AERES

Campagne d'évaluation 2011-2014
Unité de recherche

Laboratoire de l'Informatique du Parallélisme

UMR n°5668

Gilles Villard

CNRS
École Normale Supérieure de Lyon
INRIA
Université Claude Bernard Lyon 1

Bilan scientifique et annexes

Laboratoire de l'Informatique du Parallélisme

LIP, École normale supérieure de Lyon - 46, allée d'Italie, 69364 Lyon cedex 07

Sept. 2005-2009 Progress report and production list

LIP documents concerning the evaluation (extended progress report, publications, etc.)
will be available at http://www.ens-lyon.fr/LIP/Eval

Introduction

Le Laboratoire de l'Informatique du Parallélisme (LIP), crée à l'ENS en 1988, est devenu une unité mixte de recherche (UMR 5668) par association avec le CNRS en 1989, l'INRIA en 1999 et l'Université Claude Bernard Lyon 1 en 2003. Le laboratoire rassemble plus de 120 scientifiques (chercheurs, enseignants-chercheurs, doctorants et ingénieurs) organisés en six équipes de recherche et secondés par une équipe administrative et technique de 11 personnes. Les équipes sont localisées sur le site lyonnais de Gerland (bâtiments de l'ENS et de l'UCBL).

De l'informatique fondamentale et mathématique jusqu'à des développements logiciels et matériels innovants, la recherche au LIP couvre un large spectre de domaines : nouvelles architectures et infrastructures de calcul, internet du futur, calcul parallèle, protocoles de communication, réseaux dynamiques, ordonnancement, services et composants, évaluation de performance, compilation, synthèse, optimisation de code, calcul certifié, calcul algébrique, cryptographie, algèbre linéaire, réseaux Euclidiens, complexité, complexité algorithmique, systèmes à événements discrets, automates, graphes, modèles de programmation, preuve formelle, logique, etc. Plusieurs volets des activités du laboratoire sont également tournés vers d'autres disciplines scientifiques par le biais de l'institut des systèmes complexes (IXXI), d'interactions avec les sciences du vivant, et du calcul scientifique.

Les chercheurs du LIP ont une forte activité avec près de 600 livres, articles de revues ou de conférences et plus de 50 thèses et habilitations soutenues en quatre ans (près de 900 publications et interventions au total). Nos collaborations internationales se traduisent par la cosignature de chercheurs dans plus de 40 universités étrangères. La production logicielle est très diversifiée, depuis les bibliothèques de calcul jusqu'aux logiciels et outils à très large diffusion. Pendant le quadriennal, les chercheurs du LIP ont participé à plus de 100 projets nationaux et internationaux (Région, ANR, Ministère, CNRS, INRIA, Union Européenne, etc.). Ces programmes constituent la première source de financement. Le LIP est partenaire des pôles de compétitivité Minalogic et System@tic et s'implique largement en valorisation et transfert vers les industries de communication, de semi-conducteurs et de calculateurs (Alcatel-Lucent, IBM, STMicroelectronics, etc.). Trois startups émergent du laboratoire en 2009.

Une moitié approximativement des membres permanents du LIP sont enseignants-chercheurs. Avec la participation des chercheurs et des doctorants, les équipes assurent des enseignements dans un large éventail de niveaux à Lyon, et, dans une moindre mesure, à l'extérieur.

Organigramme du laboratoire, juin 2009

Laboratoire LIP

Direction : G. Villard
Direction adjointe : D. Hirschkoff
P. Vicat-Blanc

Commission des chefs d'équipe et de service

Bureau de direction

Direction
Membres invités
Réunions bimensuelles

Conseil de laboratoire

Réunions bimensuelles
Formation en assemblée générale trimestrielle

Ingénieurs, administratifs, techniciens

Correspondant : S. Torres

Ingénieurs : É. Boix, M. Imbert, J.-C. Mignot, D. Ponsard, S. Torres

Responsable administratif : C. Iafrate
J. Bresle, M. Buillas, C. Suter, (C. Desplanches)

Coordination INRIA/LIP : S. Boyer

Conseil des formations

É. Fleury

Commission des habilités

J.-M. Muller (ED) F. de Dinechin

Équipes de recherche

42 chercheurs et enseignants-chercheurs permanents 41 doctorants
15 invités, délégations, postdocs, ATER 21 ingénieurs
4 ingénieurs de recherche

Arénaire INRIA
Arithmétique des ordinateurs
G. Villard 8.5 permanents
20.5 membres

CompSys INRIA
Compilation, systèmes enfouis et calcul intensif
A. Darte 4 permanents
7 membres

Graal INRIA Algorithmique et ordonnancement pour plates-formes distribuées
F. Vivien 12.5 permanents
37 membres

MC2
Modèles de calcul et complexité
P. Koiran 5 permanents
13.5 membres

Plume
Théorie de la preuve et sémantique formelle
O. Laurent 6 permanents
12 membres

Reso INRIA Protocoles et logiciels optimisés pour réseaux haut débit
P. Vicat-Blanc 9 permanents
32 membres

Contents

1. LIP progress report

LIP, *Laboratoire de l'Informatique du Parallélisme*, is a joint research department with four heading organizations. This report summarizes its activities from Fall 2005, the date of the previous document, to September 2009. LIP has been founded at ENS—*École Normale Supérieure de Lyon*—in 1988 and associated with CNRS—*Centre National de la Recherche Scientifique*—since 1989, INRIA—*Institut National de Recherche en Informatique et Automatique*—since 1999, and UCBL—*Université Claude Bernard Lyon 1*—since 2003. About 130 people participate to its activities. There are 54 permanent members whose affiliations—12 CNRS, 17 ENS, 17 INRIA, 8 UCBL— show the balanced investment of the partners, and reflect the diversity and wealth of our research and teaching. Our main strength is the creative interaction between long-term fundamental research, innovative software and hardware design, and shorter-term projects and transfer through industrial collaborations. In most LIP thematics, this interaction fosters new trends, both theoretical and practical. Mathematical computer science and algorithmics are transverse to the lab. Along with our links to numerical and computational sciences, this illustrates our special partnership with mathematics. The last four-year policy has been to strengthen the six existing research teams, to encourage transverse and new directions, and to support collaboration with other disciplines, especially through complex system modeling. This is described and analyzed in this chapter. Subsequent chapters will give further details for research teams (Chapt. 2-7) and for the IT support team (Chapt. 8).

1.1 Scientific overview

Since our previous four-year plan report, LIP has 56% more members, reaching a total of 133. Almost one third of the permanent members in 2005 have left the lab, and 140 people have joined for at least three months. This high turnover reflects our vitality and attractiveness, and is illustrated by major progress and evolution in every team. LIP has focused on strengthening six main directions with high potential. Internally, this corresponds to an organization in six teams, among which four are joint project-teams with INRIA: Arénaire, Compsys, Reso, and Graal. These four joint projects represent about 80% of the scientific members. Hence, a large part of LIP is directly concerned by three or four heading organizations. Nevertheless, balancing the involvement of the four organizations among the teams, especially regarding MC2 and Plume who have no INRIA members, is a point that we keep in mind.

The support provided to the teams, illustrated by 13 new permanent researchers and faculty members, and one permanent engineer, has led to a qualitatively and quantitatively large production. We have produced about 600 peer-reviewed publications and over 50 doctoral and habilitation theses (about 900 publications and communications in total).

We have co-signed papers with researchers from over 40 universities abroad. We have taken responsibilities for more than 350 scientific events. About 30% of our operating (non-consolidated) budget is balanced between CNRS-ENS recurring funds, INRIA support to project-teams, and various national and international risk/specific actions funded by the four heading organizations. Hence about 70% of our income comes from external sources, the main one being the National Research Agency (ANR) with a participation to more than 30 collaborative projects. New fundamental directions emerge from this vitality, such as: HPC geometry of numbers and cryptography; applications of linear logic to implicit computational logic; algorithmics of discrete event systems; just-in-time compilation and source-to-source optimizations for high-level synthesis; static scheduling of dynamic systems described with probabilistic models; models and algorithms for Software as a Service (SaaS) platforms, etc. The quality of our doctoral program is assessed by the fact that over 70% of the PhD candidates that have defended between 2000 and 2008 have found permanent research or faculty positions.

The success of the four-year plan is also emphasized by large transverse initiatives that LIP has encouraged. The joint research around parallel computing and networking has pushed the convergence of the communication, computation, and storage aspects. Our involvement in various Grid projects and INRIA-Bell labs laboratory, together with our collaborations within the competitiveness cluster System@tic, demonstrates our R&D strengths in the domain. Our initiatives on highly-optimized arithmetic libraries and on compilation have led to a well-established partnership with STMicroelectronics, in particular through the competitiveness cluster Minalogic and the Nano2012 funding mechanism, with a quick transfer of our knowledge into STMicroelectronics compilers and tools. Beyond the rich spectrum of our software production, from widely disseminated middlewares to sophisticated and dedicated libraries, we note that three software-related start-ups (each recently awarded a national prize) have emerged from LIP.

LIP was involved in the main inter-disciplinary initiative that has been the Complex System Institute (IXXI) that is now a recognized structure, independent from LIP. LIP members, mostly from the MC2 team, have actively participated to the creation of the institute, and applying computer science tools for modeling complex phenomenons is one of our main inter-disciplinary thematics.

The increasing size of LIP has led us to find new offices outside ENS campus. We now have offices at our disposal on two new laboratory sites at Gerland, close to ENS: rented IXXI offices since 2006, and UCBL offices since 2008. These geographical aspects represent an important change in the last period that we especially want to address in the future plan. The last four years have been very successful. The thematics that have emerged, and the teams' evolution, will lead to a new LIP organization for the next plan (discussed in a separated document entitled "Project document"). Our results show that LIP is able to take advantage of its favorable environment and of the existing opportunities.

1.1.1 Teams' evolution and highlights

Since 2005 LIP has been organized in six teams, corresponding to the six main topics:

Arénaire – "Computer arithmetic". After a weakening of computer arithmetic and hardware aspects with the departures of M. Daumas and A. Tisserand (end of 2005), the team has attracted: N. Brisebarre (number theory algorithmics, CR CNRS), V. Lefèvre (computer arithmetic, CR INRIA), D. Stehlé (geometry of numbers algorithmics, CR CNRS), N. Louvet (numerical and computer arithmetic algorithmics, Mcf UCBL), and G. Hanrot (algebraic, cryptographic, and geometry of numbers algorithmics, Prof. ENS). This recruitment, along with new mathematical computer science directions, strengthens the mixing of fundamental research with software- and hardware-tools design that is a main specificity of Arénaire. Around computer arithmetic, the group studies operators together with specific application domains (cryptography, signal processing, linear algebra, lattice basis reduction, etc.), for a better understanding of the impact of the arithmetic choices on solving methods in scientific computing. The aim is to improve the quality of computation at large, and the confidence in the computed results, for example through formal proving. A significant recent orientation is toward designing synthesis tools to automatically generate libraries.

> Highlights. The new IEEE 754-2008 standard is a consequence of Arénaire's efforts since 1995 (CRLibm, MPFR). Our algorithms are used by the STMicroelectronics C compiler, and our FPGA elementary functions are used worldwide for high-performance computing. Operator synthesis prototypes, using certified libraries such as Sollya, are newly available. The algorithmic (complexity) work around matrix and lattice computing is highly recognized.

Compsys – "Compilation and embedded computing systems". The departure of T. Risset (Prof. INSA, 2005) and A. Fraboulet (Mcf INSA, 2007) have induced a more focused range of activities and objectives in compilation and circuit synthesis. The hiring of C. Alias (compilers and hardware synthesis, CR INRIA) and the ENS Emeritus Professor position of P. Feautrier should preserve the stability of the team for the upcoming years. The goal of Compsys is to adapt and to extend some code-optimization techniques, primarily developed in compilation/parallelization for high-performance computing, to the domain of compilation for embedded computing systems. Compsys focuses more particularly on: code generation for embedded processors, aggressive compilation and just-in-time compilation; high-level program analysis and transformations for high-level synthesis tools. The key direction is to help bridging the gap between compilation and architecture, and between computer science and electrical engineering. A specificity of Compsys is to tackle combinatorial optimization problems arising from actual compilation problems, and to validate these developments in compiler tools through strong industrial collaborations.

> Highlights. Compsys has established a strong collaboration with STMicroelectronics, in particular within the competitiveness cluster Minalogic, pushing the use of static single assignment (SSA) for register allocation and back-end optimizations. This has led to an international recognition. Carrying the expertise on loop transformations and polyhedral optimizations to embedded computing has led to several new and interesting concepts. In the four-year period, five publications of the team have received a best paper award in international conferences.

Graal – "Algorithms and scheduling for distributed heterogeneous platforms". The team has been strengthened with B. Tourancheau (high-performance computing, Prof. UCBL), L. Marchal (algorithmics and parallel computing, CR CNRS), B. Uçar (linear algebra, CR CNRS), G. Fedak (grid computing, CR INRIA), and C. Perez (component models, CR INRIA). From fundamental aspects to middleware developments, this (balanced) growth has reinforced Graal's three research directions: scheduling and algorithmics for heterogeneous platforms; algorithmics for parallel sparse direct solvers; high-performance servers on the grid. The Graal methodology that consists in addressing fundamental algorithmic challenges inspired by realistic platform and application models, helps bridging the gap between theoretical research and software design. Among the new themes considered we find online problems, stochastic algorithms, multi-criteria optimization, and out-of-core multifrontal solvers. The wide spectrum of the research also provides the mastering of the whole chain, from model definition to implementation. This mastering is a key point for proofs of concepts such as component models for higher abstraction levels.

> Highlights. DIET (grid middleware) has been chosen as middleware for the Décrypton grid —an AFM-CNRS-IBM initiative to support genomics and proteomics resarch— and has led to launching a startup. The MUMPS sparse linear algebra solver experiences remarkable recognition and dissemination; out-of-core factorization was a main realization. Y. Robert was Program Chair of the leading international conference in the domain: *IEEE International Symposium on Parallel Distributed Processing.*

MC2 – "Models of computation and complexity". The highly-recognized fundamental research of the team has been impacted in 2008 by the departures of five permanent researchers: 2 faculty members retired, 3 CR CNRS moved to other academic positions, and 1 ENS faculty member moved to industry. The scientific loss especially impacts discrete mathematics aspects such as graph theory, and mathematical computer science topics such as complexity theory. The investment and efforts of the team for collaborating within the Complex System Institute have also been challenged by the departure of M. Morvan (Prof. ENS). Following these departures, the MC2 members have defined a more focused range of activities namely models of computation, complexity theory and combinatorics. The initiatives on complex systems have led to boolean and gene network projects, and to an ambitious simulation platform. The creation of a startup should give an enduring legacy to our work on complex systems.

Highlights. Progress in the understanding of the relationship between some of the main problems of discrete complexity theory (e.g., "P=NP?" or "P=PSPACE ?") and the corresponding open problems of algebraic complexity theory. First group to obtain quantum lower bounds for hidden subgroup problems. New research activities: Network Calculus and game theory. Development of a general-purpose software platform for the simulation of complex systems and creation of a startup (around Spring 2010).

Plume – "Programs and proofs". P. Baillot (logic and computational complexity, CR CNRS), O. Laurent (logic and semantics, CR CNRS), A. Miquel (realizability, partial secondment), and C. Riba (term rewriting and realizability, *Chaire* CNRS/Mcf ENS) joined the team. This has a major impact on the scientific goals of Plume, with a strong new direction on proof theory. Plume research deals with methods for the formal analysis of computer programs and, more generally, of computing systems. This covers the foundations of (or of aspects of) programming languages, and static analysis of programs. The team relies on theorem-proving systems to develop rigorous formalizations and proofs about mathematical theories (mainly using the Coq system). Plume builds on the proofs-programs (Curry-Howard) correspondence to strengthen the links between proof theory in logic and computer science, with a specific focus on the integration of additional programming features such as: concurrency, mobility, modularity, and probabilistic behaviors.

Highlights. Important progress has been made on the formal tools for studying the operational equivalences of concurrent processes, such as bisimilarity. This deals with both logic for expressing equivalences (Ambient Logic) and proof techniques for establishing bisimilarity results (up-to techniques for weak bisimulation). kextraction is the first program-extraction module for Coq able to deal with classical reasoning. It is a concrete byproduct of the work on the Curry-Howard correspondence for classical logic.

Reso – "Optimized protocols and software for high performance networks". The team has attracted members coming from parent fields, specifically P. Goncalves (signal processing and statistical analysis, CR INRIA), I. Guérin Lassous (wireless and ad hoc networks, Prof. UCBL), J-P. Gelas (logistical and active networking, Mcf UCBL), and T. Begin (networking and performance evaluation, Mcf UCBL). Enriched by this new potential, Reso participated to the creation of the common research lab between INRIA and Alcatel-Lucent Bell labs on "self-organizing networks". The arrival of É. Fleury (dynamic networks, Prof. ENS) and G. Chelius (wireless and sensor networks, CR INRIA) reinforced the networking component of LIP as well as the collaboration with the Sisyphe team of the ENS's Physics laboratory on traffic and network analysis and modeling. When joining LIP, É. Fleury proposed to introduce a new research topic coupling his previous experience on sensor networks to dynamic network measurement and analysis. This will lead to a new team organization focused on dynamic networks. Reso's research on core high-speed network infrastructures shows how to efficiently aggregate network, storage, and computing resources into a virtual and integrated environment. This trend, especially investigated within the Grid and Cloud contexts, will certainly influence the design of the future Internet.

Highlights. The theoretical dimension of Reso and its collaboration with networking- and performance-evaluation communities has increased the development of a vision for the future networks integrating flow- and energy-awareness, as well as service and virtualization paradigms. A strong collaboration has been established with Alcatel-Lucent through INRIA-Bell labs laboratory and the participation to the competitiveness cluster System@tic. Reso played a key role in the Grid'5000 project (construction, validation). Innovative software will be transferred to Reso's spinoff, LinKTiss, which received the "Oseo emergence 2009" prize. Our visibility has increased as assessed by the international Young Marconi Scholarship prize received by S. Soudan.

1.1.2 Transverse and new research directions

Two main transverse directions and the complex system initiative have been successfully fostered by LIP and supported by our heading organizations.

High-performance computing and networking. The transverse Grid'5000 project between Graal and Reso is a successful operation that has been highly supported by the laboratory and various initiatives: French Ministry research grant 2004-2007, PSMN Federation FLCHP, Grid CNRS (11k€ in 2007, 38k€ in 2008), INRIA ADT Aladdin, temporary

ENS engineer positions (LIP *Fonds de roulement*). Grid'5000 was designed as a large-scale and highly-reconfigurable testbed for large-scale distributed systems. Its large scale and heterogeneity allowed researchers from the Graal team to validate the DIET middleware and several algorithms and applications (bioinformatics, cosmology). The Grid'5000 platform allows researchers from the Reso team to design and validate, at a large scale, their software protocols, services, and frameworks (e.g., the energy-aware reservation infrastructure EARI, the metrology infrastructure MetroFlux).

Software and hardware synthesis, compilation, and embedded computing systems. The collaboration of Arénaire and Compsys with STMicroelectronics has grown from informal contacts to formal joint research efforts through the Sceptre and Mediacom projects, with significant funding for PhD students, engineers, and post-docs, and many joint research papers. This has placed LIP as a significant research actor on embedded systems in Rhône-Alpes and in the formal national collaboration between INRIA and STMicroelectronics (*accord-cadre*). Compsys and STMicroelectronics organized an international Spring school on, for the first time ever, static single assignment. Almost all world experts on the topic were able to participate.

IXXI. The Complex System Institute is a main inter-disciplinary initiative that LIP has supported in particular with a 150k€ infrastructure donation in 2006. IXXI Institute targets a new kind of scientific research community, one emphasizing multi-disciplinary collaboration in pursuit of understanding the common themes that arise in natural, artificial, and social systems. Within IXXI Framework, LIP creates or reinforces collaborations to match fundamental research with the needs of society, by making its contribution to key sectors for the future: medicine, biology, and the environment.

New domains are emerging in every team, among which:

Linear logic. The arrival of P. Baillot and O. Laurent brings strong knowledge in linear logic and its variants. In particular, the applications to implicit computational complexity open the possibility of new interactions between LIP teams on questions related with resource usage.

Algorithmics of discrete event systems. The MC2 team plans to develop this research topic following its work about Network Calculus, a deterministic queuing theory for worst-case performance evaluation. Such analyses require tools from different fields like $(min, +)$ tropical algebra, linear algebra, graph algorithmics or computational geometry.

Arithmetic operator generation. Arithmetic cores get coarser, deal with new domains such as finite-field computing for cryptography, and target a wider range of technologies: from ASIC to FPGA in hardware, from embedded systems to HPC in software. To handle this increasing complexity, Arenaire works more and more on the automation of arithmetic core generation and optimization, for instance through improvements on polynomial approximation theory.

High-performance geometry of numbers and applications. With the hiring of G. Hanrot and D. Stehlé, and the existing strengths, we have an outstanding expertise around Euclidean lattices. These objects are studied for themselves, but also applied to arithmetic challenges such as as polynomial approximation or the search for correct rounding worst cases.

Compilation and embedded computing systems. Thanks to its collaboration with STMicroelectronics, Compsys has now a very strong expertise on static single assignment (SSA) that will be used to develop efficient just-in-time compilation schemes. Compsys' expertise on high-level transformations (with connections with Arénaire for FPGA platform use) is also now mature enough to get fruitful results for high-level synthesis, also with STMicroelectronics as industrial partner.

Dynamic networks. Our expertise on the deployments of large-scale sensor networks within MOSAR (Integrated Project supported for 5 years by the European Commission under the Life Science Health Priority of the Sixth Framework Program) and SensLAB (ANR Platform project) allows us to gather valuable data on dynamic networks. We want in the near future to foster our capacity of analysis and modeling of real-world graphs and complex networks.

Future Internet. We believe that self-optimization challenges will be at the core of future networks design, which will embed virtualization paradigms - deeply related to Cloud principles - traffic and energy awareness. We thus decided to strenghten our competence in performance analysis and optimisation research, notably by hiring Thomas Begin (MCF UCBL). We also aim at invigorate our interaction with local distributed systems, scheduling theory, combinatorial optimisation and signal analyses, while reinforcing our collaborations with complementary teams at INRIA, but also at EPFL and in the US. Following our motto "from the model to the real world" we will continue to transfer to industry, original, innovative and scalable softwares that implement promising concepts.

Static scheduling of dynamic systems described with probabilistic models. In the past, most of the theoretical research in Graal dealt with the design of static algorithms for static systems, that is, for systems those characteristics were supposed not to evolve during the lifespan of the studied problem (namely, the execution of an application). During the reporting period, the Graal team more and more focused on solutions for dynamic systems. We initiated our first work dealing with online problems. The hiring of A. Benoit was the opportunity for the team to acquire some knowledge on probabilistic systems. The Graal team is now considering the design of static solutions for dynamic systems that are described through probabilistic models, and has published its first works in this field.

Service Oriented Computing Platforms. The access to Grid and Cloud platforms is driven through the use of different services providing mandatory features such as security, resource discovery, virtualization, load-balancing, etc. Software as a Service (SaaS) has thus to play an important role in the future development of large scale applications. This model based on services provided and/or offered is generalized within software component models which deal with composition issues as well as with deployment issues. These software architectures need many research issues to be solved, from programming interfaces to lower middleware details in Grids, Clouds and P2P systems, and this will be the aim of the Avalon research team, started from Graal.

1.1.3 Self assessment

Our scientific objectives have largely been reealized. This is assessed by the attractivity of the teams and the evolution of the scientific points of focus of LIP. The vitality and progress of the teams keep their potential at the highest level. Following their hirings, Arénaire, Compsys, and Plume have experienced major changes. Nevertheless, the organization of these groups should be quite stable during the upcoming years. Also, thanks to various hirings, the largest teams, Graal and Reso, will give rise to smaller, more coherent, groups. The main transverse initiatives (grids and embedded systems), and the inter-disciplinary participation to the Complex System Institute, are successful. The growing importance of the latter and the departures in MC2 have induced a few team instabilities that have been brought down. In particular, having two temporary positions (researcher and faculty member) next year on quantum complexity is very satisfactory. The balanced focus on the fundamental domains (algorithmics, complexity, logic, statistics, modeling, mathematical computing) and on the software/hardware design and experimental science is at a very satisfying level and remains a clear policy of LIP.

1.2 Scientific activity and production

Publications. Our production is highly qualitative and quantitative. We have published 173 articles in peer-reviewed journals, and over 400 in peer-reviewed conferences (for an average number of 100 scientific members, among which 40 permanent members).

	2005	2006	2007	2008	2009	**Total**
ACL - International and national peer-reviewed journal	30	33	30	44	36	173
INV - Invited conferences	4	12	23	31	19	89
ACT - International and national peer-reviewed conference proceedings	53	77	83	99	90	402
COM / AFF - Short communications and posters in international and national conferences and w orkshops	3	11	10	17	10	51
OS - Scientific books and book chapters	1	3	1	3	13	21
OV - Scientific popularization	1	2	6	2	4	15
DO - Book or Proceedings editing	4	9	5	7	12	37
AP - Other Publications	3	10	6	9	8	36
TH - Doctoral Dissertations and Habilitation Theses	4	9	13	14	13	53
Total	103	166	177	226	205	877

The most represented journals are *Theoretical Computer Science* (21 times), *IEEE Transactions on Computers* (12 times), *Elsevier Future Generation Computer Systems* (11 times). For conferences, we especially publish in: *IEEE Symposium on Computer Arithmetic (ARITH)* (12 times), *International Symposium on Symbolic and Algebraic Computation (ISSAC)* (9 times), *IEEE International Symposium on Cluster Computing and the Grid* (7 times), *International Parallel and Distributed Processing Symposium (IPDPS)* (5 times).

Editorial boards, conference committees and organization. We participate to editorial boards of 17 journals, such as *Discrete Math. Theor.*, *Parallel Comput.*, *ACM TECS*, *Appl. Algebr. Eng. Comm.*, *J. Symb. Comput.*. LIP members have also been guest editors of 15 journals, such as *Ann. Telecommun*, *Par. Process. Lett.*, *IEEE T. Comput.*, *Theor. Comput. Sci.* For major international conferences, LIP members have been: chairs (ex: CCGrid'2008, SSS 2009 or Parco 2009), members/chairs of 9 steering committees (ex: CCGrid'2008, ARITH, and ISSAC); PC chair of 14 main conferences (ex: ACM Gridnets, ACM MobiHoc, IPDPS, ARITH). LIP members have been chairs or PC chairs for 56 workshops (mostly international) or tracks in conferences such as SSA'09 (Static Single-Assignment Form), E2GC2'09 (Energy Efficient Grids, Clouds and Clusters), SLSS'09 (Scheduling for large-scale systems), LCC'09 (Logic and Computational Complexity), RAIM'09 (*Rencontres Arithmétique de l'Informatique Mathématique*). We also have participated in the program comittees of about 300 various events. These significant numbers demonstrate the important LIP recognition and involvement in the national and international communities.

Recognition, prizes, and awards. F. Desprez received an IBM Faculty Award (2008). P. Feautrier received from the Euro-Par Steering Committee an award "in recognition of his outstanding contributions to parallel processing" (2009). P. Koiran is Junior Member of *Institut Universitaire de France* (2007). Y. Robert is IEEE Fellow (2006), and Senior Member of *Institut Universitaire de France* (2007). S. Soudan received the Marconi Young Scholar Prize (2009). LIP members received 16 awards for papers, presentations, or demonstrations at conferences such as ASAP, CGO, CHES, ICALP, ISSAC, JDIR, SIGMETRICS, or TLCA. The Martes ITEA project (Compsys team) received the 2008 ITEA Symposium Silver Award.

Softwares, standardization actions. We have a rich and wide software production. This goes from libraries for experimental computer science (proofs of concepts, algorithmic design), to widely-disseminated research softwares, and to industrial transfer. For example we may refer to:

- *HIPerNET:* A new framework for virtual private execution infrastructure composition and orchestration transferred to startup LinKTiss.

- *BDTS:* Software for Bulk Data Transfer Scheduling Service transferred to the LinKTiss startup and partially integrated in the CARRIOCAS-SRV prototype (Alcatel-Lucent).

- *LAO developments:* The contributions of Compsys on back-end code optimizations are transferred to STMicroelectronics, first in the prototype compiler LAO, then in the production compiler.

- *DIET:* Middleware for the deployment of computation servers over Grids and Clouds transfered to the Graal Systems startup.

- *MUMPS - a multifrontal massively parallel solver:* 200000 lines of code, 1000 downloads per year, included in lots of finite element packages and numerical simulation codes from both academia and industry.

- *IEEE-754-2008.* The specification of elementary functions in this revised floating-point standard is a direct consequence of Arénaire work on the subject.

- *FLIP - floating-point for integer processors.* Included in the standard compiler tools for the ST200 processor family.

- *FPLibrary.* First library of floating-point operators to include elementary functions for reconfigurable computers

- *Mathematical libraries* are largely disseminated through general softwares like Magma or Maple, or integrated and disseminated through open-source platforms such as Sage (ex: linbox, fplll).

- *kextraction* is a prototypical extension of the Coq proof assistant that allows one to extract certified programs from proofs in classical logic. While not software to be independently distributed, it represents a new technology, that could be included in the next versions of the Coq system.

Startups. Three start-ups are emerging from LIP, which illustrates our transfer skills. These start-ups are supported through various sources and all three of them received an Higher Education and Research Ministry Award 2009. In relation with team MC2 and the Complex System Institute, É. Boix proposes a software platform for complex systems simulation in biology. Emerging from the Graal team and the DIET software, and represented by D. Loureiro, the second project deals with consulting and software editing around access to distributed HPC resources. In relation with team Reso, P. Vicat-Blanc Primet will offer virtualization and optimization services and softwares for cloud networking.

Administration of research. Among the administrative members at the national level one may cite: J.-M. Muller, *Chargé de mission* at the *Sciences et technologies de l'information et de l'ingéniérie* CNRS institute since April 2006; N. Portier and G. Villard, members of the board of CNRS-GDR *Informatique Mathématique*; E. Fleury is member of the board of CNRS-GDR ASR *Architecture, Systèmes, Réseaux* and head of CNRS-GDR ResCom (Networking axis). E. Fleury is also in the scientific board and in the steering committee of IXXI (Complex Systems Institute), and member of the board of the regional research cluster ISLE (*Informatique, Signal, Logiciel embarqué*). P. Lescanne has been president of SPECIF (*Société des Personnels Enseignants et Chercheurs en Informatique de France*) between 2006 and 2008; F. Desprez is vice-head of the project-team committee, was a board member at the regional INRIA Research Center (2006-) and is a member of INRIA Evaluation commission (2008-); A. Darte, member of INRIA Evaluation commission, 2005-2008, and scientific referent for the national collaboration between INRIA and STMicroelectronics. LIP members are regularly elected to ENS Administration Council or Scientific Council (P. Lescanne, P. Vicat-Blanc Primet, Y. Robert, and S. Torres).

1.3 Collaborations

LIP members have a broad range of national and international collaborations that we emphasize through different aspects here. Mostly at the individual level, we have publications with 41 foreign universities in: North America (19: USA and Canada), South America (2: Brazil, Mexico), Europe (12: Denmark, Germany, Italy, Portugal, Romania, Serbia, Switzerland, UK), Russia (1), Japan (5), and Australia (1). Especially through the Explora'doc region program, we have PhD co-supervisions with universities in Australia, Germany, Portugal, Denmark, Italy, and Serbia.

Our 2005-2009 non-consolidated budget (see Section 1.7) is 8.6M€, with about 12.6% from the heading organizations (mainly CNRS and Ministry through ENS) under the control of the direction, and 7.2% from INRIA to the four project-teams. Hence about 80% of the budget come from specific project fundings. The main funding source is the National Research Agency (*ANR*), which represents 27% of the budget with 2.3M€. Here is the list of the *ANR* projects in which LIP members are involved (at various percentages):

National Research Agency (*ANR*)

Alpage *Distributed platforms*	Lego *Efficient grid*	Disc *HPC Components*	Numasis *HPC Seismology*
Solstice *Parallel solvers*	Stochagrid *Stochastic perf. models*	Gwendia *Grid workflow*	SimGrid *Simulations*
Spades *Petascale*	Cloud@Home *Cloud comput.*	Igtmd *High-speed transport*	DSLAB *DSL Internet*
Hipcal *Infrastructure virtualization*	Dmasc *Heart signal analysis*	Sycomore *Models of computation*	Mécamerge *Structure emergence*
Carpelle Virt. *Plant modeling*	MoDyFiable *Dynamic modularity*	Galapagos *Proof and geometry*	CHoCo *Curry-Howard corresp.*
Scalp *Proof and cryptography*	Complice *Implicit complexity*	Eva-Flo *Arithmetic op. synthesis*	Gecko *Geometry and complexity*
SubTile *Tilings*	MINT *Network models*	Dynanets *Dynamic networks*	GeneShape *Gene mechanisms*
LaRedA *Euclidean lattices*	Tchater *FPGA signal processing*	SensLAB *VLS sensor net. platform*	

In addition to their recurrent participation, our heading organizations, together with the region, are funders for 24% of our budget through the national and international projects listed below.

Local/Region/National projects

ARC INRIA Otaphe *Scheduling*	ARC INRIA Georep *MUMPS software*	ARC INRIA *Green-Net*
ADT INRIA Aladdin *Grid'5000*	ARC Coinc *Network calculus*	ACI Grid Explorer *Large scale exp.*
ACI GeoCal *Proof and geometry*	ACI Sécurité *Cryptography*	ACI New interfaces of mathematics
ACI Grid *Grid'5000*	Region Cluster *HPC*	Région Ragtime *Grid/Medical data*
Seiscope Consortium *MUMPS software*	ATIP CNRS *Global properties*	ATIP CNRS *Complex systems*
Competitiveness clust. System@tic *HPC network*	Competitiveness cluster Minalogic *Compilation*	Nano2012 *JIT compilation & arithmetic*
Nano2012 *Prog. transf. for high-level synthesis*	Competitiveness cluster Minalogic *Arithmetic*	Region/STMicroelectronics *Arithmetic*
Minalogic-Sporaltec ESPAD *Sportive perf. study*	CNRS PEPS *Euclidean lattices*	Region Cluster EmSoc
ADT INRIA SensTOOLS *Sensor networks*	AFFSET TubExpo *Exposure to tuberculosis study*	CNRS PEPS *Signal processing*

International projects

Exploradoc Berkeley *Out-of-core solvers*	France-Berkeley Fund *Parallel solvers*	CNRS-JST Japan *Grid platforms*
CNRS-JST Japan *Parallel solvers*	INRIA Explorer Japan	MIT-France Fund *Multicores*
CNRS Hawaii *Scheduling and bioinformatics*	INRIA Hawaii Ass. Team *Sched. and bioinf.*	INRIA Standardization Actions *Arithmetic*
PAI Fast Australia *Overlay networks*	INRIA Brazil *Java on FPGA, dyn. compilation*	INRIA Italy *Proof and concurrency*
ECOS Chili *Models of computation*	CNRS Serbia *Logic and types*	Region MIRA-Poland *Computational logic*
CNRS USA *Compilation*	PAI Romania *FPGA arithmetic*	PHC INRIA *Cryptography*
CNRS Australie *Lattices*	Australian Res. Council *Lattice enum., crypto.*	XtremeData University *FPGA computing*
French-Israeli Multicomp. *Parallel Solvers*	STIC AMSUD SAMBA *Complex Networks*	INRIA Explorer USA *Parallel solvers*

Among these projects, the investigation operations such as the Ministry ACI, CNRS PEPS, INRIA ARC, and the different INRIA/industry national joint research programmes, allow us to start new collaborations on emerging domains quickly, hence are key tools for our progress. At the international level, this is completed by 21.8% of our budget coming from European projects.

