INRIA

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Project-Team oasis

Active objects, semantics, Internet and security

Sophia Antipolis - Méditerranée

Theme : Distributed Systems and Services

Activity Report

2010

Table of contents

1. Team

Research Scientists

Eric Madelaine [Senior Researcher Inria]

Ludovic Henrio [Junior Researcher CNRS]

Faculty Members

Denis Caromel [Team Leader, Prof. UNS, HdR]

Fabrice Huet [Associate Professor UNS]

Françoise Baude [Prof. UNS, HdR]

Technical Staff

Bastien Sauvan [ADT Glaxy, to nov 30]

Elvio Borrelli [since dec 01]

Fabien Viale [FP6 S4All, FP6 NESSI-Grid, Thales]

Franca Perrina [AGOS, to jun 30]

Mathieu Schnoor [FP7 Tefis, since oct 01]

Oleg Smirnov [Microsoft - MCorePHP]

Yu Feng [ETSI, FP6 EchoGrid MCorePHP]

PhD Students

Brian Amedro [SCS-Agos, to oct 01]

Cristian Ruz [INRIA, Conicyt]

Elton Mathias [ATER UNS]

Francesco Bongiovanni [(INRIA-PACA)]

Guilherme Peretti-Pezzi [ERCIM]

Imen Filali [SCS-Agos / FP7 SOA4All, to sep 14]

Laurent Pellegrino

Marcela Rivera [ATER UNS]

Muhammad Khan [STIC-Asia, to oct 31]

Sheheryar Malik

Virginie Legrand Contes [SOA4ALL]

Paul Naoumenko [ATER UNS, to jul]

Administrative Assistants

Sandra Devauchelle [I3S]

Sylvie Lelaidier [INRIA]

Others

Elvio Borrelli [since feb 15, to nov 30]

Zhaolong Liu [since apr 13, since jun 01, to jul 31]

Zhaolong Liu [, since sep 01]

Amine Rouini [since mar 01, to aug 31]

Anemona Haegea [since jul 12, to oct 15]

Cédric Morandin [Supinfo, since apr 01, to sep 30]

Nauman Badar [since jul 05, to oct 30]

Raluca Halalai [since jul 12, to oct 30]

Xavier Bontemps [since jun 14, to sep 30]

YanWen Chen [since mar 01, to aug 31]

Zsolt Istvan [since jun 01, to sep 30]

2. Overall Objectives

2.1. Presentation

The team focuses its activities on distributed (Grid) 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): Grid 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 mobility and security. We believe that we have the skills to be significantly fruitful in this major application domain; more specifically, we aim at producing interesting results for embedded applications for mobile users, Grid computing, peer-to-peer intranet, electronic trade and collaborative applications.

2.2. Highlights

- We finished the process of standardisation of a Grid Component Model and its deployment:

  The fourth and last ESTI TC Grid standard has been aproved an published: ' "GCM Management API".

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 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 workstations (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, deployment and evolution of distributed applications ?

One of the answers we suggest relies on the concept of active object, that act, as a single entity as an abstraction of a thread, a set of objects and a location. Active object 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 [52]. 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 is intended to become a standard for Grid components, and most of our research on component models are related to it. The GCM is an extension of the Fractal model. On the practical side, ProActive/GCM is a prot otype implementation of the GCM above the ProActive library; not all GCM features are implemented in ProActive yet. ProActive/GCM is intended to become the reference implementation of the GCM, as was the goal of the European project GridCOMP.

For providing a better programming and runtime environment for object and component applications, we have developed competencies in both theoretical and applicative side fields, such as distribution, fault-tolerance.

3.2. Formal models for distributed objects

A few years ago, we designed the ASP calculus 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 generalised 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 [7].

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 models distributed asynchronous components. A functional fragment of ASP has been modelled in the Isabelle theorem prover.

3.3. Static Analysis and Verification

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, from the definition of adequate models representing programs, the mastering of state complexity through abstraction techniques or through new algorithmic approaches, to 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 [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

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, computational grids will allow new classes of applications to emerge. Another challenge is to use, for a given computation, unused CPU cycles of desktop computers in a Local Area Network. This is intranet Computational Peer-To-Peer.

There is a need for models and infrastructures for grid and peer-to-peer computing, and we promote a programming model based on communicating mobile objects and components.

4.2. Service Oriented Architectures (SOA)

Service Oriented Architectures aim at the integration of distributed services at the level of the Enterprise, or of the whole Internet. The OASIS team seeks solutions to the problems encountered here, with the underlying motivation to demonstrate the usefulness of a large-scale distributed programming approach as featured by ProActive and GCM:

- Deployment of a service on the service infrastructure: as services depend upon other services, deployment and runtime management can be eased if these dependencies are made explicit. Indeed, services required for another service to work can be instantiated or discovered more easily if the dependencies are known. The recently defined Service Component Architecture (SCA) model is gaining popularity. We are conducting research to promote the Grid Component Model as a complement to SCA. Indeed, we think that GCM is by essence well equipped for supporting at runtime SCA services that are widely distributed, and may need to be invoked in an asynchronous manner. We thus pursue works to make SCA and GCM interoperable models.

- Interoperability between services: the uniform usage of web services can provide a simple inter-operability between them. GCM components can be exposed as web services [53], and we have conducted research and development to permit a GCM component to invoke an external web service through a client interface, and thus to have GCM/SCA components be integrated in SCA-based applications relying on SCA bindings configured as web services.

- Large-scale deployment and monitoring of a set of (similar) services on a possibly large set of machines from e.g. a computing grid, a cloud of machines, etc.: such capability will really make SOA ready for the Internet scale, and we are designing some grid services, accessible as web services, in order to leverage the required functionalities for Grid/cloud deployment of components/services and monitoring of the resulting runtime infrastructure.

- Distributed and Scalable service bus: this is needed if services are composed and orchestrated through an Enterprise Service Bus. But, to scale, the ESB must itself be distributed, and must incorporate from design time, the necessary mechanisms to handle large-scale distribution and possibly huge amount of end-user or technical (service discovery, registry, orchestration and monitoring engines, etc) services. Moreover the bus itself should be deployable seamlessly on any heterogeneous combination of host machines. In this wide context, we intend to use GCM components for building the bus itself, giving it the required Internet scale capabilities (this subject is specifically addressed in the context of the SOA4ALL IP FP7 project the team is involved in, taking the OW2 PEtALS ESB from the PEtALS link SME as starting point)

- Peer-to-peer based service registry and service lookup protocols: in an Internet-based world hosting possibly billions of services, the registration and subsequent lookup of services can only be addressed along a semantic-based approach, and should allow a robust and scalable way to store and query for service descriptions. In the context of the SOA4ALL IP FP7 project, we conduct research to

contribute to the design of a semantic space where services will be stored and looked upon based on their semantic description. For scalability purposes, the space specification is organised as a peer-to-peer network, further implemented in a distributed, scalable way relying on a grid middleware as the ProActive technology.

- Self-management of the SOA infrastructure and SOA applications: this pertains to autonomic and self-management of the service infrastructure, but also of the component assemblies that constitute the Service Oriented Application. Again the use of GCM components instead of Fractal-RMI components whenever needed can be a solution to the scalability and deployment problems. For service compositions represented as component assemblies, we are exploring the use of control components put in the component membranes, acting as sensors or actuators, that can drive the self-management of composite services, e.g. according to a negotiated Service Level Agreement.

