



IN PARTNERSHIP WITH:
CNRS

**Université Nice - Sophia
Antipolis**

Activity Report 2012

Project-Team OASIS

Active objects, semantics, Internet and
security

IN COLLABORATION WITH: Laboratoire informatique, signaux systèmes de Sophia Antipolis (I3S)

RESEARCH CENTER
Sophia Antipolis - Méditerranée

THEME
Distributed Systems and Services

Table of contents

1. Members	1
2. Overall Objectives	2
2.1. Presentation	2
2.2. Highlights of the Year	2
3. Scientific Foundations	2
3.1. Programming with distributed objects and components	2
3.2. Formal models for distributed objects	3
3.3. Verification, static analysis, and model-checking	3
4. Application Domains	4
4.1. Grid and Cloud Computing	4
4.2. Service Oriented Architectures (SOA)	4
4.3. Simulation tools and methodology	5
5. Software	5
5.1. ProActive	5
5.2. Vercors platform	6
5.3. Open Simulation Architecture (OSA)	6
5.4. BtrPlace	7
6. New Results	7
6.1. Programming and Composition Models for Large-Scale Distributed Computing	7
6.1.1. Multi-active Objects	7
6.1.2. Events for Algorithmic skeletons	7
6.1.3. Behavioural models for Distributed Components	8
6.1.4. Autonomic Monitoring and Management of Components	8
6.1.5. Optimization of data transfer in SOA and EDA models	9
6.1.6. Multi-layer component architectures	9
6.2. Middleware for Grid and Cloud computing	9
6.2.1. Publish-Subscribe in Distributed Environments	9
6.2.2. Distributed algorithms for CAN-like P2P networks	10
6.2.3. Network Aware Cloud Computing	10
6.2.4. Testbed Designs from Experimenters Requirements	11
6.2.5. Energy Efficient Virtual Machines Placement in Data Centers	11
6.2.6. GPU-based High Performance Cloud Computing	12
6.3. Large-scale Simulation Platform: Techniques and methodologies	12
7. Partnerships and Cooperations	13
7.1. National Initiatives	13
7.1.1. ANR OMD2	13
7.1.2. ANR MCorePhP	13
7.1.3. ANR Soceda	13
7.1.4. ANR Songs	13
7.1.5. CPER PacaGrid	14
7.1.6. FUI CompatibleOne	14
7.1.7. FUI CloudForce (now OpenCloudWare)	14
7.1.8. Oseo-Isis Spinnaker	15
7.2. European Initiatives	15
7.2.1.1. TEFIS	15
7.2.1.2. PLAY	16
7.2.1.3. FI-WARE	16
7.3. International Initiatives	16
7.3.1. Inria Associate Teams	16

7.3.1.1.	SCADA	16
7.3.1.2.	DAESD	17
7.3.1.3.	Dissiminet	17
7.3.2.	Inria International Partners	18
7.3.3.	Participation In International Programs	18
7.4.	International Research Visitors	18
7.4.1.	Visits of International Scientists	18
7.4.1.1.	Visiting Scientists	18
7.4.1.2.	Internships	18
7.4.2.	Visits to International Teams	19
8.	Dissemination	19
8.1.	Scientific Animation	19
8.2.	Teaching - Supervision - Juries	19
8.2.1.	Teaching	19
8.2.2.	Supervision	20
8.2.3.	Juries	20
8.3.	Popularization	21
9.	Bibliography	21

Project-Team OASIS

Keywords: Safety, Component Programming, Distributed Objects, Cloud Computing, Service Oriented Architecture, Distributed System

Creation of the Project-Team: February 01, 2000 .

1. Members

Research Scientists

Eric Madelaine [Team leader, Junior Researcher Inria, HdR]
Ludovic Henrio [Junior Researcher CNRS, HdR]

Faculty Members

Françoise Baude [Vice Team leader, Prof. UNS, HdR]
Olivier Dalle [Associate Professor UNS, since Feb. 1st]
Fabien Hermenier [Associate Professor UNS]
Fabrice Huet [Associate Professor UNS]

Engineers

Amjad Alshabani [ANR Soceda, since Oct. 1st]
Iyad Alshabani [ANR Soceda]
Nilupa Bandara [FP7 TEFIS]
Semen Boikov [FUI CompatibleOne, to Nov. 30]
Elvio Borrelli [FP7 TEFIS]
Alexandre Bourdin [FP7 PLAY, to Sep. 30]
Raffaele Di Fazio [FP7 TEFIS, from Apr. 2 to Nov. 30]
Yu Feng [FUI CompatibleOne, to July 31]
Nassim Jibai [FP7 TEFIS, from June 1st to Nov.30]
Mauricio Jost [FUI CompatibleOne, to Oct. 31]
Arthur Mbonyinshuti [OJD, from Apr 2 to Oct. 31]
Cedric Morandin [OJD, to Feb. 29]
Van Dan Nguyen [OJD, to Sep. 15]
Cristian Ruz [Oseo Spinnaker, to Oct. 31]
Bastien Sauvan [FP7 PLAY]
Alexandra Savu [FP7 TEFIS, to Sep. 30]
Mathieu Schnoor [FP7 Tefis, to Oct. 31]
Oleg Smirnov [CompatibleOne, to Mar. 31]
Fabien Viale [ANR OMD2, to July 31]

PhD Students

Maeva Antoine [EDSTIC, since Oct. 1]
Michael Benguigui [Inria-PACA + Oseo Spinnaker]
Francesco Bongiovanni [Inria-PACA + FP7 PLAY, to Jan. 31]
Yanwen Chen [Cotutelle ECNU Shanghai]
Nuno Gaspar [CIFRE ActiveEon]
Sheheryar Malik [Inria]
Laurent Pellegrino [FP7 PLAY]
Alexandra Savu [FUI CloudForce, since Oct. 1]

Post-Doctoral Fellow

Emilio Mancini [ANR SONGS, to Nov. 15th]

Administrative Assistants

Christel Kozinski [Inria, since May 1st]
Sylvie Lelaidier [Inria, to April 31th]
Magali Richir [I3S]

2. Overall Objectives

2.1. Presentation

The team focuses its activities on distributed (Grid, Cloud, and more generally large-scale infrastructures) computing and more specifically on the development of secure and reliable systems using distributed asynchronous objects (active objects - OA of OASIS). From this central point of focus, other research fields are considered in the project:

- Semantics (first S of OASIS): formal specification of active objects with the definition of ASP (Asynchronous Sequential Processes) and the study of conditions under which this calculus becomes deterministic.
- Internet (I of OASIS): Large-scale, Internet-based computing with distributed and hierarchical components.
- Security (last S of OASIS): analysis and verification of programs written in such asynchronous models.

With these objectives, our approach is:

- theoretical: we study and define models and object-oriented languages (semantic definitions, equivalences, analysis);
- applicative: we start from concrete and current problems, for which we propose technical solutions;
- pragmatic: we validate the models and solutions with full-scale experiments.

Internet clearly changed the meaning of notions like locality, mobility and security. We believe that we have the skills to be significantly fruitful in this major application domain, i.e. Internet-based computing; more specifically, we aim at producing interesting results for Grid and more recently Cloud computing, peer-to-peer systems, service-based and collaborative applications.

2.2. Highlights of the Year

Ludovic Henrio defended his HdR entitled: “Formal Models for Programming and Composing Correct Distributed Systems” in September 2012.

BEST PAPER AWARD :

[22] **How to Build a Better Testbed: Lessons From a Decade of Network Experiments on Emulab in 8th International ICST Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities, Tridentcom.** F. HERMENIER, R. RICCI.

3. Scientific Foundations

3.1. Programming with distributed objects and components

The paradigm of object-oriented programming, although not very recent, is clearly still not properly defined and implemented; for example notions like inheritance, sub-typing or overloading have as many definitions as there are different object languages. The introduction of concurrency and distribution into objects also increases the complexity. It appeared that standard Java constituents such as RMI (Remote Method Invocation) do not help building, in a transparent way, sequential, multi-threaded, or distributed applications. Indeed allowing, as RMI does, the execution of the same application to proceed on a shared-memory multiprocessors architecture as well as on a network of computing units (intranet, Internet), or on any hierarchical combination of both, is not sufficient for providing a convenient and reliable programming environment.

The question is thus: how to ease the construction (i.e. programming), deployment and evolution of distributed applications ?

One of the answers we suggest relies on the concept of active object, that acts as a single entity, abstraction of a thread, a set of objects and a location. Active objects communicate by asynchronous method calls thanks to the use of futures. ProActive is a Java library that implements this notion of active objects. ProActive can also be seen as a middleware supporting deployment, runtime support, and efficient communication for large scale distributed applications.

Another answer we provide relies on component-oriented programming. In particular, we have defined parallel and hierarchical distributed components starting from the Fractal component model developed by Inria and France-Telecom [58]. We have been involved in the design of the Grid Component Model (GCM) [4], which is one of the major results produced by the CoreGrid European Network of Excellence. The GCM has been standardized at ETSI ([64] for the last published standard), and most of our research on component models are related to it. On the practical side, ProActive/GCM is the implementation of the GCM above the ProActive programming library.

We have developed over time skills in both theoretical and applicative side fields, such as distribution, fault-tolerance, verification, etc., to provide a better programming and runtime environment for object oriented and component oriented applications.

3.2. Formal models for distributed objects

A few years ago, we designed the ASP calculus [7] for modelling distributed objects. It remains to this date one of our major scientific foundations. ASP is a calculus for distributed objects interacting using asynchronous method calls with generalized futures. Those futures naturally come with a transparent and automatic synchronisation called wait-by-necessity. In large-scale systems, our approach provides both a good structure and a strong decoupling between threads, and thus scalability. Our work on ASP provides very generic results on expressiveness and determinism, and the potential of this approach has been further demonstrated by its capacity to cope with advanced issues, such as mobility, group communications, and components [6].