European Community

NoE CoreGrid *Dist.computing*	NoE Euro-NF *Networks*	NoE Artist2 *Embedded systems*
Cost Action *Energy efficiency*	Cost *Dynamic comm. networks*	Cost Once-CS *Complex systems*
STREP FP6 Euro-China GIN *Grid Networks*	FP7 AI *Autonomic Internet*	SSA FP7 OGF-europe *Standard. in Grid*
NoE Hipeac *Emb. architectures and compil.*	Eureka-ITEA Martes *Embedded systems*	NEST Morphex *Morphogenesis, compl. syst.*
Marie-Curie MetagenoGrids	Marie-Curie *Algebraic complexity*	EST Net. *Mathematical logic appl.*
IP FP6 Aeolus *Overlay theory and algorithms*	IP FP7 EDGes *Grid systems*	IP FP6 Life Science Health MOSAR

Among the above projects, several involve industrial partners, hence our direct industrial contracts, which represent 5.5% of the budget, are quite far from revealing our level of industrial collaboration. For example, the ANR grant on hardware signal processing is a collaboration with Alcatel-Lucent, the MARTES project was a collaboration with Thales, and the ANR Solstice project includes collaborations with EADS and EDF around our parallel solver MUMPS. We may also emphasize the strong partnerships developed with Alcatel-Lucent and STMicroelectronics with the support of INRIA through the competitiveness clusters Minalogic and System@tic, and INRIA - Bell-Labs Common Lab.

Associations and industry

AFM Association *Grid HPC*
Alcatel-Lucent *Services/resources*
STMicroelectronics *Compilation, arithmetic*
Orange-FT *Network load*
ADR INRIA-Bell Labs *Semantic networking*
STMicroelectronics - Thompson - Silicomps *SoC simulators*
Intel USA (computer funding)
Samtech *MUMPS software*

During the previous four-year plan, our operating budget has been 3.45M€. The spectrum and volume of our 2005-2008 fundings therefore demonstrate our vitality. This allowed us to ensure the diversity of the types of productions, between high-level publications and practical (expensive) developments.

1.4 Personnel

LIP has 133 members, representing an increase of 56% over the last four years (Sept. 1st, 2005). The distribution of permanent research staff is 17.5% of faculty members and 16.5% of researchers. The total of *ETP-chercheur* (*Équivalent Temps Plein*) is now 33.5 versus 22.5 in 2005.

		CNRS: **9.5**			ENS: **15.5**			INRIA: **15**			UCBL[1] : **8**			Temporary: **76**			
		DR	CR	IR	PR	MCF	IR	DR	CR	IR	PR	MCF	Perm.	Res.	PhD	Eng.	Tot.
Arénaire	2005	1	1			1	0.5		2				5.5	1	7	1	14.5
	2009	2	1		1	1	0.5		3			1	**9.5**	1	9	2	**21.5**
Compsys	2005	1			1				1				3		5		8
	2009	1			1				2				**4**		2	1	**7**
Graal	2005			1	1	2		1	2			1	8	4	8	1	21
	2009		2		1	2		1	4	0.5	1	1	**12.5**	8	11	5.5	**37**
MC2	2005		2		2	2					1	2	9		9		18
	2009			1	1	2						1	**5**		3	5.5	**13.5**
Plume	2005		1		1	2							4		3		7
	2009		2		1	4							**7**	2	4		**13**
Reso	2005			0.5				1	1			1	3.5	1	6		10.5
	2009			0.5	1			1	3	0.5	1	3	**10**	4	12	7	**33**
Total 2005		2	4	1.5	5	7	0.5	2	6		1	4	33	6	38	2	79
+/-		+1	+7/-6		+2/-1	+2			+7/-1	+1	+2/-1	+3/-1	+25/-10	+44/-35	+42/-39	+29/-10	**+140 / -94**
2009		3	5	1.5	6	9	0.5	2	12	1	2	6	**48**	15	41	21	**125**

The assisting permanent administrative and technical staff is 8% of the total. Non-permanent members represent 58% of the laboratory. There are 31% of PhDs, 11% of temporary researchers (ATERs, postdocs, secondments, long-term visitors), and 16% of temporary engineers. The table above gives detailed data concerning the scientific members since 2005, including a percentage for permanent research engineers, and taking into account known arrivals for Sept. 2009. The row before last emphasizes the turnover with the number of departures and arrivals during the period.

Permanent faculty members and researchers. We start in Sept. 2009 with 13 more permanent members than in 2005: 23 arrived and 10 left (mobility, promotions, and retirements)[2]. There has been one DR promotion. Apart from the main changes in MC2, this large turnover is an excellent sign. The number of CNRS researchers slightly progressed with the largest turnover in Arénaire, Graal, and Plume. The number of INRIA researchers has increased with 6 new members among the four INRIA projects. The two hired ENS professors follow departures. É. Fleury creates a new team on dynamic networks, and G. Hanrot strengthens orientations towards geometry of numbers and cryptographic applications. One *Chaire* CNRS/MCF ENS position is obtained for C. Riba in Plume, by anticipating a retirement, on logic and semantics. UCBL investment is illustrated by: the Professor position of I. Guérin-Lassous, with a new theme on wireless and ad hoc networks; the one of B. Tourancheau (after a temporary move to industry) who strengthens distributed algorithmic aspects; and three MCF positions in Arénaire and Reso. UCBL is a key support for improving in particular our teaching collaborations in Lyon. The successful LIP faculty-hiring policy has been to encourage outstanding applications and the balanced implications of the home institutions in the teams. Among the 45 permanent researchers, 18 have their *Habilitation à diriger des Recherches*. Three HDRs have been defended since 2005. The fact that most PhDs are co-supervised by junior researchers stimulates the number of HDRs and fosters the promotions.

Temporary research personnel. Taking advantage of the policies from our heading organizations and funding agencies (CNRS and INRIA postdocs and partial secondments, ANR, invited positions), LIP encourages temporary researcher hiring. This is a key point for research vitality, and for helping PhDs in finding positions. We had 23 postdocs and ATERs during the period, with 12 such positions at the beginning of 2009. Comparing that to 4 in Sept. 2005, this

[1] For the sake of simplicity, we also count here one Professor IUFM Lyon and one MCF IUT Roanne.

[2] i) With one DR promotion and a new permanent research engineer this gives 15 new permanent scientific positions. ii) A. Fraboulet and T. Risset stayed in Compsys until 2007 as LIP external collaborators, they do not appear in our (administrative) table.

reveals a growth tendency, hence a success that should continue. However, since this may also reveal a general increase of temporary positions (versus permanent ones), we have to be careful in our analysis. We also had 27 visiting junior researchers and professors for more that three months, especially on specific ENS and INRIA positions.

Doctoral students. During the four years, the number of doctoral students has been stable around 40, with 41 thesis defended compared to 26 during the previous plan. One key role of LIP is the formation of the outstanding ENS students. According to the school's policy, most of them are encouraged to leave Lyon for following doctoral programs in other cities, and only a small proportion of them prepare their thesis at LIP. Still, this has represented 29 students out of the last 80 PhDs prepared in the laboratory. The number of PhDs is slightly less than the number of permanent faculty members and researchers. On the one hand, the analysis has to take into account the fact that the growth of the number of faculty members and researchers mainly concerns juniors. On the other hand, 40 PhDs represent more than 2 PhD students per HDR, which reveals that young researchers are highly involved in student supervision.

The success of our doctoral program is clearly demonstrated by the fact that out of the last 80 defended thesis (since 2000), over 60% have permanent research or faculty positions: 18 have permanent research positions (CNRS, INRIA, Onera), 28 are faculty members in France, and 3 are faculty members abroad. Among other types of employments, 12 have engineer positions in companies such as Michelin, Google, or Intel. Regarding the most recent defenses, 15 PhDs are on postdoc or temporary faculty positions.

Attracting good doctoral students is a general difficulty in our disciplines. LIP members work actively for preserving the current supervision rate, especially by taking advantage of various sources of funding. Out of the 80 students in the last four years, 30 were ENS students with corresponding ENS/Ministry grants (Lyon or Paris), 28 (most often coming from other cities or countries) had Ministry grants (3 through INRIA), 19 had grants associated with contracts and industrial collaborations including 5 Cifre and 3 BDI CNRS, 2 had grants from Europe, and 1 from Mexico. With these various types of funding and a selective ENS master, our doctoral program keeps a very good level, and LIP is very active in the organization of the studies (see Section 1.6). The increasing involvement of the Computer Science Department at the international level, according to the ENS policy, and our increasing involvement at Lyon's level, are tools for keeping our formation at the highest level.

Administrative members. In 2005, our research was assisted by an administrative staff of 5 people, with the equivalent of 4 full-time people. We have worked in 2009 with the equivalent of 5 full-time people, and we had 3 temporary people in the meantime. We have two ITA CNRS positions. However, one of the assistants is on extended maternity leave since late 2006, and the other position is vacant since the promotion of I. Pera (Jan. 2009). Hence, considering that LIP has 56% more members than in 2005, and comparing the 8.6M€ four-year plan budget to the 3.4M€ previous one, this reveals that 2009 has been a very difficult year. This also shows how the administrative team is highly efficient and supportive. C. Iafrate (ENS) and S. Boyer (INRIA) share the role of administrative supervision. Beyond general tasks for the direction, C. Iafrate also supervises financial aspects, and a large part of the time of S. Boyer is dedicated to management tasks for INRIA Lyon, that are unrelated to LIP's activities. The three assistants, C. Suter (INRIA), J. Brajon (temporary ENS, until Aug. 2009), and M. Buillas (temporary INRIA, until Feb. 2010), are in charge of the six scientific teams. Our 2009 efforts, with the support of our heading organizations, has led to a permanent ENS assistant starting in Oct. 2009, L. Lécot, and a permanent CNRS assistant starting in Dec. 2009. We have also obtained a permanent INRIA position starting in Sept. 2009 for E. Bresle. Taking into account INRIA's tasks for projects outside LIP, by 2010 the improvement will come essentially from the replacement of temporary staff members by permanent ones. We hope that the high turnover of administrative people will stop because the training of temporary assistants induced a sustained extra load on the administrative staff, increasing its difficulties. These questions are a constant concern for the lab and the researchers. The policy of the lab has been not to hire assistants on contracts. We are not sure whether or not this policy can be kept since the support provided to researchers deteriorated between 2005 and 2009, and the complexity of the assistance tasks increased.

Technical staff. LIP has 5 permanent research and study engineers (3 CNRS, 1 ENS, 1 INRIA). M. Imbert arrived as INRIA research engineer (*Service Expérimentation et Développement*) for support on specific research projects in Graal and Reso in 2007. Our policy encourages the engineers in their research and software development activities. This policy, which keeps the engineers close to the researchers, greatly helps with the quality and the efficiency of the services, and should be preserved in the future. These research activities correspond to the equivalent of 3 full-time people and has been taken into account in the scientific members table. As in 2005, the IT support team represents the equivalent of 2 full-time engineers for everyday computer and network administration. The work of the team—J.-C. Mignot (CNRS), D. Ponsard (CNRS), and S. Torres (ENS)—has been acknowledged as a CATI CNRS in 2007 and is detailed in Chapter 8. The skills and efficiency of the group provide a continuous research support of particularly high quality. This means that most of the aspects of the infrastructures are kept up-to-date and secure (workstations and laptops; file, backup and network services; cooling; web management; collaborative and scientific softwares; special-purpose servers; video-conferencing;

tool development; etc.). The IT team also plays a key role for the administration in the Grid'5000 platform, to which no permanent engineer is assigned. LIP has provided a sustained help to the platform, which is an important research resource: obtaining a corresponding permanent position is one of our priorities. Our internal LIP organization works in tight collaboration with the IT ENS team, as well as with the numerical computing ENS pole (PSMN), and on various actions with the INRIA engineers located in Montbonnot.

Temporary engineers. During the last years, LIP has experienced a quick growth of the number of temporary engineers, who represent 16% of the scientific members (one engineer for 2.1 permanent faculty member or researcher). The equivalent of a four-year position has been provided by ENS. This help has been crucial for the administration of the Grid'5000 machines. A one-year position is provided by CNRS around the Décrypton project (Aug. 2009). Most of the other positions come from ANR and European projects, and reflect the rich activity of LIP in software development and transfer activities. Most teams are concerned, which shows that the combination of theoretical and practical aspects can follow in very different ways. This allows for instance the transfer of compilation activities in Compsys, and the transfer and dissemination of arithmetic operator libraries in Arénaire. In addition to the support of developments done by researchers, the engineer support here is crucial for transforming research prototypes into complete softwares. Larger software projects gather more engineers, typically 4 to 6 engineers per year, in MC2 for a complex system modeling platform, in Graal for the Diet software, and in Reso for optimized protocols and transport layers. The success of these activities is illustrated by the creation of three corresponding startups. The possibilities of temporary positions should not hide, however, the important need of long-term and permanent positions for most of our software production. Indeed, research softwares evolves with the fundamental research progress and, in general, the necessary involvement of the researchers themselves in code production and maintenance remains at a high level.

1.5 Governance and animation

LIP is organized in teams. The "Laboratory Council" (elected members, heads of the teams, and direction) meets roughly every two weeks. This ensures improved communications and that the large majority of strategic points are discussed by the council. General lab meetings are run several times a year. Meetings between the heads of the teams and the direction are frequent and informal. The *"Commission des habilités"* deals with questions concerning the doctoral students and CS programs, and the "Teaching council" those related to graduate CS studies.
The teams run their own scientific seminars regularly, often with outside speakers who can also be of interest for members of other teams. LIP runs a monthly seminar with two outside speakers (a junior and a senior) on a given theme, and, jointly with the CS department, we organize a student's seminar every two weeks.
F. Desprez—with vice chair G. Villard—has chaired the laboratory from July 2006 to November 2008, when he resigned. He was faced with the increasing number of administrative tasks both at LIP's and at INRIA's level. Gilles Villard then took the lab responsibility—with vice-chairs D. Hirschkoff and P. Vicat-Blanc Primet.

1.6 Teaching and training

Teaching. Through its partnership with two teaching institutions, namely École Normale Supérieure de Lyon (ENSL) and Université Claude Bernard - Lyon 1 (UCBL), LIP is naturally involved in teaching activities.
Members of the lab have teaching as well as administrative duties at UCBL, serving notably in the Masters SIR (*Systèmes d'Information et Réseaux*), RTS (*Réseaux, Télécommunications et Services*), TI (*Technologies de l'Information*), and CCI (*Compétences Complémentaires en Informatique*) and in various courses in computer science in Licence and Master 1.
ENS *Département d'Informatique* (DI) has very tight links with LIP. Administratively, it is run by members of the lab (Y. Robert was head of the DI during the period 2005-2009, É. Fleury will be the new head, starting in the academical year 2009-2010). In terms of pedagogical content, the local curriculum in computer science is defined through a close interaction between the DI and researchers in LIP, notably regarding the Master in Computer Science hosted by ENS. In 2009, a dedicated committee, the *Conseil des formations du LIP*, has been created to discuss these matters.
LIP members are also involved in the organization of the recruitment exams for admittance at ENS. Students who passed these competitive exams (*élèves normaliens*) traditionally represent an important part of the 25-30 pupils per year who follow the local curriculum. In addition to these, more students can be recruited via an application process, and we also exploit partnerships with foreign universities (typically, Erasmus programmes). The vast majority of the students attending the Master at ENSL then enroll in a PhD programme, which is consistent with the proposed courses' clear orientation towards research, and with the strong participation of LIP members in the Master.
LIP is also involved in the Doctoral School *Informatique-Mathématiques*, which gathers several doctoral programs in Lyon.

Continuing training. The members of LIP regularly participate to continuing training sessions (an average of 20 such actions per year). For administrative staff, this is training on administrative tools, preparation for civil service examination, or language training, among others. For scientific staff in general, training covers various subjects, including programming languages, computer security, as well as administrative and human resource management.

1.7 Financial resources

Consolidated budget. For the four-year plan[3], they have been of the order of 27M€, not taking into account ENS's administrative and infrastructure cost. About 70% of this budget have been salaries directly paid by our heading organizations (permanent and non-permanent members, PhD candidates). Hence, LIP members have managed an *operating budget* of 8.6M€ (see the table below) that represents almost 1/3 of the income.

Direction budget. The recurrent budget, directly managed by the lab's direction, has been 190k€/year on the average. This budget came mainly from CNRS and ENS. Indeed, INRIA is essentially funding the project-teams, directly, and UCBL's participation is essentially via salaries and offices on the University's Gerland site. Together with the income from former contracts through the ENS *Fonds de roulement* and the recent ENS *Fonds de recherche* (under the arbitrage of the direction), this totals 1.1M€ over the four years. Between 2005 and 2007, every year, a third of the funding received by the lab has been dispatched to the research teams. In the last two years, the teams' own budgets have risen, notably thanks to ANR and European grants. The budget of the direction was 18.1% of the operating budget of the lab in 2005, 14.6% in 2006, 10.5% in 2007, and 9.7% in 2008. As a consequence, the lab support to the teams has shrunk. This offers the opportunity to provide funding to specific scientific operations such as conference support and invitations. During the 2006-2008 period, around 120k€ have been devoted to the support of Grid'5000 (in addition to other fundings, as well as scientific efforts, provided by the teams involved in this initiative).

	2005	2006	2007	2008	Total	%
Direction budget	246.5k€	265.5k€	309.7k€	264k€	1085.7k€	12.6%
INRIA lab and project-teams	110k€	206k€	183.7k€	119k€	618.7k€	7.2%
CNRS/ENS/INRIA risk & specific	83.2k€	77.2k€	89.1k€	554.3k€	803.8k€	9.3%
Sub-total heading organizations	*439.7k€*	*548.7 k€*	*582.5k€*	*937.3k€*	*2508.2k€*	*29.1%*
Region	88k€	374.5k€	747.9k€	66.2k€	1276.6k€	14.8%
ANR	384k€	433.5k€	951k€	559.3k€	2327.8k€	27%
Europe	287k€	361k€	400k€	828.7k€	1876.7k€	21.8%
Sub-total contracts	*759k€*	*1169k€*	*2098k€*	*1454.2k€*	*5481.1k€*	*63.6%*
Associations and industry	85k€	106k€	12k€	333.4k€	536.4k€	6.2%
Others/mixed	72k€	0k€	16k€	8.7k€	96.7k€	1.1%
Total	**1355.7k€**	**1823.7k€**	**2709.5k€**	**2733.6.9k€**	**8622.4k€**	**100%**

Budget for each team. Once considered the budget of the direction, the remaining operating budget is 7.5M€. It is managed by the teams, with CNRS, ENS or INRIA accounting, essentially via the administrative staff of the laboratory. INRIA has provided about 0.6M€ directly to the four joint project-teams (80% of the scientific members).
The rest of the income—around 6.9M€ (8.6 − 1.1 − 0.6), that is 80% of the operating budget—comes from specific calls for proposals and fundings, from both our heading organizations and contracts (see Section 1.3). Altogether the participation of our heading organizations is close to 30% of the operating budget. This participation plays a key role in the emergence of new topics. It partly corresponds to various types of national and international small projects that remain important tools for our risk and long-term research. The remaining 70% of external fundings, with a large ANR proportion, usually applies to bigger collaborative projects, including industrial collaborations (that are not exclusively on purely industrial contracts). The numbers show that LIP members are very active in project-based research and industrial collaborations. The volume of contracts has increased the possibilities for hiring temporary researchers (mainly postdocs) and engineers, which is consistent with what we have discussed in Section 1.4. The other main source of expenses is travel, either for conference participation, for organizing dedicated workshops related to the collaborative projects, or for visiting foreign institutions in the scope of our international collaborations.

1.8 Infrastructure

Offices. The increasing number of LIP members has led us to find new offices outside ENS Gerland site. In addition to ENS main building offices used by LIP since its creation—1250m^2— we now have a total of about 500m^2 at our disposal,

[3] Since we have incomplete data for 2009 we provide details for the 2005-2008 period only. Taking into account 2009 would not modify the main part of our analysis.

distributed over two new laboratory sites. Our investment for IXXI, along with ENS's support, has led us, in 2006, to use rented IXXI offices close to ENS (200m^2). Thanks to UCBL, we also use university offices on the Gerland site since 2008 (300m^2, including a surface extension in 2009). For the Grid'5000 node, we use the infrastructures in a dedicated room in a separate ENS building. Being split among three different sites is not satisfying, especially for the MC2 team that is currently located on a site outside ENS campus, and even more so for the Graal team which is currently split over two sites. Hence we have to be very careful with both the administrative and scientific running of the laboratory, until it is possible to gather all the teams on one site only.

Technical environment. The IT support team has been acknowledged as a CATI (*Centre Automatisé du Traitement de l'Information*) by CNRS in 2007, its main objectives are: building and managing an IT environment for the laboratory activity; providing support and expertise to research teams for their projects; building and managing experimental platforms. The activities and the maintained infrastructures are detailed in Chapter 8.

1.9 Self assessment

Key elements of our strategy have been to target six strong topics and teams, foster the mixing of theoretical and practical research, train at the highest level, and balance the involvement of the four heading organizations. Our success is emphasized by the visibility, worldwide unique expertise, production, collaborations, and innovating results and future directions on the six points of focus of the four-year plan. The main criteria we use for analyzing our production are: publications, with an important role of the highest-quality conference proceedings; software and hardware developments; national and international collaborations, and community animation; industrial transfer.

The transverse aspects of our strategy has brought a strong and fruitful synergy across the teams in theoretical computer science's main trends: algorithms, complexity, logics, and semantics. As a consequence, a Coq seminar will start by Fall 2009. These transverse aspects have also led to major progress and interactions, including software development and industrial transfer, on grid computing and networking, embedded computing systems, and inter-disciplinary complex systems. Our weakening on complexity and discrete mathematics, and the difficulties for hiring in architecture and hardware/software synthesis may need a special attention.

LIP is very successful for training students, and for attracting and integrating scientists with different backgrounds. However, the productive combination of theoretical, mathematical, and practical aspects, and the growth of the lab require efforts for preserving unified projects. Our quick growth also demands an efficient training and research strategy for increasing the number of doctoral students, and preserving our supervision rate, in an international context that may not be favorable, especially in computer science.

Some of our objectives are essentially related to the technological progress that we must anticipate, and to our management of pioneering experimental facilities. The Lyon networking and computing node of Grid'5000, or our software transfer with STMicroelectronics, demonstrate the high level of our research here in various directions. Together with long-term software development and maintenance, this requires sustained infrastructure investments, and engineering support.

Our activity (such as the various nature and volume of our fundings, or the increased number of temporary people), with a deeply-modified academic system, requires us to keep a very close watch on the administrative load of the researchers. LIP researchers are so far dealing successfully with a funding policy based on short-term projects. Working tightly with four specific heading organizations clearly represent more opportunities than constraints. Nevertheless, the cost in terms of time spent for administrative duties is very high. All permanent LIP researchers agree today that this workload is becoming too heavy, and that they reach the limit of what is bearable. In a time of major French changes of research organization and funding, with impacts both at the local and national levels, we believe in the importance of fundamental and long-term research, and we invest ourselves to preserve it.

2. Arénaire - Computer Arithmetic

2.1 Team Composition

Permanent researchers Brisebarre Nicolas (CR CNRS), De Dinechin Florent (MCF ÉNS Lyon), Hanrot Guillaume (Pr ÉNS Lyon since 09/2009), Jeannerod Claude-Pierre (CR INRIA), Lefèvre Vincent (CR INRIA), Louvet Nicolas (MCF UCB Lyon1 since 11/2008), Muller Jean-Michel (DR CNRS), Revol Nathalie (CR INRIA), Villard Gilles (DR CNRS).

Post-docs, secondments and and engineers Fagbohoun Christman (Engineer, INRIA), Mayero Micaela (CR, Univ. Paris 13, secondment to INRIA), Novocin Andrew (Engineer, INRIA), Takeuming Honoré (Engineer, ÉNS Lyon/ANR TCHATER), Torres Serge (Engineer (40%), MESR), Vázquez Álvarez Álvaro (Post-doc INRIA).

Doctoral students Joldes Mioara (Région Rhône-Alpes), Jourdan Lu Jingyan (CIFRE), Martin-Dorel Érik (MESR), Morel Ivan (ÉNS Lyon), Mouilleron Christophe (ÉNS Lyon), Nguyen Hong Diep (INRIA CORDI), Panhaleux Adrien (ÉNS Lyon), Pasca Bogdan (MESR), Pujol Xavier (ÉNS Lyon), Revy Guillaume (MESR).

2.2 Research topics

Keywords computer arithmetic; floating-point arithmetic, IEEE 754 arithmetic, fixed-point arithmetic, interval arithmetic, finite field arithmetic, integer arithmetic, multiple precision, arbitrary precision; reliable computation, rounding-error analysis, numerical validation, code certification, formal proof, software synthesis; approximation theory, elementary function, global optimization, linear algebra, euclidean lattice, cryptology, pairing, computer algebra, digital signal processing; VLSI circuit, FPGA circuit, VLIW processor, embedded system.

Research area The Arénaire project-team aims at elaborating and consolidating knowledge and tools in the field of *computer arithmetic*. Computer arithmetic studies how to represent numbers in a computer and how to implement and manipulate such representations efficiently and accurately, in hardware or in software; it studies also the question of the adequation between such number representations and numerical algorithms.

Within Arénaire, we cover all of these aspects. First, we study *basic arithmetic operators in various number systems*, from their algorithmic design to their optimized implementation (in hardware or software, and on general-purpose or application-specific targets). Second, we study the *interplay between arithmetic operators and the algorithmics of several application domains* like function evaluation, polynomial approximation, linear algebra, lattice basis reduction, cryptology, and signal processing. Here, the question is twofold: How may arithmetic choices impact such application domains in terms of complexity, practical efficiency, accuracy, or reliability? Conversely, to which extent improving algorithms and implementations within those domains and within scientific computing in general, can be useful for enhancing arithmetic, that is, for making it faster and/or more reliable and/or more accurate, etc.? Finally, those two research areas are becoming increasingly tied to a third one, that of *certifying numerical computations*. Within this third field, we work not only on traditional, paper-and-pencil error analysis but also, and more and more, on automatic error analysis and formal methods.

Main goals Through computer arithmetic, we strive to improve the quality of computation at large, both in hardware and in software. We look for improvements in *performance* (frequency or throughput for hardware, cycle count for software) and *cost* (area or power consumption for hardware, code size for software, memory consumption for both). In addition, we look for improvements in *accuracy*, which is more specific to computer arithmetic. Third, we strive to increase *confidence* in the computed result: confidence that each operator is well specified and behaves as expected, confidence that the composition of such operators in a program does not entail unexpected deviations from the result to be computed. Last but not least, by demonstrating the practical feasibility of getting high-quality numerical answers, we aim at fostering and being part of *standardization* actions in the fields of computer arithmetic and numerical computing.

Methodology A specificity of computer arithmetic is that its relatively wide spectrum requires many areas of expertise, from hardware design and expert programming to formal proving and number theory. This is why Arénaire gathers people from rather different domains, whose common skills are arithmetic algorithms.

Our targeted numerical applications benefit from any progress we make at the basic operator level. In turn, our optimized designs and implementations of basic operators are nowadays made possible only through our expertise and state-of-the-art software developments in domains like linear algebra or lattice basis reduction. Consequently, we also study such domains for themselves, with an emphasis on algorithmic complexity and the design of high-performance and reliable software.

To achieve our general goal of high numerical quality, either at the basic operator level or at the application level, we need a combination of rigorous specifications and certification tools. A first approach to this is a careful study of validated arithmetics, and especially interval arithmetic, as well as the possibility of computing exact error bounds or even exact results in floating-point arithmetic. Another, complementary approach is to rely more systematically on formal verification and to use proof assistants like Coq in a way that is well-suited for our purpose.

Finally, a very significant change occurred during the recent years: while a few years ago our main realizations, besides algorithms, were expert handwritten libraries (in hardware or software, and offering operators ranging from the most basic ones to elementary functions and linear algebra kernels), we are now more and more interested in automating the development of such libraries. We now produce software tools that assist us in designing and proving optimized operator libraries. This requires new mathematical and algorithmic expertise.

This evolution from hand writing to synthesis is also pushing us towards the compilation community: If an operator can be designed, optimized, and proven automatically, it is natural that this task eventually ends up being part of the compilation process. This explains why our major industrial partner so far is STMicroelectronics' Compilation Expertise Center in Grenoble.

2.3 Research activities

2.3.1 Number Systems

The study of the number systems and, more generally, of numerical data representations is an important topic in the project. Typical examples are: the redundant number systems used inside multipliers and dividers; alternatives to floating-point representation for special purpose systems; finite field representations with a strong impact on cryptographic hardware circuits; the performance of an interval arithmetic that heavily depends on the underlying real arithmetic.

Major results Jan. 2006-Dec. 2009 P. Kornerup (Southern Danish University, Odense) and J.-M. Muller have shown that, for normal redundant digit sets with radix greater than two, a single guard digit is sufficient to determine the value of a prefix of leading nonzero digits of arbitrary length [61]. With S. Collange (then a student at ÉNS Lyon), F. de Dinechin and J. Detrey have argued that the respective merits of the floating-point and logarithmic number systems can only be assessed on a per-application basis [360], using a library of hardware operators supporting both systems [15, 76].

Self-assessment That topic used to be very important in computer arithmetic. Now, partly due to early work by our group in the late 90's, the adequate number representations for integer and real numbers, depending on the context, are well determined. However, research remains active in number representations for finite-field arithmetic.

2.3.2 Efficient Floating-Point Arithmetic and Applications

Efficient hardware for floating-point arithmetic (as specified by the IEEE-754 standards) is becoming available on an increasing number of chips. A first issue thus is to exploit the specification of such hardware operators in order to enhance the accuracy of numerical computing in general, without sacrificing efficiency. For processors without floating-point hardware, another issue is to emulate IEEE-754 arithmetic as efficiently as possible in software.

Major results Jan. 2006-Dec. 2009 We proposed several new algorithms using the fused multiply-add (FMA) floating-point instruction available on recent processors: in [287, 120] N. Brisebarre and J.-M. Muller showed how to perform multiplication by an arbitrary-precision constant very fast and how to check whether the result is correctly rounded; in [176] P. Kornerup (Southern Danish University, Odense), C. Lauter, V. Lefèvre, N. Louvet and J.-M. Muller gave algorithms that use FMAs for quickly computing correctly-rounded positive integer powers.

P. Kornerup, V. Lefèvre, N. Louvet and J.-M. Muller studied properties of summation algorithms in [673]. They showed in particular that Knuth's 2Sum algorithm is minimal (in terms of number of operations and depth of the dependency graph) among all branch-free algorithms using only rounded-to-nearest floating-point additions/subtractions. G. Melquiond, with S. Boldo (Proval, Inria Saclay), showed how to emulate an FMA [827, 116].

Error-free transforms (EFT) such as the above 2Sum algorithm have further been studied as a way of emulating extended-precision efficiently: C. Lauter gave algorithms for multi-double arithmetic in [894]; N. Louvet (jointly with Ph. Langlois from Université de Perpignan) showed in [575] how EFTs can be used to design a compensated Horner algorithm that simulates k times the precision of native floating-point arithmetic when evaluating polynomials.

Concerning software emulation of floating-point operators using integer instructions, C.-P. Jeannerod and G. Revy (jointly with H. Knochel and C. Monat from STMicroelectronics), introduced in [850] a faster algorithm for correctly-rounded square roots, based on highly-parallel bivariate polynomial evaluation and well suited to VLIW architectures. With G. Villard, they extended the approach to correctly-rounded division [669], the main difficulty here being the automatic numerical certification of the algorithm. Both algorithms were implemented in the FLIP library [934] and doubled the performance of $\sqrt{}$ and $\div$ in the code generated by the ST200 VLIW compiler.

Finally, the book "Handbook of Floating-Point Arithmetic" [769] has been published in 2009.

Self-assessment Our expertise in floating-point arithmetic is one of the strengths of Arénaire and we propose state-of-the-art algorithms and software implementations. However, we have to find a way of securing the long-term existence of our software.

2.3.3 Efficient Polynomial and Rational Approximations

Since the basic operations that are available in hardware are additions/subtractions and multiplications, the only functions of one variable one can compute are piecewise polynomials. Hence it is natural to approximate the functions we wish to evaluate by polynomials. The classical methods for computing approximations to functions cannot be used without losing accuracy, since they generate approximations with "infinitely precise" coefficients, that must be rounded to some format when we evaluate the polynomial or rational function: the corresponding "rounded approximation" may be significantly less accurate than the best approximation taken among the "rounded polynomials".

Major results Jan. 2006-Dec. 2009 We developed methods for generate very good polynomial approximations to functions, among polynomials whose coefficients follow size constraints. N. Brisebarre, J.-M. Muller and A. Tisserand, proposed to scan the integer points in a polytope [39] to obtain the optimal polynomial according to the supremum norm. N. Brisebarre and S. Chevillard then used LLL lattice reduction [438] to obtain a very good polynomial. This method, implemented in Sollya, may also accelerate the previous method. N. Brisebarre and G. Hanrot (LORIA, INRIA Lorraine) started a complete study of the problem for the L_2 norm [439], and proposed theoretical and algorithmic results that make it possible to get optimal approximants. Potential applications to hardwired operators have been studied [40] by N. Brisebarre, J.-M. Muller, A. Tisserand and S. Torres.

In [532], a joint work with M. Ercegovac (University of California at Los Angeles), N. Brisebarre, S. Chevillard, J.-M. Muller and S. Torres extend the domain of applicability of the E-method (due to M. Ercegovac), as a hardware-oriented method for evaluating elementary functions using polynomial and rational function approximations.

Self-assessment Until the papers published by the team, there was no systematic approach to obtain good approximations to a function f over an interval by polynomial or rational functions satisfying practical constraints on the coefficients. The leadership in this topic should be confirmed by future works on the several variable case.

2.3.4 Linear Algebra and Lattice Basis Reduction

The techniques for exact linear algebra evolved quickly and significantly in the last few years, substantially reducing the complexity of several algorithms.. Our main focus is on matrices whose entries are integers or univariate polynomials over a field. For such matrices, our main interest is how to relate the size of the data (integer bit lengths or polynomial degrees) to the cost of solving the problem exactly. A first goal is to design asymptotically faster algorithms, to reduce problems to matrix multiplication in a systematic way, and to relate bit complexity to algebraic complexity. Another direction is to make these algorithms fast in practice as well, especially since applications yield very large matrices that are either sparse or structured. Within the LinBox international project, we work on a software library that corresponds to our algorithmic research on matrices. LinBox is a generic library that allows to plug external components in a plug-and-play fashion. The library is devoted to sparse or structured exact linear algebra and its applications.
We recently started a direction around lattice basis reduction. Euclidean lattices provide powerful tools in various algorithmic domains. In particular, we investigate applications in computer arithmetic, cryptology, algorithmic number theory and communications theory. We work on improving the complexity estimates of lattice basis reduction algorithms and providing better implementations of them, and on obtaining more reduced bases. The above recent progress in linear algebra may provide new insights.

