



IN PARTNERSHIP WITH:
CNRS

**Institut polytechnique de
Grenoble**

**Université Joseph Fourier
(Grenoble)**

Activity Report 2014

Project-Team MESCAL

Middleware efficiently scalable

IN COLLABORATION WITH: Laboratoire d'Informatique de Grenoble (LIG)

RESEARCH CENTER
Grenoble - Rhône-Alpes

THEME
**Distributed and High Performance
Computing**

Table of contents

1. Members	1
2. Overall Objectives	2
2.1. Presentation	2
2.2. Objectives	2
3. Research Program	3
3.1. Large System Modeling and Analysis	3
3.1.1. Simulation of distributed systems	3
3.1.1.1. Flow Simulations	3
3.1.1.2. Perfect Simulation	3
3.1.2. Fluid models and mean field limits	3
3.1.3. Game Theory	3
3.2. Management of Large Architectures	4
3.2.1. Instrumentation, analysis and prediction tools	4
3.2.2. Fairness in large-scale distributed systems	4
3.2.3. Tools to operate clusters	5
3.2.4. Simple and scalable batch scheduler for clusters and grids	5
3.3. Migration resilience; Large scale data management	5
4. Application Domains	5
4.1. Cloud, Grid, Multi-core and Desktop Computing	5
4.2. Wireless Networks	6
4.3. On-demand Geographical Maps	6
4.4. Energy and Transportation	6
5. New Software and Platforms	6
5.1. Tools to visualize and analyze traces of execution of distributed applications	6
5.2. Simulation and performance evaluation tools	7
5.2.1. SimGrid	7
5.2.2. Perfect simulator	7
5.2.3. PEPS	7
5.2.4. GameSeer	7
5.3. Tools for cluster management and software development	8
5.3.1. KA-Deploy	8
5.3.2. Kameleon	8
5.4. Infrastructure Middleware and scheduler	8
5.4.1. OAR	8
5.4.2. CiGri	8
5.4.3. ComputeMode	8
5.5. Platforms	9
5.5.1. Grid'5000	9
5.5.2. The ICluster-2, the IDPot and the new Digitalis Platforms	9
5.5.3. The Bull Machine	9
6. New Results	9
6.1. Simulation of distributed architectures	9
6.2. Interactive Analysis and Visualization of Large Distributed Systems	9
6.3. Management of Parallel Architectures	10
6.4. Reproducible experiments and papers	10
6.5. Game Theory and Distributed Optimization	11
6.6. Agent-based modeling and applications to Smart Energy and Transportation Systems	11
7. Bilateral Contracts and Grants with Industry	12
7.1.1. RealTimeAtWork.com	12

7.1.2.	Alcatel Lucent-Bell	12
7.1.3.	Stimergy	12
8.	Partnerships and Cooperations	13
8.1.	Regional Initiatives	13
8.1.1.	CIMENT	13
8.1.2.	Cluster Région	13
8.2.	National Initiatives	13
8.2.1.	Inria Large Scale Initiative	13
8.2.2.	ANR	13
8.2.3.	National Organizations	14
8.3.	European Initiatives	14
8.3.1.	FP7 & H2020 Projects	14
8.3.1.1.	Mont-Blanc	14
8.3.1.2.	Mont-Blanc 2	15
8.3.1.3.	QUANTICOL	16
8.3.1.4.	NEWCOM#	16
8.3.2.	Collaborations in European Programs, except FP7 & H2020	16
8.3.3.	Collaborations with Major European Organizations	17
8.4.	International Initiatives	17
8.4.1.	Inria International Labs	17
8.4.2.	Inria Associate Teams	17
8.4.2.1.	EXASE	17
8.4.2.2.	CLOUDSHARE	18
8.4.3.	Inria International Partners	18
8.4.4.	Participation In other International Programs	18
8.5.	International Research Visitors	19
8.5.1.	Visits of International Scientists	19
8.5.2.	Visits to International Teams	19
9.	Dissemination	19
9.1.	Promoting Scientific Activities	19
9.1.1.	Scientific events organisation	19
9.1.1.1.	general chair, scientific chair	19
9.1.1.2.	member of the organizing committee	20
9.1.2.	Scientific events selection	20
9.1.3.	Journal	20
9.2.	Teaching - Supervision - Juries	20
9.2.1.	Teaching	20
9.2.2.	Supervision	21
9.2.3.	Juries	21
9.3.	Popularization	21
10.	Bibliography	21

Project-Team MESCAL

Keywords: High Performance Computing, Game Theory, Grid'5000, Scheduling, Stochastic Modeling

Creation of the Project-Team: 2006 January 01.

1. Members

Research Scientists

Bruno Gaujal [Team leader, Inria, Senior Researcher, HdR]
Nicolas Gast [Inria, Researcher]
Arnaud Legrand [CNRS, Researcher]
Panayotis Mertikopoulos [CNRS, Researcher]

Faculty Members

Yves Denneulin [INP Grenoble, Professor, HdR]
Florence Perronnin [Univ. Grenoble I, Associate Professor]
Brigitte Plateau [INP Grenoble, Professor, HdR]
Olivier Richard [Univ. Grenoble I, Associate Professor]
Jean-Marc Vincent [Univ. Grenoble I, Associate Professor]

Engineers

Maxime Boutserin [Inria]
Romain Cavagna [Univ. Grenoble I]
Augustin Degomme [CNRS]
Salem Harrache [Inria]
Franz Christian Heinrich [CNRS, from Dec 2014]
Michaël Mercier [Inria]
Pierre Neyron [CNRS]
Generoso Pagano [Inria, granted by OSEO Innovation]
Angelika Studeny [Univ. Paris VII]

PhD Students

Cristian Camilo Ruiz Sanabria [Inria]
Alexis Martin [Inria]
Erick Ramon Meneses Cuadros [Cifre, Orange Labs]
Luka Stanisic [CNRS]
Rodrigo Virote Kassick [UFRGS, Co-tutelle]
Francieli Zanon-Boito [UFRGS, Co-tutelle]

Post-Doctoral Fellow

Guillaume Massonnet [Inria, from Nov 2014]

Visiting Scientists

Fabien Chaix [Postdoc, NII, Japan, from Sep 2014 until Oct 2014]
Gilberto Diaz Toro [PhD Student, Bucaramanga, Colombie, Jun 2014]
Josu Doncel [PhD Student, LAAS, Sep 2014]
Rafael Keller Tesser [PhD Student, UFRGS, Brésil, Nov 2014]
Rhonda Righter [Professor, UC Berkeley, May 2014]
Wenjing Wu [Ass. Por. IHEP, Pekin, from Sep 2014 until Oct 2014]

Administrative Assistant

Annie Simon [Inria]

Others

Gaurav Agrawal [Intern, IIT Delhi, from May 2014 until Jul 2014]
Radhoine Ben Younes [Intern, UGA, from May 2014 until Aug 2014]
Geoffrey Danet [Intern, UGA, from Feb 2014 until Aug 2014]
Stéphane Durand [Intern, ENS of Lyon, until Mar 2014]
Vincent Mouly [Intern, ENS Paris, from Jun 2014 until Aug 2014]
Lucas Felix [Intern, UGA, from Feb 2014 until Jul 2014]
Jian Li [PhD Student, Austin, Texas, from May 2014 until Jul 2014]
Dineshkumar Rajagopal [Intern, UGA, from Feb 2014 until Aug 2014]
Yewan Wang [Intern, UGA, from Jun 2014 until Jul 2014]
Shaofan Zang [Intern, UGA, from Jun 2014 until Jul 2014]
Bruno Bzeznik [Univ. Grenoble I]
Joaquim Carvalho Assuncao [PhD Student, PUCRS, Brésil, from Oct 2014]
Vania Martin [Collaborator, UGA]
Jean François Mehaut [Collaborator, UGA]
Brice Videau [Collaborator, CNRS]

2. Overall Objectives

2.1. Presentation

MESCAL is a project-team of Inria jointly with UJF and Grenoble INP universities and CNRS, created in 2006 as an offspring of the former APACHE project-team, together with MOAIS.

MESCAL's research activities and objectives were evaluated by Inria in 2012. The MESCAL project-team received positive evaluations and useful feedback. The project-team was extended for another 4 years by the Inria evaluation commission.

2.2. Objectives

The recent evolutions in network and computer technology, as well as their diversification, go with a tremendous change in the use of these architectures: applications and systems can now be designed at a much larger scale than before. This scaling evolution concerns at the same time the amount of data, the number and heterogeneity of processors, the number of users, and the geographical diversity of the users.

This race towards *large scale* questions many assumptions underlying parallel and distributed algorithms as well as operating middleware. Today, most software tools developed for average size systems cannot be run on large scale systems without a significant degradation of their performances.

The goal of the MESCAL project-team is to design and validate efficient exploitation mechanisms (algorithms, middleware and system services) for large distributed infrastructures.

MESCAL's target infrastructures are grids obtained through sharing of available resources inside autonomous computing services, lightweight grids (such as the local CIMENT Grid), clusters of intranet resources (Condor) or aggregation of Internet resources (SETI@home, BOINC) as well as clouds (Amazon, Google clouds) and communication networks (5G, LTE and Wifi networks).

Application domains concern intensive scientific computations and low power high performance computing. We are also designing algorithms and middleware for SON (Self Organizing Networks) with implementations in wireless devices and base stations. Our range of applications also include the power grid (smart grids) as well as shared transportation systems.

MESCAL's methodology in order to ensure efficiency and scalability of proposed mechanisms is based on mathematical modeling and performance evaluation of the full spectrum of large scale systems from target architectures, software layers to applications.

3. Research Program

3.1. Large System Modeling and Analysis

Participants: Nicolas Gast, Bruno Gaujal, Arnaud Legrand, Panayotis Mertikopoulos, Florence Perronnin, Olivier Richard, Jean-Marc Vincent.