ASP provides confluence and determinism properties for distributed objects. Such results should allow one to program parallel and distributed applications that behave in a deterministic manner, even if they are distributed over local or wide area networks.

The ASP calculus is a model for the ProActive library. An extension of ASP has been built to model distributed asynchronous components. A functional fragment of ASP has been modelled in the Isabelle theorem prover [8].

3.3. Verification, static analysis, and model-checking

Even with the help of high-level libraries, distributed systems are more difficult to program than classical applications. The complexity of interactions and synchronisations between remote parts of a system increases the difficulty of analysing their behaviours. Consequently, safety, security, or liveness properties are particularly difficult to ensure for these applications. Formal verification of software systems has been active for a long time, but its impact on the development methodology and tools has been slower than in the domain of hardware and circuits. This is true both at a theoretical and at a practical level; our contributions include:

- the definition of adequate models representing programs,
- the mastering of state complexity through abstraction techniques, new algorithmic approaches, or research on advanced parallel or distributed verification methods,
- the design of software tools that hide to the final user the complexity of the underlying theory.

We concentrate on the area of distributed component systems, where we get better descriptions of the structure of the system, making the analysis more tractable, but we also find out new interesting problems. For instance, we contributed to a better analysis of the interplay between the functional definition of a component and its possible runtime transformations, expressed by the various management controllers of the component system.

Our approach is bi-directional: from models to program, or back. We use techniques of static analysis and abstract interpretation to extract models from the code of distributed applications, or from dedicated specification formalisms [3]. On the other hand, we generate “safe by construction” code skeletons, from high level specifications; this guarantees the behavioural properties of the components. We then use generic tools from the verification community to check properties of these models. We concentrate on behavioural properties, expressed in terms of temporal logics (safety, liveness), of adequacy of an implementation to its specification and of correct composition of software components.

4. Application Domains

4.1. Grid and Cloud Computing

Grid, peer-to-peer, group communication, mobile object systems, Cloud, fault tolerance, distribution, security, synchronisation

As distributed systems are becoming ubiquitous, Grid computing, and the more recent concept of Cloud computing are facing a major challenge for computer science: seamless access and use of large-scale computing resources, world-wide. It is believed that by providing pervasive, dependable, consistent and inexpensive access to advanced computational capabilities, large-scale computing infrastructures allow new classes of applications to emerge.

There is a need for models and infrastructures for grid and peer-to-peer computing, and we promote programming models based on communicating mobile and distributed objects and components that can allow to harness these infrastructures. Another challenge is to use, for a given computation, unused CPU cycles of desktop computers in a Local Area Network, or even from wide area interconnected nodes. This is local or wide area Computational Peer-To-Peer, a concept that can contribute to a global energy footprint reduction. This is a challenge that also appears to be pregnant in more stable and homogeneous environments compared to P2P systems, such as datacenters under the problematics known as Virtual Machines consolidation.

4.2. Service Oriented Architectures (SOA)

Distributed Service Composition, Distributed Service Infrastructure, Peer-to-Peer data storage and lookup, Autonomic Management, Large-Scale deployment and monitoring

Service Oriented Architectures aim at the integration of distributed services and more generally at the integration of distributed and heterogeneous data, at the level of the Enterprise or of the whole Internet.

The team seeks solutions to the problems encountered here, with the underlying motivation to demonstrate the usefulness of a large-scale distributed programming approach and runtime support as featured by ProActive and GCM:

- Interaction between services: the uniform usage of web services based *client-server* invocations, through the possible support of an Enterprise Service Bus, can provide a simple interoperability between them. GCM components can be exposed as web services [61], and we have conducted research and development to permit a GCM component to invoke an external web service through a client interface. We also have provided a Service Component Architecture (see SCA specifications at <http://www.oasis-opencsa.org/sca>) personality for GCM components (GCM/SCA) so they can be integrated in SCA-based applications relying on SCA bindings configured as web services. For more loosely coupled interactions between services (e.g. compliant to the Web Services Notification standard), we pursue efforts to support *publish-subscribe* interaction models. Scalability in terms of number of notified events per time unit, and full interoperability through the use of semantic web notations applied to these events are some of the key challenges the community is addressing and we too. Events also correspond to data that may be worth to store, for future analytics, besides being propagated to interested parties (in the form of the event content). Our research can thus also contribute to the Big Data domain, a currently hot topic in ICT.

- Services compositions on a possibly large set of machines: if service compositions can even be turned as autonomic activities, these capabilities will really make SOA ready for the Open Internet scale (because at such a scale, a global management of all services is not possible). For service compositions represented as GCM-based component assemblies, we are indeed exploring the use of control components put in the components membranes, acting as sensors or actuators, that can drive the self-deployment and self-management of composite services, according to negotiated Service Level Agreements. For service orchestrations usually expressed as BPEL like processes, and expressing the *composition in time* aspect of the composition of services, supports for deployment, management, and execution capable to support dynamic adaptations are also needed. Here again we believe a middleware based upon distributed and autonomous components as GCM can really help.

4.3. Simulation tools and methodology

simulation, component-based modeling, parallel and distributed simulation, reproducibility, architecture description language

Components are being used in simulation since many years. However, given its many application fields and its high computation needs, simulation is still a challenging application for component-based programming techniques and tools.

We have been exploring the application of Oasis programming methods to simulation problems in various areas of engineering problems, but also of financial applications.

More recently, with the arrival of O. Dalle in the team, and following a work previously started in the Mascotte project-team in 2006 [59], we are pursuing research on applying distributed component-based programming techniques to simulation. More precisely, new results are sought in three directions:

- improvement of simulation tools and methodology with techniques and tools borrowed from latest research in component-based software engineering;
- improvement of simulation scalability using high performance and distributed computing techniques;
- application of the results in the previous directions, in particular to the simulation of very large-scale distributed systems, such as peer-to-peer networks.

5. Software

5.1. ProActive

Participants: F. Baude, D. Caromel, L. Henrio, F. Huet [correspondant], F. Viale, O. Smirnov, B. Sauvan, A. Bourdin.

Proactive Parallel Suite

ProActive is a Java library (Source code under AGPL license) for parallel, distributed, and concurrent computing, also featuring mobility and security in a uniform framework. With a reduced set of simple primitives, ProActive provides a comprehensive API to simplify the programming of applications that are distributed on a Local Area Network (LAN), on cluster of workstations, Clouds, or on Internet Grids.

The library is based on an Active Object pattern that is a uniform way to encapsulate:

- a remotely accessible object,
- a thread,
- an actor with its own script,
- a server of incoming requests,
- a mobile and potentially secure agent.

and has an architecture to inter-operate with (de facto) standards such as:

- Web Service exportation (Apache Axis2 and CXF),
- HTTP transport,
- ssh, rsh, RMI/ssh tunnelling,
- Globus: GT2, GT3, GT4, gsi, Unicore, ARC (NordGrid)
- LSF, PBS, Sun Grid Engine, OAR, Load Leveler

ProActive is only made of standard Java classes, and requires **no changes to the Java Virtual Machine**, no preprocessing or compiler modification; programmers write standard Java code. Based on a simple Meta-Object Protocol, the library is itself extensible, making the system open for adaptations and optimisations. ProActive currently uses the RMI Java standard library as default portable transport layer, but others such as Ibis or HTTP can be used instead, in an adaptive way.

ProActive is particularly well-adapted for the development of applications distributed over the Internet, thanks to reuse of sequential code, through polymorphism, automatic future-based synchronisations, migration of activities from one virtual machine to another. The underlying programming model is thus innovative compared to, for instance, the well established MPI programming model.

In order to cope with the requirements of large-scale distributed and heterogeneous systems like the Grid, many features have been incorporated into ProActive, including support for many transport and job submission protocols, GCM component support, graphical visualization interface, object migration, distributed and non-functional exception handling, fault-tolerance and checkpointing mechanisms; file transfer capabilities, a job scheduler, a resource manager able to manage various hosting machines, support for JMX and OSGi capabilities, web service object exposition, an SCA personality, etc.

ProActive is a project of the former ObjectWeb, now OW2 Consortium. OW2 is an international consortium fostering the development of open-source middleware for cutting-edge applications: EAI, e-business, clustering, grid computing, managed services and more. For more information, refer to [5] [55] and to the web pages <http://www.objectweb.org> and <http://proactive.inria.fr/> which list several white papers.

ProActive management, distribution, support, and commercialisation is now ensured by the start-up company ActiveEon (<http://www.activeeon.com>), in the context of a collaboration with Inria and UNS.

5.2. Vercors platform

Participants: E. Madelaine, L. Henrio, A. Savu, M. Alexe.

The Vercors tools (<http://www-sop.inria.fr/oasis/Vercors>) include front-ends for specifying the architecture and behaviour of components in the form of UML diagrams. We translate these high-level specifications, into behavioural models in various formats, and we also transform these models using abstractions. In a final step, abstract models are translated into the input format for various verification toolsets. Currently we mainly use the various analysis modules of the CADP toolset.

- We have pursued last year experiments in distributed model-checking, and were able to generate explicit state-spaces of (sub-systems) for a new distributed use-case of several billion states. The challenges here lie in the structure of the verification workflow, and in finding strategies for separating the sub-systems in an intelligent way.
- We have also conducted intensive experiments within the Papyrus environment, aiming at the definition of a graphical specification environment combining some of the standard UML formalisms (typically class diagrams and state-machines), with a dedicated graphical formalism for the architecture of GCM components.

5.3. Open Simulation Architecture (OSA)

Participants: O. Dalle, V.D. Nguyen.