Major results Jan. 2006-Dec. 2009 I. Morel, D. Stehlé and G. Villard proposed a new floating-point LLL algorithm, relying on the Householder QR factorization algorithm [682]. This algorithm is simpler to describe than the previous ones and requires less floating-point precision. G. Hanrot and D. Stehlé investigated Kannan's algorithm for the computation of a shortest lattice vector. They improved its analysis, by providing a lower upper complexity bound [477] and showed that the new bound is tight [849]. X. Pujol and D. Stehlé showed that this algorithm can be implemented safely with floating-point arithmetic [589].

C.-P. Jeannerod, jointly with A. Bostan (EPI Algorithms, INRIA Rocquencourt) and É. Schost (University of Western Ontario), gave a new upper bound on the algebraic complexity of linear system solving for several families of structured matrices [117]. They showed further how this leads to faster algorithms for Hermite-Padé approximation and bivariate interpolation.

In order to compute a guaranteed enclosure of the solution of a linear system, H.D. Nguyen and N. Revol have devised an efficient implementation that makes use of fast existing methods for linear system solving [737].

Self-assessment The strong points here are state-of-the-art complexity results and implementations in the fields of algebraic computation and lattice reduction. Note also that the results around floating-point LLL are being used by other members of the team, e.g., for producing polynomial approximations.

2.3.5 Correct Rounding of Functions

The original IEEE Standard for Floating-Point arithmetic (1985) required that the four arithmetic operations and the square root should be correctly rounded: the returned result should always be the floating-point number nearest the exact result.However it did not require anything concerning elementary functions (sine, exponential, logarithm, etc.) because of the "Table Maker's Dilemma": an elementary function is evaluated to some internal accuracy (higher than the target precision), and then rounded to the target precision. What is the minimum accuracy necessary to ensure that rounding this evaluation is equivalent to rounding the exact result, for all possible inputs? This question cannot be answered in a simple manner, meaning that in the general case, correctly rounding functions requires arbitrary precision, which may be very slow and resource-consuming. Our work focused first on bounding this minimum accuracy for elementary functions, then on implementing efficient correctly rounded functions knowing such a bound. We have thus demonstrated code (in the CRlibm project) where the cost of correct rounding is negligible in average, and less than a factor 10 in the worst case.

Major results Jan. 2006-Dec. 2009 J.M. Muller published the second edition of his book *Elementary functions, algorithms and implementation* [753].

G. Hanrot, V. Lefèvre, D. Stehlé and P. Zimmermann [476] studied the problem of finding hard-to-round cases of a periodic function for large floating-point inputs. C. Lauter and V. Lefèvre designed and proved a new algorithm to detect the rounding boundaries for the power (x^y) function in double precision [178]. N. Brisebarre and J.-M. Muller gave bounds that allow to round algebraic functions correctly [70]. V. Lefèvre, D. Stehlé and P. Zimmermann showed that the search for worst (i.e., hardest to round) cases, started by Lefèvre and Muller in 1999, can be extended to the decimal64 format by determining the worst cases of the decimal64 exponential function [577]. In addition, the programs searching for worst cases in double precision were improved, and many new worst-case results were found.

F. de Dinechin, C. Lauter and J.-M. Muller published an algorithm for evaluating correctly rounded logarithms in double precision [74]. C. Lauter and F. de Dinechin worked on automating the generation of a machine-optimized polynomial approximation to a function [576]. The resulting polynomial is implemented as a C program resorting to a minimum amount of double-double or triple-double arithmetic if needed, and provided with a formal proof.

S. Chevillard and N. Revol designed an algorithm for computing the special function erf in arbitrary precision and with correct rounding [544].

Self-assessment Our work on correct rounding of functions persuaded the IEEE 754 revision committee to recommend that some elementary functions should be correctly rounded, and this is probably the biggest success of Arénaire in the recent years. The new standard was adopted in August 2008. CRLibm functions should now move from academic demonstrators to default functions in the GNU libc, but this requires considerable development effort.

2.3.6 Certified Computing

The certification of floating-point computations gives guarantees on the quality of the computed results. Typically, this is a bound for the error between the exact, mathematical result and the floating-point, computed result. Several levels of confidence can be reached. The first level consists in establishing an error bound, using properties of floating-point arithmetic or interval arithmetic. Another level of confidence consists in generating a certificate that can be checked using formal proof, using a theorem prover such as Coq or PVS.

Major results Jan. 2006-Dec. 2009 Arénaire is a leader for the standardization of interval arithmetic within the IEEE P1788 working group [670] and within the STL of C++ [41, 704]. *Interval arithmetic* has been used to certify efficiently (in terms of time and accuracy) the floating-point solution of a linear system [737]. Arbitrary precision interval arithmetic has been implemented in the C library MPFI. Taylor models based on this arithmetic have been developed and used to get guaranteed bounds on the error between a function and its approximating polynomial [448, 641].

Formal proof can now make use of floating-point arithmetic and interval arithmetic: interval arithmetic (and more specifically Taylor models) has been added to PVS [356, 714, 877], both arithmetics have been added to Coq [292]. This is part of the Gappa software [871] that establishes tight bounds on roundoff errors, and keeps tracks of the computations paths and transforms them so they can be checked by a theorem prover (Coq). Gappa has been successfully applied to codes of geometric computing, to check geometric certificates [702, 91], and to certify the code produced by Sollya for the CRlibm library [362, 840] and also significant parts of the code of the FLIP library [850, 669].

2.3.7 Hardware arithmetic operators

Hardware arithmetic operators are at the core of any computing system, large and small. This research domain has shifted from integer to floating-point operators, with renewed interest in FMA and decimal architectures due to the IEEE-754 standard revision. The team produces novel hardware-oriented algorithms, and also established arithmetic libraries for reconfigurable circuits (FPGAs). We also investigate operators for cryptography.

Major results Jan. 2006-Dec. 2009 M. D. Ercegovac (U.C.L.A.) and J.-M. Muller have introduced a digit-recurrence algorithm for evaluating complex square-roots [77]. They have adapted Ercegovac's E-method to the digit-recurrence solution of some tridiagonal linear systems [370] and to the evaluation of complex polynomials [459, 167]. They have designed operators for complex arithmetic [458, 657].

J. Detrey and F. de Dinechin have explored hardware algorithms specifically designed for elementary function evaluation on FPGAs. They have introduced table-based methods for exponential and logarithm [75] then sine and cosine [366, 451] up to single-precision. For larger precisions, they have introduced a recurrence algorithm well suited to the fine-grain look-up table (LUT) structure of FPGAs [452]. These operators are available as part of the FPLibrary project.

F. de Dinechin has worked with O. Creţ, I. Trestian, R. Tudoran, L. Darabant, and L. Văcariu, all from UT Cluj-Napoca, on the acceleration, using FPGAs, of a large floating-point physics simulation for a biomedical application [602, 125]. A conclusion was that efficient floating-point on FPGAs should not use the same operators as in processors, but explore original, application-specific ones [839]. This is now the subject of Bogdan Pasca's PhD and the FloPoCo project, which supersedes the FPLibrary. In addition to the previous elementary functions, operators studied so far include multipliers by a constant [534, 533], application-specific floating-point accumulators and dot-products [547], squarers and other optimized multipliers [646], square root [644].

In addition, FloPoCo is an architecture generator framework with several novel features, such as target-specific operator specialization, automatic frequency-directed pipeline, and test-bench generation [645].

With A. Tisserand (LIRMM), N. Veyrat-Charvillon used special polynomial approximations to optimize arithmetic operators [141]. With R. Michard, they also studied hardware operators for evaluating power functions [389].

J.-L. Beuchat and J.-M. Muller have exhibited a family of modular multipliers suited for FPGA applications [114]. In collaboration with T. Miyoshi and E. Okamoto (Tsukuba Univ.), they have written a survey on multiplication in GF(p) and GF(p^n).

In collaboration with E. Okamoto (Univ. Tsukuba, Japan), J.-L. Beuchat, N. Brisebarre and J. Detrey have introduced new arithmetic operators for pairing-based cryptography [430, 429, 526, 112], winning the best paper award at CHES'2007 [429].

In collaboration with R. Glabb (Calgary University), L. Imbert and A. Tisserand (LIRMM, Montpellier) and G. Jullien (Calgary University), N. Veyrat-Charvillon designed a multi-Mode operator for SHA-2 hash functions [83].

Self-assessment The strong points here is the shared Arénaire expertise in optimization, automation, and validation, enabling open-source projects such as FPLibrary/FloPoCo which are widely used worldwide. The main weak points is a lack of hand-on expertise in VLSI technology. For instance, power consumption, which is more and more important in deep sub-micron technologies, should probably be considered as a first-class objective along with cost, performance and accuracy. This weakness should be solved either by recruiting a VLSI expert or by building stronger collaboration with the VLSI community or with companies designing VLSI chips.

3. Compsys - Compilation and Embedded Computing Systems

3.1 Team Composition

Permanent researchers Christophe Alias (CR INRIA), Alain Darte (DR CNRS), Paul Feautrier (PR ENS, emeritus), Fabrice Rastello (CR INRIA).

Post-doc, engineer Quentin Colombet (INRIA, Sceptre project), Laure Gonnord-Danthony (UCBL, ATER).

Doctoral students Benoit Boissinot (ENS-Lyon), Alexandru Plesco (MESR).

3.2 Executive summary

Keywords Embedded systems, DSP, VLIW, FPGA, hardware accelerators, compilation, code & memory optimization, program analysis, high-level synthesis, parallelism, scheduling, polyhedra, graphs, regular computations.

Research area The industry sells many more embedded processors than general-purpose processors. The field of embedded systems is one of the few segments of the computer market where the European industry still has a substantial share. Compsys is primarily concerned with compilation, and more precisely code optimization, for a subclass of embedded systems, called embedded computing systems, whose applications are more compute-intensive (signal processing, multimedia, stream processing) than control-intensive or with hard real-time constraints. Two aspects are of importance: code optimization for embedded processors (DSP, VLIW, or even possibly micro-controllers) and code optimization for high-level synthesis (HLS) of hardware accelerators. The first includes aggressive static code optimization as well as just-in-time compilation. The second includes memory and communication optimizations, hardware/software co-design, power analysis, design methodologies, etc.

Main goals To adapt/extend some code optimization techniques, primarily developed in compilation/parallelization for high-performance computing (HPC), to the domain of compilation for embedded computing systems. Compsys focuses more particularly on micro-codes optimizations for embedded processors and on high-level synthesis of hardware accelerators. Our aim is to help bridging the gap between compilation and architecture, and between computer science and electrical engineering. Indeed, we are convinced that the techniques and mathematical tools developed in the past for HPC are of interest for specific optimizations in HLS and embedded programs optimizations. But we also believe that optimizing and designing embedded systems is not yet a routine effort and that a lot needs to be learned from the application side, the electrical engineering side, and the embedded architecture side.

Methodology A specificity of Compsys is to tackle combinatorial optimization problems arising from actual compilation problems (register allocation, cache optimization, memory allocation, scheduling, consumption, generation of software/hardware interfaces, etc.) and to validate these developments in compiler tools. To address relevant problems and to have more impact, we believe our research efforts should be combined with strong industrial collaborations, such as with STMicroelectronics (Grenoble and Crolles) through which we validate our results.

Highlights In the period 2006-2009, Compsys activities resulted in 6 PhD theses (4 local, 2 in Maghreb), 5 journals, 32 publications in conferences, including 5 best paper awards. The main scientific achievements were:

- We developed a strong collaboration with the compilation group at STMicroelectronics, with important results on instruction cache optimization, register allocation, and static single assingment (SSA). In particular, Compsys was the first group to push the use of SSA for register allocation. We completely deconstructed the classic view on register allocation. With our colleagues from STMicroelectronics, we are now well-identified internationally for this contribution. In addition to our results and publications, this topic created a lot of activity in seminars, tutorials, organization of workshop (for example, we organized the first seminar on SSA and will organize CGO 2011), research proposals, hiring of young researchers, etc.
- We were convinced that our past expertise on loop transformations and polyhedral optiizations could be useful for embedded computing, for many reasons (see hereafter), but that previous techniques needed to be adapted and even extended. We were successful in developing even further the foundation of high-level program transformations, including scheduling techniques for Kahn processes, lattice-based techniques for memory reuse, and more recently analysis and transformations for while loops, to go beyond standard static affine loops. Such work is generating more and more interest from the high-level synthesis community.
- We had many original contributions with partners closer to hardware constraints, including CEA, related to SoC simulation, hardware/software interfaces, power models and simulators. This activity however will be stopped due to the departure of Tanguy Risset and Antoine Fraboulet.

Compsys became an Inria project-team in 2004. The main event in 2007 was its evaluation in April. The evaluation was conducted by Erik Hagersted (Uppsala University), Vinod Kathail (Synfora, inc), and Ramanujam (Baton Rouge University). Compsys was positively evaluated and was allowed to continue for 4 more years, but in a new configuration as Tanguy Risset and Antoine Fraboulet left the project to follow research directions closer to their host laboratory at Insa-Lyon. Due to this size reduction, Compsys decided to focus on two research directions only:

- Code generation for embedded processors: aggressive compilation and just-in-time compilation.
- High-level program analysis and transformations for high-level synthesis tools.

Both activities are supported by contracts with STMicroelectronics (compilation team and HLS team, respectively) with academic partners from Inria (Alchemy and Cairn, respectively). In addition to the Inria annual funding, this will give sufficient funding for the next 3 years. Nevertheless, to make the HLS french community more tight, collaborations have started with the other HLS research french group, which will maybe lead to an ANR proposal.

The last important highlight concerns human resource. After the departure of Tanguy Risset and Antoine Fraboulet, and the predicted retiring of Paul Feautrier, Compsys was in great danger with possibly only 2 members. Fortunately, Paul Feautrier will continue as an emeritus professor, and Christophe Alias was hired as an Inria research scientist. Therefore, at least for the next 3 years, Compsys should have a stable size, with well-defined objectives.

3.3 Research activities

When creating Compsys, we were all convinced of three complementary facts: a) the mathematical tools developed in the past for manipulating programs in automatic parallelization were lacking in high-level synthesis and embedded computing optimizations and, even more, they started to be rediscovered frequently under less mature forms, b) before being able to really use these techniques in HLS and embedded program optimizations, we needed to learn a lot from the application side, from the electrical engineering side, from the embedded architecture side, c) compilation and code optimization will be at the heart of the development of embedded computing systems. Our primary goal was thus twofold: to increase our knowledge of embedded computing systems and to adapt/extend code optimization techniques, primarily designed for high performance computing, for these systems. We proposed four research directions, centered on compilation methods for embedded applications, both for software and accelerators design:

- Code optimizations for special-purpose processors;
- Platform-independent high-level code transformations;
- Hardware and software system integration;
- Development of polyhedral tools (transversal objective).

3.3.1 Objective 1: Code optimization for special-purpose processors

List of participants: B. Boissinot, F. Bouchez, Q. Colombet, A. Darte, S. Hack, F. Rastello.

Keywords Compilation, code optimization, register allocation, SSA form, VLIW and DSP processors.

Scientific issues, goals, and positioning of the team: Our goals are to develop a new understanding on optimizations and algorithms for aggressive and JIT compilation. We focus on a deep study of combinatorial problems to derive new strategies. We also pay attention to the applicability of our ideas in collaboration with STMicroelectronics.

Major results Oct. 2006-Oct. 2009: Strong collaboration with STMicroelectronics, many joint publications including 3 best papers in a row at CGO (2007-2008-2009)! Compsys is the initiator of a new and, for many, fascinating view on register allocation, in the light of SSA form, which completely deconstructs the common belief. This led to 2 tutorials on the subject and the organization of the very first seminar on SSA (4 days in Autrans).

Self-assessment Our activities with STMicroelectronics are a huge success, not only for the contracts we get. Being able to work within a complete industrial compiler makes our research relevant and more than just theoretical. We also attract young researchers with this activity (PhD students and international post-docs) and our results are getting used (or at least studied) by more and more people (including industry) in the world.

This work was done in close collaboration with STMicroelectronics, including B. Dupont de Dinechin, C. Guillon, F. de Ferrière, T. Bidault. We also have contacts (through visits) or stronger collaborations with M. Smith (Harvard), J. Palsberg (UCLA), S. Amarasinghe (MIT), P. Brisk (EPFL), S. Hack (Sarrebrücken), F. Pereira (Belo Horizonte).

Context and goals Compilation for embedded processors is either aggressive or just in time (JIT). Aggressive compilation consists in allowing more time to implement costly solutions (so, looking for complete, even expensive, studies is mandatory): the compiled program is loaded in permanent memory and the compilation time to obtain it is not so significant. For embedded systems, code size and energy consumption usually have a critical impact on the cost and the quality of the final product, hence the application is cross-compiled, i.e., compiled on a powerful platform distinct from the target processor. JIT compilation, on the other hand, corresponds to compiling applets on demand on the target processor. The code can be uploaded or sold separately on a flash memory. Compilation is performed at load time or even dynamically during execution. The heuristics, constrained by time and limited resources, cannot be too aggressive: they must be fast enough. In this context, our goal is to contribute to the understanding of *combinatorial* problems that arise in compilation for embedded processors (e.g., in opcode selection, SSA conversion, register allocation, or in code placement in the instruction cache) so as to to derive both aggressive heuristics and JIT techniques. So far, we concentrated our research on instruction cache optimization, SSA conversion, and register allocation, for *aggressive* compilation. JIT compilation will be addressed more deeply in the upcoming years. A first specificity of our research is that we try to add a theoretical value on the problems we address (using graph theory, NP-completeness, integer linear programming), even for problems that can appear "old" (such as register allocation). The second specificity is that we have the luck, thanks to our collaboration with STMicroelectronics, to be able to implement and test our techniques directly within an industrial compiler.

Instruction cache optimization In the ST220 processor, the instruction cache is "direct mapped": the cache line i can hold only instructions whose addresses, modulo L, are equal to i, where L is the number of lines. When a program starts executing a block not in the cache, it must be loaded from main memory (cache miss). This happens at the first use of a function, or after a conflict, i.e., when two functions, with interleaved executions, share the same cache line and evict each other from the cache. A cache miss costs ~ 150 cycles for the ST220, hence the interest of link-time procedure placement to minimize line sharing for conflicting functions. We revisited the approaches of Gloy-Smith and Pettis-Hansen, based on a conflict graph, and proved several complexity results related to cache conflict, code size, and locality optimization [18]. In particular, we showed that once the placement of functions is decided (i.e., modulo L), placing the functions in the code to minimize its size can be done in polynomial time (traveling salesman problem on a cyclic-metric graph). For choosing the placement modulo L, we proposed several heuristic improvements. We are still working on this problem, one of the challenges is to generate a relevant conflict graph, not from an execution trace, but from the control-flow graph.

SSA conversion and register allocation Static single assignment (SSA) is an intermediate representation in which multiplexers (called ϕ functions) are used to merge values at a "join" point in the control graph. It is now very popular in compilers due to its properties that speed up optimizations and make them easier to implement. However, ϕ functions have to be replaced, at the end of the process, by register-to-register copy instructions on control flow edges. A naive method for translating out of SSA generate many useless copies (live-range splitting) and additional goto instructions (edge splitting). A too aggressive coalescing (removal of register copies) can degrade the next compilation phase, register allocation, by increasing spilling (load/store insertions). After a first experience in out-of-SSA translation, we understood quickly that SSA could be useful for register allocation as graph coloring under SSA is polynomial because the corresponding interference graph is chordal, a property strangely never noticed nor used before in the compilation community. [1] This remark was the starting point of a deep and long research on register allocation, including the absorption of the literature on this very old problem. Before claiming anything new, we needed to study this problem on all aspects, to check that what we thought we could improve on one side was not destroyed due to another unexpected aspect.

We first revisited the initial NP-completeness proof given, in 1981, by Chaitin et al. to show that it was used wrongly to justify heuristics for problems that it does not cover. Our study shows that, somehow surprisingly, the NP-completeness of register allocation is *not* due to the coloring phase, as may be suggested by a misinterpretation of the reduction of Chaitin et al. from graph k-coloring. If live-range splitting is taken into account, deciding if k registers are enough or if some spilling is necessary is not as hard as one might think. In fact, the NP-completeness of register allocation is due to three factors: special edges (called critical or abnormal) in the control flow graph, the optimization of spilling costs (if k registers are not enough) and of coalescing costs, i.e., which live-ranges should be fused while keeping the graph k-colorable. These results have been presented at WDDD'06 [345] (the "debunking" workshop) and LCPC'06 [346]. We also made a complete study of spill-everywhere spilling problems (LCTES'07 [436]) and of coalescing problems (CGO'07 [435] for the theory, with a **best paper award**, and CASES'08 [529] for heuristics). In both cases, we analyzed the complexity in terms of graph structures, especially for chordal graphs and k-greedy graphs, those used in register allocation based on graph coloring. These studies are the basis to develop new register allocation schemes based on two phases, a first spilling phase, followed by a coloring/coalescing/splitting phase with no additional spills.

SSA and JIT compilation Liveness analysis is an important analysis in optimizing compilers, especially for code generation. Two drawbacks of conventional liveness analyzes are their fairly expensive computations and the fact that their results are easily invalidated by program transformations. We proposed a method to check liveness of SSA variables that overcomes both obstacles, by avoiding building interference graphs and liveness sets. Our analysis survives program transformations except for changes of the control-flow graph. For common program sizes, our technique is faster and consumes less memory than conventional data-flow approaches. We heavily make use of SSA-form properties to completely circumvent data-flow equation solving. This fast liveness checking received the **best paper award** at CGO'08 [528],

In SSA, many code optimizations can be performed with fast and easy-to-implement algorithms. However, some of them create situations where SSA variables corresponding to the same original variable now have overlapping live ranges. This makes the translation from SSA code to standard code more complicated. There are three issues: correctness, optimization (elimination of useless copies), speed (of algorithms). Improving previous work (Cytron et al., Briggs et al., Sreedhar et al.), we proposed a provably-correct and conceptually simpler approach, based on coalescing and a more precise view of interferences, where correctness and optimization are separated. It reduces the number of generated copies and allows us to develop simpler-to-implement algorithms, with no patches or particular cases, and fast enough for just-in-time compilation. This result received the **best paper award** at CGO'09 [631] (third year in a row!).

3.3.2 Objective 2: High-level code transformations

List of participants: C. Alias, H. Cherroun, A. Darte, P. Feautrier, O. Labbani, A. Plesco, C. Quinson, T. Risset.

Keywords FPGA, HLS, embedded systems, hardware accelerators, memory optimizations, program transformations.

Scientific issues, goals, and positioning of the team: Our goal is to extend the polyhedra-based techniques developed for loop transformation and automatic parallelization to address problems relevant for embedded systems design, such as specification, memory and communication optimization, program analysis.

Major results Oct. 2006-Oct. 2009: Two important conceptual innovations: "critical lattice" for intra-array memory reuse that subsumes all previous approaches and a modular algorithm for scheduling Kahn process networks.

[1] This fact was also independently exploited by J. Palsberg and P. Brisk (UCLA), and S. Hack (Karlsruhe, then post-doc with us).

Self-assessment Slow but strong progress for program analysis and transformations. Hard to find students who can understand simultaneously mathematics, computer science, and architecture/synthesis, also because community is small or not well-organized. Hard topic due to the interaction with HLS tools (tough developments and use).

This work was done through collaborations/contacts with STMicro (P. Urard), Thales (Martes project), and colleagues such as F. Catthoor (IMEC), S. Verdoolaege (Leiden), P. Clauss (ICPS), R. Schreiber (HP Labs), G. Villard (Arénaire).

Context and goals During the last 20 years, with the advent of parallelism in supercomputers, the bulk of research in code transformation resulted in (semi-)automatic parallelization, with techniques based on the description/manipulation of nested loops with polyhedra. Today, embedded systems generate new problems in code optimization, especially for loops, both for optimizing embedded applications and transforming programs for high-level synthesis (HLS). Data storage optimization is of prime importance as it impacts power consumption, performance, and chip area. On the application side, multimedia applications often make intensive use of multi-dimensional arrays, in sequences of (nested) loops, which make them good targets for static program analysis. Indeed, for memory optimizations, the largest gains can be obtained at source level because global program analysis/transformations are feasible. On the architecture side, the hardware, in particular memories, can be customized. When designing a programmable embedded system, the adequate parameters for cache and scratch-pad memories are selected to achieve the smallest cost for the right performance, for a given set of applications. In HLS, memories (size, topology, connections with processing elements) can even be fully customized for a given application. Embedded systems are thus good targets for memory optimizations. But powerful compile-time analysis are needed to (semi-)automatically generate a fully-customized circuit from a high-level C-like description. Our goal is to adapt/extend high-level transformations, used for automatic parallelization, to the context of HLS and embedded computing optimizations. Such techniques started to be rediscovered under various forms and we think our previous expertise can be useful both for disseminating already-known techniques and developping new ones.

Memory reduction When designing hardware accelerators, one has to solve both scheduling (when is a computation done?) and memory allocation (where is the result stored?). This is important to exploit pipelines between functional units, or between external memories and hardware accelerators, and save memory space. An example is image processing, for which it is preferable to store only a few lines and not the entire frame, data being consumed quickly after being produced. Reducing memory size can be done by remapping each array so that it reuses its memory locations when they contain a dead value (intra-array reuse). In 2003-05, we showed that all previous approaches are particular cases of a general technique, based on the construction of *admissible lattices*, whose performance can be guaranteed [12]. We then developed the algorithms needed to implement our memory reuse strategies. The resulting tool Cl@k extends our tool suite PIP/Polylib to lattice manipulations. We developed the interface with programs and the required program analysis in another tool (Bee), thanks to the source-to-source transformer ROSE (Livermore). The combination Cl@k+Bee [416] gives the first complete implementation for array contraction. As a side effect, we get the most aggressive algorithm for converting arrays into scalars. The next step is to develop good program transformations that lead to a reduced memory.

Scheduling and Kahn processes Industry is now conscious of the need for special programming models for embedded systems. Kahn process networks (KPN), introduced 30 years ago to represent parallel programs, are becoming popular again. A KPN is built from processes communicating via FIFO channels whose histories are deterministic. One can then define a semantics and talk meaningfully about the equivalence of two implementations. Also, the dataflow diagrams used by signal processing specialists can be translated on-the-fly into process networks. We extended this model into *communicating regular processes* (CRP), in which channels are represented as write-once/read-many arrays. The determinacy property still holds. As an added benefit, a communication system in which the receive operation is not destructive is closer to the expectations of system designers. Although CRP (as KPN) rely on an asynchronous execution model, we developed a modular scheduling technique (see the International Journal of Parallel Programming [53]) that synchronizes them and, in particular, enables buffer optimization. This approach promotes good programming discipline, allows reuse, and is a basic tool for the construction of libraries. The prototype scheduler Syntol is now complete.

Program analysis for Array-OL The Array-OL formalism was designed 20 years ago at Thales Research to ease the implementation of data-intensive signal processing as found in sonar and radar software. It is implemented as a design environment, SPEAR, which deals only with data movements, the data processing being hidden in black box procedures. When reusing existing software, the programmer has to abstract the movement of data from the program text, a tedious and error-prone activity. Within the Martes ITEA project, whose aim was to promote the use of UML for inter-operability between embedded system design tools, we developed, in collaboration with Thales Research, a tool for the automatic elaboration of SPEAR specifications. It uses the front-end of our tool Syntol, a common UML model, and algorithms from computer algebra for the analysis [674]. The Martes project final results won the silver award at the ITEA Symposium.

3.3.3 Objective 3: Hardware and software system integration

List of participants: H. Cherroun, P. Feautrier, N. Fournel, A. Fraboulet, P. Grosse, T. Risset, A. Scherrer.

Keywords High-level synthesis, FPGA, power consumption, simulators.

Scientific issues, goals, and positioning of the team: The goal was to continue high-level synthesis, and address it in a more general context, including hardware/software interfaces. Power was also an important issue to explore.

Major results Oct. 2006-Oct. 2009: A progressive insertion within the national and international SoC and HLS communities. Interesting models and tools for SoC simulation, sensor networks simulation, and power consumption.

Self-assessment This objective was the furthest from our initial expertise. To enter the world of hardware and software integration, we needed to learn the problems, objectives, constraints of this community, and not just stick on our high-performance computing expertise. We learned a lot, but this spread our activities on a too large spectrum with respect to our size. Also, hardware/software integration is a hard topic, with lot of developments and a lack of students with the adequate expertise if we want to increase the compilation aspects for architecture design.

This work was done through collaborations within the french HLS community, with STMicroelectronics, with CEA-LETI (work on power), and with the Ares group, at Insa-Lyon (work on sensor networks).

Context and goals The fact that Compsys was positioned as a team interested in embedded system design, even if our initial objectives were very focused compared to this large topic, pushed us to start collaborations a bit outside our central activities. We received a PhD grant from STMicroelectronics to address the problem of traffic generators for NoC simulation. Also, we were very interested in exploring how compilation techniques could have an effect on power consumption, an important issue in embedded systems. Many papers address this issue, for example with basic block scheduling techniques, but our belief was that more benefit should be found through the reduction of memory transfers and also at operating system level. To address this issue, we started a cooperation with CEA-LETI, which led us further, in the field of hardware and software power modeling and optimization. Finally, the fact that A. Fraboulet and T. Risset moved to Insa-Lyon made them more interested in topics addressed there, in the Ares team, such as sensor networks.

SoC simulation System on chip (SoC) simulation appears as a major issue in embedded computing system design. We have actively participated to SocLib (now an ANR platform project), which targets a SystemC *cycle accurate* simulation for precise performance evaluation or SoC platform design. We developed a deep understanding of SocLib simulation models, in collaboration with hardware designers (LIP6, Paris, and LabSTICC, Lorient). Our major contribution concerns the analysis and synthesis of on-chip traffic, i.e., the traffic occurring in a SoC. To prototype a network on chip (NoC) quickly, fast simulations are needed to replace the components by traffic generators. We developed a complete and flexible on-chip traffic generation environment, within SocLib, which can replay a previously recorded trace, generate a random load on the network, produce a stochastic traffic fitted to a reference trace, and take into account traffic phases. The introduction of second order statistics is an originality of the work and makes it suitable for the modeling of multi-processor on-chip traffic (see [409, 408] and A. Scherrer's PhD thesis [875]). In addition, through the Minalogic project OpenTLM, which targets a *transaction level* SystemC simulation, our experiments with SocLib showed that we need TLM simulation to prototype compilation optimization for embedded code because cycle accurate simulation is too slow.

Sensor networks Advances and miniaturization of micro-electro-mechanical systems and micro-electronic devices have pushed the development of new applications for wireless networks: wireless sensor networks (WSN). WSN are resource-constrained for computation, communication, energy. In this context, the Worldsens simulation toolchain offers an integrated platform for the design, development, prototyping, and deployment of WSN applications. Worldsens consists in two simulators, WSim and WSNet, that offer very accurate results on the application performances at the different levels of the development. WSNet is a radio medium simulator developed by Ares (Insa-Lyon) specialized in wireless network protocols and architectures. WSim is a simulator partly developed by A. Fraboulet and A. Scherrer in the GCC cross-compiler toolchain. The simulated hardware is based on the MSP430 micro-controller series from Texas Instruments, specifically designed for low power and capable of running a full operating system. For more details on WSim, see the longer version of this document. The simulation suite and performance estimation tools were presented in [372, 443, 467] and several demos at conferences [357, 715, 463]. See also the N. Fournel's PhD thesis [881].

Power models The applications in modern embedded systems (signal/audio/video processing) are increasingly complex and compute intensive. In general, high-processing power entails high power consumption. But battery capacity does not follow Moore's law. Low power consumption can be achieved at the technology level (size reduction, threshold voltage control) or at the circuit level, but also at the system and software level: this is where Compsys can contribute.

With CEA-LETI, we developed a power optimization method for a radio modem, extensible to any dataflow system. The idea is to schedule the application to deduce, from the application timing constraints, the minimum performance of each block. Supply voltages and clock frequencies are then adjusted on a per block basis to meet these performances. Detailed simulation shows that energy reduction of $\sim 50\%$ can be achieved (see [376, 172] and P. Grosse's PhD thesis [882]).

Direct physical measurements on a VLSI chip need specialized equipment, but are easier on a development platform. N. Fournel, with help from Cegely Lab (INSA-Lyon), did such experiments to model the power consumption of a processor, its cache, scratch pad, and external memory. The clock frequency influence was measured but the voltage scaling influence had to be extrapolated. The resulting model, coupled to an instruction set simulator, predicts the energy budget of quite complex softwares and of operating system services s.t. task switching or synchronization. Application to realistic image processing applications has shown that the cumulative error – power model error plus simulator approximations – is $< 10\%$. The resulting tool has been applied to the evaluation of cryptographic protocols [466] and to WSN [464, 465].

4. Graal - Algorithms and Scheduling for Distributed Heterogeneous Platforms

4.1 List of participants

Permanent researchers: Anne Benoit (MCF ENS Lyon), Yves Caniou (MCF UCBL), Eddy Caron (MCF ENSL), Frédéric Desprez (DR INRIA), Gilles Fedak (CR INRIA, since Sept. 2008), Jean-Yves L'Excellent (CR INRIA), Loris Marchal (CR CNRS, since Oct. 2007), Christian Pérez (CR INRIA, since Oct. 2008), Yves Robert (Pr ENS Lyon and Institut Universitaire de France), Bernard Tourancheau (Pr UCB Lyon 1, since Feb. 2007), Bora Uçar (CR CNRS, since Jan. 2009), Frédéric Vivien (CR INRIA).

Post-docs and engineers: Nicolas Bard (Engineer, Dec. 2006-), Laurent Bobelin (Post-doc, since March 2009), Hinde Lilia Bouziane (Post-doc, since Sep. 1, 2009), Indranil Chowdhury (Post-doc, since May 25, 2009), Haiwu He (Engineer, since Sep. 2008), MIchaël Heymann (Post-doc, since Jan. 2009), Benjamin Isnard (Engineer, since March 2008), David Loureiro (Engineer, since April 2009), Clément Mathieu (Engineer, since June 2009).

Doctoral students: Leila Ben Saad (2008-), Julien Bigot (2008-), Ghislain Charrier (2007-), Benjamin Depardon (2007-), Fanny Dufossé (2008-), Matthieu Gallet (2006-2009), Jean-Sébastien Gay (2005-), Matthias Jacquelin (2008-), Vincent Pichon (2009-), Paul Renaud-Goud (2009-), Clément Rezvoy (2007-).

4.2 Executive summary

Keywords Algorithm Design, Scheduling, Optimization, Distributed Computing, Grid Computing, Sparse Linear Solvers, Direct Methods, Network Enabled Servers, Grid Middleware.

Research area Parallel computing has spread into all fields of applications, from classical simulation of mechanical systems or weather forecast, to databases, video-on-demand servers, or search tools like Google. From the architectural point of view, parallel machines have evolved from large homogeneous machines to clusters of PCs, then to aggregations of computing or storage resources through Local Area Networks (LAN) or even Wide Area Networks (WAN). The recent progress of network technology has made it possible to use highly distributed platforms as a single parallel resource. Many research projects have been dealing with computing on these new platforms, most of them focusing on low level software details. We believe that many of these projects failed to study fundamental problems such as complexity of problems and algorithms, and guaranteed scheduling heuristics.