Markov chains, Queuing networks, Mean field approximation, Simulation, Performance evaluation, Discrete event dynamic systems.

3.1.1. Simulation of distributed systems

Since the advent of distributed computer systems, an active field of research has been the investigation of *scheduling* strategies for parallel applications. The common approach is to employ scheduling heuristics that approximate an optimal schedule. Unfortunately, it is often impossible to obtain analytical results to compare the efficiency of these heuristics. One possibility is to conduct large numbers of back-to-back experiments on real platforms. While this is possible on tightly-coupled platforms, it is unfeasible on modern distributed platforms (i.e., grids or peer-to-peer environments) as it is labor-intensive and does not enable repeatable results. The solution is to resort to *simulations*.

3.1.1.1. Flow Simulations

To make simulations of large systems efficient and trustful, we have used flow simulations (where streams of packets are abstracted into flows). SimGrid is a simulation platform that specifically targets the simulation of large distributed systems (grids, clusters, peer-to-peer systems, volunteer computing systems, clouds) from the perspective of applications. It enables to obtain repeatable results and to explore wide ranges of platform and application scenarios.

3.1.1.2. Perfect Simulation

Using a constructive representation of a Markovian queuing network based on events (often called GSMPs), we have designed perfect simulation algorithms computing samples distributed according to the stationary distribution of the Markov process with no bias. The tools based on our algorithms (ψ) can sample the stationary measure of Markov processes using directly the queuing network description. Some monotone networks with up to 10^{50} states can be handled within minutes over a regular PC.

3.1.2. Fluid models and mean field limits

When the size of systems grows very large, one may use asymptotic techniques to get a faithful estimate of their behavior. One such tool is mean field analysis and fluid limits, that can be used at a modeling and simulation level. Proving that large discrete dynamic systems can be approximated by continuous dynamics uses the theory of stochastic approximation pioneered by Michel Benaïm or population dynamics introduced by Thomas Kurtz and others. We have extended the stochastic approximation approach to take into account discontinuities in the dynamics as well as to tackle optimization issues.

Recent applications include call centers and peer to peer systems, where the mean field approach helps to get a better understanding of the behavior of the system and to solve several optimization problems. Another application concerns task brokering in desktop grids taking into account statistical features of tasks as well as of the availability of the processors. Mean field has also been applied to the performance evaluation of work stealing in large systems and to model central/local controllers as well as knitting systems.

3.1.3. Game Theory

Resources in large-scale distributed platforms (grid computing platforms, enterprise networks, peer-to-peer systems) are shared by a number of users having conflicting interests who are thus prone to act selfishly. A natural framework for studying such non-cooperative individual decision-making is game theory. In particular, game theory models the decentralized nature of decision-making.

It is well known that such non-cooperative behaviors can lead to important inefficiencies and unfairness. In other words, individual optimizations often result in global resource waste. In the context of game theory, a situation in which all users selfishly optimize their own utility is known as a *Nash equilibrium* or *Wardrop equilibrium*. In such equilibria, no user has interest in unilaterally deviating from its strategy. Such policies are thus very natural to seek in fully distributed systems and have some stability properties. However, a possible consequence is the *Braess paradox* in which the increase of resource happens at the expense of *every* user. This is why, the study of the occurrence and degree of such inefficiency is of crucial interest. Up until now, little is known about general conditions for optimality or degree of efficiency of these equilibria, in a general setting.

Many techniques have been developed to enforce some form of collaboration and improve these equilibria. In this context, it is generally prohibitive to take joint decisions so that a global optimization cannot be achieved. A possible option relies on the establishment of virtual prices, also called *shadow prices* in congestion networks. These prices ensure a rational use of resources.

Once the payoffs are fixed (using shadow prices or not), the main question is to design algorithms that allow the players to learn Nash equilibria in a distributed way, while being robust to noise and information delay as well as fast enough to outrate changing conditions of the environment.

3.2. Management of Large Architectures

Participants: Nicolas Gast, Arnaud Legrand, Olivier Richard.

Administration, Deployment, Peer-to-peer, Clusters, Grids, Clouds, Job scheduler

3.2.1. Instrumentation, analysis and prediction tools

To understand complex distributed systems, one has to provide reliable measurements together with accurate models before applying this understanding to improve system design.

Our approach for instrumentation of distributed systems (embedded systems as well as multi-core machines or distributed systems) relies on quality of service criteria. In particular, we focus on non-obtrusiveness and experimental reproducibility.

Our approach for analysis is to use statistical methods with experimental data of real systems to understand their normal or abnormal behavior. With that approach we are able to predict availability of very large systems (with more than 100,000 nodes), to design cost-aware resource management (based on mathematical modeling and performance evaluation of target architectures), and to propose several scheduling policies tailored for unreliable and shared resources.

3.2.2. Fairness in large-scale distributed systems

Large-scale distributed platforms (grid computing platforms, enterprise networks, peer-to-peer systems) result from the collaboration of many people. Thus, the scaling evolution we are facing is not only dealing with the amount of data and the number of computers but also with the number of users and the diversity of their behavior. In a high-performance computing framework, the rationale behind this joining of forces is that most users need a larger amount of resources than what they have on their own. Some only need these resources for a limited amount of time. On the opposite some others need as many resources as possible but do not have particular deadlines. Some may have mainly tightly-coupled applications while some others may have mostly embarrassingly parallel applications. The variety of user profiles makes resources sharing a challenge. However resources have to be *fairly* shared between users, otherwise users will leave the group and join another one. Large-scale systems therefore have a real need for fairness and this notion is missing from classical scheduling models.

3.2.3. Tools to operate clusters

The MESCAL project-team studies and develops a set of tools designed to help the installation and the use of a cluster of PCs. The first version had been developed for the Icluster1 platform exploitation. The main tools are a scalable tool for cloning nodes (KA-DEPLOY) and a parallel launcher based on the TAKTUK project (now developed by the MOAIS project-team). Many interesting issues have been raised by the use of the first versions among which we can mention environment deployment, robustness and batch scheduler integration. A second generation of these tools is thus under development to meet these requirements.

KA-DEPLOY has been retained as the primary deployment tool for the experimental national grid Grid'5000.

3.2.4. Simple and scalable batch scheduler for clusters and grids

Most known batch schedulers (PBS, LSF, Condor, ...) are built in a monolithic way, with the purpose of fulfilling most of the exploitation needs. This results in systems of high software complexity (150,000 lines of code for OpenPBS), offering a growing number of functions that are, most of the time, not used. In such a context, it becomes hard to control both the robustness and the scalability of the whole system.

OAR is an attempt to address these issues. Firstly, OAR is written in a very high level language (Perl) and makes intensive use of high level tools (MySQL and TAKTUK), thereby resulting in a concise code (around 5000 lines of code) easy to maintain and extend. This small code as well as the choice of widespread tools (MySQL) are essential elements that ensure a strong robustness of the system. Secondly, OAR makes use of SQL queries to perform most of its job management tasks thereby getting advantage of the strong scalability of most database management tools. Such scalability is further improved in OAR by making use of TAKTUK to manage nodes themselves.

3.3. Migration resilience; Large scale data management

Participant: Yves Denneulin.

Fault tolerance, migration, distributed algorithms.

Most propositions to improve reliability address only a given application or service. This may be due to the fact that until clusters and intranet architectures arose, it was obvious that client and server nodes were independent. This is not the case in parallel scientific computing where a fault on a node can lead to a data loss on thousands of other nodes. MESCAL's work on this topic is based on the idea that each process in a parallel application will be executed by a group of nodes instead of a single node: when the node in charge of a process fails, another in the same group can replace it in a transparent way for the application.

There are two main problems to be solved in order to achieve this objective. The first one is the ability to migrate processes of a parallel, and thus communicating, application without enforcing modifications. The second one is the ability to maintain a group structure in a completely distributed way. They both rely on a close interaction with the underlying operating systems and networks, since processes can be migrated in the middle of a communication. This can only be done by knowing how to save and replay later all ongoing communications, independently of the communication pattern. Freezing a process to restore it on another node is also an operation that requires collaboration of the operating system and a good knowledge of its internals. The other main problem (keeping a group structure) belongs to the distributed algorithms domain and is of a higher level nature.

4. Application Domains

4.1. Cloud, Grid, Multi-core and Desktop Computing

Participants: Arnaud Legrand, Olivier Richard, Jean-Marc Vincent.

Software tools were developed to carry experiments on clouds and grids (Kameleon and Expo). Other tools (Pajé, Viva, Framesoc and Ocelot) have been designed to monitor, trace and analyse applications running on multi-core and grid computers. Such traces have also been used in SIMGRID to simulate volunteer computing systems at unprecedented scale.

4.2. Wireless Networks

Participants: Bruno Gaujal, Panayotis Mertikopoulos.

MESCAL is involved in the common laboratory between Inria and Alcatel-Lucent. Bruno Gaujal is leading the Selfnets research action. This action was started in 2008 and was renewed for four more years (from 2012 to 2016). In our collaboration with Alcatel we use game theory techniques as well as evolutionary algorithms to compute optimal configurations in wireless networks (typically 3G or LTE networks) in a distributed manner. We have also been working on optimal spectrum management of MIMO systems, routing in ad-hoc works and power allocation in future 5G networks.

4.3. On-demand Geographical Maps

Participant: Jean-Marc Vincent.

This joint work involves the UMR 8504 Géographie-Cité, LIG, UMS RIATE and the Maisons de l'Homme et de la Société.

Improvements in the Web developments have opened new perspectives in interactive cartography. Nevertheless existing architectures have some problems to perform spatial analysis methods that require complex calculus over large data sets. Such a situation involves some limitations in the query capabilities and analysis methods proposed to users. The HyperCarte consortium with LIG, Géographie-cité and UMR RIATE proposes innovative solutions to these problems. Our approach deals with various areas such as spatio-temporal modeling, parallel computing and cartographic visualization that are related to spatial organizations of social phenomena.