- Distributed and agile workflow enactment: as BPMN and BPEL are the standard ways to define a service orchestration, we are considering how such a *composition in time* approach can be mapped into an architectural-based view involving (SCA) components. Besides, efficient and secured orchestration of such service compositions can benefit from distribution and parallelism. In this aim, we investigate how GCM can be successfully used to design a parallel, distributed, yet flexible orchestration engine handling a BPEL workflow description previously decomposed into sub-workflows. Deployment and management of the decomposition can also be addressed easily by having the distributed workflow relying on GCM components.

5. Software

5.1. ProActive

Participants: B. Amedro, F. Baude, D. Caromel, F. Huet [correspondant], E. Isnard, M. Khan, V. Legrand, E. Mathias, P. Naoumenko, F. Perrina, F. Viale, F. Bratu, F. Fontenoy, O. Smirnov.

ProActive is a Java library (Source code under GPLv3 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 (NorduGrid)
- 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 such as:

- The deployment framework, which has been standardised by the ETSI, allows the deployment of ProActive, native including MPI applications on almost all Grid/cluster protocol: Windows CCS, Sun Grid Engine, LSF, OAR, PBS, SSH, RSH etc. It also supports all the major virtualization products such as VMware, Xen, KVM and Virtualbox;

- The communication layer that can rely on RMI, HTTP, IBIS, RMI/ssl or RMI/ssh. In particular, this last protocol can cross firewalls in many cases;

- Two mechanisms: the ProActive message routing protocol (PAMR), and another based on the use of the ssh proxy command mechanism have been designed to allow ProActive Programming deployment in the context of limited ingoing and outgoing network connections. This is particularly useful in the following cases:

  – Network address translation devices;

  – Firewalls allowing only outgoing connections (this is the default setup of many personal firewall);

  – Virtual Machines with a virtualized network stack.

- GCM component support;

- The graphical user interface IC2D offers many views of an application such as:

  – the Monitoring view which allows better control and monitoring;

  – the TimIt view that draws graphics on time statistics;

  – the CharIt view providing charts any ProActive numerical values.

- The ability to exploit the migration capability of active objects, in network and system management;

- Object-Oriented SPMD programming model with its API;

- Distributed and Non-Functional Exceptions handling;

- Fault-Tolerance and Checkpointing mechanisms;

- File Transfer capabilities over the Grid;

- A job scheduler for scheduling many kinds of jobs like Java, ProActive, native, scripts executables, Matlab and Scilab programs, etc.

- A resource manager able to manage various hosting machines, gained through SSH, PBS, Amazon EC2, GCMD or Virtual Machine.

- ProActive connectors for remote JMX-based operations and an OSGi compliant version of the ProActive library. This involved the development of a "bundled" version of the library;

- Remote data access for accessing data stored in a remote data space;

- MPI code wrapping to deploy MPI application using the GCM deployment;

- Remote debugging support;

- Active object and component exposition as web services using Axis2 or CXF;

We have demonstrated on a set of applications the advantages of the ProActive library, and among others we are particularly proud of the following results, showing that portable and transparent Java code can compete with specific optimised approaches:

- NQueen challenge, where we equalled the world record n=24 (227 514 171 973 736 solutions) in 17 days based on ProActive's P2P infrastructure (300 machines).
- NQueen challenge, where we get the world record n=25 (2 207 893 435 808 352 solutions) in 6 months based on ProActive's P2P infrastructure using free cycles of 260 PCs.

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 [9] [50] and to the web pages http://www.objectweb.org and http://proactive.inria.fr/ which list several white papers.

The following new features have been developed in 2010:

- Multi Protocol support (.. je ne sais pas si c'a ete commence en 2009..)
- Java process builder to associate a job execution to the system user (..le travail de Zsolt Istvan)
- REST API for the job scheduler and the resource manager (fait par Arnaud Contes)
- Web Service APIs for the job scheduler
- WEB portal for job scheduler (not yet released)
- New support for Microsoft Azure Cloud Platform
- Support for MapReduce (a first prototype)

5.2. Vercors platform

Participants: E. Madelaine, R. Halalai, L. Henrio.

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.

- Our main effort this year was on experiments with a distributed infrastructure for verification activities. There are some tools in CADP that support cluster-based execution, allowing users to benefit from large CPU and memory resources, mainly at the step of state-space generation. But putting such tools in practice on a large application requires both intelligence in splitting the problem in manageable parts, and a lot of language and execution environment support for managing the verification process. We propose to extend the SVL (verification scripting) language of CADP with specific information about distributed resources and distributed execution, and to use the PacaGrid workflow definition primitives to implement SVL programs. As a first step in this direction, we have run a large case-study using various distribution and workflow strategies [22].

6. New Results

6.1. Distributed Programming Models

6.1.1. *Theory of Distributed Objects*
Participant: L. Henrio.

This work extends results published in [7]. The ς-calculus, and its semantics were published by Abadi and Cardelli [49]. In collaboration with Florian Kammüller (Technische Universität Berlin) we modelled the ς-calculus and a distributed functional calculus, based on ASP – ASP$_{fun}$ – in the Isabelle/HOL theorem prover.

- This year we have mainly focused on the representation of variables by locally nameless techniques allowing more elegant proofs and simpler theorems, at the price of a slightly more complicated theory. We encoded ASP$_{fun}$ with locally nameless notations and provided a complete comparison with the classical de Bruijn approach. This comparison has been written in a journal paper under submission.

- All the development (formalisation and proofs) in Isabelle/HOL is available at http://www-sop.inria.fr/oasis/Ludovic.Henrio/misc.html

6.1.2. *Skeleton-based parallel programming*
Participant: L. Henrio.

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 published a paper on the management of exceptions in skeletons, in order to capture helpful exception traces for the skeleton programmer [28].

This work has been realized in collaboration with Mario Leyton (NIC Labs, Universidad de Chile).

6.1.3. *Mechanised formalisation of Futures in a Distributed Component Model*
Participants: L. Henrio, M. Khan.

We formalise a component model based on the GCM [4] plus asynchronous invocations with futures between components. Our components communicate via asynchronous requests and replies where the requests are enqueued at the target component. The component making the invocation receives a future. Our model is precise enough to enable the specification of a formal semantics. Futures play a major role for such asynchronous components. ProActive/GCM is a variant of this model. This was previously published in [54]. In addition to this theoretical work, we focus on specifying and implementing the various future update strategies in ProActive with the aim of evaluating the efficiency of each strategy in a given configuration.

- We published the formalisation above extended by the precise formalisation and proof of properties on one future update strategy (Eager home) [27].

- We published a paper on the formalisation and structure of the component model itself, proving properties on the component model and the component reconfiguration [17].

- We formalised a second future update strategy (lazy).

- We have developed a semi-formal specification of future update strategies and provide their implementation in ProActive; these results are published in [37].

- Most of those results will be published in Muhammad Khan's PhD thesis.

6.1.4. *Behavioural models for Group Communication*
Participants: E. Madelaine, R. Halalai, L. Henrio.