Main goals The GRAAL team works on algorithms and scheduling for distributed heterogeneous platforms. The scope of its work goes from theoretical studies to practical developments in freely-distributed softwares. The GRAAL team tackles two main challenges for the widespread use of distributed platforms: the development of environments that will ease their use (in a seamless way) and the design and evaluation of new algorithmic approaches for applications using such platforms. The team research work is organized in three themes:

1. *Scheduling strategies and algorithm design for heterogeneous platforms*. The research work carried under this theme is rather fundamental: platform and application modeling; complexity studies; design, analysis, and evaluation of algorithms and heuristics.
2. *Scheduling for sparse direct solvers*. In this theme we only consider direct methods for sparse direct solvers, and our aim is to implement the best solutions in the freely-distributed MUMPSsoftware that we co-develop.
3. *Providing access to HPC servers on the Grid*. The aim of this theme is to propose middleware enabling to use Grid computing platforms, in particular in a client-server approach, and to study the research problems associated with them.

Methodology A typical research work in the first theme is as follows: we pick an application kernel or an application model, and we study how it can be adapted to be efficiently run on a distributed heterogeneous platform. We are especially interested in taking communications into account (accurate platform modeling, input/output data transfers, etc.). The methodology in the second theme is quite similar, but mainly targets homogeneous distributed computing platforms.

The work in the third theme goes from theoretical studies in distributed computing to development in the freely-distributed DIETmiddleware, covering aspects such as service discovery, data management and replication, performance analysis/prediction, scheduling, and middleware deployment.

Highlights

- The software package DIET, developed in the team, was chosen to be the Grid middleware of the Décrypthon Grid, a joint initiative of AFM (French association against muscular dystrophy), CNRS, and IBM, to help genomics and proteomics research. DIET ensures the load-balancing of jobs over the 6 computing centers federated by this Grid.
- 2008 saw the launch of a start-up company around the DIETmiddleware and GRAAL's expertise on the deployment of large scale applications over dedicated grids.
- Yves Robert was "Program Chair" of IPDPS 2008 (IEEE Int. Symp. Parallel Distributed Processing), the main conference in our research field.
- Yves Robert was elected IEEE Fellow (promotion 2006) and Senior Member of Institut Universitaire de France (2007).
- Frédéric Desprez received an IBM Faculty award in 2008.
- The latest version of the MUMPS solver has many new features including an out-of-core functionality that allows users to solve even larger problems than before. MUMPS is downloaded more than 3 times per day on average.
- GRAAL members have been invited to participate to the program committees of numerous conferences.
- 5 new permanent researchers and teacher-researchers joined the team during the reporting period; the team acquired new competencies and started work in new topics, like stochastic scheduling.

Three main productions/publications:

1. The DIET software.
2. The MUMPS software.
3. Anne Benoit and Yves Robert. Mapping pipeline skeletons onto heterogeneous platforms. *J. Parallel and Distributed Computing*, 68(6):790–808, 2008.

Other main productions/publications:

4. Arnaud Legrand, Alan Su, and Frédéric Vivien. Minimizing the stretch when scheduling flows of divisible requests. *Journal of Scheduling*, 11(5):381–404, 2008.
5. Olivier Beaumont, Larry Carter, Jeanne Ferrante, Arnaud Legrand, Loris Marchal, and Yves Robert. Centralized versus distributed schedulers for multiple bag-of-task applications. *IEEE Trans. Parallel Distributed Systems*, 19(5):698–709, 2008.
6. Patrick R. Amestoy, Abdou Guermouche, Jean-Yves L'Excellent, and Stephane Pralet. Hybrid scheduling for the parallel solution of linear systems. *Parallel Computing*, 32(2):136–156, 2006.
7. H. Casanova, A. Legrand, and Y. Robert. *Parallel Algorithms*. Chapman & Hall/CRC Press, 2008.
8. Eddy Caron, Frédéric Desprez, Franck Petit, and Cédric Tedeschi. Chapter DLPT: A P2P tool for Service Discovery in Grid Computing. *Handbook of Research on P2P and Grid Systems for Service-Oriented Computing: Models, Methodologies and Applications*, IGI Global, 2009.
9. Abdelkader Amar, Raphaël Bolze, Yves Caniou, Eddy Caron, Andréa Chis, F. Desprez, Benjamin Depardon, Jean-Sébastien Gay, Gaël Le Mahec, and David Loureiro. Tunable Scheduling in a GridRPC Framework. *Concurrency and Computation: Practice & Experience.* 20 (9):1051–1069, 2008.
10. Vincent Breton, Eddy Caron, Frédéric Desprez, and Gaël Le Mahec. High Performance BLAST Over the Grid. *Chapter of the book "Handbook of Research on Computational Grid Technologies for Life Sciences, Biomedicine and Healthcare"*, IGI Global, 2009.

4.3 Synthesis of Research activities

4.3.1 Scheduling Strategies and Algorithm Design for Heterogeneous Platforms

Keywords Algorithm Design, Scheduling, Complexity, Distributed platforms, Heterogeneity

Scientific issues, goals, and positioning of the team:

Scheduling sets of computational tasks on large-scale distributed platforms is quite a challenging problem. Our main objective is to enhance state-of-the-art techniques to cope with the high heterogeneity and low reliability of such platforms.

As we aim at investigating problems of practical interest, we depart from the traditional macro-data flow model and, instead, concentrate on communication-aware models. In addition, we abandon the traditional objective of scheduling algorithms: the minimization of the total execution time, or makespan. This change of objective is justified because most makespan minimization problems are NP-hard or even inapproximable, and because this metric is not relevant in the context of concurrent or online applications. Using realistic models, we concentrate on new multi-criteria approaches that aim at realizing the best possible trade-offs between performance oriented objectives (throughput, latency, stretch), and in some cases reliability or energy minimization. We investigate both offline problems, where all application and platform parameters are known beforehand, and online problems.

We target very challenging algorithmic problems. Still, our goal is to provide efficient solutions owing to the combination of realistic models, sophisticated mathematical tools, and experimental validation. At an international level, our commitment to design theoretically guaranteed solutions using realistic platform and application models helps us make the bridge between researchers in scheduling theory and designers of Grid middleware.

Major results Oct. 2005-Oct. 2009:

We have investigated multi-criteria scheduling techniques for streaming applications, mainly pipeline workflows. The considered criteria include both performance-related objectives (throughput, response-time, reliability) and cost-oriented ones (rental price, energy). Our research has been oriented to assessing new complexity results and to provide efficient heuristics (whose performance is compared to an exact linear programming formulation) for the combinatorial instances [111, 160, 156, 158].

We have worked on *online scheduling*, and especially on the minimization of the maximum-stretch (or slowdown) when scheduling flows of divisible requests. We established new theoretical results and proposed practical heuristics [138, 383]. We later extended this study by considering bag-of-tasks applications [520, 159].

We investigated a new scheduling framework, called *steady-state scheduling*, and dedicated to heterogeneous platforms with regular load. In this context, we established the complexity of a large number of scheduling problems, from collective communications [35] to series of task graphs. Recently, we proposed practical algorithms based on steady-state scheduling [560, 661].

During the reporting period, we have initiated some work on *scheduling in a stochastic context* [627, 623]. For instance, we theoretically studied, and proposed practical heuristics for, the scheduling of a divisible load on a collection of computers prone to failures, when the objective is solely the maximization of the expectation of the work completed [627].

Finally, we have provided optimal linear algebra kernels for multicore platforms [668]. This work will be extended in several directions (more precise cache and communication models, implementation on larger platforms).

Self-assessment

Strong Points.

- The project has a very good worldwide visibility and recognition. We publish in the best forums (journals such as IEEE TPDS and JPDC, conferences such as IPDPS and HPDC). We chair or are members of numerous program committees, and stand on the editorial board of several journals.
- The project is renewing its main focus at a dynamic (but reasonable) pace. In the last four years, we have investigated new topics and got acquainted with new mathematical and algorithmic techniques (online problems, stochastic algorithms, multi-criteria optimization).
- Our grant proposals received a strong acceptance rate and support from ANR, CNRS, INRIA, and other institutions.

Weak Points. We lack experimental validation for several of our research results, but the last four years have seen a steady increase of our efforts in that (time-consuming) direction. We would benefit from more industrial collaborations, but we transfer our skills and knowledge to other disciplines.

4.3.2 Scheduling for Sparse Direct Solvers

Keywords direct solvers, HPC, sparse matrices, scheduling

Scientific issues, goals, and positioning of the team:

The solution of sparse systems of linear equations (symmetric or unsymmetric, most often with an irregular structure) is at the heart of many scientific applications, usually in relation with numerical simulation. Because of their efficiency and robustness, direct methods (based on Gaussian elimination) are methods of choice to address these problems. In order to deal with numerical pivoting and adapt to the parallel architecture dynamically, we are especially interested in approaches that are intrinsically dynamic and asynchronous, leading to irregular tasks graphs that can vary dynamically. This is challenging from the point of view of parallelism and scheduling strategies. Our main goals consist in processing larger and larger problems (as required by our users) efficiently and in a scalable way, while keeping in mind numerical functionalities and accuracy. Note that our research in this field is strongly related to the software package MUMPS.

Major results Oct. 2005-Oct. 2009:

During this period, we have pursued some work on scheduling issues for parallel multifrontal solvers, taking into account both memory and performance into account [33]. Notice that the optimization of the memory usage can lead to interesting theoretical problems even in serial environments, as shown in [57]. One of our main achievements concerns the out-of-core factorization, where disks become necessary when the limits of the core memory are reached. This has led to both theoretical [414] and practical contributions [103], although software development issues appeared to be much heavier than expected. Finally, we have worked on parallelizing the symbolic and scaling techniques of parallel sparse direct solvers. As much as possible, the results of the associated research have been injected in the software package MUMPS, which implies a huge amount of software engineering work. One of our objective was to stay close to applications [185, 186], guiding our research directions with the requirements from our users.

Self-assessment

Our main strong points are:

- a good international visibility, both in academia and industry,
- a highly used software package MUMPSin various fields related to numerical simulation,
- research themes motivated by needs from applications and re-injected (as much as we can) in software.

There are also some weak points:

- the team working on this topic and on the MUMPS package is geographically distributed (Lyon, Toulouse, Bordeaux),
- software quality is difficult to maintain as there is hardly enough time to work on pure software issues; staff on temporary positions require a significant start-up on the software aspects and their developments can be difficult to integrate or maintain after their departure.

4.3.3 Providing Access to HPC Servers on the Grid

Scientific issues, goals, and positioning of the team:

Resource management is one of the key issues for the development of efficient Grid environments. Several approaches co-exist in today's middleware platforms. One efficient and simple paradigm consists in providing a semi-transparent access to computing servers by submitting jobs to dedicated servers. This model is known as the Software as a Service (SaaS) model where providers offer, not necessarily for free, computing resources (hardware and software) to clients in the same way as Internet providers offer network resources to clients. The programming granularity of this model is rather coarse. One of the advantages of this approach is that end users do not need to be experts in parallel programming to benefit from high performance parallel programs and computers. It also provides more transparency by hiding the search and allocation of computing resources.

To design middlewares implementing such paradigm, we need to address issues related to several well-known research domains. In particular, we focus on: middleware and application platforms as a base to implement the necessary "glue" to broke clients requests, find the best server available, and then submit the problem and its data, resource and service discovery, online and offline scheduling of requests, fault tolerance, link with data management and replication, distributed algorithms to manage the requests and the dynamic behavior of the platform.

Recently, we consider other programming abstractions, and in particular component based models, so as to ease the exploitation of large scale and heterogeneous platform with more abstract programming models. We also have extended our research to Desktop Grids, which use the computing, network and storage resources from idle desktop PC's distributed over multiple-LAN's or the Internet to compute a large variety of resource-demanding distributed applications.

Major results Oct. 2005-Oct. 2009:

The DIET software, first designed for our own experiments and algorithms [46, 336, 348, 440, 536] and software validation [43, 104, 343, 349, 554, 541], was chosen by IBM and AFM to be the production middleware of the Decrypthon Grid. This project was really important for the team, both around the study of bioinformatics application over Grid platform (scheduling and data replication issues [48, 450, 554, 525], middleware development and adaptation) and because it shown that our developments were mature enough to be used in production. This also let us think about starting a company around this software and our expertise in application developpement and optimization.

We have worked on service discovery on P2P environment. We designed the DLPT (Distributed Lexicographic Placement Table) architecture [71, 412, 763, 902], which proves mechanisms for load balancing and fault-tolerance. The solution centers around three parts. First, it calls upon an indexing system structured as a prefix tree, allowing multi-attribute range queries. Second, it provides the built-in mapping of such structures onto heterogeneous and dynamic networks and proposes some load balancing heuristics for it [351, 542]. Third, as our target platform is dynamic and unreliable, we worked on original fault-tolerance mechanisms, based on self-stabilization [164, 350, 352, 441, 599, 540].

We define and implement improved component models that show it is possible to offer a rich set of concepts of higher abstraction level (such as data sharing, workflow and skeletons) while preserving high performance by being able to re-use state of the art middleware systems.

Self-assessment

Strong Points.

- we are able to work on theoretical research issues linked to sofware development. Our implication in major research projects such as Grid'5000 and Decrypthon allows us to validate both theoretical and software results on large scale and production platforms.
- The DIET software has obtained a national and international visibility.
- We are able to master the whole chain, from component based model definition, to implementation so as to evaluate the benefits of our proposition.

Weak Points.

Acceptance of of new programming concepts takes times as stable enough prototypes are required.

Our software developments, while they were quite successful for our visibility, are time-consuming. We were lucky enough to have a strong support from Inria and ANR but the stability of our engineers is always linked to the success rate of our proposals.

5. MC2 - Models of Computation and Complexity

5.1 Team composition

Permanent researchers: Pascal Koiran (PR ENS Lyon), Natacha Portier (MCF ENS Lyon), Eric Thierry (MCF ENS Lyon), Eric Rémila (IUT Roanne).

Post-docs and engineers: Pablo Arrighi (Délegation CNRS, since 09/2009), Eric Boix (Engineer CNRS), Jonathan Grattage (ATER ENS Lyon, since 09/2009), Arnaud Grignard (Engineer ENS Lyon, since 05/2009), Camilo La Rota (Engineer CNRS, since 02/2007), Ricardo Uribe (Engineer CNRS, since 03/2007), Gian Chiquillo (Engineer CNRS, since 05/2008), Mathieu Malaterre (Engineer ENS Lyon, since 07/2009), Jorge Beltran (Engineer CNRS, since 05/2009).

Doctoral students: Irénée Briquel (2008-), Bruno Grenet (2009-), Laurent Jouhet (2007-), Mathilde Noual (2009-), Julien Robert (2006-).

5.2 Executive summary

Keywords theoretical computer science, algorithms, computational complexity, discrete structures, discrete dynamical systems, complex systems, gene regulatory networks, algebraic complexity, quantum computing, cellular automata, tilings, self-assembly, network calculus, game theory.

Research area The central research topics of our team are:

• *Algorithms and computational complexity.* We design and analyze efficient algorithms, and we try to understand the limitations of efficient algorithms. Understanding these limitations is perhaps the main goal of computational complexity. An early and important example is the theory of NP-completeness, which provides a way of showing that certain problems do not admit polynomial time algorithms (assuming the widely believed $P \neq NP$ conjecture).

• *Complex systems.* According to the *Réseau National des Systèmes Complexes*, complex systems from the cell to the ecosphere result from evolution and adaptation processes. They exhibit emerging properties: the underlying microscopic level gives rise to organized structures at the macroscopic level, and the macroscopic level influences the microscopic level in a feedback loop. These emerging properties are robust and can be studied from different points of views, depending on the class of complex adaptive systems under consideration.

Computational complexity and complex systems are connected through the notion of *dynamical system*. Complex systems are usually dynamic: they evolve in time. The models of computation studied in computational complexity are dynamic as well: a machine's transition function associates to the current configuration of the machine its configuration at the next time step (or the possible next configurations if the machine is not deterministic). We can therefore use our expertise in the study of such systems for the study of complex systems.

Main goals As explained above a central goal is to understand the limitations of efficient algorithms, but *efficient* can take many other meanings than worst-case polynomial time computation on a Turing machine. Algorithms can be parallel, distributed, synchronous or asynchronous, deterministic, probabilistic or quantum. Besides computation time, one can try to optimize memory requirements or volume of data exchanged. In order to focus the attention on one (or a few) of these different aspects of computation, we study various *models of computation* such as, for example, synchronous and asynchronous cellular automata or quantum models of computation.
A central goal of our research in complex systems is to understand interactions between the microscopic and macroscopic levels. We pursue this goal both for idealized mathematical models such as cellular automata or boolean networks, and "real" systems such as gene regulation networks.

Methodology Our work in algorithms and computational complexity usually follows a traditional mathematical style (definitions / theorems / proofs…). We use this mathematical methodology to study algorithms and models of computation, but also to study the underlying mathematical structures. Improved understanding of these structures often leads to better algorithms. For instance, understanding the structure of a tiling space can yield efficient algorithms for the generation of random tilings, and progress in graph theory leads to more efficient combinatorial algorithms.
Sometimes we also resort to experiments. This is especially true in our work on (low dimensional) cellular automata, whose evolution in time can be visualized easily. Rigorous mathematical results remain the ultimate goal: we use experiments to develop our intuition and formulate conjectures that will hopefully become theorems at a later stage.
In our complex systems work we analyze data (coming in particular from biological experiments), build models that explain the data and study the property of these models by mathematical analysis or by software simulation. The results

obtained can suggest further biological experiments.

5.3 Research activities

5.3.1 Algebraic complexity

Keywords: Valiant's model, Blum-Shub-Smale model, arithmetic circuits, treewidth, systems of polynomial equations.

Scientific issues, goals, and positioning of the team: Sylvain Perifel's thesis was mostly devoted to algebraic complexity theory. Connections between the complexity of decision and evaluation problems were established. Connections between algebraic and Boolean complexity were also established. We have studied the complexity of evaluating polynomials within Valiant's framework. A central open problem here is whether the permanent of a matrix can be evaluated in a polynomial number of arithmetic operations. We have studied the complexity of decision problems in the Blum-Shub-Smale model of computation. Now, the central problem is the following: is it possible to decide in a polynomial number of arithmetic operations and comparisons whether a system of multivariate polynomial equations has a solution (in the field of real or complex numbers) ?

Major results Oct. 2006-Oct. 2009: We have shown that there are strong connections between these two algebraic versions of the "P= NP ?" problem and with Boolean complexity theory. For instance, we have shown that a proof of "$P \neq NP$" over the real or complex number would imply that $P \neq NP$ in Valiant's framework, or the boolean separation $P \neq PSPACE$. In [462] we have studied the complexity of evaluating "hard" polynomials such as the permanent or the Hamiltonian for special classes of graphs. It was known that subject to these restrictions, the afore-mentioned "hard" problems become easy. We have shown that the complexity of evaluating the corresponding polynomials is captured exactly by restricted classes of arithmetic circuits, *skew circuits*.

Self-assessment: The results in Sylvain Perifel's thesis together with prior results from our group and other research groups provide a fairly complete picture of the way that some of the main open problems of algebraic and Boolean complexity are connected to each other. It remains to actually *solve* these problems, that is, to prove the corresponding lower bounds unconditionally. We also believe that the connections between Boolean complexity and Valiant's model should be studied further, and that graph-theoretic notions like tree-width will play an important role in that study.

5.3.2 Algebraic algorithms

Keywords: Computer algebra, sparse polynomials, polynomial factorization, determinants.

Scientific issues, goals, and positioning of the team: The design of algorithms for algebraic objects like polynomials or matrices has of course applications in many fields of computer science, and in particular in complexity theory where factorizing or computing determinants and permanents are cornerstone problems between different complexity classes.

Major results Oct. 2006-Oct. 2009: With E. Kaltofen, we have designed efficient algorithms for the factorization of sparse polynomials. After handling bivariate polynomials with integer coefficients and searching linear factors [309], these restrictions were lifted in [377] (factors of arbitrary degree for multivariate polynomials with coefficients in a number field).

In [569], we studied quotients of symbolic determinants: the entries of the determinants at the numerator and denominator of the quotient are constants or variables. In general, the quotient is a rational fraction. We show that when the quotient is a polynomial it can be efficiently represented by a third symbolic determinant. This situation occurs for instance in the computation of resultants of multivariate polynomial systems.

Self-assessment: Our determinantal representation of quotients is probably not efficient enough to be used in practice. It does raise the question, however, of the existence of a determinantal representation which could be used in practice.

It seems that our work on the factorization of sparse polynomials might lead to some (modest) progress on one of the central problems of complexity theory: obtaining good lower bounds on the representation of explicit polynomials by arithmetic circuits.

5.3.3 Kolmogorov complexity

Keywords: Boolean circuits, compression, mutual information.

Scientific issues, goals, and positioning of the team: We have studied the (open) question whether problems solvable in exponential time admit polynomial size circuits. It has lead to further investigations about Kolmogorov complexity which was also a active research theme for A. Romaschenko.

Major results Oct. 2006-Oct. 2009: We have shown that the answer to the first problem is negative [490] under an assumption from resource-bounded Kolmogorov complexity: polynomial-time symmetry of information. We have therefore connected two important open questions from two different areas: boolean circuit complexity and Kolmogorov complexity. We have studied another problem at the frontier of these two areas: how to extract randomness from an infinite string ? To this effect, we have studied in [507] compression of strings by pushdown automata. At the cornerstone of algorithmic information theory and in collaboration with A. Muchnik (Moscow), the stability of "'Kolmogorov type" properties to relativization was also studied with new statements about independence between words [580].

Self-assessment: Though intuitive, results are rather hard to prove. We still lack general technique suitable to attack similar problems.

5.3.4 Quantum computing

Keywords:hidden subgroup problem, lower bounds, query complexity.

Scientific issues, goals, and positioning of the team: We have focused on quantum lower bounds, and in particular on quantum lower bounds for hidden subgroup problems. It is well known that for certain problems, quantum algorithms yield an exponential speedup compared to the best known classical algorithms (e.g. Shor's algorithm and factorization). The hidden subgroup problem is an abstract framework in which most of the the known "fast" quantum algorithms such as Shor's algorithm can be cast. In this framework the algorithm has access to a *black-box function* f defined on a finite group G. This function is supposed to "hide" a subgroup $H \subseteq G$ in the sense that f is constant on each coset of H, and takes different values on different cosets. The goal of the quantum algorithm is to identify the hidden subgroup H by performing as few queries as possible to the black-box f.

Major results Oct. 2006-Oct. 2009: We have obtained optimal lower bounds on the number of queries, for Simon's problem first [311] (the case where $G = (\mathbb{Z}/2\mathbb{Z})^n$) and then for all Abelian subgroups [86].
More recently, we have begun a systematic study of lower bounds for non-adaptive quantum algorithms. In this restricted model, a quantum algorithm must perform all of its queries to the black box simultaneously. This model is especially natural for hidden subgroup problems, since many of the known algorithms for such problems turn out to be non-adaptive.

Self-Assessment: We were the first group to obtain lower bounds for hidden subgroup problems. Four year after the first version of our paper, despite intensive work on lower bounds in the quantum computing community, this is still the only lower bound available for hidden subgroup problems. It would be interesting to obtain also lower bounds for non-Abelian groups.

5.3.5 Fault-tolerant computation

Keywords: Boolean circuits, error-correcting codes.

Scientific issues, goals, and positioning of the team: We have investigated reliable computation with faulty boolean circuits in the *coded model of faulty computations* introduced in the 1990's by D. Spielman. In this model the input and the output of a computational circuit are treated as words in some error-correcting code.

Major results Oct. 2006-Oct. 2009: In [403], we investigated two models of local faults. In the random faults model we suppose that each elementary processor (a boolean gate) at each step with some small probability becomes corrupted. We proved that the complexity of a circuit resistant to random faults can be done very close to the complexity of a circuit computing the same function without fault occurences. faults. In the worst-case fulat model, we assume that at each step some arbitrary fixed fraction of elementary processors can be corrupted by an adversary. We show that one can still organize reliable computations but with an exponential blow up of the memory, with a very modest slow down.

Self-assessment: We suggested some explicit constructions for the synthesis of reliable circuits. The principal novelty of our method is a new self-correcting procedure based on a sort of mixing mapping (this rather simple mapping enjoys some useful properties of expanders).

5.3.6 Infinite words

Keywords: Topological complexity, infinite games, effective strategies.

Scientific issues, goals, and positioning of the team: We study the topological complexity of languages of infinite words by various finite machines such as finite automata, counter machines, pushdown automata, Petri nets... In particular, we try to locate these languages in the Borel or the Wadge hierarchy. To obtain decision algorithms, we considered infinite games between two players with complete information and looked for effective winning strategies. These questions are also connected to the specification and verification of systems interacting with an environment.

Major results Oct. 2006-Oct. 2009: O. Finkel has obtained with D. Lecomte new results [461] on the topological complexity of the ω-powers that appear naturally in the characterization of regular and context-free ω-languages. O. Finkel has also worked in model theory with S. Todorcevic [79] showing that some stretching theorems due to J.-P. Ressayre are in fact equivalent to the existence of n-Mahlo cardinals for all n.

Self-assessment: We have completely solved the problem of classification of ω-powers with respect to the Borel hierarchy, with effective versions of our results. Remarkably, this problem was solved by combining methods from descriptive set theory and automata theory.

5.3.7 Small-world phenomenon

Keywords: Milgram experiment, interaction networks, social networks.

Scientific issues, goals, and positioning of the team: The goal is to study the small world phenomenon emerging in real life interaction networks. The seminal work of Milgram (1967) showed that everybody is not only at about 6 handshakes from anybody else, but is also able to find these few handshakes based on his or her own local view of the huge interaction network. Kleinberg (2000) suggested a possible algorithmic nature of this phenomenon by introducing a random graph (a grid augmented with random shortcuts) where a simple greedy algorithm can find short paths between pairs of nodes.

Major results Oct. 2006-Oct. 2009: We have extended his work in several directions with N. Hanusse and P. Duchon (LaBRI). First, we have proposed a local routing algorithm that does much better on expectation than the greedy algorithm. Second, we have looked at the properties that can make it possible to turn the underlying graph (the grid in the case of Kleinberg) into a small world, such as growth [51].

Self-assessment: Our work can be considered as a first step towards a validation of the ubiquity of the small worlds in nature: these only require few resources to be built [368].

5.3.8 Networks of asynchronous automata

Keywords:cellular automata, boolean networks, stochastic automata, probabilistic analysis.

Scientific issues, goals, and positioning of the team: A natural framework to study some biological models are boolean or Hopfield networks, and cellular automata. Such model have been intensively studied for synchronous updates. We have led a study of other update regimes, like probabilistic ones to model asynchronism.

Major results Oct. 2006-Oct. 2009: We have studied two regimes: α-asynchronicity (i.i.d. updates with probability α) and full asynchronicity (1 update/step uniformly). Asynchronous behaviors differ drastically from their synchronous one in cellular automata, as observed on simulations and proved in double-quiescent 1D elementary cellular automata (classified in [52] for full asynchronism and [371] for α-asynchronism) and 2D Minority cellular automata [493, 593]. Our indicator was the relaxation time, i.e. the expected absorption time in final components of configurations. Some couplings with directed percolation were suspected by experimental studied about coaslescence [404] and have been formalized in [591].

Self-assessment: We managed to introduce some useful combinatorial tools for such analyses like tree masks and still aim at classifying distinct behaviors. Some links with other stochastic processes like Interacting Particles Systems, have not been exploited yet.

5.3.9 Blind scheduling and data broadcast

Keywords:blind scheduling, data broadcast, CPU thermal management.

Scientific issues, goals, and positioning of the team: In most operating systems, the scheduler has no precise information about processes before or during their execution. Consequently we have investigated a class of algorithms called *non-clairvoyant*, getting informed on the fly. J. Edmonds' work (STOC 1999) on the algorithm EQUI (equitably sharing the computing resources between active tasks at any instant) showed that in a general framework about the different levels of parallelism encountered (unknown to the scheduler), the total service time can be competitive with regard to the optimal clairvoyant scheduler (knowing all arrival times and process phases) up to having about twice the resources available.

Major results Oct. 2006-Oct. 2009: We have extended those results to a more general model where tasks are composed of an arbitrary graph of processes (DAG) unknown from the scheduler. We have shown [498, 595] that the algorithm EQUI$\circ$EQUI which shares equitably the resources between each job and then shares the resources allocated to each one between the active tasks, remains competitive with regard to the optimal clairvoyant scheduler up to allowing twice the resources. We have also applied our results to optimize the broadcast of a web server with dependencies between the different broadcasted pages [367, 499].

Self-assessment: This new expertise about online algorithms has led to another collaboration about schedulings ensuring some microprocessor thermal management [557].

5.3.10 Tilings

Keywords:tiling space, domino tiling, lozenge tiling, discrete rotations, self-assembly, fault-tolerance.

Scientific issues, goals, and positioning of the team: In the study of the combinatorics and the algorithmics of tilings, a emphasis has been put on a recent self-assembly model introduced by bio-computer scientist E. Winfree (STOC 2000). In this model with biochemistry and nanotechnology flavors, tiles are squares with different types of glue on their sides which guide the way they aggregate. Some external parameters like "temperature" modify the assembly. This was F. Becker's PhD research topic and it is also a joint work with I. Rapaport and N. Schabanel CMM U. Chile.

Major results Oct. 2006-Oct. 2009: Our original approach consist in constructing homothetic stable shape languages, rather than a fixed shape where the main effort is the encoding of the size, forgetting the geometry. We succeeded in obtaining size optimal and time optimal tiles sets for the construction of squares [342] and later cubes [518]. Until now, it is the only 3D non trivial construction exhibited.
In more classical settings, we have studied the flip accessibilility (sequences of local rearrangements) in lozenge tilings (joint work with O. Bodini (LIP6, Paris 6) and T. Fernique (LIF, Univ Marseille). Previouly studied for finite figures, we have obtained some characterizations for tilings of the whole plane [432, 115], later extended to domino tilings [433].

Self-assessment: When viewed as growth processes, self-assembly models may include some probabilities. Then it comes to analyzing a Markovian process with a strong combinatorial flavor, just like in our recent studies of networks of automata. One can hope that some probabilistic tools will be useful for both those models.

5.3.11 Algorithms for Network Calculus

Keywords: discrete event systems, worst-case performance analysis.

Scientific issues, goals, and positioning of the team: Network Calculus (NC) is a theory of deterministic queuing systems encountered in communications networks. Aiming at analyzing worst-case performances, it stores informations about the system in constraint functions, which are combined thanks to special operators (mainly from (min,+) algebras) to output results. Though mathematical formulas exist for small but typical networks, little is known about the way to make the formulas effective with efficient algorithms and then to aggregate the computations to analyze large and complex networks.

Major results Oct. 2006-Oct. 2009: With the perspective of offering an algorithmic toolbox and later an software, we have introduced a class of functions which is stable for the NC combinations as well as algorithms implementing them [119]. We have started to propose some algorithms to analyze complex networks with multiplexing [530] and investigated some routing issues[118, 437].

Self-assessment: It is a new topic in MC2 team. Our approach with demand for computational complexity analysis is rather new in the NC community. The challenge is to prove that this approach will be useful in the development of a concrete software tool.

5.3.12 Simulation platform for complex systems

Keywords:complex systems, modeling, simulation.

Scientific issues, goals, and positioning of the team: We are developing a software platform [918] for the simulation of complex systems. This work has been supported by the Morphex and Dynanets projects. Toan Chu-Minh from the Oslo company also participated in the development of the platform as a member of the Morphex project.

Major results Oct. 2006-Oct. 2009: The platform has reached a level of maturity which makes it possible to create a startup company. The start-up, led by Eric Boix, is currently under incubation and should begin its operations around spring 2010.

5.3.13 Modeling, simulation and analysis of gene regulatory networks

Keywords: complex systems, plant reproduction, flower development.

Scientific issues, goals, and positioning of the team: The goal of this work is to construct a predictive model for the development of the female reproductive organ of the flower. This entails the modeling, simulation and analysis of the dynamics of gene regulatory networks. This work was supported by ANR project *"carpelle virtuel"*. Several researchers outside of LIP were involved, including S. Séné (PhD student at IXXI supervised by M. Morvan and J. Demongeot), F. Monéger and J. Traas from the plant reproduction lab (RDP, ENS Lyon).

Major results Oct. 2006-Oct. 2009: With threshold networks as our modeling paradigm, each node in the network represents a gene, each state variable is Boolean and represents the concentration of the gene's product. The weight on each edge represents the type and strength of the corresponding interaction. We have obtain some robustness results about the dynamics [552]. In collaboration with RDP we have reconstructed the network corresponding to the first development stages of the flower of the Arabidopsis plant [551]. The network's parameters were optimized by mathematical programming techniques, in a collaboration with L. Liberti and F. Tarissan (LIX, Ecole Polytechnique) [574, 598].

6. Plume - Proof Theory and Formal Semantics

6.1 Team Composition

Permanent researchers Audebaud Philippe (MCF ENSL, back to Plume since 09/2006), Baillot Patrick (CR CNRS, at Plume since 09/2008), Duprat Jean (MCF ENSL), Hirschkoff Daniel (MCF ENSL), Laurent Olivier (CR CNRS, at Plume since 09/2008), Lescanne Pierre (Pr ENSL), Riba Colin (MCF ENSL, since 09/2009).

Post-docs, secondments and and engineers Miquel Alexandre (détachement ENSL), Tranquilli Paolo (Post-doc ENSL), Vecchiato Silvia (doctoral internship).

Doctoral students Demangeon Romain (ENSL), Lasson Marc (ENSL), Pardon Aurélien (ENSL), Petit Barbara (ENSL).

6.2 Executive summary

Keywords Semantics of programming languages, computer assisted proofs, proof theory.

Research area Researchers in the Plume team study methods for the formal analysis of computer programs and, more generally, of computing systems. More precisely, we work on the formal definition of (aspects of) programming languages, and of static analyses of programs. We also rely on theorem proving systems to develop rigorous formalizations and proofs about mathematical theories.

We build on the proofs-as-programs correspondence to strengthen the links between proof theory in logic and computer science. We deal in particular with the integration of additional programming features: concurrency, mobility, modularity, probabilistic behaviours, ...

Main goals The main goals of the team are:

- to strengthen the understanding of various programming constructs (control operators, probabilistic primitives, concurrency, cryptographic aspects);
- to provide logical foundations for modern programming languages;
- to build type systems controlling the behaviour of programs (termination, complexity bounds, dynamic modularity);
- to define extraction mechanisms in order to build certified software from formal proofs of specifications;
- to develop rigorous formalizations of mathematical theories with proof assistants like Coq.

Methodology Our tool-set comes from mathematical logic and the semantics of programming languages. We use the Curry-Howard (*i.e.* proofs-as-programs) correspondence as a bridge between proof-theory and computer science. In this way we are able to build mathematical reasonings about program behaviours.

Our publications often come with a mechanized verification (as Coq contributions) of the results they contain.

More recently, we have put a specific focus on the use of (variants of) linear logic and of the realizability theory as unifying tools to develop the foundations of programming languages.