4.4. Energy and Transportation

Participant: Nicolas Gast.

This work is mainly done within the Quanticol European project.

Smart urban transport systems and smart grids are two examples of collective adaptive systems. They consist of a large number of heterogeneous entities with decentralised control and varying degrees of complex autonomous behaviour. Within the QUANTICOL project, we develop an analysis tools to help to reason about such systems. Our work relies on tools from fluid and mean-field approximation to build decentralized algorithms that solve complex optimization problems. We focus on two problems: decentralized control of electric grids and capacity planning in vehicle-sharing systems to improve load balancing.

5. New Software and Platforms

5.1. Tools to visualize and analyze traces of execution of distributed applications

Participants: Jean-Marc Vincent [correspondent], Arnaud Legrand.

The Pajé (<http://paje.sourceforge.net/>) generic tool provides interactive and scalable behavioral visualizations of parallel and distributed applications, helping to capture the dynamics of their executions; because of its genericity, it can be used unchanged in a large variety of contexts. Pajé Next Generation

Pajé Next Generation (<https://github.com/schnorr/pajeng>) is a re-implementation (in C++) and direct heir of the well-known Paje visualization tool for the analysis of execution traces (in the Paje File Format) through trace visualization (space/time view). The tool is released under the GNU General Public License 3. PajeNG comprises the libpaje library, the space-time visualization tool in pajeng and a set of auxiliary tools to manage Paje trace files (such as `pj_dump` and `pj_validate`). It was started as part of the french INFRA-SONGS ANR project. Development has continued at INF/UFRGS. Viva

Viva (<https://github.com/schnorr/viva>) is an open-source tool used to analyze traces (in the Paje File Format) registered during the execution of parallel or distributed applications. The tool also serves as a sandbox to the development of new visualization techniques. Current features include: Temporal integration using dynamic time-intervals Spatial aggregation through hierarchical traces Interactive Graph Visualization with a force-directed algorithm, with viva Squarified Treemap to compare processes behavior on scale, with `vv_treemap`

Framesoc (<http://soctrace-inria.github.io/framesoc/>) is the core software infrastructure of the SoC-Trace project. It provides a graphical user environment for execution-trace analysis, featuring interactive analysis views as Gantt charts or statistics views. It provides also a software library to store generic trace data, play with them, and build other analysis tools (e.g., Ocelotl). This software is developed in partnership with Nanosim. Ocelotl

Ocelotl (<http://soctrace-inria.github.io/ocelotl/>): Multidimensional Overviews for Huge Trace Analysis is an innovative visualization tool, which provides overviews for execution trace analysis by using a data aggregation technique. This technique enables to find anomalies in huge traces containing up to several billions of events, while keeping a fast computation time and providing a simple representation that does not overload the user.

5.2. Simulation and performance evaluation tools

Participants: Arnaud Legrand [correspondent], Luka Stanisic, Augustin Degomme, Jean-Marc Vincent, Florence Perronnin.

5.2.1. SimGrid

(see <http://simgrid.gforge.inria.fr/>) is a toolkit that provides core functionalities for the simulation of distributed applications in heterogeneous distributed environments. The specific goal of the project is to facilitate research in the area of distributed and parallel application scheduling on distributed computing platforms ranging from simple network of workstations to Computational Grids.

5.2.2. Perfect simulator

– Ψ^2 (<https://gforge.inria.fr/projects/psi/>) is a simulation software of markovian models. It be able to simulate discrete and continuous time models to provide a perfect sampling of the stationary distribution or directly a sampling of functional of this distribution by using *coupling from the past* The simulation kernel is based on the CFTP algorithm, and the internal simulation of transitions on the Aliasing method.

5.2.3. PEPS

– The main objective of PEPS (<http://www-id.imag.fr/Logiciels/peps/>) is to facilitate the solution of large discrete event systems, in situations where classical methods fail. PEPS may be applied to the modelling of computer systems, telecommunication systems, road traffic, or manufacturing systems. Development has continued at INF/UFRGS.

5.2.4. GameSeer

(<http://mescal.imag.fr/membres/panayotis.mertikopoulos/publications.html>) is a tool for students and researchers in game theory that uses Mathematica to generate phase portraits for normal form games under a variety of (user-customizable) evolutionary dynamics. The whole point behind GameSeer is to provide a dynamic graphical interface that allows the user to employ Mathematica's vast numerical capabilities from a simple and intuitive front-end. So, even if you've never used Mathematica before, you should be able to generate fully editable and customizable portraits quickly and painlessly.

5.3. Tools for cluster management and software development

Participant: Olivier Richard [correspondent].

The KA-Tools (<http://ka-tools.imag.fr/>) is a software suite developed by MESCAL for exploitation of clusters and grids. It uses a parallelization technique based on spanning trees with a recursive starting of programs on nodes. Industrial collaborations were carried out with Mandrake, BULL, HP and Microsoft.

5.3.1. KA-Deploy

(<http://kadeploy3.gforge.inria.fr/>) is a fast and scalable deployment system for clusters and grids. It provides a set of tools for cloning, configuring (post installation) and managing a set of nodes. Currently it can successfully deploy Linux, *BSD, Windows and Solaris on x86 and 64 bits computers. Kameleon

5.3.2. Kameleon

(<http://kameleon.imag.fr/>) is a simple but powerful tool to generate customized appliances. With Kameleon, you make your recipe that describes how to create step by step your own distribution. At start Kameleon is used to create custom kvm, docker, VirtualBox, ..., but as it is designed to be very generic you can probably do a lot more than that.

5.4. Infrastructure Middleware and scheduler

Participant: Olivier Richard [correspondent].

5.4.1. OAR

– The OAR project (see <http://oar.imag.fr/>) focuses on robust and highly scalable batch scheduling for clusters and grids. Its main objectives are the validation of grid administration tools such as TAKTUK, the development of new paradigms for grid scheduling and the experimentation of various scheduling algorithms and policies.

The grid development of OAR has already started with the integration of best effort jobs whose purpose is to take advantage of idle times of the resources. Managing such jobs requires a support of the whole system from the highest level (the scheduler has to know which tasks can be canceled) down to the lowest level (the execution layer has to be able to cancel awkward jobs). OAR is perfectly suited to such developments thanks to its highly modular architecture. Moreover, this development is used for the CiGri grid middleware project.

The OAR system can also be viewed as a platform for the experimentation of new scheduling algorithms. Current developments focus on the integration of theoretical batch scheduling results into the system so that they can be validated experimentally.

5.4.2. CiGri

(<http://cigri.imag.fr/>) is a middleware which gathers the unused computing resource from intranet infrastructure and makes it available for the processing of large set of tasks. It manages the execution of large sets of parametric tasks on lightweight grid by submitting individual jobs to each batch scheduler. It is associated to the OAR resource management system (batch scheduler). Users can easily monitor and control their set of jobs through a web portal. CiGri provides mechanisms to identify job error causes, to isolate faulty components and to resubmit jobs in a safer context.

5.4.3. ComputeMode

(<http://computemode.imag.fr/>) is an software infrastructure that allows to extends or create a Grid through the aggregation of unused computing resources. For instance, a virtual cluster can be built using anyone's PC while not in use. Indeed, most PCs in large companies or university campus are idle at night, on weekends, and during vacations, training periods or business trips.

5.5. Platforms

5.5.1. Grid'5000

The MESCAL project-team is involved in development and management of Grid'5000 platform. The Digitalis and IDPot clusters are integrated in Grid'5000 as well as of CIMENT.

5.5.2. The ICluster-2, the IDPot and the new Digitalis Platforms

The MESCAL project-team manages a cluster computing center on the Grenoble campus. The center manages different architectures: a 48 bi-processors PC (ID-POT), and the center is involved with a cluster based on 110 bi-processors Itanium2 (ICluster-2) and another based on 34 bi-processor quad-core XEON (Digitalis) located at Inria. The three of them are integrated in the Grid'5000 grid platform.

More than 60 research projects in France have used the architectures, especially the 204 processors Icluster-2. Half of them have run typical numerical applications on this machine, the remainder has worked on middleware and new technology for cluster and grid computing. The Digitalis cluster is also meant to replace the Grimage platform in which the MOAIS project-team is very involved.

5.5.3. The Bull Machine

In the context of our collaboration with Bull the MESCAL project-team exploits a Novascale NUMA machine. The configuration is based on 8 Itanium II processors at 1.5 Ghz and 16 GB of RAM. This platform is mainly used by the Bull PhD students. This machine is also connected to the CIMENT Grid.

6. New Results

6.1. Simulation of distributed architectures

- Simgrid is a toolkit providing core functionalities for the simulation of distributed applications in heterogeneous distributed environments. It models fine-grain detail of the studied platform. In [6], we present quantitative results that show that SimGrid compares favorably to state-of-the-art domain-specific simulators in terms of scalability, accuracy, or the trade-off between the two. In [37], [29], we develop an hybrid approach of simulation and emulation of applications that use starPU. By using this approach, Simgrid calibrates the time to run specific subtasks at runtime and simulates all system calls of the application. This approach allows us to obtain performance results that are within one percent of measured results.
- In [33], [18], we study the problem of sampling the stationary distribution of a random walker in $\{0 \dots N\}^d$ using simulation. This algorithm combines the rejection method and coupling from the past of a set of trajectories of the Markov chain that generalizes the classical sandwich approach. We also provide a complexity analysis of this approach in several cases showing a coupling time in $O(N^2 d \log d)$ when no arc is forbidden and an experimental study of its performance.

6.2. Interactive Analysis and Visualization of Large Distributed Systems

- In [13], we review the methodology that we use to visualize information for large-scale data-set. Our approach uses tools from information theory to define a trade-off between the loss of information and the compactness of the representation. This methodology is applied to spatio-temporal representation of traces of execution in [30], [16], [17], [32]. In these papers, we show how to build a concise overview of the trace behavior as the result of a spatio-temporal data aggregation process. The experimental results show that this approach can help the quick and accurate detection of anomalies in traces containing up to two hundred million events.