We have extended the behavioural semantics of active objects, that was defined in [3] to handle group communication mechanisms. We now have a comprehensive semantic framework allowing for behaviour model generation, encompassing distributed active objects, handling of future proxies, and mechanisms for group communication and various kind of group synchronisation. This work was comforted by a large case-study, showing that state-of-the-art finite-state model-checking methods can be used for analysing such systems [22], [35].

6.2. Component-oriented Grid Programming

6.2.1. Designing Non-Functional Concerns as Components

Participants: F. Baude, L. Henrio, P. Naoumenko.

As part of the design of the GCM, we progressed on the research concerning the componentisation of component membranes [51]. This consists in adopting a component view of the non-functional and control aspects, in the same way the component model structures the functional concerns. The advantages of this approach are a better structuring of non-functional aspects, and better reconfiguration possibilities.

- Most of our results in this area were presented in Paul Naoumeko's PhD thesis [14] defended in July 2010.

- The Fractal reference Java implementation comes with a component factory that suffers a complex architecture and a lack of documentation. We developed a simple and powerful alternative implementation of the Fractal factory, also featuring extensions defined by ProActive/GCM. The associated software development project is called FFF [1]. This factory should allow us to exploit the new ADL for non-functional components we designed in this work.

6.2.2. Extension of the Fractal ADL for parametric topologies

Participants: E. Madelaine, A. Rouini.

We have studied an extension of the Fractal Architecture Description Language (ADL) addressing the definition of parametric topologies of components. Such an extension would be very convenient for enhancing reusability of ADL description, in many cases of classical patterns of distributed components (master/slaves, farms, rings, trees, matrices). The language extension proposal includes XML elements and tags for parameter and parameter domain definition, as well as specific syntax for defining the dataflow of parameters between components, interfaces, and bindings [43].

6.2.3. Monitoring and Management of Components

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

We have designed a monitoring and management framework that profits of the componentized membrane of GCM to separate the concerns of the MAPE (Monitoring, Analyze, Plan, Execute) loop, and provides a flexible way to add monitoring and management concerns at runtime to component-based applications and can be generalized to SCA compliant applications.

- A brief description of the monitoring capabilities is presented in [32] including an initial performance evaluation.

- A broader description of the framework and their involved components for Monitoring, SLA Analysis, Decision, and Execution is shown in [31], along with a generalization for SCA.

- We have submitted a design for a framework that covers the lifecycle of a service application from business and design level to deployment and execution concerns in a cloud environment. This work has been done in conjunction with Adrian Mos and Alain Boulze from INRIA Rhône-Alpes.

6.2.4. Enacting large-scale service composition using a component-based approach

Participants: F. Baude, V. Legrand.

---

[1] http://www-sop.inria.fr/members/Luc.Hogie/fff/

The distribution of business processes encompasses the inclusion of external service providers in the overall process as well as the usage of external infrastructures like clouds. Both of these approaches lead to decentralization and outsourcing of a part of the global workflow, resulting in a complexified management of the global orchestration. As a matter of fact, the overall data are decentralized among different domains and must, most of the time, be gathered manually. To this extent, we continue our work on agile and distributed orchestration, showing that the framework we develop eases multidomain orchestration management. Our approach extracts, gathers and digests data from the decentralized processes in order to provide an unified and global view of a distributed orchestration. This year we focus in particular on:

- **The specification and implementation of a tool** able to transform automatically a BPEL orchestration to a GCM component: Given a BPEL file, we obtain a GCM component embedding a workflow engine and managing the orchestration.

- **The test of real use-cases**.

- **The enactement of a federated workflow engine**, possibly composed of heterogeneous engines (We successfully test our framework with a federation of ActiveBPEL and Apache ODE engines).

6.2.5. *Hierarchical Grid/Cloud Programming through a component-based approach*
Participants: E. Mathias, F. Baude.

We explore GCM/ProActive component framework as the basis of a generic middleware for hierarchically programming applications or coupling legacy applications in multi-domain environments composed by private clusters, Grids and Clouds. The main idea behind this middleware is to offer a modular infrastructure that can be composed hierarchically, according to resources topology, and dynamically, according to the available resources. This middleware works as a glue between application processes running in different domains, featuring mechanisms like topology-aware point-to-point and collective communication. Our middleware grounds from the GCM (the Grid Component Model) and the ProActive Grid middleware, that we improved with features, such as: generic gathercast (Mx1) and multicast (Mx1) communication semantics, gather-multicast (MxN) component interfaces, MxN shortcuts, automated deployment, and communication tunneling and forwarding.

This middleware has been successfully used as the base of the implementation of the SOA4All fDSB (section 8.3.4) and the DiscoGrid Runtime (concluded in 2009). Main publications around this middleware are:

- a book chapter has been written on the principles, systems and applications of cloud computing [36];

- two publications in international conferences on cloud computing [18] and middleware for distributed computing [30] have been published;

- besides, the PhD thesis of Elton Mathias includes aspects related to this middleware will be defended in December 2010.

6.2.6. *Federating DSBs at Internet Scale Upon a Component-Based Approach*
Participants: F. Baude, E. Mathias, V. Legrand, C. Ruz.

The underlying solution we've pushed forward to integrate resources in a seamless manner is based on GCM components. It has already been successfully employed to couple Petals Distributed Service Buses deployed in different domains. The general architecture and concepts of the ESB federation, and associated performance results which demonstrate a low overhead are described in [21], [41], [39], [44], [13].

6.3. A Middleware for Grid and Cloud computing

6.3.1. *RDF Data Storage and Retrieval In P2P Systems*
Participants: I. Filali, F. Huet, F. Baude, F. Bongiovanni, L. Pellegrino.

We have proposed in the context of the SOA4ALL FP7-IP project the design and the implementation of a hierarchical Semantic Space infrastructure based on Structured Overlay Networks (SONS) [34]. It aims at the storage and the retrieval of the semantic description of services at the Web scale [40]. This infrastructure combines the strengths of both P2P paradigm at the architectural level and the Resource Description Framework (RDF) data model at the knowledge representation level. As it is designed, the proposed infrastructure enables the processing of simple and complex queries.

- A deep study on the existing works that have adapted the combination of RDF data model and the P2P communication model to build distributed infrastructures for RDF data storage and retrieval has been performed. This effort was the subject of a journal paper under submission at LNCS Transactions on Large-Scale Data- and Knowledge-Centered Systems. An extended version of this work spawned a research report that can be found in [38].

- An article, entitled "CAN-Based Approach for RDF Data Management in Structured P2P Systems" was written and will be submitted in forthcoming weeks. The work done in this paper introduces an implementation of a three dimensional CAN overlay network for storing and retrieving RDF triples. At the implementation level, a modular and flexible architecture for the Semantic Space infrastructure has been proposed. The implementation relies on the Pro-Active Grid middleware and provides a clear separation between its sub-components (overlay, storage, query engine, etc.). The modularity of the architecture is combined with the decentralized aspect of the infrastructure enabling the RDF data storage and retrieval at large scale. The evaluation of the infrastructure through extensive experiments on clusters and grids shows the impact of the architecture and data distribution on the performance of the storage and processing mechanisms.