6.3 Research activities

In order to guarantee safety and security, programs should be developed in a controlled way, and appropriate methods should be devised. Numerous formal methods have been introduced to help programmers in pursuing this goal. Depending on the application domain, the appropriate notion of being controlled may strongly differ (*e.g.* termination vs. liveness). The spectrum of properties that such methods make possible to address includes: security, certification, termination, resource boundedness, deadlock freedom, liveness, control on information flow, ... Inside the large body of work on formal methods, we can distinguish between static approaches (systems are analysed before execution) and dynamic approaches (analysis is done along execution). We focus on static approaches, and more precisely type systems and formal semantics, using logical means.
Researchers in the Plume team study methods for the formal analysis of computer programs and, more generally, of computing systems. More precisely, we work on the formal definition of (aspects of) programming languages, and of static analyses of programs. We also rely on theorem proving systems to develop rigorous formalizations and proofs about mathematical theories. We build on the proofs-as-programs (Curry-Howard) correspondence to strengthen the links between proof theory in logic and computer science.

The Curry-Howard correspondence which relates proofs and programs provides us with a bridge between proof-theory and computer science. It is a very powerful framework for technology transfers between these two research topics. In particular formal proofs are endowed with a computational meaning (cut-elimination is evaluation) and logic gives a crucial tool to analyze the foundations of programming languages. The meaning of this correspondence is very well understood in the intuitionistic/functional case, which is the original setting. It has been extended to classical logic and control operators (like call/cc) in the last twenty years. We are now looking for other extensions to develop logical foundations of additional programming concepts: concurrency, mobility, complexity control, ...
On the programming languages side, the team studies the formal description of expressive programming constructs (such as probabilistic operations, primitives for modularity, or concurrency mechanisms). This is done by following the approach of formal semantics of programming features. It concerns both operational semantics: the precise definition of the computational rules underlying a programming primitive, and denotational semantics: the search for invariants of computation. Our aim is to define and analyze the properties of formalisms that capture these programming constructs in a concise and clear fashion, in the same way the λ-calculus can be used to model sequential functional computation. In doing so, we mostly rely on mathematical tools, to develop the operational and denotational semantics of programs (partial orders, category theory, ...). We often use the obtained models to design type systems allowing to characterize and to analyze specific behavioural properties of programs.
Due to the very formal nature of the objects we manipulate, we are frequently turning our results into developments in the Coq proof assistant. In a more general way, we also use Coq for the formalization of various mathematical theories.

6.3.1 Curry-Howard correspondence

We work on various aspects of the Curry-Howard correspondence: from the traditional (twenty years old) question of classical logic to more recent points such as the implicit complexity point of view. Consequences in computer science are obtained by means of realizability and program extraction from proofs.

Computational interpretations of classical logic

Scientific issues, goals, and positioning of the team The Curry-Howard correspondence, originally introduced between intuitionistic logic and functional languages, is now well understood in the more general context of classical logic and control operators. This led to the development of many logical systems for presenting proofs in classical logic with a well defined computational meaning. The Plume team took part in this line of work and now focuses more on clarifying the relations between the various proposals given in the literature.

Major results Plume members worked in particular on the use of graphical syntaxes for representing proofs from classical logic.

The $\mathcal{X}$ diagrammatic calculus allows for a graphical analysis of the sequent calculus of classical logic. This belongs to the family of classical systems with non-deterministic computational behaviours.

On the other side, id-nets give a deterministic representation of classical proofs. These nets are used to provide us with a unified framework for interpreting the deterministic classical systems. This is based on a strong relation with intuitionistic logic and works both for call-by-name and for call-by-value classical systems.

Self-assessment The theory of the deterministic computational behaviour of classical logic is now well understood and appropriately unified. On the other hand there are still various proposals for non-deterministic systems. In the same way the relation between the deterministic and the non-deterministic cases are still not completely well analyzed.

Program extraction from proofs

Scientific issues, goals, and positioning of the team Not only proofs are programs, but if π is a proof of the formula F, the program P corresponding to π satisfies the specification formalized by F. This provides us with a way of developing certified software without the usual two stages approach (first write the code and then prove the correctness with respect to the specification). After a formalization of the specification of the wanted behaviour of the program through a logical formula, the developer interacts with a proof assistant to write down a proof of the formula. The code is then automatically extracted from the proof and guaranteed to satisfy the specification.

Major results Efficient program extraction from proofs can benefit from various backgrounds. Ph. Audebaud and L. Chiarabini approach aims at applying well known techniques from the programming community to the proof side; the current work concerns tail recursion and defunctionalization.

Following Krivine's work on realizability for classical logic, A. Miquel has shown how to define realizability interpretations for the calculus of inductive constructions (the type theory underlying Coq). Starting from this theoretical result, he developed an extraction module for the Coq proof assistant, addressing concrete issues related with the representation of integers for example.

Self-assessment The abstract theory of realizability is now mature for transfer towards software development. This is the line we are focusing on. The practical use of the theory will naturally bring new questions on the more abstract side.

Implicit Computational Complexity

Scientific issues, goals, and positioning of the team The linear logic (LL) approach to implicit computational complexity aims at defining variants (or restrictions) of LL for which the cut-elimination procedure encodes precisely a given complexity class. In this way, the complexity of programs becomes an additional computational behaviour accessible through logical means. We put a particular focus on two specific questions: the semantic analysis of the derived logical systems and the design of associated type systems.

This is one of the newly started research directions of the team (with the arrivals of P. Baillot and O. Laurent).

Major results P. Baillot and D. Mazza (Univ. Paris 13) have proposed a variant of linear logic called linear logic by levels (L3), which admits an elementary time complexity bound on proof normalization, and at the same time generalizes and simplifies the system ELL of Girard. A subsystem of this linear logic by levels also corresponds to polynomial time computation and generalizes the system LLL (light linear logic). One benefit of these systems with respect to ELL and LLL is that they have simpler definitions with proof-nets (the graphic representation of proofs).

Self-assessment The back and forth methodology between linear logic and type systems has been very successful in the elaboration of L3. It goes in the important general direction of weakening the constraints in light linear logics in order to be able to deal with more algorithms.

Realizability and higher-order term rewriting

Scientific issues, goals, and positioning of the team We have already seen how realizability refines typing and helps in the computational interpretation of proofs. More generally it is a key tool in the study of deep properties of computation, such as termination. Strong progress has been made on the understanding of the relations between reducibility, realizability and orthogonality. It helps in the development of applications to various extensions of the λ-calculus.

Major results Union types are required for typing various extensions of the λ-calculus but are known to interact sometimes badly with preservation under reduction, termination and realizability techniques in general. Through the introduction of a notion of value inspired by reducibility candidates, C. Riba has designed typing rules for union types which are preserved under reduction, ensure termination, ... This approach applies also to implicit existential types, since they are interpreted as infinitary unions in realizability.

The λ-calculus with constructors is designed for the abstract study of pattern matching mechanisms. In order to control the dynamics of this very rich calculus, B. Petit has used a type system based on system F and subtyping whose correction is derived through a realizability model.

Self-assessment The Plume team now has a strong expertise in realizability techniques both on the fundamental aspects, on the relation with logical systems and on applications to term rewriting. This has helped a lot in the recent results. A deeper look at relations with denotational semantics is something to address.

6.3.2 Structures of programming languages

We do not only rely on logic to study programs and in a more general way we are interested in mathematical tools for understanding various structures appearing in computing.
This line of research is related to the previous one (logic and proof theory), while being closer to programming languages, for which we try to build (semantic) models.

Modular presentations of programming languages

Scientific issues, goals, and positioning of the team We work on enhancing the available tools to present and analyze the syntax and operational semantics of calculi that are used to describe various forms of computation and interaction. These calculi include the λ-calculus, the π-calculus, as well as Milner et al.'s more recent generalization of a vast amount of formalisms, known as *bigraphs*.

We rely for this on a category-theoretic approach, the main goal being to improve the modularity of presentation, and to clarify the use of binders in such formalisms.

Major results We have been studying the use of *symmetric monoidal closed* theories to represent binding in (formal models of) programming languages. This provides a modular notion of syntax, which has allowed us to revisit the notion of context in languages with binders. We have thus been able to accommodate the λ-calculus and the π-calculus in our setting, and to propose an extended version of bigraphs.

Self-assessment In order to extend the setting we have introduced and to be able to define the operational semantics of programs, we would like to reformulate our contribution in the framework of double categories, along the lines of previous work by Melliès on models for Multiplicative Linear Logic.

Certification of randomized algorithms

Scientific issues, goals, and positioning of the team This research takes place as part of the ANR project *Scalp* which aims at providing formal models for presenting and certifying cryptographic protocols and algorithms.

More specifically, Ph. Audebaud works with Christine Paulin (LRI, Orsay) on developing a well fitted higher-order functional framework.

Major results We have designed a programing language which meets requirements with respect to expectations in the area of cryptography. It is currently delegated under implementation. For this, the *Why* verification framework has been chosen, as it already provides most of the required technology, and because a reasonable amount of modification should be needed to reach our goal. We work in parallel on enhancing the proof assistants technology (specifically, the *Coq* system) with an efficient and easy to use library in order to help users in developing formal proofs. We take advantage of recent tools such as *Class Types* (from *Haskell*) at least for internal representation of the objects involved.

We currently address the question of semantics with the help of a different language, the calculus $\lambda_\circ$, designed by Thurn, Pfenning and Park. This language is more explicit as far as randomized effects are concerned, which helps at a meta level.

Self-assessment A central aspect of this project is the close integration between its various research directions, that cover the spectrum ranging from abstract denotational models to core programming languages, and to machine mechanization of formal proofs.

Concurrency Theory

Scientific issues, goals, and positioning of the team Process calculi such as CCS or the π-calculus provide a framework where it is possible to represent many aspects of concurrent and distributed programming: message passing, dynamic reconfiguration of communications links, localized (and distant) communications, code mobility, persistent servers, ... In this setting, a program is given by a process. We develop methods for the analysis of the behaviour of processes: guarantee statically some properties at runtime, compare the behaviours of processes, compare different calculi in terms of the behaviours they can express.

Major results We have deepened the understanding of operational equivalences, which give rise to the notion of *behaviour* of processes. Behaviours provide a form of denotation, beyond the mere syntax of terms. We have studied algebraic properties and axiomatisations of behavioural equivalences, in relation with expressiveness of syntactic operators. We have also worked on proof techniques for behavioural equivalences, and studied in particular how one can combine several techniques to obtain more powerful proof tools.

Another research direction is concerned with type systems for termination in concurrent system. We have introduced such static analyses for the first- and higher-order cases (name and process passing, respectively).

Self-assessment The effort towards a deeper and more unified understanding of techniques for behavioural equivalences has led to a satisfying level of maturity. A book chapter about up-to techniques for concurrency is currently being written, and a new research activity on the mechanisation of reasoning about concurrent systems in a theorem prover has been started.

The work on termination should help in developing bridges from the theory of sequential (functional) computation, as expressed in the λ-calculus, and the world of concurrency. It seems particularly relevant for this to be able to adapt powerful λ-calculus techniques to the π-calculus, as we have started to do.

Dynamic Modularity

Scientific issues, goals, and positioning of the team We collaborate with the Sardes project (INRIA Rhône Alpes) on the design of process calculi-based models for *dynamic modularity*. This phrase stands for typical aspects of component-based programming, such as code mobility, dynamic update of modules, and reflexive programming. We want to provide a clean, mathematically defined formalism where forms of dynamic modularity can be expressed and analyzed. This formalism can serve as the basis for the design of new programming primitives.

This work has been supported by the ANR project "MoDyFiable - Modularité Dynamique Fiable" (2005-2008).

Major results We have developed a formalism where it is possible to express some forms of dynamic modularity. This process calculus, called $k\pi$, features a primitive for *passivation*, which provides the ability to freeze a running process in order to manipulate it (send it to a different location, duplicate it, replace it with a newer version, modify some of its code). We have defined an abstract machine for $k\pi$, which demonstrates how passivation can be used in a distributed setting. A prototype implementation of this machine has raised interesting questions related to the design of appropriate type systems in order to guarantee forms of runtime safety for $k\pi$ processes.

Self-assessment The primitive of passivation is very expressive, and goes beyond previously existing forms of interaction in process calculi (notably higher-order communication and forms of code mobility). Several difficult questions related to the metatheory of calculi featuring passivation deserve to be studied. Alternatively, one could also be interested in finding out whether it is possible to define controlled forms of passivation, that renounce to some of its expressiveness in favor of more tractable theoretical properties.

6.3.3 Formalizations in the Coq proof assistant

The formalization activities of the team address the machine representation of different mathematical theories through the development of computer assisted proofs (mainly with the Coq proof assistant).

Game theory

Scientific issues, goals, and positioning of the team Various concept in games theory are derived in a not formal enough way. P. Lescanne and S. Le Roux defined a notion of Feasibility-Desirability games (FD-games) for unifying the view of decision theory, strategic games and evolutionary games. This comes with the study of Nash equilibria for such games.

Major results A model of coalition games based on FD-games has been developed with F. Delaplace with applications in modelling gene regulation activities.

P. Lescanne revisited infinite sequential games using co-induction, showing that previous analyses of those games were inadequate, especially concerning "escalation" which happens to be a Nash equilibrium in the co-inductive formalization.

Self-assessment Formal approaches to game theories allow us to identify under-specified notions appearing in the literature. The introduction of logical notions such as co-induction helps in clarifying the theory by proving that behaviour usually supposed to be irrational can be shown rational in the appropriate framework.

Constructive geometry

Scientific issues, goals, and positioning of the team Starting from strong similarities between constructive proofs and proofs in Euclidean geometry, J. Duprat studies axiomatizations of geometry in the Coq proof assistant. The objective is to provide an interactive framework for describing geometric pictures and then formally proving, in the Coq proof assistant, the properties suggested and induced by the picture.

Major results Based on two primitive notions (orientation and equi-distance) and three constructions (lines, circles and intersection points), J. Duprat defined an axiomatization of constructive points of the plane. Current developments use constructive points of a given line to define distances and constructive points of the unit circle to define angles.

Self-assessment The interaction between graphical tools and a proof assistant is very promising for providing a framework for formal proofs in geometry. The current axiomatization allows one to derive all the Hilbert axioms (except of course the continuity axiom).

7. Reso - Optimized Protocols and Software for High-Performance Networks

Permanent researchers: Begin Thomas (MCF UCB Lyon 1 since 10/2009), Chelius Guillaume (CR2 INRIA since 03/2009), Fleury Eric (Pr ÉNS Lyon), Gelas Jean-Patrick (MCF UCB Lyon 1), Gluck Olivier (MCF UCB Lyon 1), Gonçalves Paulo (CR1 INRIA), Guerin-Lassous Isabelle (Pr UCB Lyon), Lefevre Laurent (CR1 INRIA), Vicat-Blanc Primet Pascale (DR2 INRIA)

Post-docs, secondments and engineers: Cheniour Abderhaman (Temp. Engineer INRIA 7/2008-), Cedeyn Aurélien (Temp. Engineer ÉNS Lyon (50%), 9/2006-08/2009), Dahal Manoj (Post-doc INRIA 11/2008-09/2009), Dias De Assunção Marcos (Post-doc INRIA, 9/2009-), Draminitos Emanouielli (Post-doc INRIA, 11/2008-), Goga Oana (Temp. Engineer INRIA, 10/2008-), Gremillet Olivier (Post-doc INRIA, 12/2008-), Ibrahim Mouhamad (Post-doc, INRIA, 1/2009-), Imbert Matthieu (Research Engineer (40%) INRIA, since 9/2007), Martinez Philippe (Temp. Engineer INRIA, 10/2008-), Mignot Jean-Christophe (Perm. Research Engineer (50%) CNRS since 9/2000), Mornard Olivier (Temp. Engineer INRIA, 3/2007-), Nussbaum Lucas (ATER UCBL Lyon 1, 10/2008-8/2009), Ragon Augustin (Temp. Engineer INRIA, 3/2009-)

Doctoral students: Anhalt Fabienne (INRIA, 10/2008-), Chis Andreea (INRIA, 10/2007-), Guichard Pierre-Solen (INRIA, 10/2008-), Guillier Romaric (INRIA, 10/2006-), Hablot Ludovic (MENRT, 10/2006-), Koslovski Guilherme (INRIA, 10/2008-), Loiseau Patrick (MENRT (AC), 10/2006-), Mon Divakaran Dinil (INRIA, 4/2007-), Orgerie Anne-Cecile (MENRT, 10/2008-), Qinna Wang (MENRT,10/2008-), Sokol Marina (INRIA, 10/09-), Soudan Sébastien (MENRT, 10/2006-), Vanier Rémi (INRIA, 10/2006-)

7.1 Executive summary

Keywords High Speed Networks, Quality of Service, Protocols, Network Services, Traffic Metrology, Energy efficiency, Virtual Infrastructures, Grid, Clouds, Future Internet.

Research area Wavelengths multiplexing and switching techniques on optical fibers allow core network infrastructures to rapidly improve their throughput, reliability and flexibility. These improvements have given the opportunity to create high-performance distributed systems that aggregate storage and computational resources into an integrated computing environment. The use of networks for on-demand computing is now gaining in the large Internet, while the optical transport layer extends to the edge (fiber to the home) enabling ultra-high-performance machine-to-machine or human-to-remote machine communications. One of the key challenges to be addressed for the large deployment of new challenging applications and services in the Internet is the dynamic provisionning of a secure, flexible, transparent and high-performance transport infrastructure. This future pervasive computing environment will rely not only on very high speed wired network but also on wireless systems. The Internet re-design raises the opportunity to better understand and assess higher-level system requirements, and use these as drivers of the lower layer architecture. In this process, mechanisms that are implemented today as part of applications, may conceivably migrate into the network itself to bring it more flexibility, intelligence and autonomy.

Main goals The goal of our research is to provide analysis of the limitations of the current communication paradigms, network software and protocols designed for standard networks and traditional usages, and to propose innovative approaches, optimizations and associated control mechanisms. RESO aims at providing software solutions but also original processes for high-performance and flexible communications on very-high-speed networking infrastructures and for an efficient exploitation of these infrastructures. These solutions must scale in increasing bandwidths, heterogeneity and number of flows and usages. During this period, it appeared crucial in the actual context that efforts should also be spent on the investigation and understanding of wireless communications and related network complexity, dynamics, and scalability.

Methodology RESO's approach relies on a tight combination of theoretical and practical work. RESO gathers expertise in the design and implementation of advanced high-performance cluster networks protocols, long-distance networks and Internet protocols architecture, distributed systems and algorithms but also scheduling theory, optimization, queuing theory, and statistical analysis. Our research framework at the interface of a specific network context (very high speed, optical technology) and a challenging application domain (Grids, Clouds) induces a close interaction with both the underlying network level and the application level. Our methodology is based on a deep evaluation of the functionalities and performance of high-speed infrastructures and on a study of the high-end and original requirements before designing and analysing new solutions. The proposals are implemented and experimented on real or emulated local or wide area testbeds, with real conditions and large-scale applications, then transfered to the industry.

Highlights During this period, the RESO team had internationally visible contributions in the following fields:

Dynamic bandwidth provisioning in Optical Networks with the work on *SRV - network-resource scheduling, virtualization and reconfiguration component in a service-oriented approach-* and participation to OGF NSI WG. (Open Grid Forum Network Service Interface working group).

Virtual private infrastructures orchestration and definition with the work on *HIPerNET and VXDL (Virtual private execution infrastructure description language)*.

Sampling methods for characterizing *heavy-tailed distributions in high-speed network traffic* and *experimental validation of the Taqqu theorem*.

Participation to the design, the development and visibility of the local and national *GRID5000 instrument* and to its international optical interconnections to Netherland (DAS3) and Japan (Naregi) in collaboration with RENATER;

Design and development of a metrology infrastructure *Metroflux* for fine-grained traffic monitoring in Grid5000 and traffic-analysis system dedicated to 10Gb/s speed links.

the strong involvement in the creation of the INRIA-Bell Labs common laboratory The motivation of our research work in the common lab is to build and to exploit the knowledge that comes along with traffic. The goal is to act in a better way and to make better decisions at the router and network levels. The main topics we are exploring in the research topic of the common laboratory ar: a) traffic identification and classification, b) traffic sampling, c) flow analysis, d), flow scheduling, e) sampling-based scheduling, e) flow-based routing.

a startup project, winner of the OSEO emergence prize, which will commercialize our solution for *virtual infrastructures orchestration*, to make this vision a reality for the Future Internet.

7.2 Research activities

7.2.1 Optimized Protocol implementations and networking equipements

Scientific issues, goals, and positioning of the team: In this research topic we focus on the implementation and optimization of the mechanisms and processes within networking devices to mainly address the following issues: flexibility, reliability and energy comsumption related to the communications in high demanding context.

Major results Oct. 2006-Oct. 2009: We proposed a new network model combining end-host, server and router virtualization to offer isolated and malleable virtual networks of different types, owners and protocols, all sharing one physical infrastructure. To have a better insight in this potential and in its limits we conducted systematic analysis and experiments of the cost of end-host, server and router virtualization.

We designed an autonomic network equipment taking into account the specific requirements of active equipments in terms of dynamic service deployment, auto-settings, self-configuration, monitoring but also in terms of hardware specification (limited resources, limited mechanical parts constraints, dimension constraints), reliability and fault tolerance.

In order to improve autonomic service deployment in large scale networks, we designed and proposed the ANPI framework (Autonomic Network Programming Interface) currently used in the Autonomic Internet European Project.

In collaboration with France Telecom RD we propose session-aware distributed network solutions integrated in a cluster-based network server which support the reliability mandatory to operators services (VOIP).

Since 2007, we are exploring the challenges associated with energy usage of large scale distributed infrastructures (Grids, clouds, Future Internet). We are developing solutions to dynamically monitor and optimize energy usage in software frameworks and validate the solutions on the Grid5000 platform.

7.2.2 Quality of Service and Transport layer for Future Networks

Scientific issues, goals, and positioning of the team: The progression of traffic over Internet links is a remarkable trend of the last years. The bandwidth increases at the end-user level, and ISPs (for home use or enterprise) has radically changed the expectations of Internet users. High bandwidth requirements, like those generated by some peer-to-peer applications, or real-time interactive communications, like voice and video over IP, will play a significant part of the next generation Internet. To offer a better quality of experience in future networks, we claim that some network-resource control has to be associated with the end-to-end flow-control approach. We proposed to investigate the temporal dimension of the Internet and to address the control timescale (control plane) in the context of Future Internet not only for performance, but also for flexibility, manageability and security purposes.

Major results Oct. 2006-Oct. 2009: During this period we activelly contributed to the collective methodological effort towards a benchmark design for high speed transport protocols comparison. We proposed scenarii, experiment deployment tools (NXE) and testbed for these studies (Grid'5000 and Metroflux).

Alternative bandwidth sharing approach have also been investigated. These solutions are intended either for wireless local networks or for grids and clouds. Flow scheduling based on the in-advance knowledge of the resource requirement of an application or online estimation of these requirements has been studied. Signaling or real-time flow analysis as well as scalability issues have been explored. Distributed and lagrangian relaxation-based solution for *bandwidth sharing* is also investigated. This approach addresses well the dynamic feature, due to node mobility or traffic variation. We study the bandwidth reservation problem for bulk data transfers in grid networks. We model grid networks as a set of distributed sites interconnected by network with potential bottlenecks, and transfer requests arrive online with specified volumes and deadlines. The model is based on network-resources reservations to support the delivery of deterministic performance to end users. We also worked on control planes to allow the automatic configuration of network equipments, in order to add or remove paths between sites or nodes and propose a model for multilayer networks with advance reservation of paths. We explored sender-side rate limitation problem.

We studied the interactions of components required to accomplish the tasks of bandwidth reservation, path computation and network signaling and propose to extend the control-plane and introduce dynamicity in network resource reservation in a network service plane.

We have proposed a new architecture for Explicit Rate Notification protocols which improves the responsiveness and the robustness of ERN flows in networks with Variable Bandwidth Environment and feedback losses. The set of mechanisms and proposed solutions have been implemented and validated on the eXplicit Control Protocol (XCP : XCP-f, XCP-r, XCP-i).

7.2.3 High Speed Network's traffic metrology and statistical analysis

Scientific issues, goals, and positioning of the team: Metrology (i.e. the deployment of a series of tools allowing the collecting of relevant information regarding the system's status) of wide-area computer networks, is a recently-introduced discipline in the context of networks, that undergoes constant developments. In a nutshell, this activity consists in measuring along time, the nature and the amount of exchanged information between the constituents of a system. It is then a matter of using the collected data to forecast the network load's evolution, so as to anticipate congestion, and more generally, to guarantee a certain Quality of Service.

From a statistical signal-processing viewpoint, collected traces correspond to (multivariate) time series principally characterized by non-properties: non-gaussianity, non-stationarity, non-linearities, absence of a characteristic time scale (scale invariance). Our research activity undertakes the development of reliable signal-analysis tools aimed at identifying these (non-)properties in the specific context of computer network traffic. Our goal is for these analyses to become in the near future a plain component not only in the study and in the development of infrastructures and computing networks, but also in real-time resources identification and management.

Major results Oct. 2006-Oct. 2009:

From 2006 to 2008, we considerably invested in the development of our metrology plateform *MetroFlux*. Our main concern was to guarantee sufficient reliability to support our future experimental research, and the great investment that has been granted to Grid5000 has been profitably used providing us with a high-performance and quite novel experimental setup to confront the proposed theoretical models with real traffic measurements. The first theorem we empirically validated with our plate-form was the relationship that bonds the tail index of a heavy-tailed flow size distribution, to the long range dependence property of the corresponding aggregated traffic. Up to our study, this theorem originally derived by Taqqu et al., had never really been corroborated on a real, large-scale network facility, under very general and fully controlled conditions.

In parallel, we concluded a theoretical work on the maximum likelihood estimation of the heavy-tail index of flow size distributions, from packet sampled traffic traces. The difficulty dwelled in the sampling process itself that yields a doubly censured data with respect to the flow size and to the flow population.
We are now studying Markov models for TCP traffic, and more specifically we revisited the multifractal analysis of these particular time series. We were led to propose a new large deviation principle adapted to the additive increase, multiplicative decrease behavior of instantaneous TCP throughputs. This is an ongoing work, in joint collaboration with J. Barral from the Sisyphe INRIA team-project.

7.2.4 Network Services for high-demanding applications

Scientific issues, goals, and positioning of the team: The purpose of Computational Grids was initially to aggregate a large collection of shared resources (computing, communication, storage, information) to build an efficient and very-high-performance computing environment for data-intensive or computing-intensive applications.But generally, the underlying communication infrastructure of these large-scale distributed environments is a complex interconnection of multi-IP domains with non-controlled performance characteristics. Consequently *the Grid Network cloud* exhibits extreme heterogeneity in performance and reliability that considerably affects the global application performance.

Major results Oct. 2006-Oct. 2009: Thanks to systematic experiments on the behavior of MPI in large-scale environments, we merge optimizations of current implementations and propose new optimizations in the communication layers in order to execute more efficiently MPI applications on the Grid. We also study the impact of using the TCP protocol for inter-site communications in the grid and its interactions with MPI applications [475]. We have proposed a new transparent layer called MPI5000 and placed between MPI and TCP allowing application composed of several tasks to be correctly distributed on available node regarding the grid topology and the application scheme.

In the CARRIOCAS project of the pôle Ile de France System@tic we explored optical resource provisionning and optimal bandwidth sharing services in the context of a 40Gb/s, dynamically-provisionable network. We developed the SRV software for dynamic network service scheduling and reconfiguration.

In the Grid 5000 testbed, responsible for the networking aspects we have adapted the Metroflux environment to deliver a network measurement service.

Through the European EC-GIN project we worked on the Bulk Data Transfer Scheduling approach. We deployed it in Grid5000 and extend it to the Internet.

Through the ANR IGTMD project, we collaborated with the LCG and real physicists. A dedicated link deployed between IN2P3 (one of the largest computing center in France) and the FermiLab laboratory in Chicago, enables us to perform transport protocol experiments as well as traffic capture.

Through the ANR HIPCAL project, RESO collaborates with biology and medical imaging applications (I3S, IBCP) to design and develop the HIPerNet framework. We propose to combine network and system virtualization with cryptographic identification and SPKI/HIP principles to help the user communities to build and share their own resource reservoirs. The HIPerNet framework enabling the creation and the management of customized confined execution environments in a large scale context implements a agile virtual network solution and a distributed security approach.

7.2.5 Wireless Communications

Scientific issues, goals, and positionning of the team:

This future pervasive computing environment will rely not only on very high speed wired network but also on two kind of wireless systems. The former will connect the world everywhere through broadband access links thanks to future 4G systems and mesh networks. The later, namely sensor networks, will make all objects able to communicate and to interact with their environment.

It appears crucial in the actual context that large efforts should be spent on the investigation and understanding of network complexity, dynamics, and scalability. Such effort should be done on the understanding of actual wireless protocols (from MAC to transport), on the behavior of such protocols in wireless networks, and hybrid architecture combining both wireless meshes and heterogeneous wire backbone. From a theoretical point of view, there is an important need of generic and powerful models in order to characterize the existing networks, in order to better understand the laws and behaviors of large scale networks.

Major results Sept 2007 -Oct 2009: We adopt a "*pragmatic*" standpoint. Whenever it's possible we want to propose a mathematical representation of the dynamics of the protocols, from which one could predict and optimize the resulting behavior and bounds on the scalability:

- **Modeling of wireless networks**. We have been deeply involved in the research and design of distributed clustering algorithm and virtual backbone in wireless networks. We used Stochastic geometry and the theory of point processes in this context to show bounds on the behavior of our proposed solutions.
- **Performance evaluation tools for wireless networks**. We have also use simulator tools and more precisely we promote the WSNet simulator developped internaly. WSNet is dedicated to wireless multi hop network s and integrates very accurate radio models (propagation, modulation).
- **Activity Scheduling and QoS**. Our research aims to save energy for general communication by scheduling nodes' duty cycles while preserving communication connectivity and bounding packet latency. Our goal is thus to ensure robustness in terms of reliability and fault-tolerance for the communications of spontaneous networks but also to ensure temporal constraints in a spontaneous network.
- **Data Aggregation & Processing**. Data aggregation has been put forward as an essential paradigm for wireless routing in sensor networks. We have study approaches designed by the P2P community by taking into account the broadcasting characteristic of the medium.
- **Sensor Network Testbed**
  We have proposed as Prime Leader SensLAB, a very ambitious project with several partners. The purpose of the SensLAB project is to deploy a very large scale open wireless sensor network platform. SensLAB's main and most important goal is to offer **an accurate and efficient scientific tool** to help in the design, development, tuning, and experimentation of real large-scale sensor network applications.

8. IT Support Team MI-LIP

8.1 Team Composition

Permanent engineers: Mignot Jean-Christophe (IR CNRS, 50%), Ponsard Dominique (IE CNRS), Torres Serge (IR ENS, 60%).

8.2 Evolution of the team

The IT support team of the laboratory — MI-LIP — was created in september 2001 with an investigator and two engineers. It's activity and membership evolved along time and it was eventually acknowledged as a CATI (Centre Automatisé du Traitement de l'Information - Automated Data Management Center) by the CNRS in 2007.
The main factor driving the team evolution is the departure and arrival of the fixed term engineers. Stéphane d'Alu has left after being hired as a permanent research engineer in the CITI laboratory at INSA Lyon. Aurélien Cedeyn has also left the LIP on July 31st 2009 and joined with the Commissariat à l'Énergie Atomique as a permanent engineer. This cyclically disrupts the team's operation.

8.3 Executive summary

Keywords
IT infrastructure management, scientific research support, scientific instruments services.

Research area
The MI-LIP does not have *per se* a research area: its activity comes as a support of scientific investigations on behalf of the laboratory research teams and, partly, of external scientists.

Main goals

- build and manage an IT environment for the laboratory activity;
- provide support and expertise to research teams for their projects;
- build and manage local nodes of distributed experimental platforms.

Methodology
The activity of the team follows the "rules of the trade" in IT support and administration. It progressively adopts, as they appear, more formal standards (e.g. ISO 2700x on security). Team's knowledge about these methods is maintained through regular training.

Highlights

- a sound and confortable IT infrastructure for the laboratory users;
- several specific and customized services to suit particular research needs;
- building and administration of local elements of large scale platforms.

The MI-LIP team does not have an independent research activity. It operates in support of the research team and activities of the laboratory and offers some services to external users.

8.4 Software production and Research Infrastructure

8.4.1 Software Descriptions

Though, as it is usual for any substantial IT administration activity, lots of code have been produced. However no pieces have been formally disseminated as software packages.

8.4.2 Contribution to Research Infrastructures

The main contributions of the MI-LIP team is its commitment to build, develop, and maintain IT research infrastructures be them of local interest or opened to a larger community.
The operational environment is widely open since all the elements have to, and actually do, interoperate with other partner services and systems: the ENS IT infrastructure (for all the elements), other laboratory affiliations (CNRS, INRIA, UCBL Lyon 1, for management information systems and electronic documentation assets) and the other distant nodes (for distributed research infrastructures such as GRID'5000 or the Décrypthon project).
Openness is also naturally emphasized by the multi-localization of the laboratory premises that span several buildings on the ENS campus, rented offices outside the campus, where a LIP team is hosted by the IXXI, offices in the Gerland branch of UCBL Lyon 1. This context is very challenging as the MI-LIP team is committed to offer the same level of service on all the laboratory locations.
Coupling between teaching and research is particularly tight in the ENS de Lyon and the relationship between LIP and the Computer Science Department must be considered along these lines. As a consequence, some infrastructure elements, managed by the MI-LIP, are used for student training and scientific investigations as well.
As a whole, the MI-LIP team has been able to follow the growth of the laboratory membership and activities in spite of a permanent shortage of IT manpower. Nevertheless the come and go of fixed term engineers shakes the team's operation that, associated with the laboratory expansion, reaches a breaking point.

Local infrastructure

It provides basic and more advanced or specialized (according to the laboratory peculiarities) services.

Basic services include:

- workstation administration (automatic installation and upgrade : Linux, for research, and Windows, for staff; over 50 workstations);
- file and backup services (over 12 Tbytes);
- printing services (spanning over all the locations of the laboratory);
- Windows terminal services (the bulk of operating systems consists in a mix of "Unices" and Mac OS/X, and scientists sometimes need access to the Microsoft Office suite);
- user support (through information services, such as the laboratory intranet, and direct interaction with users);
- LIP Web site.

The main evolution during the 2006-2009 period have been

- a net growth (storage and printing capacity, number of workstations, locations, following the overall growth of the laboratory);
- technological upgrades (Wikis and plain HTML services replaced by modern CMS tools, virtualization of servers);
- organizational consolidation (documentation, disaster recovery plans. . .);
- infrastructure consolidation by realizing the first of a two-phase cooling equipment upgrade of the laboratory computer room.

Advanced and/or specialized local services These fall under six categories:

Computing services. These are provided in two flavors:

- traditional clusters (two 24 nodes clusters, one being also used for student training) and computing servers (84 cores in 9 servers with up to 80 Gbytes RAM on one of them);
- batch jobs execution infrastructure encompassing all the computing resources of the laboratory (including user workstations) and allowing the recruitment of over 220 cores for special purposes (for instance, the quest, by Vincent Lefevre, of the worst cases for the correct rounding of elementary functions uses it extensively).

Special purpose server: servers that host specific tools such as VLSI and FPGA conception software (Xilinx, ModelSim. . .).

Videoconferencing: a permanently set up videoconferencing room, fitted with professional grade equipment (thanks to INRIA), is available for the laboratory and is also extensively used by other interested parties in the ENS de Lyon.

Multiple-architecture access: while the vast majority of equipment is based on X86 architecture (now X86-64) for some specific research needs others architectures are available, notably PowerPC and IA64.

Collaborative Development Environment: a GForge server is available for collaborative work not only between laboratory members but also with outside co-workers (more than 100 projects, 200 users).