- Trace analysis graphical user environments have to provide different views on trace data, to really help provide insights on the traced application behavior. In [22], [35], we propose an open and modular software architecture, the FrameSoC workbench, that defines clear principles for view engineering and for view consistency management. The FrameSoC workbench has been successfully applied in real trace analysis use-cases. This work has also been tested on real scenario coming from a collaboration with ST Microelectronic [25].
- In [7], we design a novel prediction method with Bayes model to predict a load fluctuation pattern over a long-term interval, in the context of Google data centers. All of the prediction methods are evaluated using Google trace with 10,000+ heterogeneous hosts. Experiments show that our Bayes method improves the long-term load prediction accuracy by up to 50%, compared to other state-of-the-art methods.

6.3. Management of Parallel Architectures

- In [12], we present a topology-aware load balancing algorithm for parallel multi-core machines and its proof of asymptotic convergence to an optimal solution. The algorithm, named HwTopoLB, takes into account the properties of current parallel systems composed of multi-core compute nodes, namely their network interconnection, and their complex and hierarchical core topology. We have implemented HwTopoLB using the Charm++ Parallel Runtime System and evaluated its performance with two different benchmarks and one application. Our experimental results confirm that HwTopoLB outperforms existing load balancing strategies on different multi-core systems.
- Large scale distributed systems typically comprise hundreds to millions of entities that have only a partial view of resources. How to fairly and efficiently share such resources between entities in a distributed way has thus become a critical question. In [31], we develop a possible answer based on Lagrangian optimization and distributed gradient descent. Under certain conditions, the resource sharing problem can be formulated as a global optimization problem, which can be solved by a distributed self-stabilizing demand and response algorithm.
- The management of resources on testbeds, including their description, reservation and verification, is a challenging issue, especially on large scale testbeds such as those used for research on High Performance Computing or Clouds. In [23], we present the solution designed for the Grid'5000 testbed in order to: (1) provide users with an in-depth and machine-parsable description of the testbed's resources; (2) enable multi-criteria selection and reservation of resources using a HPC resource manager; (3) ensure that the description of the resources remains accurate. In [24], we present Kascade, a solution for the broadcast of data to a large set of compute nodes. We evaluate Kascade using a set of large scale experiments in a variety of experimental settings, and show that Kascade: (1) achieves very high scalability by organizing nodes in a pipeline; (2) can almost saturate a 1 Gbit/s network, even at large scale; (3) handles failures of nodes during the transfer seamlessly because of its fault-tolerant design.

6.4. Reproducible experiments and papers

- In the field of large-scale distributed systems, experimentation is particularly difficult. The studied systems are complex, often nondeterministic and unreliable, software is plagued with bugs, whereas the experiment workflows are unclear and hard to reproduce. In [5], we provide an extensive list of features offered by general-purpose experiment management tools dedicated to distributed systems research on real platforms. We then use it to assess existing solutions and compare them, outlining possible future paths for improvements.
- Experiment reproducibility is a milestone of the scientific method. Reproducibility of experiments in computer science would bring several advantages such as code re-usability and technology transfer. The reproducibility problem in computer science has been solved partially, addressing particular classes of applications or single machine setups. In [26], we present our approach oriented to setup complex environments for experimentation, environments that require a lot of configuration and the

installation of several software packages. The main objective of our approach is to enable the exact and independent reconstruction of a given software environment and the reuse of code. We present a simple and small software appliance generator that helps an experimenter to construct a specific software stack that can be deployed on different available testbeds. [14],

- In [28], [45], we address the question of developing a lightweight and effective workflow for conducting experimental research on modern parallel computer systems in a reproducible way. Our workflow simply builds on two well-known tools (Org-mode and Git) and enables us to address issues such as provenance tracking, experimental setup reconstruction, replicable analysis. Although this workflow is perfectible and cannot be seen as a final solution, we have been using git for two years now and we have recently published a fully reproducible article, which demonstrates the effectiveness of our proposal.

6.5. Game Theory and Distributed Optimization

- In wireless networks, channel conditions of and user quality of service (QoS) requirements vary, often quite arbitrarily, with time (e.g. due to user mobility, fading, etc.) In this dynamic setting, static solution concepts (such as Nash equilibrium) are no longer relevant. Hence, we focus on the concept of no-regret : policies that perform at least as well as the best fixed transmit profile in hindsight. In [21], we examine the performance of the seminal Foschini–Miljanic (FM) power control scheme in a random environment. We provide a formulation of power control as an online optimization problem and we show that the FM dynamics lead to no regret in this dynamic context. We introduce an adjusted version of the FM algorithm which retains the convergence and no-regret properties of the original algorithm in this constrained setting. In [27], we examine the problem of cost / energy-efficient power allocation in uplink multi-carrier orthogonal frequency-division multiple access wireless networks. We use tools from stochastic convex programming to develop a learning scheme that retains its convergence properties irrespective of the magnitude of the observational errors. In [11], we consider a cognitive radio network where wireless users with multiple antennas communicate over several non-interfering frequency bands. We draw on the method of matrix exponential learning and online mirror descent techniques to derive a no-regret policy that relies only on local channel state information.
- In game theory, the best-response strategy of a player is a strategy that maximizes the selfish payoff of this player. A natural and popular question is, when players update their strategy over time, do they converge to a Nash equilibrium. In [15], we characterize the revision sets in different variants of the best response algorithm that guarantee convergence to pure Nash Equilibria in potential games. We prove that if the revision protocol is separable, then the greedy version as well as smoothed versions of the algorithm converge to pure Nash equilibria. If the revision protocol is not separable, then convergence to Nash Equilibria may fail in both cases. In [43], we investigate a class of reinforcement learning dynamics in which each player plays a "regularized best response" to a score vector consisting of his actions' cumulative payoffs. Our main results extend several properties of the replicator dynamics such as the elimination of dominated strategies, the asymptotic stability of strict Nash equilibria and the convergence of time-averaged trajectories to interior Nash equilibria in zero-sum games.

6.6. Agent-based modeling and applications to Smart Energy and Transportation Systems

- Renewable energy sources, such as wind, are characterized by non-dispatchability, high volatility, and non-perfect forecasts. Energy storage or electric loads that have a flexible consumption are viewed as a way to mitigate these effects. In [9], [19], we study centralized and distributed algorithms for solving this problem. We provide theoretical bounds on the trade-off between energy loss and the use of reserves. We develop a centralized algorithm that attains this bound in [9]. In [19], we study a distributed optimization problem by modeling a two-stage electricity market. We show that the

market is efficient: the players' selfish responses to prices coincide with a socially optimal policy. We develop a distributed solution technique based on the Alternating Direction Method of Multipliers (ADMM) and trajectorial forecasts to compute the Nash-equilibrium.

- Bike-sharing systems are becoming important for urban transportation. In these systems, users arrive at a station, pick up a bike, use it for a while, and then return it to another station of their choice. In [8], we propose a stochastic model of an homogeneous bike-sharing system and study the effect of the randomness of user choices on the number of problematic stations. Even in a homogeneous city, the system exhibits a poor performance: the minimal proportion of problematic stations is of the order of the inverse of the capacity. We show that simple incentives, such as suggesting users to return to the least loaded station among two stations, improve the situation by an exponential factor.
- In [10], we discuss the validation of an agent-based model of emergent city systems with heterogeneous agents. We transform our model into an analytically tractable discrete Markov model, and we examine the city size distribution. We show that the Markov chains lead to a power-law distribution when the ranges of migration options are randomly distributed across the agent population. We also identify sufficient conditions under which the Markov chains produce the Zipf's Law, which has never been done within a discrete framework. The conditions under which our simplified model yields the Zipf's Law are in agreement with, and thus validate, the configurations of the original heterogeneous agent-based model.

7. Bilateral Contracts and Grants with Industry

7.1. Bilateral Contracts with Industry

7.1.1. *RealTimeAtWork.com*

is a startup from Inria Lorraine created in December 2007. Bruno Gaujal is a founding partner and a scientific collaborator of the startup. Its main target is to provide software tools for solving real time constraints in embedded systems, particularly for superposition of periodic flows. Such flows are typical in automotive and avionics industries who are the privileged potential users of the technologies developed by RealTimeAtWork.com

7.1.2. *Alcatel Lucent-Bell*

A common laboratory between Inria and the Alcatel Lucent-Bell Labs was created in early 2008 and consists on three research groups (ADR). MESCAL leads the ADR on self-optimizing networks (SELFNET). The researchers involved in this project are Bruno Gaujal and Panayotis Mertikopoulos.

7.1.3. *Stimergy*

Stimergy is a startup that aims at developing a distributed data center built by connecting mini data centers embedded in digital boilers installed in multi-unit residential buildings. Each boiler contains several servers and the dissipated power can thus be used to cover a large part of the annual energy requirements for preparing domestic hot water for a building. Such infrastructure drastically reduces the energy required to operate data centers, while reducing total cost of infrastructure and ownership. Mescal (Olivier Richard, and Michael Mercier, full-time Inria engineer) provides the necessary expertise for the realization and implementation of software infrastructure allowing the coordination of operating such mini data center.

8. Partnerships and Cooperations

8.1. Regional Initiatives

8.1.1. CIMENT

The CIMENT project (Intensive Computing, Numerical Modeling and Technical Experiments, <http://ciment.ujf-grenoble.fr/>) gathers a wide scientific community involved in numerical modeling and computing (from numerical physics and chemistry to astrophysics, mechanics, bio-modeling and imaging) and the distributed computer science teams from Grenoble. Several heterogeneous distributed computing platforms were set up (from PC clusters to IBM SP or alpha workstations) each being originally dedicated to a scientific domain. More than 600 processors are available for scientific computation. The MESCAL project-team provides expert skills in high performance computing infrastructures. The members of MESCAL involved in this project are Pierre Neyron and Olivier Richard.