6.3.2. *Matlab/Scilab parallel programming*
Participant: F. Viale.

Matlab & Scilab, with millions of users around the world are industry standard for numerical computing. They both lack a powerful and modern parallel computing framework to meet the industry's growing demand in terms of parallel processing. This project is intended to integrate into both softwares a toolbox for parallel processing, based on ProActive.

- This year, we extended ProActive Matlab's toolbox with dataspace capabilities (automatic data files, source code and workspace environment transfering). We extended the compatibility and reliability of the toolbox to work on Linux and Windows both 32 and 64 bits. We also allowed it to work across heterogeneous platforms (Linux to Windows and Windows to Linux).

6.3.3. *Losing rendezvous*
Participant: B. Amedro.

To manage the concurrency on the request content, we extended the programming model with the ForgetOn-Send primitive and the extension of the Wait-by-necessity. This way we permit to define some candidates to delegate the rendezvous in concurrence with the computation. We also distinguish between the different types of requests, according to their characteristics and behavior. In particular, we defined the sterility, as a particular characteristic of request from which a causal ordering disruption can be controlled. Then, we are able to combine all these solutions to enable the delegation of the rendezvous in order to have non-blocking request sendings.

- This work has been presented as an invited talk at the 4th INRIA-Illinois Joint laboratory workshop [46].
- The PhD thesis of Brian Amedro includes those aspects and will be defended in December 2010.

6.3.4. *HPC on clouds*
Participants: B. Amedro, F. Baude, F. Huet, E. Mathias, O. Smirnov.

We explores in details the relevance of using hybrid Grid/Cloud environments and the technical challenges that such mixing raise at the middleware level. In particular, we analyses performance results we obtained on commercial clouds such as Amazon EC2 running some of the NAS parallel benchmarks. We also work on two complementary and relevant concepts, Cloud bursting and Cloud seeding.

Most realizations of this year include publication of this work:

- a book chapter has been written on the principles, systems and applications of cloud computing [36].
- a publication to an international conference on cloud computing has been published [19]

6.3.5. *Grid Computing for Computational Finance*
Participants: F. Baude, V. D. Doan.

- In 2010, we obtained new results regarding the valuation of American option pricing, relying upon up to 40 underlying assets. The prices of such options are the first one published in the public domain.
- More importantly, the use of PicsouGrid master/slave framework to execute the billions of needed MonteCarlo simulations prove invaluable. Results using clusters using up to 256 nodes demonstrated near to optimal speedup, thus making it now possible to envision to price such very complex options within a handful of hours instead of days/nights.
- All details are exposed in the PhD of Viet Dung Doan [12].
- Paper [25] has been definitely accepted in the final proceedings of the conference held in 2009, to be published in 2010, and the article submitted to Mathematics and Computer in Simulation has been finally accepted and is under press [15].

6.3.6. *Network Aware Cloud Computing*
Participants: F. Huet, S. Malik, Denis Caromel.

We are working on a cloud scheduler module named Network Awareness Module (NAM), which helps the scheduler to make the more efficient scheduling decisions on the basis of resource characteristics. The scheduling decision is done on the basis of resource characteristics like reliability, network latency, bandwidth, error rate, topology, proximity, processing power, fault tolerance, memory availability, library availability, environment compatibility, and monetary cost of cloud service. The primary objective is to achieve this for the network latency, reliability, topology, environment compatibility and monetary cost issues. NAM has different sub-modules, which have implementation for models for different resource characteristics. It also has a repository, which contains information about cloud services and their characteristics available from different cloud operators.

- We have designed and implemented an algorithm for grouping of nodes on the basis of inter-node latencies. This algorithm works with the incomplete information about node latency and make groups heuristically.It features the following characteristics:
  – Dynamic grouping of nodes on the basis of incomplete latency information
  – Reformation of groups due to change in latency at run time
  – Updation of node latency by a mechanism similar to force based embedding

7. Contracts and Grants with Industry

7.1. Contracts with Industry

7.1.1. *AGOS*

"Architecture de Grille Orientée Services" is a project labelled by the *pôle de compétitivité* SCS ("Solutions Communicantes Sécurisées"), and financed by *FCE Ministère de l'Industrie* (from October 2007 to March 2010).

AGOS is a development project integrating and standardising a scientific approach (INRIA) and an industrial approach (HP and Oracle) of two innovative technologies: Grid computing and service oriented architecture. AGOS defines such a generic functional integration architecture. AGOS delivers also a secured software platform providing the following:

- A library of services based on standards;
- A set of tools to build comprehensive applications both Grid and SOA compliant, with their associated operational and business process monitoring in real time;
- A methodology expertise to build on or migrate to this architecture.

7.1.2. Compatible One

The OASIS Team is involved in the COMPATIBLE ONE national project. The topic of this project is the Cloud Computing. It aims to provide a set of tools and integrated software components for the IaaS, PaaS and SaaS layers, in order to offer low cost solutions, to create, deploy, administrate and to evolve private and public clouds. The keyword basis of this Cloud Computing framework are openness, interoperability and expandability. The team is mainly involved in the work packages 2.1 (PaaS Services, Scheduling and distributed computation), 2.3 (Storage and caching services), 2.4 (OSGI Runtime), 4.3 (3D distributed rendering use case). The project duration is 2 years and started on November 1st, 2010.

7.2. Grants with Industry

First, as a small collaboration involving an industrial partner, the thesis of Paul Naoumenko was situated in the context of a collaboration with France Telecom, interactions mainly occur on the domain of the M2M (Machine to Machine architecture) and service oriented computing.

Also the PhD thesis of Guilherme Peretti-Pezzi is financed as a CIFRE funding with the "Canal de Provence" company.

8. Other Grants and Activities

8.1. Regional Initiatives

8.1.1. Contrat Plan État Région Grille et Calcul Pair-à-Pair

This contract aims at building a regional computing platform. This is achieved by mixing desktop machines with dedicated ones like clusters. Users willing to submit a job will do so by accessing a web-page and uploading their program. It will then be scheduled and executed on a free machine. The scheduler is currently under development.

In the first part of the project, the access to the platform will be restricted to Inria members. Once most of the tools have been developed, the access will be open to industrial partners.

A convention has been signed with Microsoft to provide a specific cluster with Microsoft Compute Cluster Server.

The members of this project are the Inria and the Eurecom institute (Télécom Paris Ecole Polytechnique Fédérale de Lausanne).

The total budget for this project is 500kEuros for Inria and 100kEuros for Eurecom.

8.2. National Initiatives

8.2.1. *Soceda*

ANR Plateform project, ARPEGE, *SocEDA* "Plate-forme EDA sociale, largement distribuée pour l'informatique dans les nuages. Partners: EBM WebSourcing (lead), ActiveEon, Ecole des Mines Albi, I3S CNRS (OASIS), CNRS/LIG (SARDES), INSA Lyon/LIRIS, France Telecom, INRIA Lille (ADAM), THALES communications. The goal of SocEDA is to build a flexible, elastic and efficient platform for handling events generated by services deployed on top of federated and distributed enterprise service buses. To scale and be highly configurable, the platform will be designed as a GCM based application built along Peer-to-peer principles for scalability and robustness, and deployed on Grid/Cloud hybrid environments. Publications/subscriptions will be organized along social links that may exist between the deployed services.