Technical expertise: beyond basic user support, the team engineers are often called to study and give advice on the technology used in research projects. This covers hardware, development tools, elements integration as well as compatibility with the ENS network.

The main evolution during the 2006-2009 period have been:

- the extension of the available computing power, mainly by the deployment of the batch infrastructure and by the addition, as well, of some specific computing servers (e.g. 16 cores/80 Gbytes of RAM server);
- the upgrade and consolidation of the other services
- upgrade, always with the support of INRIA, of videoconferencing equipment;

Large Scale Research Instruments

The team is engaged in two national initiatives of unequal importance (as far as MI-LIP is concerned):

- GRID'5000 (http://www.grid5000.fr/);
- The Décrypthon project (http://www.decrypthon.fr/).

GRID'5000 The team manages the Lyon node of the instrument. It's computing power relies on 260 processors. While this is a rather modest size, even at the scale of GRID'5000 , it's operation represents a large part of the team activity (one full time engineer). Time has been dedicated to five main realms:

- day-to-day management;
- dealing with hardware problems;
- deploying research equipments and resources (i.e. energy sensors, GNet network control boxes. . .);
- going with network evolutions of GRID'5000;
- improving monitoring services;
- organizing the technical team at national level.

A lot of time has been spent (lost?) to fix technical problems that where irrelevant to research issues. Most of them where related to data center cooling malfunctions that have forced us to a time and money costly moving from one computer room to another inside the ENS. In spite of this shift, cooling problems have continued to plague the operation of our GRID'5000 equipment. Other problems arose from the inability of hardware providers to comply with their commitments for the remote operation of the servers. Last but not least, it took a lengthy and painfull debugging processes to deal with issues in the Lyon metropolitan network when we have switched from 1 Gbits to 10 GBits links. Nevertheless, the Lyon node has been able to keep pace with the whole of GRID'5000 and busily take it's share of the experimental burden.

Many improvements have taken place in the management and services area. First, a standardization of the organization has been realized, coupled with a better integration of services. This allows engineers from different locations to safely and usefully intervene on a distant node should the in-charge administrator be absent (this can be a long time when a fixed term engineer leaves!). Second, extensive monitoring of the platform has been deployed to allow proactive action to be taken instead of mere reaction to failures. Third, following a nowadays common trend, all the new and most of the old services have been virtualized. In the networking realm, the main change was the switch to 10 Gbits. If this upgrade was rather transparent in the core of the network, up to the local RENATER NOC, things were more complicated downstream, in the metropolitan network (LYRES). Eventually, after a lot of busy work and debugging, the Lyon joined the pack of 10 Gbits GRID'5000 nodes.

Décrypthon The MI-LIP team implication is only a small part of the LIP activity in the Décrypthon project. Software developments and research have taken place in the GRAAL team. Yet the basic management of the Lyon node and provision of good operational condition have been supplied by the MI-LIP.

8.5 Team organization, infrastructure management, and adequacy to laboratory needs

Above realizations were made possible thanks to team and laboratory efforts in several directions.

Team organization: The process of acknowledgement of the team as a CATI, has helped to formalize and strengthen it is organization (meeting planning, work plan, cross training and information sharing). A training policy has been defined to procure the new skills and keep the level of proficiency needed to carry on the team's missions.

Infrastructure management: The team has internally developed some tools needed to manage the plateform it operates (e.g. automatic workstation and server installation, backup tools to supplement the ENS wide infrastructure). Rules and multiannual planning have been set up to standardize and optimize equipment expenditures so has to relieve the burden of an anarchic proliferation of equipment configurations and maximize the efficiency/cost ratio. The maintenance and upgrade of basic facilities such as computer room have been undertaken. Expenditure and operational disturbance are planned over several years.

Adequacy to laboratory needs Laboratory needs have been more rigorously taken into account. Some elements of the platform management have been discussed by laboratory council. Efforts have been concentrated on coping with the mostly wanted services for which proximity is essential. Other services are delegated to external providers (e.g. the ENS IT services or high performance computing centers). A reflexion is undertaken to refocus the laboratory platform policy and find a better articulation between the activity in the research teams and the MI-LIP.

8.6 Consulting Activities

As mentioned earlier, the team members have been internally consulted for technology selection for research. They have also participated, as experts, in several workgroups inside the ENS for government contracts or for infrastructure projects.

9. Production list

We give here our 2005-2009 list of publications, softwares and patents. Productions noted as "external" correspond to productions by LIP members before their arrival in the department, and are not counted in the table on page 5.

International and national peer reviewed journals [ACL]

2005

[1] Alessandro Bassi, Micah Beck, Fabien Chanussot, Jean-Patrick Gelas, Robert Harakaly, Laurent Lefèvre, Terry Moore, James Plank, and Pascale Vicat-Blanc Primet. Active and logistical networking for grid computing: the e-toile architecture. *The International Journal of Future Generation Computer Systems (FGCS) - Grid Computing: Theory, Methods and Applications*, 21(1):199–208, January 2005. Elsevier B.V (ed),ISSN 0167-739X.

[2] Cédric Bastoul and Paul Feautrier. Adjusting a program transformation for legality. *Parallel Processing Letters*, 15(1-2):3–17, March-June 2005.

[3] V. Blondel, E. Jeandel, P. Koiran, and N. Portier. Decidable and undecidable problems about quantum automata. *SIAM Journal on Computing*, 34(6):1464–1473, 2005.

[4] Faycal Bouhafs, Jean-Patrick Gelas, Laurent Lefèvre, Moufida Maimour, Cong-Duc Pham, Pascale Vicat-Blanc Primet, and Bernard Tourancheau. Designing and evaluating an active grid architecture. *The International Journal of Future Generation Computer Systems (FGCS) - Grid Computing: Theory, Methods and Applications*, 21(2):315–330, February 2005.

[5] Nicolas Boullis and Arnaud Tisserand. Some optimizations of hardware multiplication by constant matrices. *IEEE Transactions on Computers*, 54(10):1271–1282, October 2005.

[6] Nicolas Brisebarre, David Defour, Peter Kornerup, Jean-Michel Muller, and Nathalie Revol. A new range reduction algorithm. *IEEE Transactions on Computers*, 54(3):331–339, 2005.

[7] Nicolas Brisebarre and Georges Philibert. Effective lower and upper bounds for the Fourier coefficients of powers of the modular invariant j. *J. Ramanujan Math. Soc.*, 20(4):255–282, 2005.

[8] Franck Cappello, Samir Djilali, Gilles Fedak, Thomas Herault, Frédéric Magniette, Vincent Néri, and Oleg Lodygensky. Computing on Large Scale Distributed Systems: XtremWeb Architecture, Programming Models, Security, Tests and Convergence with Grid. *Future Generation Computer Systems*, 21(3):417–437, mar 2005. External publication.

[9] Arnaud Carayol, Daniel Hirschkoff, and Davide Sangiorgi. On the representation of McCarthy's amb in the pi-calculus. *Theor. Comput. Sci.*, 330(3):439–473, 2005.

[10] F. Chavanon, M. Latapy, M. Morvan, E. Rémila, and L. Vuillon. Configurations induced by discrete rotations: periodicity and quasiperiodicity properties. *Theoretical Computer Science*, 346:226–253, 2005.

[11] Alain Darte and Guillaume Huard. New complexity results on array contraction and related problems. *Journal of VLSI Signal Processing-Systems for Signal, Image, and Video Technology*, 40(1):35–55, May 2005.

[12] Alain Darte, Robert Schreiber, and Gilles Villard. Lattice-based memory allocation. *IEEE Transactions on Computers*, 54(10):1242–1257, October 2005. Special Issue, Tribute to B. Ramakrishna (Bob) Rau.

[13] Florent de Dinechin and Arnaud Tisserand. Multipartite table methods. *IEEE Transactions on Computers*, 54(3):319–330, 2005.

[14] H. Derksen, E. Jeandel, and P. Koiran. Quantum automata and algebraic groups. *Journal of Symbolic Computation*, 39 (3-4):357–371, 2005. special issue on the occasion of MEGA 2003.

[15] Jérémie Detrey and Florent de Dinechin. Outils pour une comparaison sans a priori entre arithmétique logarithmique et arithmétique flottante. *Technique et science informatiques*, 24(6):625–643, 2005.

[16] N. Fatès and M. Morvan. An experimental study of robustness to asynchronism for elementary cellular automata. *Complex Systems*, 16(1):1–27, 2005.

[17] Paulo Gonçalves and Rudolf H. Riedi. Diverging moments and parameter estimation. *Journal of American Statistical Association*, 100(472):1382–1393, December 2005.

[18] Christophe Guillon, Fabrice Rastello, Thierry Bidault, and Florent Bouchez. Procedure placement using temporal-ordering information: Dealing with code size expansion. *Journal of Embedded Computing*, 1(4):437–459, December 2005.

[19] Tom Hirschowitz and Xavier Leroy. Mixin modules in a call-by-value setting. *ACM Trans. Program. Lang. Syst.*, 27(5):857–881, 2005.

[20] Claude-Pierre Jeannerod and Gilles Villard. Essentially optimal computation of the inverse of generic polynomial matrices. *Journal of Complexity*, 21(1):72–86, 2005.

[21] E. Kaltofen and G. Villard. On the complexity of computing determinants. *Computational Complexity*, 13:91–130, 2005.

[22] Olivier Laurent. Classical isomorphisms of types. *Mathematical Structures in Computer Science*, 15(5):969–1004, October 2005. External publication.

[23] E. Lebhar and N. Schabanel. Close to optimal decentralized routing in long-range contact networks. *Theoretical Computer Science*, 348(2-3):294–310, 2005. special issue on ICALP 2004.

[24] Étienne Lozes. Elimination of spatial connectives in static spatial logics. *Theor. Comput. Sci.*, 330(3):475–499, 2005.

[25] Jean-Philippe Martin-Flatin and Pascale Vicat-Blanc Primet. Editorial of the special issue high performance networking and services in grids: the datatag project. *International Journal of Future Generation Computer System, FGCS*, 21(Issue 4):439–623, April 2005.

[26] Jean-Philippe Martin-Flatin and Pascale Vicat-Blanc Primet. Special issue high performance networking and services in grids: the datatag project. *International Journal of Future Generation Computer System, FGCS*, 21(Issue 4):439–442 (, April 2005.

[27] B. Nouvel and E. Rémila. Configurations induced by discrete rotations: periodicity and quasiperiodicity properties. *Discrete Applied Mathematics*, 147:325–343, 2005.

[28] Jose-Alejandro Pineiro, Stuart F. Oberman, Jean-Michel Muller, and Javier D. Bruguera. High-speed function approximation using a minimax quadratic interpolator. *IEEE Transactions on Computers*, 54(3):304–318, March 2005.

[29] E. Rémila. Tiling a polygon with two kinds of rectangles. *Discrete and Computational Geometry*, 34:313–330, 2005.

[30] N. Revol, K. Makino, and M. Berz. Taylor models and floating-point arithmetic: proof that arithmetic operations are validated in COSY. *Journal of Logic and Algebraic Programming*, 64:135–154, 2005.

[31] N. Revol and F. Rouillier. Motivations for an arbitrary precision interval arithmetic and the MPFI library. *Reliable Computing*, 11(4):275–290, 2005.

[32] Pascale Vicat-Blanc Primet, F. Echantillac, and M. Goutelle. Experiments of the equivalent differentiated service model in grids. *in International Journal Future Generation Computer Systems FGCS, special issue on High Performance Networking and Services in Grids*, 21(Issue 4):512–524 (, April 2005.

2006

[33] Patrick R. Amestoy, Abdou Guermouche, Jean-Yves L'Excellent, and Stephane Pralet. Hybrid scheduling for the parallel solution of linear systems. *Parallel Computing*, 32(2):136–156, 2006.

[34] Anestis Antoniadis, Andrey Feuerverger, and Paulo Gonçalves. Wavelet based estimation for univariate stable laws. *Annals of the Institute of Mathematical Statistics*, 58(4):779–807, 2006.

[35] Olivier Beaumont, Loris Marchal, and Yves Robert. Complexity results for collective communications on heterogeneous platforms. *Int. Journal of High Performance Computing Applications*, 20(1):5–17, 2006.

[36] Bernhard Beckermann, George Labahn, and Gilles Villard. Normal forms for general polynomial matrices. *Journal of Symbolic Computation*, 41(6):708–737, 2006.

[37] Anne Benoit, Brigitte Plateau, and William J. Stewart. Memory Efficient Kronecker algorithms with applications to the modelling of parallel systems. *Future Generation Computer Systems, special issue on System Performance Analysis and Evaluation*, 22(7):838–847, August 2006.

[38] Raphaël Bolze, Franck Cappello, Eddy Caron, Michel Daydé, Frederic Desprez, Emmanuel Jeannot, Yvon Jégou, Stéphane Lanteri, Julien Leduc, Noredine Melab, Guillaume Mornet, Raymond Namyst, Pascale Primet, Benjamin Quetier, Olivier Richard, El-Ghazali Talbi, and Touché Irena. Grid'5000: a large scale and highly reconfigurable experimental grid testbed. *International Journal of High Performance Computing Applications*, 20(4):481–494, November 2006.

[39] Nicolas Brisebarre, Jean-Michel Muller, and Arnaud Tisserand. Computing machine-efficient polynomial approximations. *ACM Transactions on Mathematical Software*, 32(2):236–256, June 2006.

[40] Nicolas Brisebarre, Jean-Michel Muller, Arnaud Tisserand, and Serge Torres. Hardware operators for function evaluation using sparse-coefficient polynomials. *IEE Electronics Letters*, 42(25):1441–1442, December 2006.

[41] Hervé Brönnimann, Guillaume Melquiond, and Sylvain Pion. The design of the Boost interval arithmetic library. *Theoretical Computer Science*, 351:111–118, 2006.

[42] Yves Caniou and Emmanuel Jeannot. Multi-criteria scheduling heuristics for GridRPC systems. *Special edition of The International Journal of High Performance Computing Applications (IJHPCA)*, 20(1):61–76, 2006.

[43] Eddy Caron and Frédéric Desprez. Diet: A scalable toolbox to build network enabled servers on the grid. *International Journal of High Performance Computing Applications*, 20(3):335–352, 2006.

[44] F. Chavanon and E. Rémila. Rhombus tilings: decomposition and space structure. *Discrete and Computational Geometry*, 35:329–358, 2006.

[45] Pushpinder Kaur Chouhan, Holly Dail, Eddy Caron, and Frédéric Vivien. Automatic middleware deployment planning on clusters. *International Journal of High Performance Computing Applications*, 20(4):517–530, November 2006.

[46] Holly Dail and Frédéric Desprez. Experiences with hierarchical request flow management for network-enabled server environments. *International Journal of High Performance Computing Applications*, 20(1):143–157, February 2006.

[47] Florent de Dinechin and Gilles Villard. High precision numerical accuracy in physics research. *Nuclear Inst. and Methods in Physics Research, A*, 559(1):207–210, 2006.

[48] Frédéric Desprez and Antoine Vernois. Simultaneous scheduling of replication and computation for data-intensive applications on the grid. *Journal of Grid Computing*, 4(1):19–31, March 2006.

[49] S. Desreux and E. Rémila. An optimal algorithm to generate tilings. *Journal of Discrete Algorithms*, 4:168–180, 2006.

[50] Daniel J. Dougherty, Pierre Lescanne, and Luigi Liquori. Addressed term rewriting systems: application to a typed object calculus. *Mathematical Structures in Computer Science*, 16(4):667–709, 2006.

[51] P. Duchon, N. Hanusse, E. Lebhar, and N. Schabanel. Could any graph be turned into a small world? *Theoretical Computer Science*, 355(1):96–103, 2006. Special issue on complex networks.

[52] N. Fatès, M. Morvan, N. Schabanel, and E. Thierry. Fully asynchronous behavior of double-quiescent elementary cellular automata. *Theoretical Computer Science*, 362(1-3):1–16, 2006.

[53] Paul Feautrier. Scalable and structured scheduling. *International Journal of Parallel Programming*, 34(5):459–487, October 2006.

[54] N. Garnier, A. Friedrich, R. Bolze, E. Bettler, L. Moulinier, C. Geourjon, J. D. Thompson, G. Deleage, and O. Poch. MAGOS: multiple alignment and modelling server. *Bioinformatics*, 22(17):2164–2165, 2006.

[55] A. Gebremedhin, J. Gustedt, I. Guérin Lassous, M. Essaïdi, and J. A. Telle. Pro: A model for the design and analysis of efficient and scalable parallel algorithms. *Nordic Journal of Computing*, 13(4), 2006.

[56] Arnaud Giersch, Yves Robert, and Frédéric Vivien. Scheduling tasks sharing files on heterogeneous master-slave platforms. *Journal of Systems Architecture*, 52(2):88–104, 2006.

[57] Abdou Guermouche and Jean-Yves L'Excellent. Constructing memory-minimizing schedules for multifrontal methods. *ACM Transactions on Mathematical Software*, 32(1):17–32, 2006.

[58] Daniel Hirschkoff, Étienne Lozes, and Davide Sangiorgi. On the expressiveness of the ambient logic. *Logical Methods in Computer Science*, 2(2), 2006.

[59] Claude-Pierre Jeannerod and Gilles Villard. Asymptotically fast polynomial matrix algorithms for multivariable systems. *International Journal of Control*, 79(11):1359–1367, 2006.

[60] Peter Kornerup and Jean-Michel Muller. Choosing starting values for certain Newton-Raphson iterations. *Theoretical Computer Science*, 351(1):101–110, February 2006.

[61] Peter Kornerup and Jean-Michel Muller. Leading guard digits in finite-precision redundant representations. *IEEE Transactions on Computers*, 55(5):541–548, May 2006.

[62] Pierre Lescanne. Mechanizing common knowledge logic using coq. *Ann. Math. Artif. Intell.*, 48(1-2):15–43, 2006.

[63] Loris Marchal, Yang Yang, Henri Casanova, and Yves Robert. Steady-state scheduling of multiple divisible load applications on wide-area distributed computing platforms. *Int. Journal of High Performance Computing Applications*, 20(3):365–381, 2006.

[64] Hélène Renard, Yves Robert, and Frédéric Vivien. Data redistribution algorithms for heterogeneous processor rings. *Int. Journal of High Performance Computing Applications*, 20(1):31–43, 2006.

[65] S. Sarr, C. Chaudet, G. Chelius, and I. Guérin Lassous. A node-based available bandwidth evaluation in ieee 802.11 ad hoc networks. *International Journal of Parallel, Emergent and Distributed Systems*, 21(6), 2006.

2007

[66] Gabriel Antoniu, Hinde Bouziane, Mathieu Jan, Christian Pérez, and Thierry Priol. Combining data sharing with the master-worker paradigm in the common component architecture. *Cluster Computing*, 10(3):265 – 276, 2007. External publication.

[67] V. Atassi, P. Baillot, and K. Terui. Verification of Ptime reducibility for system F terms: type inference in Dual Light Affine Logic. *Logical Methods in Computer Science*, 3(4:10):1–32, 2007. External publication.

[68] Michele Benzi and Bora Uçar. Block triangular preconditioners for M-matrices and Markov chains. *Electronic Transactions on Numerical Analysis*, 26:209–227, 2007. External publication.

[69] Valérie Berthé and Bertrand Nouvel. Discrete rotations and symbolic dynamics. *Theoretical Computer Science*, 380(3):276–285, 2007.

[70] Nicolas Brisebarre and Jean-Michel Muller. Correct rounding of algebraic functions. *Theoretical Informatics and Applications*, 41:71–83, January 2007.

[71] Eddy Caron, Frédéric Desprez, and Cédric Tedeschi. Enhancing computational grids with peer-to-peer technology for large scale service discovery. *Journal of Grid Computing*, 5(3):337–360, September 2007.

[72] Eddy Caron, Vincent Garonne, and Andreï Tsaregorodtsev. Definition, modelling and simulation of a grid computing scheduling system for high throughput computing. *Future Generation Computer Systems*, 23(Issue 8):968–976, November 2007. ISSN:0167-739X.

[73] Hadda Cherroun, Alain Darte, and Paul Feautrier. Reservation table scheduling: Branch-and-bound based optimization vs. integer linear programming techniques. *RAIRO-OR*, 41(4):427–454, December 2007.

[74] Florent de Dinechin, Christoph Quirin Lauter, and Jean-Michel Muller. Fast and correctly rounded logarithms in double-precision. *Theoretical Informatics and Applications*, 41:85–102, 2007.

[75] Jérémie Detrey and Florent de Dinechin. Parameterized floating-point logarithm and exponential functions for FPGAs. *Microprocessors and Microsystems, Special Issue on FPGA-based Reconfigurable Computing*, 31(8):537–545, December 2007.

[76] Jérémie Detrey and Florent de Dinechin. A tool for unbiased comparison between logarithmic and floating-point arithmetic. *Journal of VLSI Signal Processing*, 49(1):161–175, 2007.

[77] Milos Ercegovac and Jean-Michel Muller. Complex square root with operand prescaling. *Journal of VLSI Signal Processing*, 49(1):19–30, October 2007.

[78] O. Finkel. An example of $\mathbf{\Pi}_3^0$-complete infinitary rational relation. *Computer Science Journal of Moldova*, 15 (1):3–21, 2007.

[79] O. Finkel and S. Todorcevic. Local sentences and Mahlo cardinals. *Mathematical Logic Quaterly*, 53 (6):558–563, 2007.

[80] Laurent Fousse, Guillaume Hanrot, Vincent Lefèvre, Patrick Pélissier, and Paul Zimmermann. MPFR: A multiple-precision binary floating-point library with correct rounding. *ACM Trans. Math. Softw.*, 33(2), 2007.

[81] Antoine Fraboulet and Tanguy Risset. Master interface for on-chip hardware accelerator burst communications. *Journal of VLSI Signal Processing*, 29(1):73–85, 2007.

[82] A. Gajardo and J. Mazoyer. One head machines from a symbolic approach. *Theoretical Computer Science*, 370 (1-3):34–47, 2007.

[83] Ryan Glabb, Laurent Imbert, Graham Jullien, Arnaud Tisserand, and Nicolas Veyrat-Charvillon. Multi-mode operator for SHA-2 hash functions. *Journal of Systems Architecture, Special issue on "Embedded Cryptographic Hardware"*, 53(2-3):127–138, 2007.

[84] M. Habib, D. Kelly, E. Lebhar, and C. Paul. Can transitive orientation make sandwich problems easier? *Discrete Mathematics*, 307(16):2030–2041, 2007. special issue on EuroComb'03.

[85] Daniel Hirschkoff, Damien Pous, and Davide Sangiorgi. An efficient abstract machine for safe ambients. *J. Log. Algebr. Program.*, 71(2):114–149, 2007.

[86] P. Koiran, V. Nesme, and N. Portier. The quantum query complexity of the abelian subgroup problem. *Theoretical Computer Science*, 380 (1-2):115–126, 2007. special issue on the occasion of ICALP'2005.

[87] P. Koiran and S. Perifel. The complexity of two problems on arithmetic circuits. *Theoretical Computer Science*, 389:172–181, 2007.

[88] Derrick Kondo, Gilles Fedak, Franck Cappello, Andrew A. Chien, and Henri Casanova. Resource Availability in Enterprise Desktop Grids. *Future Generation Computer Systems*, 23(7):888–903, August 2007. External publication.

[89] E. Lebhar and N. Schabanel. Routage dans les petits mondes. *)i(Interstice (INRIA)*, 2007. Article de popularisation.

[90] Loris Marchal, Veronika Rehn, Yves Robert, and Frédéric Vivien. Scheduling algorithms for data redistribution and load-balancing on master-slave platforms. *Parallel Processing Letters*, 17(1):61–77, 2007.

[91] Guillaume Melquiond and Sylvain Pion. Formally certified floating-point filters for homogeneous geometric predicates. *Theoretical Informatics and Applications*, 41(1), 2007.

[92] M. Morvan, E. Rémila, and E. Thierry. A note on the structure of spaces of domino tilings. *Discrete Mathematics*, 307:769–773, 2007.

[93] Stéphane Operto, Jean Virieux, Patrick Amestoy, Jean-Yves L'Excellent, Luc Giraud, and Hafedh Ben Hadj Ali. 3d finite-difference frequency-domain modeling of visco-acoustic wave propagation using a massively parallel direct solver: A feasibility study. *Geophysics*, 72(5):SM195–SM211, 2007.

[94] Edmundo Pereira de Souza Neto, Patrice Abry, Patrick Loiseau, Jean-Christophe Cejka, Marc-Antoine Custaud, Jean Frutoso, Claude Gharib, and Patrick Flandrin. Empirical mode decomposition to assess cardiovascular autonomic control in rats. *Fundamental & Clinical Pharmacology*, 21(5):481–496, October 2007.

[95] Damien Pous. New up-to techniques for weak bisimulation. *Theor. Comput. Sci.*, 380(1-2):164–180, 2007.

[96] Clément Rezvoy, Delphine Charif, Laurent Guéguen, and Gabriel A.B. Marais. MareyMap: an R-based tool with graphical interface for estimating recombination rates. *Bioinformatics*, 23(16):2188–2189, 2007.

[97] Gabriel Rilling, Patrick Flandrin, Paulo Gonçalves, and Jonathan M. Lilly. Bivariate empirical mode decomposition. *IEEE, Signal Processing Letters*, 14(12):936–939, 2007.

[98] Antoine Scherrer, Nicolas Larrieu, Pierre Borgnat, Philippe Owezarski, and Patrice Abry. Non Gaussian and long memory statistical characterisations for internet traffic with anomalies. *IEEE Transactions on Dependable and Secure Computing (TDSC)*, 4(1):56–70, 2007.

[99] William Thies, Frédéric Vivien, and Saman Amarasinghe. A step towards unifying schedule and storage optimization. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 29(6):45 p., 2007.

[100] Bora Uçar and Cevdet Aykanat. Partitioning sparse matrices for parallel preconditioned iterative methods. *SIAM Journal on Scientific Computing*, 29(4):1683–1709, 2007. External publication.

[101] Bora Uçar and Cevdet Aykanat. Revisiting hypergraph models for sparse matrix partitioning. *SIAM Review*, 49(4):595–603, 2007. External publication.

2008

[102] P. Koiran A. Chistov, H. Fournier and S. Perifel. On the construction of a family of transversal subspaces over finite fields. *Linear Algebra and its Applications*, 429 (2-3):589–600, 2008.

[103] Emmanuel Agullo, Abdou Guermouche, and Jean-Yves L'Excellent. A parallel out-of-core multifrontal method: Storage of factors on disk and analysis of models for an out-of-core active memory. *Parallel Computing, Special Issue on Parallel Matrix Algorithms*, 34(6-8):296–317, 2008.

[104] Abelkader Amar, Raphaël Bolze, Yves Caniou, Eddy Caron, Benjamin Depardon, Frédéric Desprez, Jean-Sébastien Gay, Gaël Le Mahec, and David Loureiro. Tunable scheduling in a GridRPC framework. *Concurrency and Computation: Practice and Experience*, 20(9):1051–1069, 2008.

[105] Narjess Ayari, Denis Barbaron, Laurent Lefèvre, and Pascale Vicat-Blanc Primet. Fault tolerance for highly available internet services: Concepts, approaches, and issues. *IEEE Communications Surveys and Tutorials*, 10(2):34–46, July 2008.

[106] Cevdet Aykanat, Berkant B. Cambazoglu, and Bora Uçar. Multi-level direct K-way hypergraph partitioning with multiple constraints and fixed vertices. *Journal of Parallel and Distributed Computing*, 68:609–625, 2008. External publication.

[107] Zoltan Balaton, Zoltan Farkas, Gabor Gombas, Peter Kacsuk, Robert Lovas, Attila Csaba Marosi, Ad Emmen, Gabor Terstyanszky, Tamas Kiss, Ian Kelley, Ian Taylor, Oleg Lodygensky, Miguel Cardenas-Montes, Gilles Fedak, and Filipe Araujo. EDGeS: the Common Boundary Between Service and Desktop Grids. *Parallel Processing Letters*, 18(3):433–453, September 2008. External publication.

[108] Olivier Beaumont, Larry Carter, Jeanne Ferrante, Arnaud Legrand, Loris Marchal, and Yves Robert. Centralized versus distributed schedulers for multiple bag-of-task applications. *IEEE Trans. Parallel Distributed Systems*, 19(5):698–709, 2008.

[109] Elyes Ben Hamida, Guillaume Chelius, Anthony Busson, and Eric Fleury. Neighbor discovery in multi-hop wireless networks: evaluation and dimensioning with interferences considerations. *Discrete Mathematics & Theoretical Computer Science*, 2008.

[110] Anne Benoit, Veronika Rehn-Sonigo, and Yves Robert. Replica placement and access policies in tree networks. *IEEE Trans. Parallel Distributed Systems*, 19(12):1614–1627, 2008.

[111] Anne Benoit and Yves Robert. Mapping pipeline skeletons onto heterogeneous platforms. *J. Parallel and Distributed Computing*, 68(6):790–808, 2008.

[112] Jean-Luc Beuchat, Nicolas Brisebarre, Jérémie Detrey, Eiji Okamoto, Masaaki Shirase, and Tsuyoshi Takagi. Algorithms and arithmetic operators for computing the η_t pairing in characteristic three. *IEEE Transactions on Computers*, 57(11):1454–1468, November 2008.

[113] Jean-Luc Beuchat, Takanori Miyoshi, Jean-Michel Muller, and Eiji Okamoto. Horner's rule-based multiplication over GF(p) and GF(p^n): A survey. *International Journal of Electronics*, 95(7):669–685, 2008.

[114] Jean-Luc Beuchat and Jean-Michel Muller. Automatic generation of modular multipliers for FPGA applications. *IEEE Transactions on Computers*, 57(12), December 2008.

[115] O. Bodini, T. Fernique, and E. Rémila. A characterization of flip-accessibility for lozenge tilings of the whole plane. *Information and Computation*, 206):1065–1073, 2008.

[116] Sylvie Boldo and Guillaume Melquiond. Emulation of a FMA and correctly rounded sums: Proved algorithms using rounding to odd. *IEEE Transactions on Computers*, 57(4):462–471, April 2008.

[117] Alin Bostan, Claude-Pierre Jeannerod, and Éric Schost. Solving structured linear systems with large displacement rank. *Theoretical Computer Science*, 407(1:3):155–181, November 2008.

[118] A. Bouillard, B. Gaujal, S. Lagrange, and E. Thierry. Optimal routing for end-to-end guarantees using network calculus. *Performance Evaluation*, 65(11-12):883–906, 2008.

[119] A. Bouillard and E. Thierry. An algorithmic toolbox for network calculus. *Discrete Event Dynamic Systems*, 18(1):3–49, 2008.

[120] Nicolas Brisebarre and Jean-Michel Muller. Correctly rounded multiplication by arbitrary precision constants. *IEEE Transactions on Computers*, 57(2):165–174, February 2008.

[121] Eddy Caron, Andréea Chis, Frédéric Desprez, and Alan Su. Design of plug-in schedulers for a GridRPC environment. *Future Generation Computer Systems*, 24(1):46–57, January 2008.

[122] Hugo Carrao, António Araújo, Paulo Gonçalves, and Mário Caetano. Multitemporal meris images for land cover mapping at national scale: the case study of portugal. *International Journal of Remote Sensing*, 2008. To appear.

[123] Hugo Carrao, Paulo Gonçalves, and Mário Caetano. Contribution of multispectral and multitemporal information from modis images to land cover classification. *Elsevier, Remote Sensing of Environment*, 112(3):986–997, 2008.

[124] P. Charbit, E. Jeandel, P. Koiran, S. Périfel, and S. Thomassé. Finding a vector orthogonal to roughly half a collection of vectors. *Journal of Complexity*, 24(1):39–53, 2008.

[125] Octavian Creţ, Ionut Trestian, Radu Tudoran, Laura Darabant, Lucia Văcariu, and Florent de Dinechin. Accelerating the computation of the physical parameters involved in transcranial magnetic stimulation using FPGA devices. *Romanian Journal of Information, Science and Technology*, 10(4):361–379, 2008.

[126] M. Delorme and J. Mazoyer. Examples of cellular automata classes. *RAIRO – Theoretical Informatics and Applications*, 42 (1):37–54, 2008.

[127] Jérémie Detrey and Florent de Dinechin. Fonctions élémentaires en virgule flottante pour les accélérateurs reconfigurables. *Technique et Science Informatiques*, 27(6):673–698, 2008.

[128] Jack Dongarra, Jean-François Pineau, Yves Robert, Zhiao Shi, and Frédéric Vivien. Revisiting matrix product on master-worker platforms. *International Journal of Foundations of Computer Science*, 19(6):1317–1336, 2008.

[129] Daniel J. Dougherty, Silvia Ghilezan, and Pierre Lescanne. Characterizing strong normalization in the Curien-Herbelin symmetric lambda calculus: Extending the Coppo-Dezani heritage. *Theor. Comput. Sci.*, 398(1-3):114–128, 2008.

[130] O. Finkel. Topological complexity of locally finite omega languages. *Archive for Mathematical Logic*, 47(6):625–651, 2008.

[131] Matthieu Gallet, Yves Robert, and Frédéric Vivien. Comments on "Design and performance evaluation of load distribution strategies for multiple loads on heterogeneous linear daisy chain networks". *J. Parallel and Distributed Computing*, 68(7):1021–1031, 2008.

[132] Jean-Sébastien Gay and Yves Caniou. Étude de la précision de Simbatch, une API pour la simulation de systèmes batch. *RSTI - Techniques et Sciences Informatiques*, 27(3-4):373–394, 2008.

[133] Brice Goglin, Olivier Gluck, and Pascale Vicat-Blanc Primet. Interaction efficace entre les réseaux rapides et le stockage distribué dans les grappes de calcul. *Technique et Science Informatiques*, 27(7/2008):911–940, September 2008.

[134] G. Hanrot, G. Tenenbaum, and J. Wu. Moyennes de certaines fonctions arithmétiques sur les entiers friables, ii. *Proc. London Math. Soc.*, 96:107–135, 2008. External publication.

[135] Daniel Hirschkoff and Damien Pous. A Distribution Law for CCS and a New Congruence Result for the Pi-calculus. *Logical Methods in Computer Science*, 4(2), 2008.

[136] P. Jensen, J.-B. Rouquier, P. Kreimer, and Y. Croissant. Scientists who engage with society perform better academically. *Science and Public Policy*, 35(7):527–541, Aug 2008.

[137] Laurent Lefèvre and Jean-Patrick Gelas. Ian2 : Industrial autonomic network node architecture for supporting personnalized network services in the industrial context. *The International Journal of Future Generation Computer Systems (FGCS) - Grid Computing: Theory, Methods and Applications*, 24(1):58–65, January 2008.

[138] Arnaud Legrand, Alan Su, and Frédéric Vivien. Minimizing the stretch when scheduling flows of divisible requests. *Journal of Scheduling*, 11(5):381–404, 2008.

[139] Étienne Lozes, Daniel Hirschkoff, and Davide Sangiorgi. Separability in the ambient logic. *Logical Methods in Computer Science*, 4(3:4), September 2008.

[140] G. Malod and N. Portier. Characterizing valiant's algebraic complexity classes. *Journal of Complexity*, 24(1), 2008. journal version of [387].

[141] Romain Michard, Arnaud Tisserand, and Nicolas Veyrat-Charvillon. Optimisation d'opérateurs arithmétiques matériels à base d'approximations polynomiales. *Technique et science informatiques*, 27(6):699–718, June 2008.