8.1.2. Cluster Région

Partners: the Inria GRAAL project-team, the LSR-IMAG and IN2P3-LAPP laboratories.

The MESCAL project-team is a member of the regional "cluster" project on computer science and applied mathematics, the focus of its participation is on handling large amount of data large scale architecture.

8.2. National Initiatives

8.2.1. Inria Large Scale Initiative

- *HEMERA, 2010-2014* Leading action "Completing challenging experiments on Grid'5000 (Methodology)" (see <https://www.grid5000.fr/Hemera>).

Experimental platforms like Grid'5000 or PlanetLab provide an invaluable help to the scientific community, by making it possible to run very large-scale experiments in controlled environment. However, while performing relatively simple experiments is generally easy, it has been shown that the complexity of completing more challenging experiments (involving a large number of nodes, changes to the environment to introduce heterogeneity or faults, or instrumentation of the platform to extract data during the experiment) is often underestimated.

This working group explores different complementary approaches, that are the basic building blocks for building the next level of experimentation on large scale experimental platforms.

8.2.2. ANR

- *ANR GAGA (2014-2017)*

GAGA is a "Young Researchers" project funded by the French National Research Agency (ANR) to explore the Geometric Aspects of GAMES. The GAGA teams spread over three different locations in France (Paris, Toulouse and Grenoble), and is coordinated by Vianney Perchet, assistant professor (Maître de Conférences) in the Probabilities and Random Models laboratory in Université Paris VII.

As the name suggests, our project's focus is game theory, a rapidly developing subject with growing applications in economics, social sciences, computer science, engineering, evolutionary biology, etc. As it turns out, many game theoretical topics and tools have a strong geometrical or topological flavor: the structure of a game's equilibrium set, the design of equilibrium-computing algorithms, Blackwell approachability, the geometric character of the replicator dynamics, the use of semi-algebraicity concepts in stochastic games, and many others. Accordingly, our objective is to perform a systematic study of these geometric aspects of game theory and, by so doing, to establish new links between areas that so far appeared unrelated (such as Hessian-Riemannian geometry and discrete choice theory).

- *ANR MARMOTE, 2013-2016*. Partners: Inria Sophia (MAESTRO), Inria Rocquencourt (DIOGEN), PRiSM laboratory from University of Versailles-Saint-Quentin, Telecom SudParis (SAMOVAR), University Paris-Est Créteil (*Spécification et vérification de systèmes*), Université Pierre-et-Marie-Curie/LIP6.

The project aims at realizing a software prototype dedicated to Markov chain modeling. It gathers seven teams that will develop advanced resolution algorithms and apply them to various domains (reliability, distributed systems, biology, physics, economy).

- *ANR NETLEARN, 2013-2015*. Partners: PRiSM laboratory from University of Versailles-Saint-Quentin, Telecom ParisTech, Orange Labs, LAMSADE/University Paris Dauphine, Alcatel-Lucent, Inria (MESCAL).

The main objective of the project is to propose a novel approach of distributed, scalable, dynamic and energy efficient algorithms for managing resources in a mobile network. This new approach relies on the design of an orchestration mechanism of a portfolio of algorithms. The ultimate goal of the proposed mechanism is to enhance the user experience, while at the same time to better utilize the operator resources. User mobility and new services are key elements to take into account if the operator wants to improve the user quality of experience. Future autonomous network management and control algorithms will thus have to deal with a real-time dynamicity due to user mobility and to traffic variations resulting from various usages. To achieve this goal, we focus on two central aspects of mobile networks (the management of radio resources at the Radio Access Network level and the management of the popular contents users want to get access to) and intend to design distributed learning mechanisms in non-stationary environments, as well as an orchestration mechanism that applies the best algorithms depending on the situation.

- *ANR SONGS, 2012-2015*. Partners: Inria Nancy (Algorille), Inria Sophia (MASCOTTE), Inria Bordeaux (CEPAGE, HiePACS, RunTime), Inria Lyon (AVALON), University of Strasbourg, University of Nantes.

The last decade has brought tremendous changes to the characteristics of large scale distributed computing platforms. Large grids processing terabytes of information a day and the peer-to-peer technology have become common even though understanding how to efficiently exploit such platforms still raises many challenges. As demonstrated by the USS SimGrid project funded by the ANR in 2008, simulation has proved to be a very effective approach for studying such platforms. Although even more challenging, we think the issues raised by petaflop/exaflop computers and emerging cloud infrastructures can be addressed using similar simulation methodology.

The goal of the SONGS project (Simulation of Next Generation Systems) is to extend the applicability of the SimGrid simulation framework from grids and peer-to-peer systems to clouds and high performance computation systems. Each type of large-scale computing system will be addressed through a set of use cases and led by researchers recognized as experts in this area.

Any sound study of such systems through simulations relies on the following pillars of simulation methodology: Efficient simulation kernel; Sound and validated models; Simulation analysis tools; Campaign simulation management.

8.2.3. National Organizations

Jean-Marc Vincent is member of the scientific committees of the CIST (Centre International des Sciences du Territoire).

8.3. European Initiatives

8.3.1. FP7 & H2020 Projects

8.3.1.1. Mont-Blanc

Program: FP7 Programme

Project acronym: Mont-Blanc

Project title: Mont-Blanc: European scalable and power efficient HPC platform based on low-power embedded technology

Duration: October 2011 - October 2014

Coordinator: Alex Ramirez

Other partners: BSC (Barcelone), Bull, ARM (UK), Julich (Germany), Genci, CINECA (Italy), CNRS (LIRMM, LIG)

Abstract: There is a continued need for higher compute performance: scientific grand challenges, engineering, geophysics, bioinformatics, etc. However, energy is increasingly becoming one of the most expensive resources and the dominant cost item for running a large supercomputing facility. In fact, the total energy cost of a few years of operation can almost equal the cost of the hardware infrastructure. Energy efficiency is already a primary concern for the design of any computer system and it is unanimously recognized that Exascale systems will be strongly constrained by power.

The analysis of the performance of HPC systems since 1993 shows exponential improvements at the rate of one order of magnitude every 3 years: One petaflops was achieved in 2008, one exaflops is expected in 2020. Based on a 20 MW power budget, this requires an efficiency of 50 GFLOPS/Watt. However, the current leader in energy efficiency achieves only 1.7n GFLOPS/Watt. Thus, a 30x improvement is required.

In this project, the partners believe that HPC systems developed from today's energy-efficient solutions used in embedded and mobile devices are the most likely to succeed. As of today, the CPUs of these devices are mostly designed by ARM. However, ARM processors have not been designed for HPC, and ARM chips have never used in HPC systems before, leading to a number of significant challenges.

8.3.1.2. *Mont-Blanc 2*

Program: FP7 Programme

Project acronym: Mont-Blanc 2

Project title: Mont-Blanc: European scalable and power efficient HPC platform based on low-power embedded technology

Duration: October 2013 - September 2016

Coordinator: BSC (Barcelone)

Other partners: BULL - Bull SAS (France), STMicroelectronics - (GNB SAS) (France), ARM - (United Kingdom), JUELICH - (Germany), BADW-LRZ - (Germany), USTUTT - (Germany), CINECA - (Italy), CNRS - (France), Inria - (France), CEA - (France), UNIVERSITY OF BRISTOL - (United Kingdom), ALLINEA SW LIM - (United Kingdom)

Abstract: Energy efficiency is already a primary concern for the design of any computer system and it is unanimously recognized that future Exascale systems will be strongly constrained by their power consumption. This is why the Mont-Blanc project has set itself the following objective: to design a new type of computer architecture capable of setting future global High Performance Computing (HPC) standards that will deliver Exascale performance while using 15 to 30 times less energy. Mont-Blanc 2 contributes to the development of extreme scale energy-efficient platforms, with potential for Exascale computing, addressing the challenges of massive parallelism, heterogeneous computing, and resiliency. Mont-Blanc 2 has great potential to create new market opportunities for successful EU technology, by placing embedded architectures in servers and HPC.

The Mont-Blanc 2 proposal has 4 objectives:

1. To complement the effort on the Mont-Blanc system software stack, with emphasis on programmer tools (debugger, performance analysis), system resiliency (from applications to architecture support), and ARM 64-bit support.

2. To produce a first definition of the Mont-Blanc Exascale architecture, exploring different alternatives for the compute node (from low-power mobile sockets to special-purpose high-end ARM chips), and its implications on the rest of the system.
3. To track the evolution of ARM-based systems, deploying small cluster systems to test new processors that were not available for the original Mont-Blanc prototype (both mobile processors and ARM server chips).
4. To provide continued support for the Mont-Blanc consortium, namely operations of the Mont-Blanc prototype, and hands-on support for our application developers

8.3.1.3. *QUANTICOL*

Program: The project is a member of Fundamentals of Collective Adaptive Systems (FOCAS), a FET-Proactive Initiative funded by the European Commission under FP7.

Project acronym: QUANTICOL

Project title: A Quantitative Approach to Management and Design of Collective and Adaptive Behaviours

Duration: 04 2013 – 03 2017

Coordinator: Jane Hillston (University of Edinburgh, Scotland)

Other partners: University of Edinburgh (Scotland); Istituto di Scienza e Tecnologie della Informazione (Italy); IMT Lucca (Italy) and University of Southampton (England).

Abstract: The main objective of the QUANTICOL project is the development of an innovative formal design framework that provides a specification language for collective adaptive systems (CAS) and a large variety of tool-supported, scalable analysis and verification techniques. These techniques will be based on the original combination of recent breakthroughs in stochastic process algebras and associated verification techniques, and mean field/continuous approximation and control theory. Such a design framework will provide scalable extensive support for the verification of developed models, and also enable and facilitate experimentation and discovery of new design patterns for emergent behaviour and control over spatially distributed CAS.