This project started November 2010 for 3 years

8.2.2. *ADT Galaxy (INRIA transversal action)*

The ADT Galaxy contributes to make INRIA a value-added player in the SOA arena, by designing and developing an Open Framework for Agile and Dynamic Software Architecture. This ADT will work for INRIA and INRIA's research project-teams direct benefit, and aims at pre-assembling technological bricks from various teams, projects and preparing them to be transferred through the open source software channel.

The ADT aims at providing an IT agile platform, built on dynamic software architecture principles, and fitting for flexibility, dynamical reconfiguration, adaptability, continuity, and autonomic computing. Fractal, SCA-Tinfi and GCM/ProActive are the major technological drivers of this ADT. The different usage scenarios as well as the different tools developed at infrastructure, application, and business levels demonstrate that this platform is able to support the design, modelling, deployment, and execution of business processes. In the same time, the ADT targets the definition of a new common language to manipulate dynamically adaptive distributed SOA-based systems, encompassing application, and middleware layers. This common language will take different forms, inherited from works done by several project-teams with their distinct skills, and illustrates a new kind of collaboration between teams, coupling research and development works.

Contributors to this ADT are mainly research project-teams, including OASIS, ADAM (Lille), ECOO (Nancy), ASCOLA (Rennes), ObjectWeb/TUVALU (Grenoble), SARDES (Grenoble) and TRISKELL (Rennes), and the ADT Galaxy is led and managed by the TUVALU team.

The duration of this ADT is over 28 months : the kickoff meeting has been held on July 3rd, 2008 and ended at the end of October, 2010

8.2.3. *ANR OMD2*

The aim of this project is to define, develop and experiment a collaborative platform of mutidisciplinary optimization - As "platform" we consider here a software environment hosting heterogeneous code and data, geographically distributed in equally heterogeneous machines. These codes can be sequential or parallel. These machines can be data servers, supercomputers, PC farms, etc. - As "collaborative" we consider this environment to be able to host, control and allow communication of these codes transparently for the user, according to their own work habits. - As "multidisciplinary optimization", we consider the collection of methods and numerical tools, objects of the other tasks in OMD2 and previously realized during project OMD. It has been decided that ProActive will be used as middleware for communication between machines, and that the ProActive Scheduler will be used as a basis for the collaborative platform. Similarly, Scilab will be used as a common language to describe optimization strategies.

The project started January 1st 2009 and will last 3 years, with a total budget of 214 kEuros.

8.3. European Initiatives

8.3.1. *IP TEFIS*

TEstbed for Future Internet Services (TEFIS) is a Large-Scale Integrating Project funded by the European Seventh Framework Programme, under the ICT research area and, more precisely, around the theme of Future Internet Experimental Facility and Experimentally-driven Research.

Internet is more and more used by services and applications as the common communication infrastructure. However, the Internet and Software industry is facing two main challenges. The first one is that Internet is progressively reaching a saturation point in meeting an ever increasing variety of user expectations. The second one is that the increasing complexity of ICT environments (user communities, network heterogeneity, multiplicity of terminals and platforms) is paralysing testing processes for the assessment of Internet-based services at large-scale. Product and business developers must be able to rapidly satisfy and anticipate user requirements while testing and sizing their applications and services, being sure to constantly have the right amount of reactivity with regards to market demands.

To address these challenges, TEFIS will provide an open platform to enable the design, dimensioning and user-centric validation of innovative applications and services on top of upcoming Future Internet technologies. The TEFIS platform will be a central access point to heterogeneous and complementary experimental facilities addressing the full development lifecycle of innovative services with the appropriate tools and testing methodologies.

The OASIS team is involved in work packages 2 (Platform architecture and engineering), 3 (TEFIS portal), 4 (Core services), 6 (Experimental Data Lifecycle Management), 8 (Dissemination, Exploitation and Community Engagement), 9 (FireStation coordination).

The project will have a duration of 30 months and it started on June 2010.

8.3.2. STREP PLAY

STREP FP7, Call 5, *PLAY* " Pushing dynamic and ubiquitous interaction between services Leveraged in the Future Internet by ApplYing complex event processing" . Partners: FZI (lead), EBM WebSourcing, INRIA (OASIS and SARDES), France Telecom, ICCS, Ecole des Mines Albi, CIM. The goal of PLAY is to develop and validate an elastic and reliable federated SOA architecture for dynamic and complex, event-driven interaction in large highly distributed and heterogeneous service systems. Such architecture will enable exchange of contextual information between heterogeneous services, providing the possibilities to optimize/personalize the execution of them, resulting in the so called situational-driven adaptivity.

This project started October 2010 for 3 years

8.3.3. EIT ICT Labs

EIT ICT Lab is a new lab, under construction, at the European scale. OASIS is being involved in an RTD activity part of the theme "Future Media and Content Delivery", entitled End-to-End Cloud Infrastructure for Media and Content Delivery. Request funding and planned dissemination actions have strong complementarities with those conducted by the team within the PLAY and TEFIS, recently started FP7 projects.

8.3.4. IP SOA4All

Service Oriented Architectures for All (SOA4All) is a Large-Scale Integrating Project funded by the European Seventh Framework Programme, under the Service and Software Architectures, Infrastructures and Engineering research area.

Computer science is entering a new generation. The emerging generation starts by abstracting from software and sees all resources as services in a service-oriented architecture (SOA). In a world of services, it is the service that counts for a customer and not the software or hardware components which implement the service. Service-oriented architectures are rapidly becoming the dominants computing paradigm. However, current SOA solutions are still restricted in their application context to being in-house solutions of companies. A service Web will have billions of services. While service orientation is widely acknowledged for its potential to revolutionize the world of computing by abstracting from the underlying hardware and software layers, its success depends on resolving a number of fundamental challenges that SOA does not address today.

SOA4All will help to realize a world where billions of parties are exposing and consuming services via advanced Web technology: the main objective of the project is to provide a comprehensive framework and infrastructure that integrates complementary and evolutionary technical advances (i.e., SOA, context management, Web principles, Web 2.0 and Semantic Web) into a coherent and domain-independent service delivery platform.

OASIS is involved in work packages 1 (SOA4All Runtime), 2 (Service Deployment and Use) and 6 (Service Construction).

We strongly collaborate with the ObjectWeb/TUVALU EPI and also ADAM, from which Philippe Merle is co-leading with Françoise Baude the PhD thesis of Virginie Legrand-Contes.

8.3.5. *IP Bionets*

The OASIS team is involved in the European project called BIONETS (BIOlogically-inspired autonomic NETworks and Services)

The motivation for BIONETS comes from emerging trends towards pervasive computing and communication environments, where myriads of networked devices with very different features will enhance our five senses, our communication and tool manipulation capabilities. The complexity of such environments will not be far from that of biological organisms, ecosystems, and socio-economic communities. Traditional communication approaches are ineffective in this context, since they fail to address several new features: a huge number of nodes including low-cost sensing/identifying devices, a wide heterogeneity in node capabilities, high node mobility, the management complexity, and the possibility of exploiting spare node resources. BIONETS aims at a novel approach able to address these challenges. BIONETS overcomes device heterogeneity and achieves scalability via an autonomic and localised peer-to-peer communication paradigm. Services in BIONETS are also autonomic, and evolve to adapt to the surrounding environment, like living organisms evolve by natural selection. Biologically-inspired concepts permeate the network and its services, blending them together, so that the network moulds itself to the services it runs, and services, in turn, become a mirror image of the social networks of users they serve.