[142] Jean-François Pineau, Yves Robert, and Frédéric Vivien. The impact of heterogeneity on master-slave scheduling. *Parallel Computing*, 34(3):158–176, 2008.

[143] Damien Pous. Using bisimulation proof techniques for the analysis of distributed abstract machines. *Theor. Comput. Sci.*, 402(2-3):199–220, 2008.

[144] T. Razafindralambo and I. Guérin Lassous. Increasing fairness and efficiency using the madmac protocol in ad hoc networks. *Ad Hoc Networks*, 6(3), May 2008.

[145] T. Razafindralambo, I. Guérin Lassous, L. Iannone, and S. Fdida. Dynamic packet aggregation to solve performance anomaly in 802.11 wireless networks. *Computer Networks*, 52(1), January 2008.

[146] Cheikh Sarr, Claude Chaudet, Guillaume Chelius, and Isabelle Guerin Lassous. Bandwidth Estimation for IEEE 802.11-based Ad Hoc Networks. *IEEE Transactions on Mobile Computing*, 7:1228–1241, 2008. External publication.

[147] S. Sarr, C. Chaudet, G. Chelius, and I. Guérin Lassous. Bandwidth estimation for IEEE 802.11-based ad hoc networks. *IEEE Transactions on Mobile Computing*, 7(10), 2008.

[148] Antoine Scherrer, Pierre Borgnat, Eric Fleury, Jean-Loup Guillaume, and Céline Robardet. Description and simulation of dynamic mobility networks. *Computer Networks (1976)*, 2008. This work is partially financed by the European Commission under the Framework 6 HealthCare Project LSH PL037941 Mastering hOSpital Antimicrobial Resistance and its spread into the community (MOSAR) and AEOLUS project IST IP-FP6-015964. The views given herein represent those of the authors and may not necessarily be representative of the views of the project consortium as a whole.

[149] Steffen van Bakel and Pierre Lescanne. Computation with classical sequents. *Mathematical Structures in Computer Science*, 18(3):555–609, 2008.

2009

[150] Philippe Audebaud and Christine Paulin-Mohring. Proofs of randomized algorithms in Coq. *Science of Computer Programming*, 74(8):568–589, June 2009.

[151] O. Audouin, D. Barth, M. Gagnaire, C. Mouton, P. Vicat-Blanc Primet, D. Rodrigues, L. Thual, and D. Verchère. Carriocas project: Towards converged internet infrastructures supporting high performance distributed applications. *IEEE/OSA Journal of Lightwave Technology*, 2009. accepted.

[152] P. Baillot and K. Terui. Light types for polynomial time computation in lambda calculus. *Information and Computation*, 207(1):41–62, 2009. External publication.

[153] Patrick Baillot and Damiano Mazza. Linear logic by levels and bounded time complexity. *Theoretical Computer Science*, 2009. 49 pages, to appear.

[154] Françoise Baude, Caromel Denis, Cédric Dalmasso, Marco Danelutto, Vladimir Getov, Ludovic Henrio, and Christian Pérez. Gcm: A grid extension to fractal for autonomous distributed components. *Special Issue of Annals of Telecommunications: Software Components – The Fractal Initiative*, 64(1):5–24, 2009.

[155] Elyes Ben Hamida, Guillaume Chelius, and Jean-Marie Gorce. Impact of the Physical Layer Modeling on the Accuracy and Scalability of Wireless Network Simulation. *Simulation*, 2009. External publication.

[156] Anne Benoit, Mourad Hakem, and Yves Robert. Contention awareness and fault tolerant scheduling for precedence constrained tasks in heterogeneous systems. *Parallel Computing*, 2009. To appear.

[157] Anne Benoit, Mourad Hakem, and Yves Robert. Multi-criteria scheduling of precedence task graphs on heterogeneous platforms. *The Computer Journal*, 2009. To appear.

[158] Anne Benoit, Harald Kosch, Veronika Rehn-Sonigo, and Yves Robert. Multi-criteria scheduling of pipeline workflows (and application to the JPEG encoder). *Int. Journal of High Performance Computing Applications*, 23(2):171–187, 2009.

[159] Anne Benoit, Loris Marchal, Jean-François Pineau, Yves Robert, and Frédéric Vivien. Scheduling concurrent bag-of-tasks applications on heterogeneous platforms. *IEEE Transactions on Computers*, 2009. To appear.

[160] Anne Benoit and Yves Robert. Complexity results for throughput and latency optimization of replicated and data-parallel workflows. *Algorithmica*, 2009. To appear.

[161] Anne Benoit, Eric Thierry, and Yves Robert. On the complexity of mapping linear chain applications onto heterogeneous platforms. *Parallel Processing Letters*, 19(3):383–397, 2009.

[162] V. Bertis, R. Bolze, F. Desprez, and K. Reed. From Dedicated Grid to Volunteer Grid: Large Scale Execution of a Bioinformatics Application. *Journal of Grid Computing*, 2009. To appear.

[163] V. Blondel, S. Gaubert, and N. Portier. The set of realizations of a max-plus sequence is semi-polyhedral. *Journal of Computer and System Sciences*, 2009. To appear.

[164] Eddy Caron, Frédéric Desprez, Cédric Tedeschi, and Franck Petit. Snap-stabilizing prefix tree for peer-to-peer systems. *Parallel Processing Letters*, 2009. To appear.

[165] Hugo Carrao, Paulo Gonçalves, and Mário Caetano. A nonlinear model for satellite images time series: analysis and prediction of land cover dynamics. *IEEE Transactions on Geoscience and Remote Sensing*, 2009. To appear.

[166] R. Demangeon, D. Hirschkoff, and D. Sangiorgi. Mobile processes and termination. *Festschrift in the Honour of Peter Mosses*, 2009. to appear.

[167] Milos Ercegovac and Jean-Michel Muller. An efficient method for evaluating complex polynomials. *Journal of VLSI Signal Processing Systems*, 2009. To appear.

[168] Gilles Fedak, Haiwu He, and Franck Cappello. A Data Management and Distribution Service with Multi-Protocol and Reliable File Transfer. *Journal of Network and Computer Applications*, 2009. To appear.

[169] Eric Fleury and Céline Robardet. Communities detection and the analysis of their dynamics in collaborative networks. *International Journal of Web Based Communities (IJWBC)*, 2009.

[170] Stef Graillat, Jean-Luc Lamotte, and Diep Nguyen Hong. Extended precision with a rounding mode toward zero environment. application on the CELL processor. *International Journal of Reliability and Safety, Special issue on "Reliable Engineering Computing*, 2009. To appear.

[171] Stef Graillat, Philippe Langlois, and Nicolas Louvet. Algorithms for accurate, validated and fast polynomial evaluation. *Japan Journal of Industrial and Applied Mathematics*, 2009. To appear.

[172] Philippe Grosse, Yves Durand, and Paul Feautrier. Methods for power optimization in SOC-based data flow systems. *ACM Transactions on Design Automation of Electronic Systems*, 14(3):1–20, 2009.

[173] P. Jensen, J.-B. Rouquier, and Y. Croissant. Testing bibliometric indicators by their prediction of scientists promotions. *Scientometrics*, 78(3), March 2009.

[174] Kamer Kaya and Bora Uçar. Exact algorithms for a task assignment problem. *Parallel Processing Letters*, 19(3):451–465, 2009.

[175] P. Koiran, J. Landes, N. Portier, and P. Yao. Adversary lower bounds for nonadaptive quantum algorithms. *Journal of Computer and System Sciences*, 2009. Special issue on Wollic'08, to appear.

[176] Peter Kornerup, Christoph Quirin Lauter, Vincent Lefèvre, Nicolas Louvet, and Jean-Michel Muller. Computing correctly rounded integer powers in floating-point arithmetic. *ACM Transactions on Mathematical Software*, 37(1), 2009. To appear.

[177] Olivier Laurent. On the categorical semantics of elementary linear logic. *Theory and Applications of Categories*, 22(10):269–301, June 2009. External publication.

[178] Christoph Quirin Lauter and Vincent Lefèvre. An efficient rounding boundary test for `pow(x,y)` in double precision. *IEEE Transactions on Computers*, 58(2):197–207, February 2009.

[179] Laurent Lefèvre and Anne-Cécile Orgerie. Towards energy aware reservation infrastructure for large-scale experimental distributed systems. *Parallel Processing Letters*, 2009.

[180] Patrick Loiseau, Paulo Gonçalves, Guillaume Dewaele, Pierre Borgnat, Patrice Abry, and Pascale Vicat-Blanc Primet. Investigating self-similarity and heavy-tailed distributions on a large scale experimental facility. *IEEE/ACM Transactions on Networking*, 2009. To appear.

[181] Ivan Morel, Damien Stehlé, and Gilles Villard. Analyse numérique et réduction des réseaux. *Technique et Science Informatiques*, 2009. To appear.

[182] Phong Q. Nguyen and Damien Stehlé. An LLL algorithm with quadratic complexity. *SIAM Journal on Computing*, 2009. To appear.

[183] Phong Q. Nguyen and Damien Stehlé. Low-dimensional lattice basis reduction revisited. *ACM Transactions on Algorithms*, 2009.

[184] Sebastien Soudan, Binbin Chen, and Pascale Vicat-Blanc Primet. Flow scheduling and endpoint rate control in gridnetworks. *Future Generation Computer Systems*, 25(8):904–911, 2009.

[185] Florent Sourbier, Stéphane Operto, Jean Virieux, Patrick R. Amestoy, and Jean-Yves L'Excellent. FWT2D: a massively parallel program for frequency-domain full-waveform tomography of wide-aperture seismic data – part 1: Algorithm. *Computer and Geosciences*, 35(3):487–495, 2009.

[186] Florent Sourbier, Stéphane Operto, Jean Virieux, Patrick R. Amestoy, and Jean-Yves L'Excellent. FWT2D: a massively parallel program for frequency-domain full-waveform tomography of wide-aperture seismic data – part 2: numerical examples and scalability analysis. *Computer and Geosciences*, 35(3):496–514, 2009.

[187] Piero Spinnato, Pascale Vicat-Blanc Primet, Chris Edwards, and Michael Welzl. Editorial special section on networks for grid applications. *International Journal on Future Generation Computer Systems*, april 2009.

[188] Pascale Vicat-Blanc Primet, Sebastien Soudan, and Dominique Verchere. Virtualizing and scheduling optical network infrastructure for emerging it services. *Optical Networks for the Future Internet (special issue of Journal of Optical Communications and Networking (JOCN))*, 1(2):A121–A132, 2009.

Invited conferences [INV], international seminars, and tutorials

2005

[189] Alain Darte. Lattice-based memory allocation. In *Seminar on Sphere Packings*, Mathematisches Forschungsinstitut Oberwolfach (Germany), November 2005. Invited talk.

[190] Nicolas Schabanel. Asynchronous randomized automata: how does randomness affect decentralized algorithms execution?, Juin 2005. Invited talk at “Approximation and randomized algorithms", Dagstuhl seminar 05201.

[191] Arnaud Tisserand. Algorithms and number systems for hardware computer arithmetic. In *International Symposium on Symbolic and Algebraic Computation (ISSAC)*, Beijing, China, July 2005. Invited tutorial.

[192] Gilles Villard. Efficient algorithms in linear algebra, May 2005. 33rd Theoretical Computer Science Spring School, Computational Complexity, Montagnac-les-truffes.

2006

[193] Paulo Gonçalves. Empirical Mode Decomposition: Definition and applications to non linear time series. Nonlinear Dynamical methods and time series analysis workshop, Udine (Italy), August 2006.

[194] P. Koiran. Algebraic versions of the P=NP problem, 2006. Invited talk at the midterm conference of the european network MODNET (model theory network). Antalya, Turkey, Novembre 2006.

[195] P. Koiran. Algebraic versions of the P=NP problem, 2006. Invited lecture at CSR 2006 (International Computer Science Symposium in Russia, Saint Petersburg, june 2006).

[196] P. Koiran. Decision versus evaluation in algebraic complexity theory, 2006. Invited talk at the session on “Complexity and Computability in Analysis, Geometry, and Dynamics” of the winter 2006 CMS meeting (Toronto, December 2006).

[197] Laurent Lefèvre. Designing high performance autonomic gateways for large scale grids and distributed environments. CCGSC 2006 : Clusters and Computational Grids for Scientific Computing Workshop, Flat Rock, USA, September 2006.

[198] Laurent Lefèvre. High performance programmable network support for grid infrastructures. Kolloquium of Computer Science, Linz University, Austria, January 2006.

[199] Yves Robert. Static scheduling for large-scale platforms: can one hope for efficiency? Invited conference at *IPDPS’06, the IEEE Int. Parallelism and Distributed Processing Symposium*, 2006. http://www.ipdps.org/ipdps2006.

[200] Pascale Vicat-Blanc Primet. High speed transport protocol evaluation in grid5000. National ARRU workshop organised by RENATER, 2006.

[201] Pascale Vicat-Blanc Primet. High speed transport protocol evaluation in grid5000. 4th International TERENA conference "NRENs and Grids", 2006.

[202] Pascale Vicat-Blanc Primet. Qos and security issues in grids. First International workshop ITU/OGF (International Telecommunication Union/Open Grid Forum) on Grids and New Generation Networks, 2006.

[203] Pascale Vicat-Blanc Primet. Research on high speed networks and transport protocols for grid applications. AIST booth - SC06, 2006.

[204] Frédéric Vivien. Minimizing the stretch when scheduling flows of divisible requests. Invited conference at the *Scheduling Algorithms for new Emerging Applications* workshop, May 30, 2006. http://www-id.imag.fr/SAEA06/index.php.

2007

[205] Anne Benoit. Mapping skeleton workflows onto heterogeneous platforms. Invited presentation at the Workshop on Parallelism Oblivious Programming, July 2007. http://www.ipl.t.u-tokyo.ac.jp/~kmatsu/pop07/programme.html.

[206] Anne Benoit. Mapping skeleton workflows onto heterogeneous platforms. Invited conference at the *9th High Performance Computing Conference*, June 2007. http://www.hpcc.unical.it/hpc2007/.

[207] Yves Caniou. Tunable scheduling and middleware deployment for a scalable gridrpc middleware. Invited presentation at the first workshop Grid@Mons, May 4 2007. http://grid.umh.ac.be/2007/menu2007.

[208] Alain Darte. Revisiting register coalescing in the light of SSA form. In *Workshop Compilers for Parallel Computing*, Lisboa, Portugal, July 2007. Invited talk.

[209] Paul Feautrier. Embedded systems and compilation. In *MSR (Modélisation des systèmes réactifs) conference*, 2007. Invited talk.

[210] I. Guérin Lassous. Mac protocols for ad hoc networks. IRAMUS workshop, Val Thorens, France, January 2007.

[211] P. Koiran. Decision versus evaluation in algebraic complexity theory, 2007. Invited talk at the MCU 2007 Conference (Modèles de Calculs Universels). Orléans, September 2007.

[212] P Koiran. Decision versus evaluation in algebraic complexity theory, 2007. Invited talk at the IMA workshop on Complexity, Coding, and Communications (Minneapolis, avril 2007).

[213] P Koiran. Interpolation in Valiant's theory, 2007. Invited talk at the Oberwolfach complexity theory workshop (june 2007).

[214] P. Koiran. Problèmes de décision et d'évaluation en complexité algébrique, 2007. Journées du GDR Informatique Mathématique (IHP, Paris, February 2007).

[215] Laurent Lefèvre. Autonomic and programmable networks approach for supporting long latency (inter-planetary) grids. Otago University, Seminar of New Zealand Distributed Information Systems group, New Zealand, August 2007.

[216] Laurent Lefèvre. Next generation router-assisted transport protocols for high performance grids : interoperability and fairness issues. Ho Chi Minh Ville University, Vietnam, May 2007.

[217] Laurent Lefèvre. Towards new services and capabilities for next generation grids. Otago University, Seminar of New Zealand Distributed Information Systems group, New Zealand, August 2007.

[218] Jean-Michel Muller. Exact computations with approximate arithmetic. Invited conference at the "Journees en l'honneur de Donald Knuth", for the Doctor Honoris Causa diploma given to D. Knuth, Bordeaux, France, 29-31 Oct. 2007.

[219] Yves Robert. Algorithms and scheduling techniques for clusters and grids. Invited conference at the *9th High Performance Computing Conference*, 2007. http://www.hpcc.unical.it/hpc2007/.

[220] Yves Robert. Algorithms and scheduling techniques for heterogeneous platforms. Invited conference at the *Workshop on Parallelism Oblivious Programming*, 2007. http://www.ipl.t.u-tokyo.ac.jp/~kmatsu/pop07/programme.html.

[221] Yves Robert. Algorithms for heterogeneous platforms. Invited conference at *HeteroPar'2007*, 2007. http://www.labri.fr/perso/obeaumon/heteropar07.

[222] Yves Robert. Scheduling bags of tasks. Invited conference at *Arny's fest, Amherst*, 2007. http://www.cs.umass.edu/~immerman/TheorySeminar/arnyFest.pdf.

[223] Damien Stehlé. Floating-point LLL: Theoretical and practical aspects. LLL+25 conference, in honour of the 25th anniversary of the LLL algorithm, Caen, June 2007.

[224] Damien Stehlé. Numerical analysis and LLL reduction. Oberwolfach bi-annual workshop on Explicit Methods in Number Theory, Germany, July 2007.

[225] Pascale Vicat-Blanc Primet. High speed cluster virtualisation. STIC07, 2007.

[226] Pascale Vicat-Blanc Primet. Self-management and grid networks. IM07, May 2007.

[227] Gilles Villard. Some recent progress in symbolic linear algebra and related questions *(invited tutorial)*. In *Proc. International Symposium on Symbolic and Algebraic Computation, Waterloo, Canada*, pages 391–392. ACM Press, August 2007.

2008

[228] Anne Benoit. Multi-criteria scheduling of workflow applications. Invited presentation at the Workshop on Clusters and Computational Grids for Scientific Computing, September 2008.

[229] Eddy Caron. Diet: un intergiciel de grille pour l'ordonnancement d'applications. Invited conference at *JDIR'08*. 9èmes Journées Doctorales en Informatique et Réseaux, January 16, 2008. http://www2.lifl.fr/JDIR2008/.

[230] Olivier Finkel and Dominique Lecomte. On the topological complexity of ω-powers, 2008. Dagstuhl Seminar on "Topological and Game-Theoretic Aspects of Infinite Computations", Dagstuhl, Allemagne, 29.06.08 - 04.07.08.

[231] I. Guérin Lassous. Resources allocation for multihop wireless networks. IFI, Hanoi, Vietnam, October 2008.

[232] P. Koiran. Decision versus evaluation in algebraic complexity theory, 2008. Invited talk at the DAIMI workshop on algebraic complexity theory (Aarhus, Danemark, September 2008).

[233] P. Koiran. On the expressive power of planar perfect matching and permanents of bounded treewidth matrices, 2008. Invited talk at the DAIMI workshop on algebraic complexity theory (Aarhus, Danemark, septembre 2008).

[234] P. Koiran. On the expressive power of planar perfect matching and permanents of bounded treewidth matrices, 2008. Invited talk at FoCM 2008 (Foundations of Computational Mathematics), workshop on real number algorithms (Hong-Kong, june 2008).

[235] Laurent Lefèvre. Green-* : Towards energy efficient solutions for next generation large scale distributed systems. ACOMP 2008 : International Workshop on Advanced Computing and Applications, Ho Chi Minh City, Vietnam, March 2008.

[236] Laurent Lefèvre. Proposing inter-operable router-assisted transport protocol for transferring large volume of data in high performance grids. University of Sevilla, Spain, March 2008.

[237] Laurent Lefèvre. Towards energy aware resource infrastructure for large scale distributed systems. CCGSC 2008 : Clusters and Computational Grids for Scientific Computing Workshop, Flat Rock, North Carolina, USA, September 2008.

[238] Laurent Lefèvre. Towards new services and capabilities for next generation grids. University of Sevilla, Spain, March 2008.

[239] Laurent Lefèvre and Anne Cecile Orgerie. Energy efficiency challenges for large scale distributed systems. Deakin University, Australia, December 2008.

[240] Loris Marchal. Offline and online scheduling of concurrent bags-of-tasks on heterogeneous platforms. Invited presentation at the *New Challenges on Scheduling Theory* workshop, May 12-16, 2008. http://www-id.imag.fr/NCST08/index.php.

[241] Jean-Michel Muller. Some algorithmic improvements due to the availability of an FMA, invited talk at the Dagstuhl seminar 08021, numerical validation in current hardware architectures, January 2008.

[242] Natacha Portier. Algorithmes quantiques : est-ce que faire des requêtes en parallèle fait perdre du temps ?, 2008. exposé invité, Réunion finale du GDR Information et communication quantique du 6 au 8 octobre 2008 à Paris.

[243] Natacha Portier. Bornes inférieures pour algorithmes quantiques non adaptatifs, 2008. journée information quantique, 23 janvier 2008, IHP.

[244] Fabrice Rastello. SSA-based register allocation. In *ESWeek'08 (Embedded Systems Week)*, Atlanta, October 2008. Tutorial with P. Brisk, S. Hack, J. Palsberg, and F. Pereira.

[245] Eric Rémila. Quelques constructions sur les pavages auto-assemblants, 2008. Invited conference at the "JOurnées Rouennaises de Combinatoire et Algorithmique ", (JORCAD'08), Rouen, France, 15-19 Sept. 2008.

[246] Nathalie Revol. Automatic adaptation of the computing precision, January 2008. Dagstuhl seminar 08021 on Numerical Validation in Current Hardware Architectures.

[247] Nathalie Revol. Introduction to interval analysis and to some interval-based software systems and libraries. In *ECMI 2008, The European Consortium For Mathematics In Industry*, July 2008.

[248] Yves Robert. Algorithm design and scheduling techniques for clusters and grids. Invited conference at the *Workshop on Parallel Routines Optimization and Applications*, 2008. http://dis.um.es/~domingo/08/CD/programa.html.

[249] Yves Robert. Static strategies for worksharing with unrecoverable interruptions. Invited presentation at the Workshop on Clusters and Computational Grids for Scientific Computing, September 2008.

[250] Eric Thierry. Algorithmique du network calculus, 2008. Invited talk, 9ème Atelier d'Evaluation de Performances, Aussois, juillet 2008.

[251] Pascale Vicat-Blanc Primet. CARRIOCAS architecture for coordination of very high speed network and high end applications. CARRIOCAS-PHOSPHORUS workshop meeting, 2008.

[252] Pascale Vicat-Blanc Primet. The CARRIOCAS project: Orchestrating dynamic network service deliveries over ultra high capacity optical networks. CCGrid2008, May 2008.

[253] Pascale Vicat-Blanc Primet. Dynamic bandwidth provisioning in the CARRIOCAS project. AIST, 2008.

[254] Pascale Vicat-Blanc Primet. Flow scheduling and network virtualization. Tokyo Technical University (TiTech), 2008.

[255] Pascale Vicat-Blanc Primet. High speed transport and flow scheduling. Osaka University, 2008.

[256] Pascale Vicat-Blanc Primet. Network virtualization. ITU-T's System and Network Operation group meeting at Beijing, 2008.

[257] Pascale Vicat-Blanc Primet. Network virtualization: perspectives in grids. GridNets2008, Beijing, 2008.

[258] Pascale Vicat-Blanc Primet. Presentation of the CARRIOCAS project. Optical Networks workshop of GridNets2008, Beijing, 2008.

2009

[259] Anne Benoit. Scheduling pipelined applications: models, algorithms and complexity. Invited presentation at the Workshop on Algorithms and Techniques for Scheduling on Clusters and Grids, June 2-5, 2009. http://www-id.imag.fr/ASTEC09/index.php.

[260] Benoit Boissinot. Fast liveness checking in SSA. In *SSA form seminar*, Autrans, April 2009. Invited talk.

[261] Alain Darte. Out-of-SSA translation. In *SSA form seminar*, Autrans, April 2009. Invited talk.

[262] Alain Darte and Fabrice Rastello. SSA-based register allocation. In *CGO 2008*, Seattle, March 2009. Tutorial with P. Brisk and J. Palsberg.

[263] Iain S. Duff and Bora Uçar. Combinatorial problems in solving linear systems. Invited presentation at Dagstuhl Seminar on Combinatorial Scientific Computing delivered by Iain S. Duff, February 1–6, 2009.

[264] Paulo Gonçalves. Empirical Mode Decomposition: Definition and applications to non linear time series. IEEE, 17th Signal Processing and Communications - Application Conference, Antalya, Turkey, April 2009.

[265] Paulo Gonçalves. Fractal dimension estimation: EMD versus wavelets. Istanbul Technical University, Elec. and Comp. Eng. Dept., April 2009.

[266] I. Guérin Lassous. Quality of service issues in multihop wireless networks. Polytechnic University of Catalonia, Department of Telematic Engineering, 10h, June 2009.

[267] P Koiran, 2009. Complexity theory workshop (Oberwolfach, Germany, November 2009).

[268] P. Koiran. Près de 40 ans après le théorème de Cook, où en est la complexité algorithmique ? (Une petite histoire du problème "P=NP ?"), 2009. Colloquium Jacques Morgenstern (Sophia-Antipolis, may 2009).

[269] Laurent Lefèvre. Energy efficiency issues for large scale distributed systems : the green-net initiative. OGF 25 : Open Grid Forum during "OGF-EU: Using IT to reduce Carbon Emissions and Delivering the Potential of Energy Efficient Computing" sesion, Catania, Italy, March 2009.

[270] Laurent Lefèvre. Pourquoi chasser les watts dans les systèmes distribués à grande échelle ? les approches green-*. Keynote Talk Renpar 2009 Conference, Toulouse, France, September 2009.

[271] Laurent Lefèvre and Anne Cecile Orgerie. Towards energy aware resource infrastructure for large scale distributed systems. University of Melbourne, Australia, January 2009.

[272] Loris Marchal. Efficient scheduling of task graph collections on heterogeneous resources. Invited presentation at the Workshop on Algorithms and Techniques for Scheduling on Clusters and Grids, June 2-5, 2009. http://www-id.imag.fr/ASTEC09/index.php.

[273] Jean-Michel Muller. Exact computations with an arithmetic known to be approximate. Invited lectures (3 hours) at the conference *Numeration: Mathematics and Computer Science*, CIRM, Marseille, Mar. 2009.

[274] Jean-Baptiste Rouquier. Cartographie des pratiques du vélo'v : le regard de physiciens et d'informaticiens, 2009. Invited talk, conference “Le métier de cartographe hier et aujourd'hui”, ENSSIB, Lyon, France, May 28th 2009.

[275] Pascale Vicat-Blanc Primet. Virtualizing and scheduling network resource for emerging it services: the CARRIOCAS approach. Optical Forum and Conference /NFOEC Workshop on Grid vs Cloud/Utility Computing and Optical Networks, 2009.

[276] Pascale Vicat-Blanc Primet. VXDL: Virtual execution infrastructures description language. OGF 25 NML working group, 2009.

[277] Pascale Vicat-Blanc Primet. Where are we going with bandwidth on demand. 8th International TERENA conference, 2009.

International and national peer-reviewed conference proceedings [ACT]

2005

[278] Christophe Alias. Tema: an efficient tool to find high-performance library patterns in source code. In *International Workshop on Patterns in High-Performance Computing*, 2005. External publication.

[279] Christophe Alias and Denis Barthou. On domain specific languages re-engineering. In *International ACM Conference on Generative Programming and Component-based Engineering (GPCE'05)*, September 2005. External publication.

[280] Narjess Ayari, Denis Barbaron, Laurent Lefèvre, and Pascale Vicat-Blanc Primet. A survey on high availability mechanisms for ip services. In *HAPCW2005 : High Availability and Performance Computing Workshop*, Santa Fe, New Mexico, USA, October 2005.

[281] P. Baillot and K. Terui. A Feasible Algorithm for Typing in Elementary Affine Logic. In *Proceedings of 7th International Conference on Typed Lamb da Calculi and Applications (TLCA 2005)*, volume 3461 of *LNCS*, pages 55–70. Springer, 2005. External publication.

[282] Rachid Beguenane, Jean-Luc Beuchat, Jean-Michel Muller, and Stéphane Simard. Modular multiplication of large integers on FPGA. In *39th Asilomar Conference on Signals, Systems & Computers*. IEEE, November 2005.

[283] Jean-Luc Beuchat and Jean-Michel Muller. Multiplication algorithms for radix-2 RN-codings and two's complement numbers. In *Proceedings of the 16th IEEE International Conference on Application-Specific Systems, Architectures, and Processors (ASAP 2005)*, pages 303–308. IEEE, 2005.

[284] Jean-Luc Beuchat and Jean-Michel Muller. RN-codes : algorithmes d'addition, de multiplication et d'élévation au carré. In *SympA'2005: 10e édition du SYMPosium en Architectures nouvelles de machines*, pages 73–84, April 2005.

[285] Sylvie Boldo and Guillaume Melquiond. When double rounding is odd. In *Proceedings of the 17th IMACS World Congress on Computational and Applied Mathematics*, Paris, France, 2005.

[286] Sylvie Boldo and Jean-Michel Muller. Some functions computable with a fused-MAC. In *Proc. 17th IEEE Symposium on Computer Arithmetic (ARITH-17)*, pages 52–58, Cape Cod, USA, 2005.

[287] Nicolas Brisebarre and Jean-Michel Muller. Correctly rounded multiplication by arbitrary precision constants. In *Proc. 17th IEEE Symposium on Computer Arithmetic (ARITH-17)*. IEEE Computer Society Press, June 2005.

[288] Franck Cappello, Frederic Desprez, Michel Dayde, Emmanuel Jeannot, Yvon Jegou, Stephane Lanteri, Nouredine Melab, Raymond Namyst, Pascale Vicat-Blanc Primet, Olivier Richard, Eddy Caron, Julien Leduc, and Guillaume Mornet. Grid'5000: A large scale, reconfigurable, controlable and monitorable grid platform. In *Grid2005 6th IEEE/ACM International Workshop on Grid Computing*, 2005.

[289] Martine Chaudier, Jean-Patrick Gelas, and Laurent Lefèvre. Towards the design of an autonomic network node. In *IWAN2005 : Seventh Annual International Working Conference on Active and Programmable Networks*, Nice, France, November 2005.

[290] Alain Darte, Steven Derrien, and Tanguy Risset. Hardware/software interface for multi-dimensional processor arrays. In *IEEE International Conference on Application-Specific Systems, Architecture, and Processors (ASAP'05)*, pages 28–35. IEEE Computer Society Press, July 2005.

[291] Alain Darte and Robert Schreiber. A linear-time algorithm for optimal barrier placement. In *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP'05)*, pages 26–35, Chicago, IL, USA, June 2005.

[292] Marc Daumas, Guillaume Melquiond, and Cesar Muñoz. Guaranteed proofs using interval arithmetic. In *Proceedings of the 17th IEEE Symposium on Computer Arithmetic*, pages 188–195, Cape Cod, Massachusetts, USA, 2005.

[293] Florent de Dinechin, Alexey Ershov, and Nicolas Gast. Towards the post-ultimate libm. In *17th Symposium on Computer Arithmetic*, pages 288–295. IEEE, 2005.

[294] Jérémie Detrey and Florent de Dinechin. A parameterizable floating-point logarithm operator for FPGAs. In *39th Asilomar Conference on Signals, Systems & Computers*. IEEE, November 2005.

[295] Jérémie Detrey and Florent de Dinechin. A parameterized floating-point exponential function for FPGAs. In *IEEE International Conference on Field-Programmable Technology (FPT'05)*. IEEE, December 2005.

[296] Jérémie Detrey and Florent de Dinechin. Table-based polynomials for fast hardware function evaluation. In *16th Intl Conference on Application-specific Systems, Architectures and Processors (ASAP 2005)*. IEEE, July 2005.

[297] Dan J. Dougherty, Pierre Lescanne, Luigi Liquori, and Frédéric Lang. Addressed term rewriting systems: Syntax, semantics, and pragmatics: Extended abstract. *Electr. Notes Theor. Comput. Sci.*, 127(5):57–82, 2005.

[298] Daniel J. Dougherty, Silvia Ghilezan, Pierre Lescanne, and Silvia Likavec. Strong normalization of the dual classical sequent calculus. In *Proc. of Logic for Programming, Artificial Intelligence, and Reasoning, 12th International Conference, LPAR 2005, Montego Bay, Jamaica, December 2-6, 2005, Proceedings*, volume 3835 of *Lecture Notes in Computer Science*, pages 169–183. Springer, 2005.

[299] P. Duchon, N. Hanusse, E. Lebhar, and N. Schabanel. Could any graph be turned into a small world? In *actes d'AlgoTel 2005*, 2005.

[300] Jean Duprat. About constructive vectors. In P. Lescanne, R. David, and M. Zaionc, editors, *Proceedings of the Second Workshop on Computational Logic and Applications*, volume 140 of *Electr. Notes Theor. Comput. Sci.*, pages 93–100, 2005.

[301] Milos D. Ercegovac, Jean-Michel Muller, and Arnaud Tisserand. Simple seed architectures for reciprocal and square root reciprocal. In *Proc. 39th Asilomar Conference on Signals, Systems and Computers*, Pacific Grove, California, U.S.A., October 2005.

[302] N. Fatès, M. Morvan, N. Schabanel, and E. Thierry. Fully asynchronous behavior of double-quiescent elementary cellular automata. In *Proceedings of MFCS 2005 (30th International Symposium on Mathematical Foundations of Computer Science)*, Lecture Notes in Computer Science 3618, pages 316–327, 2005.

[303] F. Fomin, F. Mazoit, and I. Todinca. Computing branchwidth via efficient triangulations and blocks. In *Proceedings of WG'05 (31st Workshop on Graph-Theoretic Concepts in Computer Science)*, Lecture Notes in Computer Science 5344, pages 374–384, 2005.

[304] Brice Goglin, Olivier Gluck, and Pascale Vicat-Blanc Primet. An efficient network api for in-kernel applications in clusters. In *Proceedings of the IEEE International Conference on Cluster Computing*, Boston, Massachussets, September 2005. IEEE Computer Society Press.

[305] Brice Goglin, Olivier Gluck, Pascale Vicat-Blanc Primet, and Jean-Christophe Mignot. Accès optimisés aux fichiers distants dans les grappes disposant d'un réseau rapide. In *Actes de RenPar'16, CFSE'4, SympAAA'2005*, pages 37–46, Le Croisic, Presqu'ile de Guérande, France, April 2005.

[306] Daniel Hirschkoff, Tom Hirschowitz, Damien Pous, Alan Schmitt, and Jean-Bernard Stefani. Component-oriented programming with sharing: Containment is not ownership. In *Generative Programming and Component Engineering, 4th International Conference, GPCE 2005, Tallinn, Estonia, September 29 - October 1, 2005, Proceedings*, volume 3676 of *Lecture Notes in Computer Science*, pages 389–404. Springer, 2005.

[307] Daniel Hirschkoff, Damien Pous, and Davide Sangiorgi. A correct abstract machine for safe ambients. In *Coordination Models and Languages, 7th International Conference, COORDINATION 2005, Namur, Belgium, April 20-23, 2005, Proceedings*, volume 3454 of *Lecture Notes in Computer Science*, pages 17–32. Springer, 2005.

[308] E. Jeandel. Topological automata. In *Proceedings of STACS 2005*, Lecture Notes in Computer Science 3404. Springer Verlag, 2005.