8.3.1.4. *NEWCOM#*

Program: FP7-ICT-318306

Project acronym: NEWCOM#

Project title: Network of Excellence in Wireless Communications

Duration: 11 2012 – 10 2015

Coordinator: Consorzio Nazionale Interuniversitario per le Telecomunicazioni (Italy)

Other partners: Aalborg Universitet (AAU). Denmark; Bilkent Üniversitesi (Bilkent). Turkey; Centre National de la Recherche Scientifique (CNRS). France; Centre Tecnològic de Telecomunicacions de Catalunya (CTTC). Spain; Institute of Accelerating Systems and Applications (IASA). Greece; Inesc Inovacao; Instituto de Novas Tecnologias (INOV). Portugal; Poznan University of Technology (PUT). Poland; Technion - Israel Institute of Technology (TECHNION). Israel; Technische Universität Dresden (TUD). Germany; University of Cambridge (UCAM). United Kingdom; Université Catholique de Louvain (UCL). Belgium; Oulun Yliopisto (UOULU). Finland

Abstract: NEWCOM# is a project funded under the umbrella of the 7th Framework Program of the European Commission (FP7-ICT-318306). NEWCOM# pursues long-term, interdisciplinary research on the most advanced aspects of wireless communications like Finding the Ultimate Limits of Communication Networks, Opportunistic and Cooperative Communications, or Energy- and Bandwidth-Efficient Communications and Networking.

8.3.2. *Collaborations in European Programs, except FP7 & H2020*

8.3.2.1. *CROWN*

Program: European Community and Greek General Secretariat for Research and Technology

Project acronym: CROWN

Project title: Optimal Control of Self Organized Wireless Networks

Duration: 2012-2015

Coordinator: Tassiulas Leandros

Other partners: Thales, University of Thessaly, National and Kapodistrian University of Athens, Athens University of Economics and Business

Abstract: Wireless networks are rapidly becoming highly complex systems with large numbers of heterogeneous devices interacting with each other, often in a harsh environment. In the absence of central control, network entities need to self-organize to reach an efficient operating state, while operating in a distributed fashion. Depending on whether the operating criteria are individual or global, nodes interact in an autonomic or coordinated way. Despite recent progress in autonomic networks, the fundamental understanding of the operational behaviour of large-scale networks is still lacking. This project will address these emergent network properties, by introducing new tools and concepts from other disciplines.

We will first analyze how imperfect network state information can be harvested and distributed efficiently through the network using machine learning techniques. We will design flexible methodologies to shape the competition between autonomous nodes for resources, with aim to maintain robust social optimality. Both cooperating and non-cooperating game-theoretic models will be used. We also consider networks with nodes coordinating to achieve a joint task, e.g., global optimization. Using algorithms inspired from statistical physics, we will address two representative paradigms in the context of wireless ad hoc networks, namely connectivity optimization and the localization of a network of primary sources from a sensor network.

Finally, we will explore delay tolerant networks as a case study of an emerging class of networks that, while sharing most of the characteristics of traditional autonomic or coordinated networks, they present unique challenges, due to the intermittency and constant fluctuations of the connectivity. We will study tradeoffs involving delay, the impact of mobility on information transfer, and the optimal usage of resources by using tools from information theory and stochastic evolution theory.

8.3.3. Collaborations with Major European Organizations

University of Athens: Panayotis Mertikopoulos was an invited professor for 3 months.

EPFL: Laboratoire pour les communications informatiques et leurs applications 2, Institut de systèmes de communication ISC, Ecole polytechnique fédérale de Lausanne (Switzerland). We collaborate with Jean-Yves Leboudec (EPFL) and Pierre Pinson (DTU) on electricity markets.

University of Antwerp: we collaborate with Benny Van Houdt on caching problems.

TU Wien: Research Group Parallel Computing, Technische Universität Wien (Austria). We collaborate with Sascha Hunold on experimental methodology and reproducibility of experiments in HPC.

8.4. International Initiatives

8.4.1. Inria International Labs

8.4.1.1. North America

- JLESC (former JLPC) (Joint Laboratory for Extreme-Scale Computing) with University of Illinois Urbana Champaign, Argonne Nat. Lab and BSC. Several members of MESCAL are partners of this laboratory, and have done several visits to Urbana-Champaign or NCSA.
- Associated Team with Berkeley. MESCAL is thus involved in the Inria@SiliconValley program.

8.4.2. Inria Associate Teams

8.4.2.1. EXASE

Title: Exascale Computing Scheduling and Energy

International Partner (Institution - Laboratory - Researcher):

Universidade Federal do Rio Grande do Sul (Brazil)

Duration: 2014 -

See also: <https://team.inria.fr/exase/>

The main scientific goal of this collaboration for the three years is the development of state-of-the-art energy-aware scheduling algorithms for exascale systems. Three complementary research directions have been identified : (1) Fundamentals for the scaling of schedulers: develop new scheduling algorithms for extreme exascale machines and use existing workloads to validate the proposed scheduling algorithms (2) Design of schedulers for large-scale infrastructures : propose energy-aware schedulers in large-scale infrastructures and develop adaptive scheduling algorithms for exascale machines (3) Tools for the analysis of large scale schedulers : develop aggregation methodologies for scheduler analysis to propose synthesized visualizations for large traces analysis and then analyze schedulers and energy traces for correlation analysis

8.4.2.2. CLOUDSHARE

Title: Guaranteed Application Performance on Idle Data Center Resources

International Partner (Institution - Laboratory - Researcher):

University of California Berkeley (United States)

Duration: 2009 - 2014

See also: <http://mescal.imag.fr/membres/derrick.kondo/ea/ea.html>

Data centers are often 85% idle as they must over-provision to ensure service level agreements. At the same time, high data center utilization is essential for efficient resource usage and optimal revenue. One way to improve utilization is for low-priority applications to use the idle resources of data centers, allowing high-priority applications to preempt them at any time. While users benefit from the lower costs of using these idle resources, parallel applications such as Map-Reduce can suffer severe overheads and unpredictable performance due to unexpected preemption and unavailability. The goal of this project is to enable complex applications to utilize idle data center resources with guaranteed performance. Our approach will be as follows. First, we will investigate novel statistical methods to predict the execution time of complex batch applications. Second, we will apply machine learning methods to predict idleness in data centers. Third, we will craft fair scheduling algorithms for multiple applications that compete for idle data center resources. The collaboration bridges experts in statistical modeling and simulation from the Inria MESCAL team with system and scheduling experts in the Berkeley BOINC team and the Google Infrastructure team.

8.4.3. Inria International Partners

8.4.3.1. Declared Inria International Partners

- MESCAL has strong connections with both UFRGS (Porto Alegre, Brazil) and USP (Sao Paulo, Brazil). The creation of the LICIA common laboratory (see next section) has made this collaboration even tighter.
- MESCAL has strong bounds with the University of Illinois Urbana Champaign, within the (Joint Laboratory on Petascale Computing, see previous section).
- MESCAL also has long lasting collaborations with University of California in Berkeley.

8.4.4. Participation In other International Programs

8.4.4.1. South America

- LICIA. The CNRS, Inria, the Universities of Grenoble, Grenoble INP and Universidade Federal do Rio Grande do Sul have created the LICIA (*Laboratoire International de Calcul intensif et d'Informatique Ambiante*). Jean-Marc Vincent is the director of the laboratory, on the French side.
The main themes are high performance computing, language processing, information representation, interfaces and visualization as well as distributed systems.
More information can be found at <http://www.inf.ufrgs.br/licia/>.

8.5. International Research Visitors

8.5.1. Visits of International Scientists

- Rhonda Righter (UC Berkeley), two weeks in May.
- Mario Bravo (University of Chile), one week in March.
- Josu Donsel (LAAS), two weeks in September.
- William H. Sandholm (University of Wisconsin), 4 days in September.
- Jian Li (Texas-A&M University) visited as a PhD intern for two months.
- Wenjing Wu (Chinese Academy of Science) (one month, Sept.- Oct.)
- Rafael Tesser (UFRGS) visited as a PhD intern for one month.
- Philippe Navaux (UFRGS), Nicolas Maillard (UFRGS) and Alexandre Carissimi (UFRGS) and Lucas Schnorr (UFRGS) visited Mescal for two weeks in Jan. and Oct. 2014.

8.5.2. Visits to International Teams

8.5.2.1. Research stays abroad

- Panayotis Mertikopoulos visited the University of Athens for one trimester to the Department of Physics and the Department of Economics (invited by Aris L. Moustakas and Andreas Polydoros).
- Panayotis Mertikopoulos visited the University of Neuchâtel for one week (Department of Mathematics, invited by Michel Benaïm).
- Panayotis Mertikopoulos visited the University of Wisconsin–Madison for one week (Department of Economics, invited by William H. Sandholm).
- Arnaud Legrand, Luka Stanisic and Augustin Degomme visited the Barcelona Supercomputer Center in November 2014.
- Jean-Marc Vincent visited UFRGS for two weeks in Feb. - Mar. 2014.

9. Dissemination

9.1. Promoting Scientific Activities

9.1.1. Scientific events organisation

9.1.1.1. general chair, scientific chair

- Panayotis Mertikopoulos was technical program co-chair of WiOpt '14: the 12-th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks.
- Bruno Gaujal was the technical program chair of EAP 2014, Sophia-Antipolis.
- Bruno Gaujal, Florence Perronnin and Jean-Marc Vincent organized a Marmote workshop, July 9-10, 2014.
- Nicolas Gast organized a Quanticol workshop, April 14, 2014.

- Arnaud Legrand organized the Summer School on Performance Metrics, Modeling and Simulation of Large HPC Systems, Sophia-Antopolis, June 2014.
- Jean-Marc Vincent organized the Licia Workshop in Grenoble, Oct. 2014.

9.1.1.2. *member of the organizing committee*

- Arnaud Legrand was a workshop organizer of the 1st workshop on reproducibility in parallel computing.
- Olivier Richard was an organizer of the 2st event “Realis 2014 : Reproductibilité expérimentale pour l’informatique en parallélisme, architecture et système” during Compas 14.