OSA stands for Open Simulation Architecture. OSA is primarily intended to be a federating platform for the simulation community: it is designed to favor the integration of new or existing contributions at every level of its architecture. The platform core supports discrete-event simulation engine(s) built on top of the ObjectWeb Consortium's Fractal component model. In OSA, the systems to be simulated are modeled and instrumented using Fractal components. In OSA, the event handling is mostly hidden in the controller part of the components, which alleviates noticeably the modeling process, but also eases the replacement of any part of the simulation engine. Apart the simulation engine, OSA aims at integrating useful tools for modeling, developing, experimenting, and analysing simulations. OSA is also a platform for experimenting new techniques and approaches in simulation, such as aspect oriented programming, separation of concerns, innovative component architectures, and so on.

5.4. BtrPlace

Participant: F. Hermenier.

Btrplace (<http://btrp.inria.fr>) is an open source virtual machine (VM) placement algorithm for datacenters. BtrPlace has been designed to be extensible. It can be customized by plugins from third party developers to address new SLAs or optimization objectives. Its extensibility is possible thanks to a composable core reconfiguration algorithm implemented using Constraint Programming.

Btrplace is a part of the OW2 project Entropy. It has been originally developed by Fabien Hermenier during its PhD at the Ecole des Mines of Nantes. BtrPlace is now a standalone project that is currently used to address fault tolerance, isolation, infrastructure management, performance, and energy efficiency concerns inside the national project OpenCloudWare (<http://opencloudware.org/>) and the European project Fit4Green.

This year, our development has been guided by our collaborations. The Fit4Green project chose to rely on BtrPlace to compute an energy-efficient placement for their VMs while some partners inside OpenCloudWare required new placement constraints. The inferring capabilities of BtrPlace and its catalog of placement constraint have then been upgraded accordingly.

6. New Results

6.1. Programming and Composition Models for Large-Scale Distributed Computing

6.1.1. Multi-active Objects

Participants: L. Henrio, F. Huet, A. Bourdin.

The active object programming model is particularly adapted to easily program distributed objects: it separates objects into several *activities*, each manipulated by a single thread, preventing data races. However, this programming model has its limitations in terms of expressiveness – risk of deadlocks – and of efficiency on multicore machines. We proposed to extend active objects with *local multi-threading*. We rely on declarative *annotations* for expressing potential concurrency between requests, allowing easy and high-level expression of concurrency. This year we realized the following:

- improvement on the model and its formalisation
- use of the new model in our CAN P2P network (see below); this was also the opportunity to improve our implementation.

This year, we also spent considerable efforts to publish this work; a conference paper is currently under review.

6.1.2. Events for Algorithmic skeletons

Participant: L. Henrio.

In the context of the SCADA associated team, we worked on the algorithmic skeleton programming model. The structured parallelism approach (skeletons) takes advantage of common patterns used in parallel and distributed applications. The skeleton paradigm separates concerns: the distribution aspect can be considered separately from the functional aspect of an application.

- This year we focused on the handling of events in algorithmic skeletons: adding the possibility for a skeleton to output an event should increase the control and monitoring capabilities of algorithmic skeletons. The ultimate goal is to improve autonomy for algorithmic skeletons.

6.1.3. Behavioural models for Distributed Components

Participants: E. Madelaine, N. Gaspar, A. Savu, L. Henrio.

In the past [3], we defined the behavioural semantics of active objects and components. This year we extended this work to address group communications. On the practical side, this work contributes to the Vercors platform; the overall picture being to provide tools to the programmer for defining his application, including its behavioural specification. Then some generic properties like absence of deadlocks, but also application specific properties can be validated on the composed model using an existing model-checker. We mainly use the CADP model-checker, that also supports distributed generation of state-space. This year our main achievements are the following:

- We entirely formalised the specification of the behavioural model generation for component systems. This should provide us both a stronger formal background for our works in this area, and a specification for the automatic generation of behavioural models for our component systems.
- We additionally have put considerable efforts on the improvement of the Vercors platform and its integration with the Papyrus framework (see Section 5.2).

The formal work has been published as a research report [40]. A journal version is under submission. This work was done in collaboration with Rabéa Ameur-Boulifa from Télécom-Paristech.

In parallel with core developments of the behavioural specification environment, our collaborations led us to the study of the following application domain. In the context of the Spinnaker project, we are interested in developing a component-based distributed application to manage and monitor some pre-existing component-based distributed application - and hence, we called it The HyperManager. Our in-house component model (GCM) provides all the means to define, compose and dynamically reconfigure such applications. However, special care must be taken for this kind of undertaking. To this end, this year:

- We made the first steps towards a platform for the mechanized specification and verification, in the Coq Proof Assistant, of GCM applications. This work was published in [33], and is progressively being updated ¹ to cope with behavioural specification, and to seamlessly combine deductive and model-checking techniques.
- We studied a real-life application scenario for our HyperManager prototype using distributed model-checking techniques in order to cope with the huge space state generated from reconfigurable applications.

6.1.4. Autonomic Monitoring and Management of Components

Participants: F. Baude, C. Ruz, B. Sauvan.

We have completed the design of a framework for autonomic monitoring and management of component-based applications. We have provided an implementation using GCM/ProActive taking advantage of the possibility of adding components in the membrane. For this purpose, we finalized the implementation of a factory which, from any GCM ADL description can instantiate the requested non functional components of a GCM application.

¹<http://www-sop.inria.fr/members/Nuno.Gaspar/Mefresa.php>

The framework for autonomic computing allows the designer to describe in a separate way each phase of the MAPE autonomic control loop (Monitoring, Analysis, Planning, and Execution), and to plug them or unplug them dynamically. We have demonstrated how such a control loop can be relevant to drive the dynamic reconfiguration of services part of a SOA application, considering as in the SCA standard, that services are components [15].

Our objective now is to exemplify such autonomic and structured approach in the management of any distributed middleware or application, e.g. in the Spinnaker industrial context.

6.1.5. *Optimization of data transfer in SOA and EDA models*

Participants: I. Alshabani, F. Baude, L. Pellegrino, B. Sauvan, Q. Zagarese.

Traditional client-server interactions rely upon method invocations with copy of the parameters. This can be useless in particular if the receiver does not effectively uses them. On the contrary, copying and transferring parameters lazily so to allow the receiver to proceed even without all of them is a meaningful idea that we proved to be effective for active objects in the past [56]. This idea wasn't so far realized in the context of the web services technology, the most popular one used today for client-server SOAP-based interactions.

- To such an aim, we contributed to the offloading of objects representing parameters of the web service Java Apache CXF API [29]. It is innovative notably in the way the offloading of parameters for on-demand access can be delegated from services to services, which resembles the concept of first-class futures.
- Relying upon such an effective approach, we have applied a similar idea of “lazy copying and transfer” to the data parts of events in the context of event-driven architecture applications [28]. The middleware dynamically off-loads data (generally of huge size) attached to an event, according to some user-level policy expressed as annotation in the Java code at the subscriber side. The event itself, without its attachments, gets forwarded into the publish/subscribe brokering system (in our case, the event cloud middleware, that is the subject of section 6.2.1) and its attachments are transferred to the subscriber only on-demand. Compared to some existing propositions geared towards a data centric publish-subscribe pattern (e.g. the DDS OMG standard), ours is more user-friendly as it does not require the user code to explicitly program when to get the data attached to notified events.

Overall, this work opens the way towards a strong convergence between service oriented and event-driven technologies.

6.1.6. *Multi-layer component architectures*

Participant: O. Dalle.

Since a few years, we have been investigating the decomposition of a simulation application into multiple layers corresponding to the various concerns commonly found in a simulation: in addition to the various modeling domains that may be found in a single simulation application (e.g. telecommunications networks, road-networks, power-grids, and so on), a typical simulation includes various orthogonal concerns such as system modelling, simulation scenario, instrumentation and observation, distribution, and so on. This large number of concerns has put in light some limits of the traditional hierarchical component-based architectures and their associated ADL, as found in the FCM and GCM. In order to cope with these limitations, we started a new component architecture model called Binding Layers centered on the binding rather than the component, with no hierarchy but advanced layering capabilities, and offering advanced support for dynamic structures[32].

6.2. Middleware for Grid and Cloud computing

6.2.1. *Publish-Subscribe in Distributed Environments*

Participants: F. Baude, F. Huet, F. Bongiovanni, L. Pellegrino, B. Sauvan, I. Alshabani, A. Bourdin, M. Antoine, A. Alshabani.

In the context of the SOA4ALL FP7-IP project, we designed and implemented a hierarchical Semantic Space infrastructure based on Structured Overlay Networks (SONS) [62], [63]. It originally aimed at the storage and the retrieval of the semantic description of services at the Web scale [57]. This infrastructure combines the strengths of both the P2P paradigm at the architectural level and the Resource Description Framework (RDF) data model at the knowledge representation level. The achievements of this year are the following:

- In the context of the FP7 Strep PLAY and French ANR SocEDA research projects, we have been extending the aforementioned work with a content-based Publish/Subscribe abstraction in order to support asynchronous queries for RDF-based events in large scale settings, which raises some interesting challenges [26]. The goal is to build a platform for large scale distributed reasoning[25]. Such an integrated working platform [39], [38] has been presented in two tutorials [27], [54].
- We have also investigated the Publish/Subscribe paradigm in the MapReduce programming model. We have proposed the concept of continuous job which allows MapReduce jobs to be re-executed when new data are added to the system. To maintain the correctness of the execution, we have introduced the notion of carried data, i.e. data which are kept between subsequent executions. An implementation has been written on top of Hadoop and a paper submitted.

6.2.2. *Distributed algorithms for CAN-like P2P networks*

Participants: L. Henrio, F. Bongiovanni, F. Huet.

