Secure Chemical Programs with Aspects*, Institute on Knowledge and Data Management , CoreGRID - Network of Excellence, August 2008, n^{o} TR-0166, http://www.coregrid.net/mambo/images/stories/TechnicalReports/tr-0166.pdf.

[3] 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/.

[4] 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.

[5] 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.

[6] E. JEANVOINE. *Intergiciel pour l'exécution efficace et fiable d'applications distribuées dans des grilles dynamiques de très grande taille*, Université de Rennes 1, IRISA, Rennes, France, November 2007, In french., ftp://ftp.irisa.fr/techreports/theses/2007/jeanvoine.pdf.

[7] Z. LAHJOMRI, T. PRIOL. *KOAN: A Shared Virtual Memory for iPSC/2 Hypercube*, in "Proc. of the 2nd Joint Int'l Conf. on Vector and Parallel Processing (CONPAR'92)", Lecture Notes in Computer Science, Springer Verlag, September 1992, vol. 634, p. 441–452, http://hal.inria.fr/inria-00074927.

[8] R. LOTTIAUX. *Gestion globale de la mémoire physique d'une grappe pour un système à image unique : mise en œuvre dans le système Gobelins*, IRISA, Université de Rennes 1, December 2001.

[9] T. PRIOL. *Efficient support of MPI-based parallel codes within a CORBA-based software infrastructure*, in "Response to the Aggregated Computing RFI from the OMG, Document orbos/99-07-10", July 1999.

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

[11] 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.

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

Publications of the year

Articles in International Peer-Reviewed Journal

[13] J. GALLARD, A. LÈBRE. *Managing Virtual Resources: Fly through the Sky*, in "ERCIM News", 10 2010, p. 36–37, http://hal.inria.fr/inria-00536994/en/.

[14] U. SCHWIEGELSHOHN, R. BADIA, M. BUBAK, M. DANELUTTO, S. DUSTDAR, F. GAGLIARDI, A. GEIGER, L. HLUCHY, D. KRANZLMUELLER, E. LAURE, T. PRIOL, A. REINEFELD, M. RESCH, A. REUTER, O. RIENHOFF, T. RUETER, P. SLOOT, D. TALIA, K. ULLMANN, R. YAHYAPOUR, G. VON VOIGT. *Perspectives on grid computing*, in "Future Generation Computer Systems", May 2010, vol. 26, n^{o} 8, p. 1104-1115 [*DOI :* 10.1016/J.FUTURE.2010.05.010], http://hal.inria.fr/inria-00508409/en.

International Peer-Reviewed Conference/Proceedings

[15] F. ANDRÉ, E. DAUBERT, G. GAUVRIT. *Towards a Generic Context-Aware Framework for Self-Adaptation of Service-Oriented Architectures*, in "5th International Conference on Internet and Web Applications and Services", Spain Barcelona, May 2010, p. 309-314 [*DOI :* 10.1109/ICIW.2010.52], http://hal.inria.fr/inria-00470487/en.

[16] 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)", India New Dehli, November 2010, http://hal.inria.fr/inria-00526101/en.

[17] E. FELLER, J. MEHNERT-SPAHN, M. SCHOETTNER, C. MORIN. *Independent Checkpointing in a Heterogeneous Grid Environment*, in "First International Conference on Utility and Cloud Computing (UCC 2010)", India Chennai, December 2010, http://hal.inria.fr/inria-00532912/en.

[18] 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.

[19] H. FERNANDEZ, T. PRIOL, C. TEDESCHI. *Decentralized Approach for Execution of Composite Web Services Using the Chemical Paradigm*, in "IEEE 8th International Conference on Web Services (ICWS 2010)", United States Miami, July 2010, http://hal.inria.fr/inria-00508346/en.

[20] J. GALLARD, A. LÈBRE, C. MORIN. *Saline: Improving Best-Effort Job Management in Grids*, in "PDP 2010: The 18th Euromicro International Conference on Parallel, Distributed and Network-Based Computing – Special Session: Virtualization in Distributed Systems", Pisa Italy, 2010, http://hal.inria.fr/inria-00426373/en/.

[21] J. GALLARD, C. MORIN, G. VALLÉE, T. NAUGHTON, S. SCOTT, A. LÈBRE. *Architecture for the Next Generation System Management Tools.*, in "First International Conference on Utility and Cloud Computing (UCC 2010)", Chennai Inde, 2011, http://hal.inria.fr/inria-00548605/en/.