9.1.2. *Scientific events selection*

9.1.2.1. *member of the conference program committee*

- Panayotis Mertikopoulos was member of the program committee of Valuetools 2014 and NetGCoop 2014.
- Bruno Gaujal was member of the program committee of of Sigmetrics 2014, Wodes 2014 and Performance 2014.
- Nicolas Gast was member of the program comittee of E-Energy 2014 and ValueTools 2014.
- Arnaud Legrand served on the committee of Compas 2014, ICPP and HiPC.
- Jean-Marc Vincent served in the committee of ASMTA 2014, SimulTech 2014, Epew 2014.
- Olivier Richard served on the committe of Compas 2014, CARLA 2014/First HPCLATAM, Grid’5000 School, Euro-Par2014WS, HPCS-14.

9.1.3. *Journal*

9.1.3.1. *reviewer*

- Team members served as reviewers for: IEEE Transactions on Information Theory, IEEE Transactions on Automatic Control, IEEE Transactions on Signal Processing, IEEE Journal on Selected Areas in Communications, Advances in Applied Probability, Games and Economic Behavior, Journal of Economic Theory, Theoretical Economics, Journal of Optimization Theory and Applications, Mathematics of Operations Research, SIAM Journal on Control and Optimization, Automatica, CCPE, IEEE Transactions on Communication, Journal of Grid Computing.

9.2. Teaching - Supervision - Juries

9.2.1. *Teaching*

Master/PhD: Panayotis Mertikopoulos, "Game Theory for the Working Economist", 35 Eq. TD, University of Athens, Department of Economic Sciences PhD program (UADPhilEcon), Athens, Greece.

Master : Bruno Gaujal, Discrete Event Systems, 18 h, (M2R), MPRI, Paris.

Master : Bruno Gaujal, Advanced Peformance evaluation, 9 h, (M2), Ensimag, Grenoble.

Master: Arnaud Legrand, Parallel Systems, 21 h, (M2R), Mosig.

Master: Arnaud Legrand and Jean-Marc Vincent, Performance Evaluation, 15 h, (M2R), Mosig.

Master: Arnaud Legrand and Jean-Marc Vincent, Probability and simulation, performance evaluation 72 h, (M1), RICM, Polytech Grenoble.

Master: Jean-Marc Vincent, Mathematics for computer science, 18 h , (M1) Mosig.

Master: Jean-Marc Vincent, workshop on the methodology in computer science research, 15 h, (M2R) Mosig.

Master : Olivier Richard, Networking, 40 Eq. TD, (M1), RICM, Polytech Grenoble

Master : Olivier Richard, Physical Computing Eq. 60 TD, L1 and M1, Joseph Fourier University and Polytech Grenoble

DU: Jean-Marc Vincent, Informatique et sciences du numérique, 20 h, (Professeurs de lycée).

E-learning

SPOC: Jean-Marc Vincent, Informatique et sciences du numérique, 6 mois, plate-forme pairformnce, rectorat de Grenoble, Professeurs de lycée, formation continue, environ 50 inscrits.

9.2.2. Supervision

PhD : Christian Ruiz, *Méthodes et outils pour des expériences difficiles sur Grid 5000: un cas d'utilisation sur une simulation hybride en électromagnétisme* Defended Dec. 15, 2014.

9.2.3. Juries

- Panayotis was a member of the jury eximanor for the Ph.D. thesis of Tatiana Seregina: "Applications of Game Theory to Distributed Routing and Delay Tolerant Networking", University of Toulouse, November 2014.
- Bruno Gaujal served as a reviewer in the PhD thesis defense of William Mangoua Sofack (Onera), Stefano Iellamo (Telecom Paris tech) and Simon Stuker (University of Toulouse).
- Olivier Ricahrd was a member of the jury eximanor for the Ph.D. thesis of George Markomanolis : "Performance evaluation of parallel applications and performance prediction through simulation", Inria/École Normale Supérieure de Lyon, January 2014.

9.3. Popularization

- fête de la science: atelier sciences manuelles du numérique.
- MathC2+: sciences manuelles du numérique.
- Interventions dans des classes de seconde sur la simulation simcity.
- participation à l'IREM de Grenoble, groupe algorithmique.

10. Bibliography

Major publications by the team in recent years

- [1] H. CASANOVA, A. GIERSCH, A. LEGRAND, M. QUINSON, F. SUTER. *Versatile, Scalable, and Accurate Simulation of Distributed Applications and Platforms*, in "Journal of Parallel and Distributed Computing", June 2014, vol. 74, n^o 10, pp. 2899-2917 [DOI : 10.1016/j.jpdc.2014.06.008], <https://hal.inria.fr/hal-01017319>
- [2] N. GAST, B. GAUJAL, J.-Y. LE BOUDEC. *Mean field for Markov decision processes: from discrete to continuous optimization*, in "Automatic Control, IEEE Transactions on", 2012, vol. 57, n^o 9, pp. 2266–2280
- [3] B. JAVADI, D. KONDO, J.-M. VINCENT, D. P. ANDERSON. *Discovering Statistical Models of Availability in Large Distributed Systems: An Empirical Study of SETI@home*, in "IEEE Transactions on Parallel and Distributed Systems", 2010
- [4] P. MERTIKOPOULOS, E. V. BELMEGA, A. L. MOUSTAKAS, S. LASAULCE. *Distributed Learning Policies for Power Allocation in Multiple Access Channels*, in "IEEE JSAC", January 2012, vol. 30, n^o 1, pp. 96-106

Publications of the year

Articles in International Peer-Reviewed Journals

- [5] T. BUCHERT, C. RUIZ, L. NUSSBAUM, O. RICHARD. *A survey of general-purpose experiment management tools for distributed systems*, in "Future Generation Computer Systems", 2015, vol. 45, pp. 1 - 12 [DOI : 10.1016/J.FUTURE.2014.10.007], <https://hal.inria.fr/hal-01087519>
- [6] H. CASANOVA, A. GIERSCH, A. LEGRAND, M. QUINSON, F. SUTER. *Versatile, Scalable, and Accurate Simulation of Distributed Applications and Platforms*, in "Journal of Parallel and Distributed Computing", June 2014, vol. 74, n^o 10, pp. 2899-2917 [DOI : 10.1016/J.JPDC.2014.06.008], <https://hal.inria.fr/hal-01017319>
- [7] S. DI, D. KONDO, W. CIRNE. *Google hostload prediction based on Bayesian model with optimized feature combination*, in "Journal of Parallel and Distributed Computing", 2014, vol. 74, n^o 1, pp. 1820-1832 [DOI : 10.1016/J.JPDC.2013.10.001], <https://hal.inria.fr/hal-00936829>
- [8] C. FRICKER, N. GAST. *Incentives and redistribution in homogeneous bike-sharing systems with stations of finite capacity*, in "EURO Journal on Transportation and Logistics", June 2014, 31 p. [DOI : 10.1007/s13676-014-0053-5], <https://hal.inria.fr/hal-01086009>
- [9] N. GAST, D.-C. TOMOZEI, J.-Y. LE BOUDEC. *Optimal Generation and Storage Scheduling in the Presence of Renewable Forecast Uncertainties*, in "IEEE Transactions on Smart Grid", 2014, 12 p. [DOI : 10.1109/TSG.2013.2285395], <https://hal.inria.fr/hal-01086022>
- [10] B. GAUJAL, L. GULYAS, Y. MANSURI, E. THIERRY. *Validating an agent-based model of the Zipf's law: a discrete Markov chain approach*, in "Journal of Economic Dynamics & Control", April 2014, pp. 38-49 [DOI : 10.1016/J.JEDC.2014.02.002], <https://hal.inria.fr/hal-00787998>
- [11] P. MERTIKOPOULOS, E. V. BELMEGA. *Transmit without regrets: online optimization in MIMO-OFDM cognitive radio systems*, in "IEEE Journal on Selected Areas in Communications", 2014, vol. 32, n^o 11, pp. 1987-1999, 13 pages, 6 figures, <https://hal.inria.fr/hal-01073500>
- [12] L. PILLA, C. POUSA RIBEIRO, P. COUCHENEY, F. BROQUEDIS, B. GAUJAL, P. NAVAU, J.-F. MEHAUT. *A Topology-aware Load Balancing Algorithm for Clustered Hierarchical Multi-core Machines*, in "Future Generation of Computing Systems (FGCS)", 2014, vol. 30, n^o 1, pp. 191-201 [DOI : 10.1016/J.FUTURE.2013.06.023], <https://hal.inria.fr/hal-00953132>

Articles in National Peer-Reviewed Journals

- [13] R. LAMARCHE-PERRIN, L. M. SCHNORR, J.-M. VINCENT, Y. DEMAZEAU. *Agrégation de traces pour la visualisation de grands systèmes distribués*, in "Technique et Science Informatiques", 2014, forthcoming, <https://hal.inria.fr/hal-00918432>

Invited Conferences

- [14] L. NUSSBAUM, O. RICHARD. *Realis'2014: Reproductibilité expérimentale pour l'informatique en parallélisme, architecture et système*, in "ComPAS : Conférence en Parallélisme, Architecture et Système", Neuchâtel, Switzerland, April 2014, <https://hal.inria.fr/hal-01011401>