The nature of some large-scale applications, such as content delivery systems or publish/subscribe systems, built on top of SONS, demands application-level dissemination primitives which do not overwhelm the overlay, i.e. efficient, and which are also reliable. Building such communication primitives in a reliable manner on top of such networks would increase the confidence regarding their behavior prior to deploying them in real settings. In order to come up with real efficient primitives, we take advantage of the underlying geometric topology of the overlay network and we also model the way peers communicate with one another. Our objective is to design and prove an efficient and reliable broadcast algorithm for CAN-like P2P networks. To this aim, in 2012 we:

- Improved the formalisation in Isabelle/HOL of a CAN-like P2P system, devised formalised tools to reason on CAN topologies, and on communication protocols on top of CANs. We designed and proved the efficiency of a first naive algorithm.
- Sketched on paper the proof of completeness and efficiency for the algorithm we designed and implemented last year.

Part of this work was done in the PhD thesis of F. Bongiovanni [10]

We are also investigating the new algorithms to efficiently build a SONS in the presence of existing data. Most of the work on SONS assume that new peers joining the network will arrive without data or fail to take into account the cost of distributing these data. Indeed, depending on the key subspace given to the new peer, some or all its data will have to be distributed in the network. In 2012:

- We proposed a first version of new join algorithms which try to allocate key sub-spaces to peers so that the amount of data that needs to be moved is minimal. An expected benefit of this work is that it should allow for fast and efficient reconstruction of a SON in case of a crash, without having to use distributed snapshots.

6.2.3. *Network Aware Cloud Computing*

Participants: S. Malik, F. Huet.

We have worked on the Resource Aware Cloud Computing project. Its primary purpose is to address different issues which can help the scheduler to make more efficient scheduling decisions. These issues are related to the resource characteristics.

- We introduce a framework, which increases the performance of the application and ensures high level of reliability during the scheduling of application onto the cloud. It is a cloud scheduler module named as Resource Aware Cloud Scheduling (RACS) module. It helps the scheduler in making the scheduling decisions on the basis of different characteristics of cloud resources. These characteristics are reliability, network latency, and monetary cost. RACS consists of multiple sub modules, which are responsible for their corresponding tasks. In RACS, we have done the implementation for the different issues.
- We worked on a model for the reliability assessment of the cloud's computing nodes. This reliability assessment mechanism helps to do the scheduling on cloud infrastructure and perform fault tolerance on the basis of the reliability values acquired during reliability assessment. The model has different algorithms for different types of applications. Thus it has multiple reliability values for each computing node. For real time applications, the model has time based reliability assessment algorithms.

This work is part of S. Malik's PhD thesis [12]

6.2.4. *Testbed Designs from Experimenters Requirements*

Participant: F. Hermenier.

The physical design of the Emulab facility, and many other testbeds like it, has been based on the facility operators' expectations regarding user needs and behavior. If operators' assumptions are incorrect, the resulting facility can exhibit inefficient use patterns and sub-optimal resource allocation.

- We have collaborated with Robert Ricci from the University of Utah on the study of the Utah' Emulab facility to provide better testbed designs. Our study gained insight into the needs and behaviors of networking researchers by analyzing more than 500,000 topologies from 13,000 experiments submitted to Emulab.
- Using this dataset, we re-visited the assumptions that went into the physical design of the Emulab facility and considered improvements to it. Through extensive simulations with real workloads, we evaluated alternative testbeds designs for their ability to improve testbed utilization and reduce hardware costs.

The results have been published to TridentCom [22], the reference conference related to testbeds and research infrastructures, and the article received the best paper award.

6.2.5. *Energy Efficient Virtual Machines Placement in Data Centers*

Participant: F. Hermenier.

Data centres are powerful ICT facilities which constantly evolve in size, complexity, and energy consumption. At the same time, tenants' and operators' requirements become more and more complex. The data centre operators may target different energy-related objectives while the workload volatility may alter the data centre capacity at supporting load spikes. Finally, clients of data centres are looking for dependable infrastructures that can comply with their SLA requirements.

To stay attractive, a data centre should then support these expectations. These constraints are however very specific to each of the tenants but also to the infrastructure. They also cover a large range of concerns (hardware requirements, performance, security ...) that are continuously evolving according to new trends and new technologies. Existing solutions are however ad-hoc and can not be updated easily to fit the data centres and the workload specificities.

We proposed a flexible energy-aware framework to address the multiple facets of an energy-aware consolidation of VMs in a cloud data centre.[21] This framework extended BtrPlace to make it able to address specific energy concerns. We integrated a fine grain energy model reducing either gas emissions or power consumption. We also proposed constraints to control the aggressiveness of these objectives to let the data centre reactive when a load spike occurs. We finally proposed various constraints to satisfy the hardware and the resource requirements of the tenants. The evaluation on a testbed running an industrial workload validated the practical benefits provided by the usage of our framework.

6.2.6. GPU-based High Performance Cloud Computing

Participants: M. Benguigui, F. Baude, F. Huet.

To address HPC, GPU devices are now considered as unavoidable cheap, energy efficient and very efficient alternative computing units. The barrier to handle such devices is the programming model: it is both very fine grained and synchronous.

Our long term goal is to devise some generic solutions in order to incorporate GPU-specific code whenever relevant into a parallel and distributed computation. The first step towards this objective is to gain some insight on how to efficiently program a non trivial but well known algorithm. We selected the American basked option pricing non embarrassingly parallel problem that was previously parallelized and distributed using ProActive master-slave approach [60], achieving an almost linear speedup and good performances (64 CPUs based computation allowed us to solve the problem in about 8 hours). The same algorithm has been reorganized for running on **a single GPU** [17] and achieved the same option pricing computation in about 9 hours. The current work is to succeed to take advantage of GPUs, even if non homogeneous, hired from a Cloud or a federation of clouds at once, orchestrated by an active object acting as a GPU task delegator. The goal is to drastically lower the overall computation time for such highly time consuming stochastic simulation problems.

6.3. Large-scale Simulation Platform: Techniques and methodologies

Participants: O. Dalle, E. Mancini.

In the domain of simulation techniques and methodologies, this year, we conducted research in the three following areas:

Simulation in the Cloud In recent years, numerous applications have been deployed into mobile devices. However, until now, there have been no attempts to run simulations on handheld devices. In the context of the DISSIMINET Associate Team, we work in collaboration with our partners at the Carleton University to investigate different architectures for running and managing simulations on handheld devices, and putting the simulation services in the Cloud[24]. We propose a hybrid simulation and visualization approach, where a dedicated mobile application is running on the client side and the RISE simulation server is hosted in the Cloud.

Simulation Methodology In the context of the ANR INFRA SONGS project, we are involved (as coordinators) in a Work-Package called “Open Science” whose aim is to investigate and contribute means to ensure the long term visibility and reproductibility of simulation results obtained using the SimGrid simulation platform. Our preliminary work in this direction consisted in identifying the issues, trends and potential solutions to ensure the long-term reproducibility of simulations[16].

Peer-to-peer Simulation In order to evaluate the performance and estimate the resource usage of peer-to-peer backup systems, it is important to analyze the time they spend in storing, retrieving and keeping the redundancy of the stored files. The analysis of such systems is difficult due to the random behavior of the peers and the variations of network conditions. In the context of the ANR USS-SIMGRID and INFRA-SONGS projects, we investigated means for reproducing such varying conditions in a controlled way. We worked on the design of a general simulation meta-model for peer-to-peer backup systems and a tool-chain, based on SimGrid, to help in their analysis[20]. We validated the meta-model and tool-chain through the analysis of a common scenario, and verified that they can be used, for example, for retrieving the relations between the storage size, the saved data fragment sizes and the induced network workload. We also started to investigate a new simulation technique for very-large scale distributed simulation of peer-to-peer systems based on the decomposition of a simulation into many micro-simulation steps[31] in order to optimize the overlap between communications and computations.

7. Partnerships and Cooperations

7.1. National Initiatives

7.1.1. ANR OMD2

Title: MultiDisciplinary Distributed Optimization

Program: Conception and Simulation 2008

Duration: July 2009 - September 2012

Coordinator: Renault

Others partners: SMEs: CD-adapco, SIREHNA, ACTIVEEON, academics: Inria, ENSM-SE, UTC, ECP, IRCCyN, ENS CACHAN, and consortium DIGITEO.

See also: <http://omd2.scilab.org/>

Abstract: OMD2 (MultiDisciplinary Distributed Optimization) is a national research project led by Renault and gathering several academics and industrial partners which aims at developing methods and tools to generalize the use of optimization on large scale engineering problems. Scilab is the chosen generic programming tool to gather the different developments in a unique optimization environment. ProActive Parallel Suite is used to execute the Workflows in parallel, and to manage the Grid and Cloud resources

7.1.2. ANR MCorePhP

Title: Multi-Core Parallel Heterogeneous Programming

Program: Blanc international

Duration: January 2010 - December 2012

Coordinator: Inria Oasis

Others partners: Tsinghua University Beijing (China)

See also: http://www-sop.inria.fr/oasis/mcorephp_home.htm

Abstract: McorePhP is dedicated to programming models and middleware for large-scale, multilevel infrastructures including multi-core, clusters, and large scale grid/cloud resources. We will ensure the compatibility of the new programming model with the China Grid specifications, and will assess the viability and efficiency of the approach on a large example from the area of bioinformatics.

7.1.3. ANR Soceda

Title: SOCial Event Driven Architecture

Program: Platform

Duration: July 2009 - September 2012

Coordinator: Linagora (ex EBM Web Sourcing)

Others partners: SMEs: ACTIVEEON, industry: Thales, OrangeLabs, academics: Inria, CNRS IMAG, LIRIS, ARMINES

See also: <http://www.soceda.org/display/soceda/>

Abstract: SocEDA is an ANR project of type Platform, also labelled by two competitiveness clusters, PEGASE and SCS. The aim is to provide a "Cloud based platform for large scale social aware Event-Driven Architecture (EDA)". OASIS is in charge of managing the storage and publication/subscription of events on the cloud.