[22] G. GAUVRIT, E. DAUBERT, F. ANDRÉ. *SAFDIS: A Framework to Bring Self-Adaptability to Service-Based Distributed Applications*, in "36th EUROMICRO Conference on Software Engineering and Advanced Applications (SEAA)", France Lille, September 2010 [*DOI :* 10.1109/SEAA.2010.25], http://hal.inria.fr/inria-00480221/en.

[23] A. LAGE FREITAS, N. PARLAVANTZAS, J.-L. PAZAT. *A QoS Assurance Framework for Distributed Infrastructures*, in "3rd International Workshop on Monitoring, Adaptation and Beyond (MONA+)", Cyprus Ayia Napa, December 2010, http://hal.inria.fr/inria-00533377/en.

[24] A. LAGE FREITAS, J.-L. PAZAT. *A Self-Adaptable Approach for Easing the Development of Grid-Oriented Services*, in "International Conference on Computer and Information Technology (CIT 2010)", United Kingdom Bradford, June 2010 [*DOI :* 10.1109/CIT.2010.56], http://hal.inria.fr/inria-00474992/en.

[25] A. LAGE FREITAS, J.-L. PAZAT, N. PARLAVANTZAS. *Ensuring Resource-Level Quality for Services on Grids*, in "2010 6th World Congress on Services", United States Miami, July 2010, p. 168-169 [*DOI :* 10.1109/SERVICES.2010.75], http://hal.inria.fr/inria-00530191/en.

[26] C. LEORDEANU, T. ROPARS, Y. JÉGOU, V. CRISTEA, C. MORIN. *Using a Failure History Service for Reliable Grid Node Information*, in "Fifth International Conference on P2P, Parallel, Grid, Cloud and Internet Computing", Japan Fukuoka, November 2010, http://hal.inria.fr/inria-00534838/en.

[27] T. ROPARS, C. MORIN. *Improving Message Logging Protocols Scalability through Distributed Event Logging*, in "16th International Euro-Par Conference", Italy Ischia, August 2010, http://www.springerlink.com/content/m7127166764674422/, http://hal.inria.fr/inria-00526097/en.

[28] M. TSUGAWA, P. RITEAU, A. MATSUNAGA, J. FORTES. *User-level Virtual Networking Mechanisms to Support Virtual Machine Migration Over Multiple Clouds*, in "The 2nd IEEE International Workshop on Management of Emerging Networks and Services (IEEE MENS 2010)", United States Miami, December 2010, http://hal.inria.fr/inria-00535569/en.

[29] C. WANG, J.-L. PAZAT. *Using Chemical Metaphor to Express Workflow and Service Orchestration*, in "The 10th IEEE International Conference on Computer and Information Technology", United Kingdom Bradford, June 2010, http://hal.inria.fr/inria-00534135/en.

[30] M. ZOUARI, M.-T. SEGARRA, F. ANDRÉ. *A framework for distributed management of dynamic self-adaptation in heterogeneous environments*, in "The 10th IEEE International Conference on Computer and Information Technology", United Kingdom Bradford, June 2010, http://hal.inria.fr/inria-00471892/en.

Scientific Books (or Scientific Book chapters)

[31] 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, December 2011, http://hal.inria.fr/hal-00519596/en.

Books or Proceedings Editing

[32] B. CHAPMAN, F. DESPREZ, G. JOUBERT, A. LICHNEWSKY, F. PETERS, T. PRIOL (editors). *Parallel Computing: From Multicores and GPU's to Petascale. Proceedings of ParCO 2009*, Advances in Parallel Computing, IOS Press, 2010, vol. 19, http://hal.inria.fr/inria-00484902/en.

[33] F. DESPREZ, V. GETOV, T. PRIOL, R. YAHYAPOUR (editors). *Grids, P2P and Services Computing*, CoreGRID book series, Springer, July 2010, vol. 10, http://hal.inria.fr/inria-00508331/en.

Research Reports

[34] V. ACRETOAIE, E. CARON, C. TEDESCHI. *A Practical Study of Self-Stabilization for Prefix-Tree Based Overlay Networks*, INRIA, April 2010, n° RR-7252, http://hal.inria.fr/inria-00474376/en.