International Conferences with Proceedings

- [15] P. COUCHENEY, S. DURAND, B. GAUJAL, C. TOUATI. *General Revision Protocols in Best Response Algorithms for Potential Games*, in "Network Games, Control and Optimization (NetGCoop)", Trento, Italy, I. EXPLORE (editor), October 2014, <https://hal.inria.fr/hal-01085077>
- [16] D. DOSIMONT, R. LAMARCHE-PERRIN, L. MELLO SCHNORR, G. HUARD, J.-M. VINCENT. *A Spatiotemporal Data Aggregation Technique for Performance Analysis of Large-scale Execution Traces*, in "IEEE Cluster 2014", Madrid, Spain, September 2014, <https://hal.inria.fr/hal-01065093>
- [17] D. DOSIMONT, G. PAGANO, G. HUARD, V. MARANGOZOVA-MARTIN, J.-M. VINCENT. *Efficient Analysis Methodology for Huge Application Traces*, in "HPCS 2014 - The 2014 International Conference on High Performance Computing & Simulation", Bologna, Italy, July 2014, <https://hal.inria.fr/hal-01065783>
- [18] S. DURAND, B. GAUJAL, F. PERRONNIN, J.-M. VINCENT. *A perfect sampling algorithm of random walks with forbidden arcs*, in "QEST 2014 - 11th International Conference on Quantitative Evaluation of Systems", Florence, Italy, Springer, September 2014, vol. 8657, pp. 178-193 [DOI : 10.1007/978-3-319-10696-0_15], <https://hal.inria.fr/hal-01069975>
- [19] N. GAST, J.-Y. LE BOUDEC, D.-C. TOMOZEI. *Impact of Demand-Response on the Efficiency and Prices in Real-Time Electricity Markets*, in "ACM e-Energy 2014", Cambridge, United Kingdom, June 2014 [DOI : 10.1145/2602044.2602052], <https://hal.inria.fr/hal-01086036>
- [20] B. GIACOMO, E. V. BELMEGA, P. MERTIKOPOULOS, S. LUCA. *Energy-aware competitive link adaptation in small-cell networks*, in "WiOpt '14: Proceedings of the 12th International Symposium and Workshops on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks", Hammamet, Tunisia, 2014, <https://hal.inria.fr/hal-01073501>
- [21] S. IOANNIS, P. MERTIKOPOULOS, C. TOUATI. *No regrets: Distributed power control under time-varying channels and QoS requirements*, in "Allerton '14: Proceedings of the 51st Annual Allerton Conference on Communication, Control, and Computing", Allerton, United States, 2014, <https://hal.inria.fr/hal-01073503>
- [22] V. MARANGOZOVA-MARTIN, G. PAGANO. *The FrameSoC Software Architecture for Multiple-View Trace Data Analysis*, in "ACM SIGCHI Symposium on Engineering Interactive Computing Systems", Rome, Italy, June 2014, 6 p. [DOI : 10.1145/2607023.2610274], <https://hal.archives-ouvertes.fr/hal-01017743>
- [23] D. MARGERIE, E. MOREL, L. NUSSBAUM, O. RICHARD, C. ROHR. *Resources Description, Selection, Reservation and Verification on a Large-scale Testbed*, in "TRIDENTCOM - 9th International Conference on Testbeds and Research Infrastructures for the Development of Networks & Communities", Guangzhou, China, May 2014, <https://hal.inria.fr/hal-00965708>
- [24] S. MARTIN, T. BUCHERT, P. WILLEMET, O. RICHARD, E. JEANVOINE, L. NUSSBAUM. *Scalable and Reliable Data Broadcast with Cascade*, in "HPDIC - International Workshop on High Performance Data Intensive Computing, in conjunction with IEEE IPDPS 2014", Phoenix, United States, May 2014, <https://hal.inria.fr/hal-00957671>
- [25] A. MARTIN, G. PAGANO, J. CORRENOZ, V. MARANGOZOVA-MARTIN. *Analyse de systèmes embarqués par structuration de traces d'exécution*, in "ComPAS : Conférence en Parallélisme, Architecture et Système",

Neuchâtel, Switzerland, P. FELBER, L. PHILIPPE, E. RIVIERE, A. TISSERAND (editors), April 2014, <https://hal.inria.fr/hal-00989532>

- [26] C. C. RUIZ SANABRIA, O. RICHARD, J. EMERAS. *Reproducible Software Appliances for Experimentation*, in "TRIDENTCOM - 9th International Conference on Testbeds and Research Infrastructures for the Development of Networks & Communities (2014)", Guangzhou, China, May 2014, <https://hal.inria.fr/hal-01064825>
- [27] D. SALVATORE, P. MERTIKOPOULOS, M. ARIS L., P. SERGIO. *Adaptive transmit policies for cost-efficient power allocation in multi-carrier systems*, in "WiOpt '14: Proceedings of the 12th International Symposium and Workshops on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks", Hammamet, Tunisia, 2014, <https://hal.inria.fr/hal-01073502>
- [28] L. STANISIC, A. LEGRAND. *Effective Reproducible Research with Org-Mode and Git*, in "1st International Workshop on Reproducibility in Parallel Computing", Porto, Portugal, August 2015, <https://hal.inria.fr/hal-01083205>
- [29] L. STANISIC, S. THIBAUT, A. LEGRAND, B. VIDEAU, J.-F. MÉHAUT. *Modeling and Simulation of a Dynamic Task-Based Runtime System for Heterogeneous Multi-Core Architectures*, in "Euro-par - 20th International Conference on Parallel Processing", Porto, Portugal, Euro-Par 2014, LNCS 8632, Springer International Publishing Switzerland, August 2014, pp. 50-62, <https://hal.inria.fr/hal-01011633>

National Conferences with Proceedings

- [30] D. DOSIMONT, G. HUARD, J.-M. VINCENT. *Agrégation temporelle pour l'analyse de traces volumineuses*, in "10ème Atelier en Evaluation de Performances", Sophia Antipolis, France, June 2014, <https://hal.inria.fr/hal-01065862>

Conferences without Proceedings

- [31] A. LEGRAND. *On the Use of Lagrangian Optimization For Designing Distributed Self-Stabilizing Protocols*, in "ROADEF - 15ème congrès annuel de la Société française de recherche opérationnelle et d'aide à la décision", Bordeaux, France, Société française de recherche opérationnelle et d'aide à la décision, February 2014, Ce talk est un des trois "exposés phares" invités par Olivier Beaumont, <https://hal.archives-ouvertes.fr/hal-00946262>

Research Reports

- [32] D. DOSIMONT, L. M. SCHNORR, G. HUARD, J.-M. VINCENT. *A Trace Macroscopic Description based on Time Aggregation*, April 2014, n^o RR-8524, Trace visualization; trace analysis; trace overview; time aggregation; parallel systems; embedded systems; information theory; scientific computation; multimedia application; debugging; optimization, <https://hal.inria.fr/hal-00981020>
- [33] S. DURAND, B. GAUJAL, F. PERRONNIN, J.-M. VINCENT. *A perfect sampling algorithm of random walks with forbidden arcs*, March 2014, n^o RR-8504, 23 p. , <https://hal.inria.fr/hal-00964098>
- [34] K. GEORGIEV, V. MARANGOZOVA-MARTIN. *Deterministic Partial Replay for MPSoC Debugging*, April 2014, n^o RR-8515, 30 p. , <https://hal.inria.fr/hal-00969478>
- [35] A. MARTIN, V. MARANGOZOVA-MARTIN. *Analyse de traces d'exécutions pour les systèmes embarqués : détection d'anomalies par corrélation temporelle*, Inria, October 2014, n^o RT-0450, 20 p. , <https://hal.inria.fr/hal-01073315>

- [36] G. PAGANO, V. MARANGOZOVA-MARTIN. *FrameSoC Workbench: Facilitating Trace Analysis through a Consistent User Interface*, Inria, April 2014, n^o RT-0447, 26 p. , <https://hal.inria.fr/hal-00977887>
- [37] L. STANISIC, S. THIBAUT, A. LEGRAND, B. VIDEAU, J.-F. MÉHAUT. *Modeling and Simulation of a Dynamic Task-Based Runtime System for Heterogeneous Multi-Core Architectures*, March 2014, n^o RR-8509, <https://hal.inria.fr/hal-00966862>

Other Publications

- [38] G. BACCI, E. V. BELMEGA, P. MERTIKOPOULOS, L. SANGUINETTI. *Energy-Aware Competitive Power Allocation for Heterogeneous Networks Under QoS Constraints*, August 2014, 30, pages, 10 figures, submitted to IEEE J. Select. Areas Commun. (Special Issue on HetNets), <https://hal.inria.fr/hal-01073494>
- [39] P. COUCHENEY, B. GAUJAL, P. MERTIKOPOULOS. *Penalty-regulated dynamics and robust learning procedures in games*, April 2014, 33 pages, 3 figures, <https://hal.inria.fr/hal-01073497>
- [40] D. DOSIMONT, R. LAMARCHE-PERRIN, L. M. SCHNORR, G. HUARD, J.-M. VINCENT. *Combining Data and Visual Aggregation Techniques to Build a Coherent Spatiotemporal Overview*, October 2014, <https://hal.inria.fr/hal-01065853>
- [41] K. GEORGIEV, V. MARANGOZOVA-MARTIN. *MPSoC Zoom Debugging: A Deterministic Record-Partial Replay Approach*, June 2014, 8 p. , Accepté à EUC'2014, <https://hal.archives-ouvertes.fr/hal-01006231>
- [42] J. KWON, P. MERTIKOPOULOS. *A continuous-time approach to online optimization*, January 2014, <https://hal.inria.fr/hal-00937400>
- [43] P. MERTIKOPOULOS, W. H. SANDHOLM. *Regularized Best Responses and Reinforcement Learning in Games*, July 2014, 34 pages, 6 figures, <https://hal.inria.fr/hal-01073491>
- [44] S. PERKINS, P. MERTIKOPOULOS, S. LESLIE. *Game-theoretical control with continuous action sets*, 2014, <https://hal.inria.fr/hal-01100277>
- [45] L. STANISIC, A. LEGRAND. , P. FELBER, L. PHILIPPE, E. RIVIERE, A. TISSERAND (editors) *Writing a Reproducible Article*, April 2014, ComPAS : Conférence en Parallélisme, Architecture et Système, <https://hal.inria.fr/hal-00994575>