7.1.4. ANR Songs

Title: Simulation of Next Generation Systems

Program: Infra 13

Duration: January 2012 - December 2015

Coordinator: Inria (Nancy, Grenoble, Bordeaux)

Others partners: IN2P3 Villeurbanne, LSIIT Strasbourg, I3S Sophia-Antipolis, LINA Nantes

See also: <http://infra-songs.gforge.inria.fr/>

Abstract: SONGS (2012-2015) is the continuity of SIMGRID project (2009-2012), in the ANR INFRA program. The aim of SONGS is to continue the development of the SimGrid simulation platform for the study of large distributed architectures, including data grids, cloud computing facilities, peer-to-peer applications and HPC/exascale architectures.

7.1.5. CPER PacaGrid

Duration: January 2010 - December 2012

See also: <http://www-sop.inria.fr/oasis/pacagrid/>

Abstract: ProActive PacaGrid is a set of machines deployed at Inria Sophia Antipolis (1400 cores, 150 TB storage) accessible via Graphical Interactive interfaces based on ProActive Parallel Suite. This Grid is available for Inria, UNS, and PACA (regional) labs, as well as for SMEs for R&D purpose, and international partners in R&D projects. It has been funded by EU FEDER, PACA and Alpes Maritimes Landers, and EIT ICT Labs (about 1.7 MEuros in total). Users include for instance INRA (Institut de recherche en Agronomie), IPMC INSERM (Institut de Pharmacologie Moléculaire et Cellulaire), LCMBA (Laboratoire de Chimie des Molécules Bioactives et des Arômes), IGS (Laboratoire Information Génomique et Structurale, Marseille), LIFM (Laboratoire d'Informatique Fondamentale de Marseille), K-Epsilon SME, Renault, Sirehna DCNS, Poznań Supercomputing and Networking Center (Poland), National University of Singapore.

7.1.6. FUI CompatibleOne

Title: The Open Source Cloud Broker

Program: Conception and Simulation 2008

Duration: July 2009 - September 2012

Coordinator: OW2

Others partners: industry: ActiveEon, Bull, CityPassenger, eNovance, Eureva, Mandriva, Nexedi, Nuxeo, XWiki, Prologue; academic: Inria, Institut Telecom

See also: <http://www.compatibleone.org>

Abstract: CompatibleOne is an open source project which provides a model, CORDS (CompatibleOne Resource Description System), and a platform, ACCORDS (Advanced Capabilities for CORDS), for the description and federation of different clouds comprising resources provisioned by heterogeneous cloud service providers. CompatibleOne's flexible service architecture makes it independent from any Cloud Service Provider (from OpenStack to OpenNebula, from Azure to Vcloud) and can address all types of cloud services (IaaS, PaaS, SaaS, XaaS, BaaS, ...) and any type of cloud service deployment (public, private, community and hybrid).

7.1.7. FUI CloudForce (now OpenCloudWare)

Program: FSN, labelled by Minalogic, Systematic and SCS.

Duration: January 2012 - December 2014

Coordinator: France-Telecom Research

Others partners: ActiveEon, Armines, Bull, eNovance, eXo Platform, France Telecom (coordinator), Inria, IRIT – INP Toulouse, Linagora, OW2, Peergreen, Télécom Paris Tech, Télécom Saint Etienne, Thales Communications, Thales Services, Université Joseph Fourier, Université de Savoie – LISTIC, UShareSoft

See also: <http://www.opencloudware.org/>

Abstract: The OpenCloudware project aims at building an open software engineering platform, for the collaborative development of distributed applications to be deployed on multiple Cloud infrastructures.

The results of OpenCloudware will contain a set of software components to manage the lifecycle of such applications, from modelling (Think), developing and building images (Build), to a multi-IaaS compliant PaaS platform (Run) for their deployment, orchestration, performance testing, self-management (elasticity, green IT optimisation) and provisioning. Applications will be deployed potentially on multi IaaS (supporting either one IaaS at a time, or hybrid scenarios). The results of the project will be made available as open source components through the OW2 Open Source Cloudware initiative.

7.1.8. Oseo-Isis Spinnaker

Duration: June 2011 - May 2014

Coordinator: Tagsys-RFID

Others partners: SMEs: Inside-Secure, STIC, Legrand; Academic: IPG, ENS des Mines de St Etienne, Un. du Maine, Un. F. Rabelais Tours, AETS ESEO Angers, Un. Marne la Vallée, Un. Paris 6, Un. Rennes 1, Inria.

See also: <http://www.spinnaker-rfid.com/>

Abstract: The objective of Spinnaker is to really allow RFID technology to be widely and easily deployed. The role of the OASIS team in this project is to allow the wide scale deployment and management of the specific RFID application servers in the cloud, so to build an end-to-end robust and flexible solution using GCM technology.

7.2. European Initiatives

7.2.1. FP7 Projects

7.2.1.1. TEFIS

Title: TEstbed for Future Internet Services

Type: COOPERATION (ICT)

Defi: Future Internet Experimental Facility and Experimentally-driven Research

Instrument: Integrated Project (IP)

Duration: June 2010 - November 2012

Coordinator: THALES (France)

Others partners: Engineering Ingegneria Informatica S.p.A. (It); IT Innovation (UK); Fundação de Apoio à Universidade de São Paulo (Br); Thales Communications (Fr); ActiveEon (Fr); Lulea University of Technology (Se); Software Quality System S.A. (Es); Fraunhofer Institute FOKUS (De)

See also: <http://www.tefisproject.eu/>

Abstract: TEFIS will support Future Internet of Services Research by offering a single access point to different testing and experimental facilities for communities of software and business developers to test, experiment, and collaboratively elaborate knowledge.

The project develops an open platform to access heterogeneous and complementary experimental facilities addressing the full development lifecycle of innovative services with the appropriate tools and testing methodologies. Through the TEFIS platform users will be supported throughout the whole experiment lifecycle by access to different testing tools covering most of the software development-cycle activities such as software build and packaging, compliance tests, system integration, SLA dimensioning, large-scale deployment, and user evaluation of run-time services. The platform will provide the necessary services that will allow the management of underlying testbeds resources. In particular, it will handle generic resource management, resource access scheduling, software deployment, matching and identification of resources that can be activated, and measurement services for a variety of testbeds.

7.2.1.2. PLAY

Title: Pushing dynamic and ubiquitous interaction between services Leveraged in the Future Internet by ApplYing complex event processing

Type: COOPERATION (ICT)

Defi: Internet of Services, Software & Virtualisation

Instrument: Specific Targeted Research Project (STREP)

Duration: October 2010 - September 2013

Coordinator: FZI (Germany)

Others partners: EBM WebSourcing (Fr), Inria (OASIS and SARDES) (Fr), France Telecom (Fr), ICCS (Gr), Ecole des Mines Albi (Fr), CIM (Serbia).

See also: <http://www.play-project.eu/>

Abstract: The PLAY project will develop and validate an elastic and reliable architecture for dynamic and complex, event-driven interaction in large highly distributed and heterogeneous service systems. Such an architecture will enable ubiquitous exchange of information between heterogeneous services, providing the possibilities to adapt and personalize their execution, resulting in the so-called situational-driven process adaptivity. The OASIS Team is in charge of designing the key element of the PLAY Platform: the event cloud that is a publish/subscribe P2P based system, developed using the GCM technology.

7.2.1.3. FI-WARE

Title: Morphus

Type: COOPERATION (ICT)

Defi: PPP FI: Technology Foundation: Future Internet Core Platform

Instrument: Integrated Project (IP)

Duration: September 2011 - May 2014

Coordinator: Telefonica (Spain)

Others partners:Thales, SAP, Inria

See also: <http://www.fi-ware.eu/>

Abstract: FI-WARE will deliver a novel service infrastructure, building upon elements (called Generic Enablers) which offer reusable and commonly shared functions making it easier to develop Future Internet Applications in multiple sectors. This infrastructure will bring significant and quantifiable improvements in the performance, reliability and production costs linked to Internet Applications ? building a true foundation for the Future Internet.

7.3. International Initiatives

7.3.1. Inria Associate Teams

7.3.1.1. SCADA

Title: Safe Composition of Autonomic Distributed Applications

Inria principal investigator: Ludovic Henrio

International Partner (Institution - Laboratory - Researcher):

University of Chile (Chile) - NIC Chile Research Labs - Mario Leyton

Duration: 2012 - 2014

See also: <http://team.inria.fr/scada>

The SCADA project aims at promoting the collaboration between NIC LABS (Santiago - Chile) and OASIS team (Inria Sophia Antipolis - France) in the domain of the safe composition of applications. More precisely the project will extend existing composition patterns dedicated to parallel or distributed computing to ease the reliable composition of applications. The strong interactions between formal aspects and practical implementation are a key feature of that projects, where formal methods, and language theory will contribute to the practical implementation of execution platforms, development and debugging tools, and verification environments. The composition models we focus on are algorithmic skeletons, and distributed components; and we will particularly focus on the programming and verification of non-functional features. Overall, from formal specification and proofs, this project should lead to the implementation of tools for the design and execution of distributed and parallel applications with a guaranteed behavior.

7.3.1.2. DAESD

Title: Distributed/Asynchronous, Embedded/synchronous System Development

Inria principal investigator: Eric Madelaine

International Partner (Institution - Laboratory - Researcher):

East China Normal University (ECNU) Shanghai - SEI - Yixiang Chen

Duration: 2012 - 2014

See also: <http://team.inria.fr/DAESD>

The development of concurrent and parallel systems has traditionally been clearly split in two different families; distributed and asynchronous systems on one hand, now growing very fast with the recent progress of the Internet towards large scale services and clouds; embedded, reactive, or hybrid systems on the other hand, mostly of synchronous behaviour. The frontier between these families has attracted less attention, but recent trends, e.g. in industrial systems, in “Cyber-Physical systems”, or in the emerging “Internet of Things”, give a new importance to research combining them.