The team is involved in work packages 3.1 (Requirement Analysis and Architecture),3.2 (Autonomic Service Life-Cycle and Service Ecosystems), and 3.4 (Probes for Service Framework).

The project started in 2006, for 48 months, for a total budget of 127 kEuros. Project terminated in March 2010.

8.4. International Initiatives

8.4.1. *Stic Asia*

Stic Asia is a multilateral project with universities of BUPT (Beijing, China), Tsinghua (Beijing, China), SCUT (Ghuanzhou, China), and NUST (Pakistan). *Experiments and Dissemination on Grid Standard: ProActive GCM*, is a collaborative research and academic exchanges project on Grid standard between Inria and the partners from Asia. It is partially funded by French ministry of Foreign affairs starting from July 2007 and will finish in December 2010.

The main objective of this project is to foster international scientific cooperation in Grid research between France and Asian partners, share experiments and disseminate ProActive and Grid Component Model (GCM) standard for Grid Middleware and applications interoperability. Furthermore, it is intended to support and establish partnership from mobility programs in a short and long term.

8.4.2. *ANR Blanc International MCore-PhP*

This is a 3 years bilateral project started in march 2010, between the Oasis EPI and Tsinghua university in Beijing (China). The full title is "Multi-Core Parallel Heterogeneous Programming". In this project we shall study new programming models for heterogeneous infrastructures, from multicore and GPUs to clouds, support for such programming libraries on the large chinese grid infrastructures (China-Grid), and application to large use-cases in bioinformatics.

The project started in March 2010, for 36 months, with a budget of 166 KEuros for the Oasis EPI.

9. Dissemination

9.1. Animation of the scientific community

9.1.1. Program Committees and Conference Organisation

- **Ludovic Henrio**
  Program Committee member of FESCA'10, FOCLASA'10, FESCA'10.
- **Eric Madelaine**
  Steering committee member of FACS'10 and FMCO'10
  Program committee member of conferences FACS'10 and FMCO'10, and reviewer for the journal SCP (Science of Computer Programming)
- **Françoise Baude**
  Reviewer for IEEE Transactions on Parallel and Distributed Systems
  PC member for Workshop for Component-Based High-Performance Computing (CBHPC2010) joint with GRID 2010
- **Fabrice Huet**
  Program Committee member of PCGrid 2010, Workshop on Component-Based High Performance Computing 2010

9.1.2. Short Visits and external collaborators

- Florian Kammüller, Technische Universität Berlin, now at Middlesex University, London. Florian works with us in the context of several common developments and joint works in the domain of mechanised formalisation of programming languages and framework.
- WU Nuo from Tshinghua Un., Bejing, China, 3 months visit
- SUN Dou from Beihang Un., Bejing, China, 3 months visit

9.1.3. Theses

The following theses are in preparation:

- Muhammad Uzair Khan: "Supporting First Class Futures in a Fault-Tolerant Java Middleware" (Since Oct 2007), director Ludovic Henrio.
- Marcela Rivera: "Reconfiguration and Life-cycle of Distributed Components: Asynchrony, Coherence and Verification" (since Dec 2006), director Ludovic Henrio.
- Imen Filali: "Peer to Peer computational Grids with reservation in service oriented architectures" (Since Oct 2007), director Fabrice Huet.
- Virginie Legrand-Contes: "Large Scale and Distributed Services Orchestration" (Since Mar 2008), directors Françoise Baude and Philippe Merle
- Cristian Ruz: "Autonomic Service Deployment and Management of grid-based enterprise services" (Since Mar 2008), director Françoise Baude.
- Guilherme Peretti-Pezzi: "ProActive Parallel Hydraulic Simulations for Grid and SOA Environments" (Since October 2008), director Denis Caromel, CIFRE funding with the "Canal de Provence" company.
- Francesco Bongiovanni: "Self organizing overlay networks and generic overlay computing systems" (Since Oct 2008, joined OASIS in Sept 2009), directors Françoise Baude and Fabrice Huet.
- Laurent Pellegrino "Pushing dynamic and ubiquitous event-based interaction in the Internet of services: a middleware for event clouds" (since Sept 2010), director Françoise Baude.
- Sheheryar Malik "Network Aware Cloud Computing" (Since Oct 2009), Supervisors: Denis CAROMEL, Fabrice HUET

9.1.4. Member of thesis Committees

- **Fabrice Huet**
  Reviewer and member of PhD thesis committee of Niels Drost (25/11/10), "Real-World Distributed Supercomputing", under the supervision of Prof. Henri Bal, Amsterdam.

- **Françoise Baude**
  PhD reviewer and defense jury member : Jean Arnaud, SARDES, Univ. de Grenoble, Septembre 2010, and Ghislain Charrier, GRAAL, ENS Lyon, December 2010

9.1.5. Invited Talks

- **Brian Amedro** Invited talk to the 4th INRIA-Illinois Joint Lab workshop: " Improving Asynchrony in an Active Object Model"

- **Eric Madelaine** "Specification and Verification for distributed applications running on heterogeneous infrastructures" invited talk at East China Normal University, Dep. of Software Engineering, Shanghai, China.

- **Denis Caromel**
  "Cloud Computing Revolution", Invited Keynote at Exposition Universal 2010, Oct. 18, Shanghai, China.
  "From Multi-Cores to Clouds with ProActive Parallel Suite: Use cases in Biotech, IT, Finance, and Engineering". Invited keynote at conference GCC, The 9th International Conference on Grid and Cloud Computing, Nov. 3rd 2010, Nanjing, Jiangsu, China.
  "Cloud Computing", Invited Seminar at Nanjing University, Key laboratory of Software technology, October 20th 2010.
  "ProActive Parallel Suite", Invited Seminar at Tongji University, Shanghai, October 22nd 2010.
  "ProActive Parallel Suite and OW2 OSCi initiative: from Multi-Cores to Multi-Clouds", Invited keynote at OW2 Annual Conference, November 24, Paris.
  "CLOUD Computing for Digital Cities", invited presentation at ICT Labs, Helsinki, Findland, September 2nd, 2010.
  "Bridging Multi-Core and Distributed Computing: All the Way Up To the Cloud", Invited keynote at conference 10th WSEAS International Conference on Applied Informatics and Communications (AIC '10), Taipei, Taiwan, August 20-22, 2010.
  "Cloud Open Source", invited presentation at Journée Rencontres Pole SCS, June 18th 2010, Gardanne, France.
  "Clouds with ProActive Parallel Suite on Windows Together with Legacy Linux Clusters", invited Presentation at Microsoft Booth, June 2nd 2010, ISC Conference, Hamburg, Germany.
  "Bridging Multi-Core and Distributed Computing: all the way up to the Cloud", Invited Talk at Tsinghua UNiversity, Beijing, China, May 25th 2010.
  "Open Source pour simulations parallèles et réparties", invited presentation at Journée CSMA, may 2010, Université de Technologie de Compiègne, Centre de Recherches de Royallieu.