[35] S. COSTACHE, N. PARLAVANTZAS, C. MORIN, S. KORTAS. *Use Cases of Virtualization in the Management of Distributed and Parallel Systems*, INRIA, December 2010, n° RT-0399, http://hal.inria.fr/inria-00547804/en.

[36] S. COSTACHE, T. ROPARS, C. MORIN. *Towards Highly Available and Self-Healing Grid Services*, INRIA, April 2010, n° RR-7264, http://hal.inria.fr/inria-00476276/en.

[37] E. FELLER, J. MEHNERT-SPAHN, M. SCHOETTNER, C. MORIN. *Independent Checkpointing in a Heterogeneous Grid Environment*, INRIA, September 2010, n° 7399, http://hal.inria.fr/inria-00521443/en.

[38] E. FELLER, C. MORIN, D. LEPRINCE. *State of the Art of Power Saving in Clusters and Results from the EDF Case Study*, INRIA, December 2010, n° RR-7473, http://hal.inria.fr/inria-00542889/en.

[39] E. FELLER, L. RILLING, C. MORIN, R. LOTTIAUX, D. LEPRINCE. *Snooze: A Scalable, Fault-Tolerant and Distributed Consolidation Manager for Large-Scale Clusters*, INRIA, September 2010, n° 7398, http://hal.inria.fr/inria-00521390/en.

[40] H. FERNANDEZ, T. PRIOL, C. TEDESCHI. *Decentralized Approach for Execution of Composite Web Services using the Chemical Paradigm*, INRIA, April 2010, n° RR-7267, http://hal.inria.fr/inria-00476991/en.

[41] J. GALLARD, G. VALLÉE, T. NAUGHTON, A. LÈBRE, S. SCOTT, C. MORIN. *Architecture for the Next Generation System Management Tools for Distributed Computing Platforms*, INRIA, May 2010, n° RR-7325, http://hal.inria.fr/inria-00494328/en.

[42] S. JEULAND, C. MORIN, Y. JÉGOU. *On-the-fly Dynamic Virtual Organizations in a Secured Grid Environment*, INRIA, September 2010, http://hal.inria.fr/inria-00515841/en.

[43] Y. JÉGOU. *Installing XtreemOS on a Virtual Machine*, INRIA, October 2010, n° XtreemOS Technical Report # 6, http://hal.inria.fr/inria-00545871/en.

[44] L. LIQUORI, C. TEDESCHI, L. VANNI, F. BONGIOVANNI, V. CIANCAGLINI, B. MARINKOVIC. *Synapse: a Scalable Protocol for Interconnecting Heterogeneous Overlay Networks*, INRIA, April 2010, n° RR-7255, http://hal.inria.fr/inria-00474529/en.

[45] P. RITEAU, C. MORIN, T. PRIOL. *Shrinker: Efficient Wide-Area Live Virtual Machine Migration using Distributed Content-Based Addressing*, INRIA, February 2010, n° RR-7198, http://hal.inria.fr/inria-00454727/en.

[46] E.-D. TIRSA, P. RITEAU, J. GALLARD, C. MORIN, Y. JÉGOU. *Towards XtreemOS in the Clouds – Automatic Deployment of XtreemOS Resources in a Nimbus Cloud*, INRIA, October 2010, n° RT-0395, http://hal.inria.fr/inria-00524843/en.

Other Publications

[47] F. ANDRÉ, E. DAUBERT, G. NAIN, B. MORIN, O. BARAIS. *F4Plan: An Approach to build Efficient Adaptation Plans*, 2010, http://hal.inria.fr/hal-00548736/en.

[48] XTREEMOS CONSORTIUM. , Y. JÉGOU (editor)*Installing XtreemOS on a Virtual Machine*, October 2010, XtreemOS Technical Report #6, http://xtreemos.eu/project/publications/VMInstall.pdf.

References in notes

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

[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] 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/.

[52] 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/.

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

[54] XTREEMOS CONSORTIUM. , M. COPPOLA (editor)*XtreemOS Advanced Guide: Installation and Administration*, October 2009, XtreemOS Technical Report #9, http://www.xtreemos.eu/software/adminguide.pdf.

[55] XTREEMOS CONSORTIUM. , M. COPPOLA (editor)*XtreemOS User's Guide*, October 2009, XtreemOS Technical Report #10, http://www.xtreemos.eu/software/userguide.pdf.