The aim of the DAESD associate team is to combine the expertise of the Oasis and Aoste teams at Inria, the SEI-Shone team at ECNU-Shanghai, and to build models, methods, and prototype tools inheriting from synchronous and asynchronous models. We plan to address modelling formalisms and tools, for this combined model; to establish a method to analyze temporal and spatial consistency of embedded distributed real-time systems; to develop scheduling strategies for multiple tasks in embedded and distributed systems with mixed constraints.

7.3.1.3. Dissiminet

Title: Web-Service approaches for simulation

Inria principal investigator: Olivier Dalle

International Partner (Institution - Laboratory - Researcher):

Carleton University (Ottawa, Canada) - Advanced Real-Time Simulation Laboratory - Gabriel Wainer

Duration: 2011 - 2013

See also: <http://www.inria.fr/en/teams/dissiminet>

This Franco-Canadian team will advance research on the definition of new algorithms and techniques for component-based simulation using a web-services based approach. On one hand, the use of web-services is expected to solve the critical issues that pave the way toward the simulation of systems of unprecedented complexity, especially (but not exclusively) in the studies involving large networks such as Peer-to-peer networks. Web-Service oriented approaches have numerous advantages, such as allowing the reuse of existing simulators, allowing non-computer experts to merge their respective knowledge, or seamless integration of complementary services (eg. on-line storage and repositories, weather forecast, traffic, etc.). One important expected outcome of this approach is to significantly

enhance the simulation methodology in network studies, especially by enforcing the seamless reproducibility and traceability of simulation results. On the other hand, a net-centric approach of simulation based on web-services comes at the cost of added complexity and incurs new practices, both at the technical and methodological levels. The results of this common research will be integrated into both teams' discrete-event distributed simulators: the CD++ simulator at Carleton University and the simulation middle-ware developed in the MASCOTTE EPI, called OSA, whose developments are supported by an Inria ADT starting in December 2011.

7.3.2. Inria International Partners

Fit4Green (<http://fit4green.eu>) is a FP7 project that aimed at creating an energy-aware layer of plug-in on top of the current data centres' management tools to orchestrate the placement of VMs with regards to energy-efficiency concerns. In 2012, the consortium decided to rely on Btrplace to compute the VM placement. Accordingly, Fabien Hermenier collaborated with them to integrate their work with BtrPlace.

7.3.3. Participation In International Programs

7.3.3.1. CIRIC Chili

Ciric research line: Telecommunications

Inria principal investigator: Eric Madelaine

Duration: 2012 - 2021

Our activities with CIRIC have slowly been starting during this year, while CIRIC and Inria-Chile set-up their local organisations. We took the opportunity of our visit in July in Santiago de Chile (workshop if the SCADA associated Team), to discuss with Ciric, and to setup our plans. Later in November Tomas Barrós (PI on the Ciric side) visited us in Siophia-Antipolis, and we were able to pursue our plans.

The current state is that we have listed two chilean software companies, one in the area of telecommunications, the other in the area of banks, that have an interest in method for the development of safe large and complex applications. The role of Ciric in a first step is to set-up a first technical contact with these companies, discuss the use-cases, the common interests, and a preliminar workplan. The next step (in 2013) will involve the work of Ciric engineers on the case-study definition, and a longer visit of E. Madelaine (and possibly other Inria people) in Santiago to start concrete work on this line.

7.4. International Research Visitors

7.4.1. Visits of International Scientists

7.4.1.1. Visiting Scientists

Min Zhang Sep. 15th to Dec. 15th. This visit is in the framework of our "DAESD" Associated Team with ECNU Shanghai. The subject is on contextual/parametric bisimulations for the pNets (Parameterized Networks of Synchronized automata).

Gabriel Wainer Jun. 15th - July 7th. This visit is in the context of the DISSIMINET Associate Team between Inria and Carleton University. The subject is on simulation in the Cloud and Handheld devices.

7.4.1.2. Internships

Yanwen CHEN: Cotutelle with ECNU Shanghai, visits in Inria planned 6 to 9 month each year.

Subject: Programmation d'applications hétérogènes embarquées et distribuées

Institution: UNS & East China Normal University (China)

Quirino ZAGARESE (from Jan 2012 until Aug 2012)

Subject: Lazy loading of data in service oriented and event oriented interaction software architecture models

Institution: University Sannio (Italy)

Michel Jackson DE SOUZA (from Jul 2012 until Aug 2013)

Subject: Distributed coherent snapshot solution for the P2P CAN-based Event Cloud

Institution: UFBA Federal University of Bahia (Brasil), Science sans Frontière brazilian mobility program

7.4.2. Visits to International Teams

- Fabien Hermenier visited the Flux Team at the University of Utah from September to December 2012. This visit allowed us to enhance our collaboration on the study of the Utah' Emulab in order to improve testbed designs. [22]
- Ludovic Henrio, Eric Madelaine, and Cristian Ruz visited NIC-Labs and CIRIC center in Santiago de Chile in July 2012 (1 week visit); a workshop was also held during the week.

8. Dissemination

8.1. Scientific Animation

- Eric Madelaine is member of the steering committees of the FACS and FMCO symposia, and has acted as a program committee of FACS'12, FMCO'12, and of the Science of Computer Programming (SCP) journal.
- Françoise Baude is member of the editorial board of the Techniques et Sciences Informatique French journal from Feb. 2011, was program committee member of ISPA 2012, of RenPar'21 to be held in January 2013. She is University of Nice official representative within the KIC EIT ICT Labs, including the Master school, from 2012.
- Ludovic Henrio: program committee member of FESCA'12, FOCLASA'12, FESCA'13 and reviewer for the journal SCP (Science of Computer Programming) and MSCS.
- Fabien Hermenier: program committee member of VTDC 2012 and PDP 2012 (energy efficiency track) and reviewer for IEEE Transactions on Network and Service Management and the Journal of System Architecture. He also presented btrPlace during the OW2Con, the OW2 Annual Conference.
- Fabrice Huet : program committee member of HPDC 2012, AP2PS 2012, reviewer for Transactions on Services Computing, Concurrency and Computation: Practice and Experience.

8.2. Teaching - Supervision - Juries

8.2.1. Teaching

Licence : Olivier Dalle, Introduction à la Programmation Objet, 15 H eqTD, niveau L1, Département Informatique, UNS, France

Licence : Olivier Dalle, Introduction à la Programmation Objet, 15 H eqTD, niveau L1, Département Informatique, UNS, France

Licence : Fabrice Huet, Informatique Générale, 48 H eqTD, niveau L1, Département Informatique, UNS, France

Licence : Françoise Baude, Algorithmique et Programmation en Java, 70 H eqTD, niveau L2, Polytech'Nice Sophia, UNS, France

Licence : Fabien Hermenier, Introduction à Internet, 103 H eqTD, niveau L2, Polytech'Nice Sophia, UNS, France

Licence : Olivier Dalle, Algorithmes et Structures de Données, 36 H eqTD, niveau L3, Département Informatique, UNS, France

Licence : Olivier Dalle, Harmonisation Informatique, 36 H eqTD, niveau L3, IUP Miage, UNS, France

Licence : Fabien Hermenier, Outils pour le Génie Logiciel, 26 H eqTD, niveau L3, Polytech’Nice Sophia, UNS, France

Licence : Fabrice Huet, Outils de Génie Logiciel, 24 H eqTD, niveau L3, IUT, UNS, France

Licence : Fabrice Huet, Programmation de Jeux Réseau, 18 H eqTD, niveau L3, IUT, UNS, France

Master : Françoise Baude, Applications Réparties, 45 H eqTD, niveau M1, Polytech’Nice Sophia, UNS, France

Master : Olivier Dalle, Ingénierie des Protocoles, 21 H eqTD, niveau M1, Département Informatique, UNS, France

Master : Fabrice Huet, Distribution et parallélisme, 49H eqTD, niveau M1, Département Informatique, UNS, France

Master : Fabrice Huet, Syst èmes Distribués, 42H eqTD, niveau M1, IUP Miage, UNS, France

Master : Françoise Baude et Ludovic Henrio, Distributed systems: an Algorithmic approach, 17 H + 17 h eqTD, niveau M2, Polytech’Nice Sophia/UFR Sciences, UNS, France

Master: Ludovic Henrio, Sémantique des Systèmes Distribués et Embarqués, 11 H eqTD, niveau M1, Département Informatique, UNS, France

8.2.2. Supervision

PhD in progress : Laurent Pellegrino “Pushing dynamic and ubiquitous event-based interaction in the Internet of services: a middleware for event clouds”, since Sept 2010, advisor Françoise Baude.

PhD in progress : Nuno Gaspar “Integrated, Autonomic, and Reliable Deployment and Management for SaaS composite applications”, since Nov 2011, advisor: Eric Madeleine

PhD in progress : Michaël Benguigui, “Modèles de programmation pour l’analyse de données sur machines multi-cœurs, grilles et clouds”, since Dec 2011, advisor Françoise Baude.