- **Ludovic Henrio** "Typing and Exceptions for Algorithmic Skeletons" Journée du GdT LaMHA (GDR GPL)

9.2. Teaching

- **Ludovic Henrio** Semantics of Distributed and Embedded Systems (Master 1), and involvement in a few Master 2 courses.

- **Brian Amedro** "Distributed Software Architectures" (Master 1); "Parallel and Distributed Computing" (Master 1); "Java Programming" (Licence 3)

- **Virginie Legrand-Contes** Course of "Web Services and Service Oriented Architectures" (Master 2 MIAGE, Nouvelles Technologies et Direction de Projets)

- **Fabrice Huet**
  Coordinator of the 1st year of Master of Computer Science,
  Course convenor of "Advanced Operating Syste"m (Master 1),
  "Parallel and Distributed Programming" (Master 1),
  "Distributed Systems" (Master 1) and "Network Game Programming".

- **Françoise Baude**
  Responsible for the "Distributed Algorithmics" course, Master 2 IFI, CSSR
  Co-director of the computer science department of UFR Sciences, UNS, till June 2010

- **Eric Madelaine** Course Convenor of Semantics of Distributed and Embedded Systems (Master 1)

- **Elton Mathias**
  Attaché Temporaire d'Enseignement et Recherche (ATER) at IUT - Université de Nice: "Programmation Réseaux, Internet et Web" (Licence Professionelle 3)

- **Marcela Rivera**
  full-time ATER at Université de Nice in 2009/2010. course in "Architecture", "système et réseaux", "Outils génie logiciel", and "système et réseaux". She is still ATER in 2010-2011.

- **Denis Caromel**
  Course convenor of Distributed Programming and Multi-Tiers Architectures (Master 1) with the involvement of Brian Amedro. Compute and data grids: large scale distributed systems, Master 2 Ubinet, Master of Science in Ubiquitous Networking and Computing.

10. 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. *Programming, Composing, Deploying for the Grid (chapter 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. *Towards a Method of Object-Oriented Concurrent Programming*, in "Communications of the ACM", September 1993, vol. 36, n^o 9, p. 90–102.

[7] D. CAROMEL, L. HENRIO. *A Theory of Distributed Object*, Springer-Verlag, 2005.

[8] 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 [*DOI :* 10.1145/964001.964012].

[9] D. CAROMEL, W. KLAUSER, J. VAYSSIÈRE. *Towards Seamless Computing and Metacomputing in Java*, in "Concurrency Practice and Experience", November 1998, vol. 10, n^o 11–13, p. 1043–1061.

[10] 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

[11] B. AMEDRO. *Optimizing an Active Object Framework: From Multi-Cores to Multi-Clouds*, University of Nice-Sophia Antipolis, Dec 2010.

[12] 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/.

[13] E. MATHIAS. *Hierarchical Multi-Domain computing based upon a component-oriented approach*, University of Nice-Sophia Antipolis, Dec 2010.

[14] P. NAOUMENKO. *Designing Non-Functional Aspects with Components*, University of Nice-Sophia Antipolis, July 2010.

Articles in International Peer-Reviewed Journal

[15] V. D. DOAN, M. BOSSY, F. BAUDE, I. STOKES-REES. *Grid computing of Monte Carlo based American option pricing: Analysis of two methods*, in "Mathematics and Computers in Simulation", 2010, doi: 10.1016/j.matcom.2010.08.005.

Articles in Non Peer-Reviewed Journal

[16] F. BAUDE, V. LEGRAND, L. HENRIO, P. NAOUMENKO, H. PFEFFER, L. BASSBOUSS, D. LINNER. *Mixing Workflows and Components to Support Evolving Services*, in "International Journal of Adaptive, Resilient and Autonomic Systems (IJARAS)", 2010, To appear.

Invited Conferences

[17] L. HENRIO, F. KAMMÜLLER, M. U. KHAN. *A Framework for Reasoning on Component Composition*, in "FMCO 2009", Lecture Notes in Computer Science, Springer, 2010.

International Peer-Reviewed Conference/Proceedings

[18] B. AMEDRO, F. BAUDE, F. HUET, E. MATHIAS. *"Middleware level solutions for SPMD Applications in Grids and Clouds"*, in ""2nd IEEE International Conference on Cloud Computing Technology and Science"", Dec. 2010.

[19] B. AMEDRO, F. BAUDE, F. HUET, E. MATHIAS. *Middleware level solutions for SPMD Applications in Grids and Clouds*, in "2nd International Conference on Cloud Computing Technology and Science", Indianapolis, IN United States, 11 2010, http://salsahpc.indiana.edu/CloudCom2010/papers.html.

[20] B. AMEDRO, D. CAROMEL, F. HUET, V. BODNARTCHOUK, C. DELBÉ, G. TABOADA. *HPC in Java: Experiences in Implementing the NAS Parallel Benchmarks*, in "New Aspects of Applied Informatics, Biomedical Electronics and Informatics and Communications", Taipei, Taiwan, WSEAS, 08 2010, http://hal.inria.fr/inria-00504630/PDF/document.pdf.

[21] 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.

[22] R. BOULIFA, L. HENRIO, E. MADELAINE. *Behavioural Models for Group Communications*, in "WCSI-10: International Workshop on Component and Service Interoperability", 2010, to appear.

[23] A. CANSADO, L. HENRIO, E. MADELAINE, P. VALENZUELA. *Unifying Architectural and Behavioural Specifications of Distributed Components*, in "International Workshop on Formal Aspects of Component Software (FACS'08)", Malaga, Electronic Notes in Theoretical Computer Science (ENTCS), Sept 2010, vol. 260, p. 25-45.

[24] I. FILALI, F. HUET. *Dynamic TTL-Based Search In Unstructured Peer-to-Peer Networks*, in "IEEE International Symposium on Cluster Computing and the Grid (CCGRID 2010)", 2010, to appear, http://hal.archives-ouvertes.fr/inria-00485790/en/.

[25] A. GAIKWAD, V. D. DOAN, M. BOSSY, F. ABERGEL, F. BAUDE. *SuperQuant Financial-Benchmark Suite for Performance Analysis of Grid Middlewares*, in "4th Int. Conference on High Performance Scientific Computing, Modeling, Simulation and Optimization of Complex Processes", Hanoi, Vietnam, Institute of Mathematics, Vietnamese Academy of Science and Technology, March 2010, 70, Full version accepted for post-proceeding publication by Springer.

[26] A. GAIKWAD, V. D. DOAN, M. BOSSY, F. ABERGEL, F. BAUDE. *SuperQuant Financial Benchmark Suite for Performance Analysis of Grid Middlewares*, in "Post-proceedings of the 4th Int. Conference on High Performance Scientific Computing, Modeling, Simulation and Optimization of Complex Processes", Hanoi, Vietnam, March 2-6 2011, To appear.

[27] L. HENRIO, M. U. KHAN. *Asynchronous Components with Futures: Semantics and Proofs in Isabelle/HOL*, in "Proceedings of the Seventh International Workshop, FESCA 2010", ENTCS, 2010, http://hal.archives-ouvertes.fr/inria-00462325/en/.