PhD in progress : Yanwen Chen, “Formal model and Scheduling for cyberphysical systems”, since Dec 2011, advisors Eric Madelaine and Yixiang Chen (ECNU Shanghai)

PhD in progress : Alexandra Savu, “Design, specification, and verification of parameterized topologies of distributed components”, since Oct. 2012, advisor Eric Madelaine

PhD in progress : Maeva Antoine, “Plateforme élastique pour le stockage et la notification d’Evènements”, since Oct. 2012, advisors Eric Madelaine, Fabrice Huet

8.2.3. Juries

Françoise Baude participated in PhD evaluations as

- jury supervisor for the PhD of Marcela Rivera held on December 15th 2011, UNS; PhD of Imen Ben Lhamar, on November 15th 2012, at Telecom Sud-Paris / Univ. Evry Val d’Essone
- PhD reviewer and jury member for the PhD of Gabriel Hermosilio on 5th June 2012, University of Lille (ADAM EPI); PhD of Xavier Etchevers on December 12th 2012, at LIG / FranceTelecom R&D (SARDES EPI)

She participated in the recruitment committees for

- Maitre assistant position at Ecole des Mines Albi-Carnaux, in January 2012
- Professor at Univ. of Henri Poincaré, Ecole Supérieure d’Informatique et Applications de Lorraine, in May 2012

Fabrice Huet participated to the mid-term evaluation jury of Giuseppe Reina (April 2012) and Mario Pastorelli (October 2012), Eurecom PhD students.

8.3. Popularization

Françoise Baude is in charge since September 2011 of the set up and animation of an educational program *Cordées de la Réussite*. She leads as Polytech'Nice Sophia member, and in connection with MASTIC activities at Sophia-Antipolis, two *Cordées pour la Science à Sophia* groups (one with Lycée Vinci Antibes, the other with Lycée Tocqueville Grasse).

Fabrice Huet gave a seminar to high school students about the inner working of computers. He also leads the ISN *Informatiques et Sciences du Numérique* courses for the academia of Nice. These courses were given to high school teachers who volunteered to offer a Computer Science option to their students.

9. Bibliography

Major publications by the team in recent years

- [1] L. BADUEL, F. BAUDE, D. CAROMEL. *Asynchronous Typed Object Groups for Grid Programming*, in "International Journal of Parallel Programming", 2007, vol. 35, n^o 6, p. 573-614.
- [2] L. BADUEL, F. BAUDE, D. CAROMEL, A. CONTES, F. HUET, M. MOREL, R. QUILICI. 9, in "Grid Computing: Software Environments and Tools", Springer, 2006, ISBN: 1-85233-998-5.
- [3] T. BARROS, R. BOULIFA, A. CANSADO, L. HENRIO, E. MADELAINE. *Behavioural Models for Distributed Fractal Components*, in "Annals of Telecommunications", jan 2009, vol. 64, n^o 1–2, also Research Report Inria RR-6491 [DOI : 10.1007/s12243-008-0069-7], <http://hal.inria.fr/inria-00268965>.
- [4] F. BAUDE, D. CAROMEL, C. DALMASSO, M. DANELUTTO, V. GETOV, L. HENRIO, C. PÉREZ. *GCM: A Grid Extension to Fractal for Autonomous Distributed Components*, in "Annals of Telecommunications", 2009, vol. 64, n^o 1, p. 5–24 [DOI : 10.1007/s12243-008-0068-8], <http://hal.archives-ouvertes.fr/inria-00323919/en/>.
- [5] F. BAUDE, D. CAROMEL, F. HUET, L. MESTRE, J. VAYSSIÈRE. *Interactive and Descriptor-based Deployment of Object-Oriented Grid Applications*, in "11th IEEE International Symposium on High Performance Distributed Computing", 2002, p. 93–102.
- [6] D. CAROMEL, L. HENRIO. *A Theory of Distributed Object*, Springer-Verlag, 2005.
- [7] D. CAROMEL, L. HENRIO, B. SERPETTE. *Asynchronous and Deterministic Objects*, in "Proceedings of the 31st ACM Symposium on Principles of Programming Languages", ACM Press, 2004, p. 123–134, <http://doi.acm.org/10.1145/964001.964012>.
- [8] L. HENRIO, F. KAMMÜLLER, B. LUTZ. *ASPfun : A typed functional active object calculus*, in "Science of Computer Programming", 2011, vol. In Press, Corrected Proof [DOI : DOI: 10.1016/J.SCICO.2010.12.008], <http://www.sciencedirect.com/science/article/B6V17-51XR3D4-1/2/52153869e8c4df60a65b2d52bbf9828c>.
- [9] F. HUET, D. CAROMEL, H. E. BAL. *A High Performance Java Middleware with a Real Application.*, in "Proceedings of the ACM/IEEE SC2004 Conference on High Performance Networking and Computing, 6-12 November 2004, Pittsburgh, PA, USA", 2004.

Publications of the year

Doctoral Dissertations and Habilitation Theses

- [10] F. BONGIOVANNI. *Design, formalization and implementation of overlay networks ; application to RDF data storage*, University of Nice-Sophia Antipolis, Apr 2012.
- [11] L. HENRIO. *Formal Models for Programming and Composing Correct Distributed Systems*, Université de Nice Sophia-Antipolis, July 2012, Habilitation à Diriger des Recherches.
- [12] S. MALIK. *Network Aware Cloud Computing*, University of Nice-Sophia Antipolis, Dec 2012.

Articles in International Peer-Reviewed Journals

- [13] L. HENRIO, F. KAMMÜLLER, B. LUTZ. *ASPfun : A typed functional active object calculus*, in "Science of Computer Programming", 2012, vol. 77, Issue 7-8, p. 823 - 847 [DOI : DOI: 10.1016/J.SCICO.2010.12.008], <http://www.sciencedirect.com/science/article/B6V17-51XR3D4-1/2/52153869e8c4df60a65b2d52bbf9828c>.
- [14] E. MATHIAS, F. BAUDE. *A Component-Based Middleware for Hybrid Grid/Cloud Computing Platforms*, in "Concurrency and Computation: Practice and Experience", 2012, vol. 24, n^o 13, p. 1461–1477.
- [15] C. RUZ, F. BAUDE, B. SAUVAN. *Using Components to Provide a Flexible Adaptation Loop to Component-based SOA Applications*, in "International Journal on Advances in Intelligent Systems", July 2012, vol. 5, n^o 1 & 2, p. 32–50, ISSN: 1942-2679, http://www.ariajournals.org/intelligent_systems/.

Invited Conferences

- [16] O. DALLE. *On Reproducibility and Traceability of Simulations*, in "Winter Simulation Conference (WSC'12)", C. LAROQUE, J. HIMMELSPACH, R. PASUPATHY, O. ROSE, A. M. UHRMACHER (editors), December 2012, To appear.

International Conferences with Proceedings

- [17] M. BENGUIGUI, F. BAUDE. *Towards Parallel and Distributed Computing on GPU for American Basket Option Pricing*, in "GPU Computing in the Clouds (GPU-Cloud 2012) workshop in conjunction with 4th IEEE international conference on Cloud Computing Technology and Science (CloudCom 2012)", IEEE (editor), Dec 2012.
- [18] F. BONGIOVANNI, L. HENRIO. *A Mechanized Model for CAN Protocols*, in "16th International Conference on Fundamental Approaches to Software Engineering (FASE'13)", LNCS, Springer, 2013, To appear.
- [19] Y. CHEN, Y. CHEN, E. MADELAINE. *Timed-pNets: A formal communication behavior model for real-time CPS systems*, in "Workshop on Trustworthy Cyber Physical Systems", TCPS, sep 2012.
- [20] O. DALLE, E. MANCINI. *Integrated Tools for the Simulation Analysis of Peer-To-Peer Backup Systems*, in "Proceedings of the 2012 Intl Conference on Simulation Tools and Techniques (SIMUTOOLS 2012)", Sirmione, Italy, F. QUAGLIA, J. HIMMELSPACH (editors), March 2012, p. 178–183 [DOI : 10.4108/ICST.SIMUTOOLS.2012.247783], <ftp://ftp-sop.inria.fr/mascotte/Publications/DaMa12.pdf>.

- [21] C. DUPONT, G. GIULIANI, F. HERMENIER, T. SCHULZE, A. SOMOV. *An Energy Aware Framework for Virtual Machine Placement in Cloud Federated Data Centres*, in "3rd IEEE/ACM International Conference on Future Energy Systems, e-energy", May 2012.
- [22] *Best Paper*
F. HERMENIER, R. RICCI. *How to Build a Better Testbed: Lessons From a Decade of Network Experiments on Emulab*, in "8th International ICST Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities, Tridentcom", Jun 2012, Best Paper award.
- [23] S. MALIK, F. HUET, D. CAROMEL. *Latency Based Dynamic Grouping Aware Cloud Scheduling*, in "26th IEEE International Conference on Advanced Information Networking and Applications Workshops (AINA 2012)", Fukuoka, Japan, IEEE, Mar. 2012.
- [24] E. P. MANCINI, G. WAINER, K. AL-ZOUBI, O. DALLE. *Simulation in the Cloud Using Handheld Devices*, in "MSGC@CCGRID - Workshop on Modeling and Simulation on Grid and Cloud Computing - 2012", Ottawa, Canada, IEEE, May 2012, <http://hal.inria.fr/hal-00691248>.
- [25] L. PELLEGRINO, I. ALSHABANI, F. BAUDE, R. STÜHMER, N. STOJANOVIC. *An approach for efficiently combining real-time and past events for ubiquitous business processing*, in "7th International Workshop on Semantic Business Process Management in conjunction with 9th Extended Semantic Web Conference", LNCS, Springer, 2012.
- [26] L. PELLEGRINO, F. BAUDE, I. ALSHABANI. *Towards a scalable cloud-based RDF storage offering a pub/sub query service*, in "CC 2012, The Third International Conference on Cloud Computing, GRIDs, and Virtualization", July 2012.
- [27] N. STOJANOVIC, R. STÜHMER, F. BAUDE, P. GIBERT. *Tutorial: Where Event Processing Grand Challenge meets Real-time Web: PLAY Event Marketplace*, in "DEBS'12, the 6th ACM International conference on Distributed Event-based system", ACM, 2012, p. 341–352.
- [28] Q. ZAGARESE, G. CANFORA, E. ZIMEO, I. ALSHABANI, L. PELLEGRINO, F. BAUDE. *Efficient Data-Intensive Event-Driven Interaction in SOA*, in "SAC 2013, 28th ACM Symposium On Applied Computing", March 2013, to appear.
- [29] Q. ZAGARESE, G. CANFORA, E. ZIMEO, F. BAUDE. *Enabling Advanced Loading Strategies for Data-Intensive Web Services*, in "19th Int. Conference on Web Services, ICWS", IEEE, 2012, p. 480–487.