[28] M. LEYTON, L. HENRIO, J. M. PIQUER. *Exceptions for Algorithmic Skeletons*, in "16th Int. European Conference on Parallel and Distributed Computing (Euro-Par 2010)", 2010, http://hal.archives-ouvertes.fr/hal-00486108/en/.

[29] L. LIQUORI, C. TEDESCHI, L. VANNI, F. BONGIOVANNI, V. CIANCAGLINI, B. MARINKOVIC. *Synapse: A Scalable Protocol for Interconnecting Heterogeneous Overlay Networks*, in "9th International IFIP-TC 6 Networking Conference (Networking 2010)", 2010.

[30] E. MATHIAS, F. BAUDE. *"Multi-domain Grid/Cloud Computing Through a Hierarchical Component-Based Middleware"*, in ""8th International Workshop on Middleware for Grids, Clouds and e-Science - MGC 2010. In conjunction with ACM/IFIP/USENIX 11th International Middleware Conference 2010"", Dec. 2010.

[31] C. RUZ, F. BAUDE, B. SAUVAN. *"Component-based generic approach for reconfigurable management of component-based SOA applications"*, in ""3rd Workshop on Service Monitoring, Adaptation, and Beyond (MONA+), in conjunction with ECOWS 2010 - The 8th IEEE European Conference on Web Services "", Dec. 2010, to appear.

[32] C. RUZ, F. BAUDE, B. SAUVAN. *"Enabling SLA Monitoring for Component-Based SOA Applications – A Component Based Approach"*, in ""SEAA'10: 36th Euromicro Conference, Software Engineering and Advanced Applications, Work in Progress Session"", Sep. 2010.

[33] C. RUZ, F. BAUDE, B. SAUVAN. *Enabling SLA Monitoring for Component-Based SOA Applications – A Comp onent-Based Approach*, in "36th Euromicro Conf. on Software Engineering and Advanced Applicati ons, SEAA, Work In Progress Session", 2010.

National Peer-Reviewed Conference/Proceedings

[34] I. FILALI, F. HUET, L. PELLEGRINO. *RDF Data Based Storage and Retrieval In Structured P2P Networks*, in "11èmes Journées Doctorales en Informatique et Réseaux", 2010.

[35] L. HENRIO, E. MADELAINE. *Experiments with distributed Model-Checking of group-based applications*, in "Sophia-Antipolis Formal Analysis Workshop", France Sophia-Antipolis, Oct 2010, 3p., http://hal.inria.fr/inria-00538499/en.

Scientific Books (or Scientific Book chapters)

[36] B. AMEDRO, F. BAUDE, D. CAROMEL, C. DELBÉ, I. FILALI, F. HUET, E. MATHIAS, O. SMIRNOV. *An efficient framework for running applications on clusters, grids and clouds*, in "Cloud Computing: Principles, Systems and Applications", Springer, 2010, ISBN: 978-0-387-47656-8.

[37] L. HENRIO, M. U. KHAN, N. RANALDO, E. ZIMEO. *First Class Futures: Specification and implementation of Update Strategies*, Coregrid Series, Springer, August 2010.

Research Reports

[38] I. FILALI, F. BONGIOVANNI, F. HUET, F. BAUDE. *RDF Data Indexing and Retrieval: A survey of Peer-to-Peer based solutions*, INRIA, Nov. 2010, n° RR-7457, http://hal.archives-ouvertes.fr/inria-00540314/en/.

[39] C. HAMERLING, E. MATHIAS. *SOA4All Runtime v2*, SOA4All IP Project Delivrable from the SOA4All Runtime WP (WP1), Aug 2010, n° D1.4.2B.

[40] R. KRUMMENACHER, F. HUET, M. FRIED, L. PELLEGRINO, I. PEIKOV, A. SIMOV. *A Distributed Semantic Marketplace*, SOA4All IP Project Delivrable from the SOA4All Runtime WP (WP1), Feb 2010, n° D1.3.3A.

[41] R. KRUMMENACHER, J.-P. LORRÉ, C. HAMERLING, A. DUJE, M. VILLA, F. BAUDE, E. MATHIAS, V. LEGRAND, C. RUZ, D. LIU, C. PEDRINACI, T. P. LOBO, M. DIMITROV, P. MERLE. *Final SOA4All Reference Architecture Specification*, SOA4All IP Project Delivrable from the SOA4All Runtime WP (WP1), Feb 2010, n° D1.4.2A.

[42] L. LIQUORI, C. TEDESCHI, L. VANNI, F. BONGIOVANNI, V. CIANCAGLINI, B. MARINKOVIC. *Synapse: a Scalable Protocol for Interconnecting Heterogeneous Overlay Networks*, INRIA, 04 2010, n[o] RR-7255, http://hal.inria.fr/inria-00474529/PDF/RR-7255.pdf.

[43] A. ROUINI. *Parametric Component Topologies: language extension and implementation*, Master Ubinet, Univ. od Nice Sophia Antipolis, 2010.

[44] B. SCHREDER, R. KRUMMENACHER, J. P. MARTINEZ, M. VILLA, G. D. MATTEO, F. HUET, E. MATHIAS. *Testbeds Validation*, SOA4All IP Project Delivrable from the SOA4All Runtime WP (WP1), Dec 2010, n[o] D1.5.3.

Patents and standards

[45] 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.

Other Publications

[46] B. AMEDRO. *Improving Asynchrony in an Active Object Model*, November 2010, Invited talk at 4th INRIA-Illinois Joint Laboratory for Petascale Computing Workshop, Nov 22-24, NCSA/University of Illinois, USA.

[47] OASIS TEAM. *ProActive parallel suite*, 2010, http://proactive.inria.fr/.

[48] OASIS. TEAM. *The VERCORS platform: VERification of models for distributed communicating COmponants, with safety and Security*, 2010, http://www-sop.inria.fr/oasis/index.php?page=vercors.

References in notes

[49] M. ABADI, L. CARDELLI. *A Theory of Objects*, Springer-Verlag, New York, 1996.

[50] L. BADUEL, F. BAUDE, D. CAROMEL, A. CONTES, F. HUET, M. MOREL, R. QUILICI. *Programming, Composing, Deploying for the Grid (chapter 9)*, in "Grid Computing: Software Environments and Tools", O. F. CUNHA (editor), Springer, 2006, ISBN: 1-85233-998-5.

[51] F. BAUDE, L. HENRIO, P. NAOUMENKO. *A Component Platform for Experimenting with Autonomic Composition*, in "First International Conference on Autonomic Computing and Communication Systems (Autonomics 2007)", ACM Digital Library, Oct 2007, Invited Paper.

[52] 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.

[53] 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, http://www.coregrid.net/mambo/images/stories/TechnicalReports/tr-0003.pdf.

[54] L. HENRIO, F. KAMMÜLLER, M. RIVERA. *An Asynchronous Distributed Component Model and Its Semantics*, in "FMCO 2008", F. S. DE BOER, M. M. BONSANGUE, E. MADELAINE (editors), Lecture Notes in Computer Science, Springer, dec 2009, vol. 5751, p. 159-179, http://www.springer.com/978-3-642-04166-2.