Conferences without Proceedings

- [30] M. ANTOINE, F. BAUDE, F. HUET. *Évaluation d'une architecture de stockage RDF distribuée*, in "Atelier Ontologies et Jeux de Données pour évaluer le web sémantique (OJD 2012) conjoint à IC 2012, 23èmes journées francophones de l'ingénierie des connaissances", 2012, http://www.irit.fr/OJD/OJD_3.pdf.
- [31] O. DALLE. *Activity-based parallel and distributed simulation: application to Peer-to-Peer systems*, in "Activity-based Modeling & Simulation Workshop, Créteil, France", 2012, <http://www.spatial-computing.org/actims-paris/start>.

- [32] O. DALLE. *Measuring activity at structure-change level: Questions and ideas*, in "Activity-based Modeling & Simulation Workshop, Créteil, France", 2012, <http://www.spatial-computing.org/actims-paris/start>.
- [33] N. GASPARD, E. MADELAINE. *Fractal à la Coq*, in "Conférence en Ingénierie du Logiciel", CIEL 2012, june 2012.
- [34] L. HENRIO. *Behavioural specification and verification of GCM component*, in "Journées communes LTP - LAC - LAMHA, Orléans, France", 2012, <http://www.spatial-computing.org/actims-paris/start>.
- [35] L. HENRIO. *GCM/ProActive: a distributed component model, its implementation, and its formalisation*, in "Dagstuhl Seminar 12511 - Divide and Conquer: the Quest for Compositional Design and Analysis", 2012.
- [36] L. HENRIO. *Multi-active Objects*, in "Journées communes LTP - LAC - LAMHA, Orléans, France", 2012.

Books or Proceedings Editing

- [37] S. STRASSBURGER, O. DALLE, G. F. RILEY (editors). *Themed Issue: Recent advances in parallel and distributed simulation*, Palgrave MacMillan, August 2012, vol. 6, p. 141–213, <http://www.palgrave-journals.com/jos/journal/v6/n3/index.html>.

Research Reports

- [38] I. ALSHABANI. *Assessment of the PLAY integrated platform V1, Deliverable D5.2.1*, FP7 Play EU project, Sep. 2012.
- [39] I. ALSHABANI, F. BAUDE, L. PELLEGRINO, B. SAUVAN, P. GIBERT, C. HAMERLING. *PLAY federated middleware specification and implementation V1: State of the Art, Deliverable D.2.5.1*, FP7 Play EU project, May 2012.
- [40] R. AMEUR-BOULIFA, L. HENRIO, E. MADELAINE, A. SAVU. *Behavioural Semantics for Asynchronous Components*, Inria, December 2012, n^o RR-8167, 58, <http://hal.inria.fr/hal-00761073>.
- [41] E. BORRELLI. *D2.3.2 - Integration plan*, FP7 TEFIS EU project, Jun. 2012, final version.
- [42] E. BORRELLI. *D4.1.2 - Resource management implementation*, FP7 TEFIS EU project, Jun. 2012, final version.
- [43] E. BORRELLI. *D4.3.1 - Code provisioning management system*, FP7 TEFIS EU project, Apr. 2012, final version.
- [44] E. BORRELLI. *D4.3.2 - Code provisioning management system*, FP7 TEFIS EU project, Jun. 2012, final version.
- [45] L. HENRIO, F. HUET, Z. ISTVÁN. *A Language for Multi-threaded Active Objects*, Inria, July 2012, n^o RR-8021, <http://hal.inria.fr/hal-00720012>.
- [46] E. MADELAINE, F. BAUDE, C. RUZ, N. GASPARD, M. BENGUIGUI, A. CONTES, L. BELISSARD, N. MERLE, C. MOUTOT. *Architecture de la plateforme Cloud/RFID*, OSEO Projet Spinnaker, Livrable 1.2.2, February 2012, n^o LIV-0102-02-0212.

- [47] E. MADELAINE, F. BAUDE, C. RUZ, N. GASPAR, M. BENGUIGUI, A. CONTES, L. BELISSARD, C. MOUTOT. *Caractérisation des Technologies Cloud existantes*, OSEO Projet Spinnaker, Livrable 1.2.1, January 2012, n^o LIV-0102-01-0112.
- [48] F. PERRINA, K. CAMPOWSKY, B. PICKERING, G. GIAMMATTEO, J. ROBERTO, U. ACOSTA, J. GONZÁLEZ, I. FERNÁNDEZ, A. GRACIA. *D3.3.2 - User tools implementation*, FP7 TEFIS EU project, May 2012, final version.
- [49] F. PERRINA, B. PICKERING, G. GIAMMATTEO, K. CAMPOWSKY, U. ACOSTA, J. GONZÁLEZ, I. FERNÁNDEZ, A. GRACIA. *D3.2.2 - Building blocks integrated into TEFIS portal*, FP7 TEFIS EU project, Jun. 2012, final version.
- [50] F. PERRINA, J. ROBERTO, J. LEGUAY, G. GIAMMATTEO, A. SÄLLSTRÖM, U. ACOSTA, J. GONZÁLEZ, I. FERNÁNDEZ, A. GRACIA. *D4.4.2 - Test toolkit implementation*, FP7 TEFIS EU project, Jun. 2012, final version.
- [51] P. PHAM. *Optimizing the performance of parallel simulator using transactional memory*, Master Ubinet, Univ. of Nice Sophia Antipolis, 2012.
- [52] B. PICKERING, S. TAYLOR, E. BORRELLI. *D6.3.2 - Data services implementation - updated*, FP7 TEFIS EU project, Apr. 2012, final version.
- [53] E. SALAGEANU, C. DALMASSO, F. PERRINA. *D4.2.2 - Experiment and Workflow Scheduler prototype update*, FP7 TEFIS EU project, May 2012, final version.

Other Publications

- [54] O. ETZION, F. BAUDE, N. STOJANOVIC, L. STOJANOVIC. *Scalable semantic processing of huge, distributed real-time streams: Semantics between Event Processing and Cloud Computing*, Nov 2012, Tutorial at 11th International Semantic Web Conference (ISWC 2012), <http://tutorial-iswc12.fzi.de>.

References in notes

- [55] L. BADUEL, F. BAUDE, D. CAROMEL, A. CONTES, F. HUET, M. MOREL, R. QUILICI. 9, in "Grid Computing: Software Environments and Tools", Springer, 2006, ISBN: 1-85233-998-5.
- [56] F. BAUDE, D. CAROMEL, N. FURMENTO, D. SAGNOL. *Optimizing remote method invocation with communication-computation overlap*, in "Future Generation Computer Systems, Elsevier", 2002, vol. 18, p. 769–778, Selected article from PaCT 2001.
- [57] F. BAUDE, I. FILALI, F. HUET, V. LEGRAND, E. MATHIAS, P. MERLE, C. RUZ, R. KRUMMENACHER, E. SIMPERL, C. HAMERLING, J.-P. LORRÉ. *ESB Federation for Large-Scale SOA*, in "SAC '10: Proceedings of the 2010 ACM Symposium on Applied Computing", ACM, March 2010, p. 2459–2466.
- [58] E. BRUNETON, T. COUPAYE, M. LECLERCP, V. QUEMA, JEAN-BERNARD. STEFANI. *An Open Component Model and Its Support in Java.*, in "7th Int. Symp. on Component-Based Software Engineering (CBSE-7)", LNCS 3054, may 2004.

-
- [59] O. DALLE. *OSA: an Open Component-based Architecture for Discrete-Event Simulation*, in "20th European Conference on Modeling and Simulation (ECMS)", Bonn, Germany, May 2006, p. 253–259.
- [60] V. D. DOAN. *Grid computing for Monte Carlo based intensive calculations in financial derivative pricing applications*, University of Nice-Sophia Antipolis, March 2010, http://www-sop.inria.fr/oasis/personnel/Viet_Dung.Doan/thesis/.
- [61] J. DÜNNWEBBER, F. BAUDE, V. LEGRAND, N. PARLAVANTZAS, S. GORLATCH. *Towards Automatic Creation of Web Services for Grid Component Composition*, in "Integrated Research in GRID computing - Invited papers from the 1st CoreGRID Integration Workshop 2005", Springer, 2006, ISBN: 978-0-387-47656-8. Also as the CoreGrid TR003 report.
- [62] I. FILALI. *Improving Resource Discovery in P2P Systems*, University of Nice-Sophia Antipolis, February 2011.
- [63] I. FILALI, L. PELLEGRINO, F. BONGIOVANNI, F. HUET, F. BAUDE. *Modular P2P-Based Approach for RDF Data Storage and Retrieval*, in "AP2PS 2011, The Third International Conference on Advances in P2P Systems", November 2011, p. 39-46, Best paper award.
- [64] TC-GRID, ETSI. *ETSI TS 102 830: GRID; Grid Component Model; Part 4: GCM Fractal JAVA API*, European Telecommunications Standards Institute (ETSI), Sophia-Antipolis, France, 2010, http://webapp.etsi.org/workprogram/Report_WorkItem.asp?wki_id=28857.