[309] E. Kaltofen and P. Koiran. On the complexity of factoring bivariate supersparse (lacunary) polynomials. In *Proceedings of ISSAC 2005*. ACM Press, 2005. Distinguished paper award.

[310] P. Koiran. The limit theory of generic polynomials. In *Logic Colloquium 2001*, volume Lecture notes in Logic 20, pages 242–254. Association for Symbolic Logic, 2005.

[311] P. Koiran, V. Nesme, and N. Portier. A quantum lower bound for the query complexity of Simon's problem. In *Proceedings of ICALP 2005*. Springer Verlag, 2005.

[312] Peter Kornerup and Jean-Michel Muller. RN-coding of numbers: definition and some properties. In *Proceedings of the 17th IMACS World Congress on Scientific Computation, Applied Mathematics and Simulation*, Paris, July 2005.

[313] Dieter Kranzlmuller and Laurent Lefèvre. A record and replay mechanism on programmable network card. In *The IASTED International Conference on Parallel and Distributed Computing and Networks (PDCN 2005)*, Innsbruck, Austria, February 2005.

[314] Julien Laganier and Pascale Vicat-Blanc Primet. Hipernet: a decentralized security infrastructure for large scale grid environments. In *6th IEEE/ACM International Conference on Grid Computing (GRID 2005), November 13-14, 2005, Seattle, Washington, USA, Proceedings*, pages 140–147. IEEE, 2005.

[315] Laurent Lefèvre. Heavy and lightweight dynamic network services : challenges and experiments for designing intelligent solutions in evolvable next generation networks. In IEEE Society, editor, *Workshop on Autonomic Communication for Evolvable Next Generation Networks - The 7th International Symposium on Autonomous Decentralized Systems*, pages 738–743, Chengdu, Jiuzhaigou, China, April 2005. ISBN : 0-7803-8963-8.

[316] Laurent Lefèvre and Aweni Saroukou. Active network support for deployment of java-based games on mobile platforms. In IEEE Computer Society, editor, *The First International Conference on Distributed Frameworks for Multimedia Applications (DFMA'2005)*, pages 88–95, Besancon, France, February 2005.

[317] Dino M Lopez-Pacheco and Congduc Pham. Robust transport protocol for dynamic high-speed networks: enhancing the xcp approach. In *Proceedings of IEEE International Conference on Networks*, volume 1, pages 404–409, Kuala Lumpur, Malaysia, November 2005.

[318] Renaud Lottiaux, Benoit Boissinot, Pascal Gallard, Geoffroy Vallee, and Christine Morin. OpenMosix, OpenSSI, and Kerrighed: A comparative study. In *5th IEEE International Symposium on Cluster Computing and the Grid (CCGrid '05)*, volume 2, pages 1016–1023. IEEE Computer Society, May 2005.

[319] Loris Marchal, Pascale Vicat-Blanc Primet, Yves Robert, and Jingdi Zeng. Optimizing network resource sharing in grids. In *IEEE GLOBECOM'05, USA*, November 2005.

[320] Guillaume Melquiond and Sylvain Pion. Formal certification of arithmetic filters for geometric predicates. In *Proceedings of the 17th IMACS World Congress on Computational and Applied Mathematics*, Paris, France, 2005.

[321] Romain Michard, Arnaud Tisserand, and Nicolas Veyrat-Charvillon. Divgen: a divider unit generator. In *Proc. Advanced Signal Processing Algorithms, Architectures and Implementations XV*, volume 5910, page 59100M, San Diego, California, U.S.A., August 2005. SPIE.

[322] Romain Michard, Arnaud Tisserand, and Nicolas Veyrat-Charvillon. Étude statistique de l'activité de la fonction de sélection dans l'algorithme de E-méthode. In *5es journées d'études Faible Tension Faible Consommation (FTFC)*, pages 61–65, Paris, May 2005.

[323] Romain Michard, Arnaud Tisserand, and Nicolas Veyrat-Charvillon. Évaluation de polynômes et de fractions rationnelles sur FPGA avec des opérateurs à additions et décalages en grande base. In *10^{e} SYMPosium en Architectures nouvelles de machines (SYMPA)*, pages 85–96, Le Croisic, April 2005.

[324] Romain Michard, Arnaud Tisserand, and Nicolas Veyrat-Charvillon. Small FPGA polynomial approximations with 3-bit coefficients and low-precision estimations of the powers of x. In *Proc. 16th International Conference on Application-specific Systems, Architectures and Processors (ASAP 2005)*, pages 334–339, Samos, Greece, July 2005. IEEE. Best Paper Award.

[325] Jean-Michel Muller, Arnaud Tisserand, Benoît Dupont de Dinechin, and Christophe Monat. Division by constant for the ST100 DSP microprocessor. In *Proc. 17th Symposium on Computer Arithmetic (ARITH-17)*, pages 124–130, Cape Cod, MA., U.S.A, June 2005. IEEE Computer Society.

[326] V. Poupet. Cellular automata: Real-time equivalence between one-dimensional neighborhoods. In *Proceedings of STACS 2005*, Lecture Notes in Computer Science 3404, page 133144. Springer Verlag, 2005.

[327] Damien Pous. Up-to techniques for weak bisimulation. In *Automata, Languages and Programming, 32nd International Colloquium, ICALP 2005, Lisbon, Portugal, July 11-15, 2005, Proceedings*, volume 3580 of *Lecture Notes in Computer Science*, pages 730–741. Springer, 2005.

[328] P. Vicat-Blanc Primet and J. Zeng. An overlay infrastructure for bulk data transfer in grids. In *Proceedings of 3rd International Workshop on Protocols for Fast Long-Distance Networks (PFLDnet'05)*, Lyon, France, February 2005.

[329] Arne Storjohann and Gilles Villard. Computing the rank and a small nullspace basis of a polynomial matrix. In *Proc. International Symposium on Symbolic and Algebraic Computation, Beijing, China*, pages 309–316. ACM Press, July 2005.

[330] G. Theyssier. How common can be universality in cellular automata? In *Proceedings of STACS 2005*, Lecture Notes in Computer Science 3404, pages 121–132. Springer Verlag, 2005.

[331] Steffen van Bakel, Stéphane Lengrand, and Pierre Lescanne. The language chi: Circuits, computations and classical logic. In *Proc. of Theoretical Computer Science, 9th Italian Conference, ICTCS 2005, Siena, Italy, October 12-14, 2005, Proceedings*, volume 3701 of *Lecture Notes in Computer Science*, pages 81–96. Springer, 2005.

[332] Pascale Vicat-Blanc Primet and Jingdi Zeng. Traffic isolation and network resource sharing for performance control in grids. In *Proceedings of the Joint International Conference on Autonomic and Autonomous Systems and International Conference on Networking and Services (ICAS/ICNS 2005) 0-7695-2450-8/05 2005 IEEE*, April 2005.

[333] Jingdi Zeng and Pascale Vicat-Blanc Primet. An overlay infrastructure for bulk data transfers in grids. In *in Proceedings of the Third International Workshop on Protocol For Very Long Distance Fat Networks, Pfldnet2005*, volume Lyon, February 2005.

2006

[334] Emmanuel Agullo, Abdou Guermouche, and Jean-Yves L'Excellent. A preliminary out-of-core extension of a parallel multifrontal solver. In *EuroPar'06 Parallel Processing*, volume 4128 of *Lecture Notes in Computer Science*, pages 1053–1063, 2006.

[335] Emmanuel Agullo, Abdou Guermouche, and Jean-Yves L'Excellent. Étude mémoire d'une méthode multifrontale parallèle hors-mémoire. In *17e Rencontres Francophones en Parallélisme (RenPar'17)*, pages 220–227, Perpignan, France, 2006.

[336] Abelkader Amar, Raphaël Bolze, Aurélien Bouteiller, Pushpinder Kaur Chouhan, Andréea Chis, Yves Caniou, Eddy Caron, Holly Dail, Benjamin Depardon, Frédéric Desprez, Jean-Sébastien Gay, Gaël Le Mahec, and Alan Su. Diet: New developments and recent results. In Lehner et al. (Eds.), editor, *CoreGRID Workshop on Grid Middleware (in conjunction with EuroPar2006)*, number 4375 in LNCS, pages 150–170, Dresden, Germany, August 28-29 2006. Springer.

[337] Pierre Amiranoff, Albert Cohen, and Paul Feautrier. Beyond iteration vectors: Instancewise relational abstract domains. In *Static Analysis Symposium (SAS'06)*, volume 4134 of *Lecture Notes In Computer Science*, pages 161–180, Seoul, Corea, August 2006. Springer.

[338] Philippe Audebaud and Christine Paulin-Mohring. Proofs of randomized algorithms in Coq. In Tarmo Uustalu, editor, *Mathematics of Program Construction, MPC'2006*, volume 4014 of *LNCS*, pages 49–68, Kuressaare, Estonia, 2006. Springer Verlag.

[339] Olivier Beaumont, Larry Carter, Jeanne Ferrante, Arnaud Legrand, Loris Marchal, and Yves Robert. Centralized versus distributed schedulers for multiple bag-of-task applications. In *International Parallel and Distributed Processing Symposium IPDPS'2006*. IEEE Computer Society Press, 2006.

[340] Olivier Beaumont, Anne-Marie Kermarrec, Loris Marchal, and Étienne Rivière. Voronet, un réseau objet-à-objet sur le modèle petit-monde. In *CFSE'5: Conférence Française sur les Systèmes d'Exploitation*, 2006.

[341] Olivier Beaumont, Loris Marchal, Veronika Rehn, and Yves Robert. FIFO scheduling of divisible loads with return messages under the one-port model. In *HCW'2006, the 15th Heterogeneous Computing Workshop*. IEEE Computer Society Press, 2006.

[342] F. Becker, I. Rapaport, and E. Rémila. Self-assembling classes of shapes, fast and with a minimal number of tiles. In *Proceedings of FSTTCS 2006*, Lecture Notes in Computer Science 4337. Springer Verlag, 2006.

[343] Raphaël Bolze, Eddy Caron, Frédéric Desprez, Georg Hoesch, and Cyril Pontvieux. A monitoring and visualization tool and its application for a network enabled server platform. In M. Gavrilova, editor, *Computational Science and Its Applications - ICCSA 2006*, volume 3984 of *LNCS*, pages 202–213, Glasgow, UK., May 8-11 2006. Springer.

[344] P. Borgnat, N. Larrieu, P. Owezarski, P. Abry, J. Aussibal, L. Gallon, G. Dewaele, N. Nobelis, L. Bernaille, A. Scherrer, Y. Zhang, Y. Labit, and et al. Détection d'attaques de dénis de service par un modèle non gaussien multirésolution. In *Colloque francophone sur l'ingénierie des protocoles (CFIP)*, Tozeur, Tunisie, November 2006.

[345] Florent Bouchez, Alain Darte, Christophe Guillon, and Fabrice Rastello. Register allocation: What does the NP-completeness proof of Chaitin et al. really prove? In *5th Annual Workshop in Duplicating, Deconstructing, and Debunking (WDDD'06), held in conjunction with the 33rd International Symposium on Computer Architecture (ISCA-33)*, Boston, USA, July 2006.

[346] Florent Bouchez, Alain Darte, Christophe Guillon, and Fabrice Rastello. Register allocation: What does the NP-completeness proof of Chaitin et al. really prove? Or revisiting register allocation: Why and how. In *19th International Workshop on Languages and Compilers for Parallel Computing (LCPC'06)*, New Orleans, USA, November 2006.

[347] M. Brahma, M. Chaudier, E. Garcia, J.P. Gelas, H. Guyennet, F. Hantz, L. Lefèvre, P. Lorenz, and H. Tobiet. Temic: a new cooperative platform for industrial tele-maintenance. In *DFMA06 : International Conference on Distributed Framework for Multimedia Applications*, Penang, Malaysia, May 2006.

[348] Eddy Caron, Andréea Chis, Frédéric Desprez, and Alan Su. Plug-in scheduler design for a distributed grid environment. In ACM/IFIP/USENIX, editor, *4th International Workshop on Middleware for Grid Computing - MGC 2006*, Melbourne, Australia, November 2006. In conjunction with ACM/IFIP/USENIX 7th International Middleware Conference 2006.

[349] Eddy Caron, Pushpinder Kaur Chouhan, and Holly Dail. GoDIET: A deployment tool for distributed middleware on grid'5000. In IEEE, editor, *EXPGRID workshop. Experimental Grid Testbeds for the Assessment of Large-Scale Distributed Apllications and Tools. In conjunction with HPDC-15.*, pages 1–8, Paris, France, June 2006.

[350] Eddy Caron, Frédéric Desprez, Charles Fourdrignier, Franck Petit, and Cédric Tedeschi. A repair mechanism for fault-tolerance for tree-structured peer-to-peer systems. In Yves Robert, Manish Parashar, Ramamurthy Badrinath, and Viktor K. Prasanna, editors, *HiPC'2006. 13th International Conference on High Performance Computing.*, volume 4297 of *LNCS*, pages 171–182, Bangalore. India, December 18-21 2006. Springer-Verlag Berlin Heidelberg.

[351] Eddy Caron, Frédéric Desprez, and Cédric Tedeschi. A dynamic prefix tree for the service discovery within large scale grids. In A. Montresor, A. Wierzbicki, and N. Shahmehri, editors, *The Sixth IEEE International Conference on Peer-to-Peer Computing, P2P2006*, pages 106–113, Cambridge, UK., September 2006. IEEE.

[352] Eddy Caron, Charles Fourdrignier, Franck Petit, and Cédric Tedeschi. Mécanisme de réparations pour un système P2P de découverte de services. In *Perpi'2006 - Conférences conjointes RenPar'17 / SympA'2006 / CFSE'5 / JC'2006*, pages 252–259, Canet en Roussillon, October 4-6 2006.

[353] Hugo Carrao, Mário Caetano, and Paulo Gonçalves. Meris based land cover characterization: a comparative study. In *Proceedings of the ASPRS international conference*, Reno (NV, US), May 2006.

[354] Hugo Carrao, Paulo Gonçalves, and Mário Caetano. Use of intra-annual satellite imagery time-series for land cover characterization purpose. In *Proceedings of the 26th EARSeL Symposium*, Warsaw (Poland), 2006.

[355] C. Chaudet and I. Guérin Lassous. Etat des lieux sur la qualité de service dans les réseaux ad hoc. In *Colloque Francophone sur l'Ingénierie des Procoles (CFIP)*, Tozeur, Tunis, November 2006.

[356] Francisco Cháves and Marc Daumas. A library to Taylor models for PVS automatic proof checker. In *Proceedings of the NSF workshop on reliable engineering computing*, pages 39–52, Savannah, Georgia, 2006.

[357] Guillaume Chelius, Antoine Fraboulet, and Eric Fleury. Demonstration of Worldsens: A fast prototyping and performance evaluation tool for wireless sensor network applications & protocols. In *Second International Workshop on Multi-hop Ad Hoc Networks: From Theory to Reality (REALMAN)*, pages 131–133, Firenze, Italia, May 2006. ACM.

[358] Hadda Cherroun, Alain Darte, and Paul Feautrier. Scheduling under resource constraints using dis-equalities. In *Design Automation and Test in Europe (DATE'06)*, pages 1067–1072, Munich, Germany, March 2006.

[359] Chafika Chettaoui, Franck Delaplace, Pierre Lescanne, Mun'delanji Vestergaard, and René Vestergaard. Rewriting game theory as a foundation for state-based models of gene regulation. In *Computational Methods in Systems Biology, International Conference, CMSB 2006, Trento, Italy, October 18-19, 2006, Proceedings*, volume 4210 of *Lecture Notes in Computer Science*, pages 257–270. Springer, 2006.

[360] Sylvain Collange, Jérémie Detrey, and Florent de Dinechin. Floating point or LNS: choosing the right arithmetic on an application basis. In *9th Euromicro Conference on Digital System Design: Architectures, Methods and Tools (DSD'2006)*, pages 197–203, Dubrovnik, Croatia, August 2006. IEEE.

[361] B. David, R. Chalon, O. Delotte, G. Masserey, and M. Imbert. ORCHESTRA: formalism to express mobile cooperative applications. In *Proceedings of the 12th International Workshop, CRIWG*, volume 4154 of *Lecture Notes in Computer Science*, 2006. External publication.

[362] Florent de Dinechin, Christoph Quirin Lauter, and Guillaume Melquiond. Assisted verification of elementary functions using Gappa. In *Proceedings of the 2006 ACM Symposium on Applied Computing*, pages 1318–1322, 2006.

[363] Florent de Dinechin and Sergey Maidanov. Software techniques for perfect elementary functions in floating-point interval arithmetic. In *Real Numbers and Computers*, July 2006.

[364] Florent de Dinechin, Eric McIntosh, and Franck Schmidt. Massive tracking on heterogeneous platforms. In *9th International Computational Accelerator Physics Conference (ICAP)*, October 2006.

[365] Jérémie Detrey. TutoRISC : un RISC dans mon FPGA. In *9es Journées Pédagogiques du CNFM*, pages 153–158, Saint-Malo, France, November 2006.

[366] Jérémie Detrey and Florent de Dinechin. Opérateurs trigonométriques en virgule flottante sur FPGA. In *RenPar'17, SympA'2006, CFSE'5 et JC'2006*, pages 96–105, Perpignan, France, October 2006.

[367] Sandeep Dey and Nicolas Schabanel. Customized newspaper problem: Data broadcast with dependancies. In *Proc. of the 7th Latin American Theoretical Informatics conference (LATIN)*, volume LNCS 3887, pages 362–373, Mar. 2006.

[368] Philippe Duchon, Nicolas Hanusse, Emmanuelle Lebhar, and Nicolas Schabanel. Towards small world emergence. In *Proc. of ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, pages 225–232, Jul. 2006.

[369] Wayne Eberly, Mark Giesbrecht, Pascal Giorgi, Arne Storjohann, and Gilles Villard. Solving sparse integer linear systems. In *Proc. International Symposium on Symbolic and Algebraic Computation, Genova, Italy*, pages 63–70. ACM Press, July 2006.

[370] Milos Ercegovac and Jean-Michel Muller. Arithmetic processor for solving tri-diagonal systems of linear equations. In *proc. 40th Conference on Signals, Systems and Computers*, Pacific Grove, CA, November 2006.

[371] Nazim Fatès, Damien Regnault, Nicolas Schabanel, and Éric Thierry. Asynchronous behavior of double-quiescent elementary cellular automata. In *Proc. of the 7th Latin American Theoretical Informatics conference (LATIN)*, volume LNCS 3887, pages 455–466, Mar. 2006.

[372] Antoine Fraboulet, Guillaume Chelius, and Eric Fleury. Worldsens: System tools for embedded sensor networks. In *Real-Time Systems Symposium (RTSS 2006) (Work in Progress)*, Rio de Janeiro, Brasil, December 2006. IEEE.

[373] Jean-Sébastien Gay and Yves Caniou. Simbatch : une API pour la simulation et la prédiction de performances de systèmes batch. In *17ieme Rencontres Francophones du Parallélisme, des Architectures et des Systèmes*, pages 180–187, Perpignan, October 2006.

[374] Ryan Glabb, Laurent Imbert, Graham Jullien, Arnaud Tisserand, and Nicolas Veyrat-Charvillon. Multi-mode operator for SHA-2 hash functions. In *Proc. International Conference on Engineering of Reconfigurable Systems and Algorithms (ERSA)*, pages 207–210, Las Vegas, Nevada, U.S.A., June 2006.

[375] Paulo Gonçalves, Hugo Carrao, and Mário Caetano. Parametric model for intra-annual reflectance time series. In *Proceedings of the 26th EARSeL Symposium*, Warsaw (Poland), 2006.

[376] Philippe Grosse, Yves Durand, and Paul Feautrier. Power modeling of a NoC based design for high-speed telecommunication systems. In *16th PATMOS Workshop - International Workshop on Power And Timing Modeling, Optimization and Simulation*, Montpellier, France, September 2006.

[377] E. Kaltofen and P. Koiran. Finding small degree factors of multivariate supersparse (lacunary) polynomials over algebraic number fields. In *Proceedings of ISSAC 2006*. ACM Press, 2006.

[378] P. Koiran and S. Perifel. Valiant's model: from exponential sums to exponential products. In *Proceedings of MFCS 2006*, Lecture Notes in Computer Science 4708. Springer Verlag, 2006.

[379] Olivier Laurent and Lorenzo Tortora de Falco. Obsessional cliques: a semantic characterization of bounded time complexity. In *Proceedings of the twenty-first annual IEEE symposium on Logic In Computer Science*, pages 179–188. IEEE Computer Society Press, August 2006. External publication.

[380] Laurent Lefèvre and Jean-Patrick Gelas. Towards interplanetary grids. In *Workshop on Next Generation Communication Infrastructure for Deep-Space Communications held in conjunction with the Second International Conference on Space Mission Challenges for Information Technology (SMC-IT)*, Pasadena, California, July 2006.

[381] Laurent Lefèvre and Jean-Marc Pierson. Just in time entertainment deployment on mobile platforms. In *ICIW'06 : International Conference on Internet and Web Applications and Services*, Guadeloupe, French Caribbean, February 2006.

[382] Laurent Lefèvre and Paul Roe. Improving the flexibility of active grids through web services. In Australian Computer Society, editor, *Fourth Australasian Symposium on Grid Computing and e-Research (AusGrid2006)*, volume 28, pages 3–8, Hobart, Australia, January 2006. ISBN 1-920-68236-8.

[383] Arnaud Legrand, Alan Su, and Frédéric Vivien. Minimizing the stretch when scheduling flows of biological requests. In *SPAA '06: Proceedings of the eighteenth annual ACM symposium on Parallelism in algorithms and architectures*, pages 103–112, Cambridge, Massachusetts, USA, 2006. ACM Press.

[384] Dino M. Lopez Pacheco, Cong-Duc Pham, and Laurent Lefèvre. Xcp-i : explicit control protocol for heterogeneous inter-networking of high-speed networks. In *Globecom 2006*, San Francisco, California, USA, November 2006.

[385] Dino M. Lopez Pacheco, Cong-Duc Pham, and Laurent Lefèvre. Xcp-i: explicit control protocol pour l'interconnexion de réseaux haut-débit hétérogènes. In *CFIP 2006: Colloque Francophone sur l'Ingénierie des Protocoles*, Tozeur, Tunisia, October 2006.

[386] Dino M Lopez-Pacheco and Congduc Pham. Enabling large data transfers on dynamic, very high-speed network infrastructures. In *Proceedings of ICN/ICONS/MCL 2006*, Mauritius, april 2006.

[387] G. Malod and N. Portier. Characterizing valiant's algebraic complexity classes. In *Proceedings of MFCS 2006 (31st International Symposium on Mathematical Foundations of Computer Science)*, Lecture Notes in Computer Science 4162. Springer Verlag, 2006.

[388] Loris Marchal, Pascale Primet, Yves Robert, and Jingdi Zeng. Optimal bandwidth sharing in grid environment. In *HPDC'2006, the 15th International Symposium on High Performance Distributed Computing*. IEEE Computer Society Press, 2006.

[389] Romain Michard, Arnaud Tisserand, and Nicolas Veyrat-Charvillon. New identities and transformations for hardware power operators. In *Proc. Advanced Signal Processing Algorithms, Architectures and Implementations XVI*, San Diego, California, U.S.A., August 2006. SPIE.

[390] Romain Michard, Arnaud Tisserand, and Nicolas Veyrat-Charvillon. Optimisation d'opérateurs arithmétiques matériels à base d'approximations polynomiales. In *11^{e} SYMPosium en Architectures nouvelles de machines (SYMPA)*, Perpignan, October 2006.

[391] Jean-Michel Muller. Generating functions at compile-time. In *Proc. 40th Conference on Signals, Systems and Computers*, Pacific Grove, CA, November 2006.

[392] Pablo Neira Ayuso, Laurent Lefèvre, and Rafael M. Gasca. High availability support for the design of stateful networking equipments. In *ARES'06 : The First International Conference on Availability, Reliability and Security*, Vienna, Austria, April 2006.

[393] B. Nouvel and E. Rémila. Incremental and transitive discretized rotations. In *Proceedings of IWCIA 2006*, Lecture Notes in Computer Science 4040. Springer Verlag, 2006.

[394] Stéphane Operto, Jean Virieux, Patrick Amestoy, Luc Giraud, and Jean-Yves L'Excellent. 3D frequency-domain finite-difference modeling of accoustic wave propagation using a masively parallel direct solver: a feasable study. In *The Society of Exploration Geophysicists (SEG 2006), New Orleans (USA), 01/10/06-06/10/06*, pages 2265–2269, 2006.

[395] Jean-François Pineau, Yves Robert, and Frédéric Vivien. The impact of heterogeneity on master-slave on-line scheduling. In *HCW'2006, the 15th Heterogeneous Computing Workshop*. IEEE Computer Society Press, 2006.

[396] Jean-François Pineau, Yves Robert, and Frédéric Vivien. Off-line and on-line scheduling on heterogeneous master-slave platforms. In *PDP'2006, 14th Euromicro Workshop on Parallel, Distributed and Network-based Processing*. IEEE Computer Society Press, 2006.

[397] André Pinheiro, Paulo Gonçalves, Hugo Carrao, and Mário Caetano. A toolbox for multi-temporal analysis of satellite imagery. In *Proceedings of the 26th EARSeL Symposium*, Warsaw (Poland), 2006.

[398] Damien Pous. On bisimulation proofs for the analysis of distributed abstract machines. In *Trustworthy Global Computing, Second Symposium, TGC 2006, Lucca, Italy, November 7-9, 2006, Revised Selected Papers*, volume 4661 of *Lecture Notes in Computer Science*, pages 150–166. Springer, 2006.

[399] Damien Pous. Weak bisimulation up to elaboration. In *CONCUR 2006 - Concurrency Theory, 17th International Conference, CONCUR 2006, Bonn, Germany, August 27-30, 2006, Proceedings*, volume 4137 of *Lecture Notes in Computer Science*, pages 390–405. Springer, 2006.

[400] R. Razafindralambo, I. Guérin Lassous, L. Iannone, and S. Fdida. Aggrégation dynamique de paquets pour résoudre l'anomalie de performance des réseaux sans fil ieee 802.11. In *Colloque Francophone sur l'Ingénierie des Procoles (CFIP)*, Tozeur, Tunis, November 2006.

[401] T. Razafindralambo, I. Guérin Lassous, L. Iannone, and S. Fdida. Dynamic packet aggregation to solve performance anomaly in 802.11 wireless networks. In *9th ACM/IEEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, Terremolinos, Spain, October 2006.

[402] D. Regnault. Abrupt behaviour changes in cellular automata under asynchronous dynamics. In *Proceedings of 2nd European Conference on Complex Systems (ECCS)*, Oxford, UK, 2006.

[403] A. Romashchenko. Reliable computations based on locally decodable codes. In *Proceedings of STACS 2006*, Lecture Notes in Computer Science 3884. Springer Verlag, 2006.

[404] Jean-Baptiste Rouquier and Michel Morvan. Coalescing cellular automata. In Vassil N. Alexandrov, G. Dick van Albada, Peter M. A. Sloot, and Jack Dongarra, editors, *Proceedings of the 6th International Conference on Computational Science (ICCS)*, volume 3993 of *Lectures Notes in Computer Science*, pages 321–328. Springer, May 2006.

[405] Silvius Rus, Guobin He, Christophe Alias, and Lawrence Rauchwerger. Region array SSA. In *15th ACM/IEEE International Conference on Parallel Architectures and Compilation Techniques (PACT'06)*, pages 43–52, Seattle, WA, USA, September 2006. ACM.

[406] Silvius Rus, Guobin He, Christophe Alias, and Lawrence Rauchwerger. Region array SSA. In *International ACM/IEEE Conference on Parallel Architectures and Compilation Techniques (PACT'06)*, September 2006. External publication.

[407] C. Sarr, C. Chaudet, G. Chelius, and I. Guérin Lassous. Improving accuracy in available bandwidth estimation for 802.11-based ad hoc networks. In *Third IEEE International Conference on Mobile Ad-hoc and Sensor Systems (MASS)*, Hong-Kong, China, December 2006.

[408] Antoine Scherrer, Antoine Fraboulet, and Tanguy Risset. Automatic phase detection for stochastic on-chip traffic generation. In *Proceedings of the 4th International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS)*, pages 88–93, seoul, South Corea, October 2006. ACM Press.

[409] Antoine Scherrer, Antoine Fraboulet, and Tanguy Risset. A generic multi-phase on-chip traffic generation environment. In *IEEE 17th International Conference on Application-Specific Systems, Architectures and Processors (ASAP'06)*, Steamboat Springs, Colorado, USA, September 2006.

[410] Antoine Scherrer, Nicolas Larrieu, Pierre Borgnat, Philippe Owezarski, and Patrice Abry. Non Gaussian and long memory statistical modeling of internet traffic. In *4th International Workshop on Internet Performance, Simulation, Monitoring and Measurement (IPS MOME)*, Salzbourg, Austria, March 2006.

[411] Antoine Scherrer, Nicolas Larrieu, Pierre Borgnat, Philippe Owezarski, and Patrice Abry. Une caractérisation non gaussienne et longue mémoire du trafic internet et de ses anomalies. In *5th Conference on Security and Network Architectures (SAR)*, Seignosse, France, June 2006.

[412] Cédric Tedeschi. Découverte de services pour les grilles de calcul dynamiques large échelle. In *Perpi'2006 - Conférences conjointes RenPar'17 / SympA'2006 / CFSE'5 / JC'2006*, pages 52–59, Canet en Roussillon, France, October 4-6 2006.

[413] Pascale Vicat-Blanc Primet, R. Takano, Y. Kodama, T. Kudoh, Olivier Gluck, and C. Otal. Large scale gigabit emulated testbed for grid transport evaluation. In *Proceedings of The Fourth International Workshop on Protocols for Fast Long-Distance Networks, PFLDnet'2006*, Nara, Japan, February 2006.

2007

[414] Emmanuel Agullo, Abdou Guermouche, and Jean-Yves L'Excellent. On reducing the I/O volume in a sparse out-of-core solver. In *HiPC'07 14th International Conference On High Performance Computing*, number 4873 in Lecture Notes in Computer Science, pages 47–58, Goa, India, December 17-20, 2007.

[415] Christophe Alias, Fabrice Baray, and Alain Darte. Bee+cl@k: An implementation of lattice-based memory reuse in the source-to-source translator ROSE. In *International ACM Conference on Languages, Compilers, and Tools for Embedded Systems (LCTES'07)*, June 2007. External publication.

[416] Christophe Alias, Fabrice Baray, and Alain Darte. An implementation of lattice-based memory reuse in the source-to-source translator ROSE. In *ACM SIGPLAN/SIGBED Conference on Languages, Compilers, and Tools for Embedded Systems (LCTES'07)*, pages 73–82, San Diego, USA, June 2007.

[417] Gabriel Antoniu, Eddy Caron, Frédéric Desprez, Aurélia Fèvre, and Mathieu Jan. Towards a transparent data access model for the aGridRPC paradigm. In S. Aluru et al. (Eds), editor, *HiPC'2007. 14th International Conference on High Performance Computing.*, number 4873 in LNCS, pages 269–284, Goa. India, December 2007. Springer Verlag Berlin Heidelberg.

[418] Hrachya Astsatryan, Michel Daydé, Aurélie Hurault, Marc Pantel, and Eddy Caron. On defining a web interface for linear algebra tasks over computational grids. In *International Conference on Computer Science and Information Technologies (CSIT'07)*, Yerevan (Arménie), September 2007.

[419] Narjess Ayari, Denis Barbaron, Laurent Lefèvre, and Pascale Vicat-Blanc Primet. Sara: A session aware infrastructure for high performance next generation cluster-based servers. In *ATNAC 2007 : Australasian Telecommunication Networks and Applications Conference*, pages 648–655, Christchurch, New Zealand, December 2007.

[420] Narjess Ayari, Denis Barbaron, Laurent Lefèvre, and Pascale Vicat-Blanc Primet. Session awareness issues for next-generation cluster-based network load balancing frameworks. In *AICCSA07 : ACS/IEEE International Conference on Computer Systems and Applications*, pages 180–186, Amman, Jordan, May 2007.

[421] Narjess Ayari, Denis Barbaron, Laurent Lefèvre, and Pascale Vicat-Blanc Primet. T2cp-ar: A system for transparent tcp active replication. In *AINA-07 : The IEEE 21st International Conference on Advanced Information Networking and Applications*, pages 648–655, Niagara Falls, Canada, May 2007.

[422] P. Baillot, P. Coppola, and U. Dal Lago. Light Logics and Optimal Reduction: Completeness and Complexity. In *Proceedings of Symposium on Logic in Computer Science (LICS'07)*, pages 421–430. IEEE Computer Society, 2007. External publication.

[423] Olivier Beaumont, Anne-Marie Kermarrec, Loris Marchal, and Etienne Riviére. Voronet: A scalable object network based on voronoi tessellations. In *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS 2007)*, pages 1–10. IEEE Computer Society Press, 2007.

[424] Anne Benoit, Veronika Rehn, and Yves Robert. Strategies for replica placement in tree networks. In *HCW'2007, the 16th Heterogeneous Computing Workshop*. IEEE Computer Society Press, 2007.

[425] Anne Benoit, Veronika Rehn-Sonigo, and Yves Robert. Impact of QoS on replica placement in tree networks. In *ICCS'2007, the 7th 2007 International Conference on Computational Science*, LNCS 4487, pages 366–373. Springer Verlag, 2007.

[426] Anne Benoit, Veronika Rehn-Sonigo, and Yves Robert. Multi-criteria scheduling of pipeline workflows. In *HeteroPar'2007: International Conference on Heterogeneous Computing, jointly published with Cluster'2007*. IEEE Computer Society Press, 2007.

[427] Anne Benoit and Yves Robert. Complexity results for throughput and latency optimization of replicated and data-parallel workflows. In *HeteroPar'2007: International Conference on Heterogeneous Computing, jointly published with Cluster'2007*. IEEE Computer Society Press, 2007.

[428] Anne Benoit and Yves Robert. Mapping pipeline skeletons onto heterogeneous platforms. In *ICCS'2007, the 7th International Conference on Computational Science*, LNCS 4487, pages 591–598. Springer Verlag, 2007.

[429] Jean-Luc Beuchat, Nicolas Brisebarre, Jérémie Detrey, and Eiji Okamoto. Arithmetic operators for pairing-based cryptography. In *Cryptographic Hardware and Embedded Systems – CHES 2007*, number 4727 in Lecture Notes in Computer Science, pages 239–255. Springer, 2007. Best paper award.

[430] Jean-Luc Beuchat, Nicolas Brisebarre, Masaaki Shirase, Tsuyoshi Takagi, and Eiji Okamoto. A coprocessor for the final exponentiation of the η_T pairing in characteristic three. In *Proceedings of Waifi 2007*, volume 4547 of *Lecture Notes in Computer Science*, pages 25–39. Springer, 2007.

[431] Julien Bigot and Christian Pérez. Enabling collective communications between components. In *CompFrame'07: Proceedings of the 2007 symposium on Component and framework technology in high-performance and scientific computing*, pages 121–130, New York, NY, USA, 21-22 October 2007. ACM Press.

[432] O. Bodini, T. Fernique, and E. Rémila. A characterization of flip-accessibility for lozenge tilings of the whole plane. In *Proceedings of LATA 2007*, 2007.

[433] O. Bodini, T. Fernique, and E. Rémila. Characterizations of flip-accessibility for domino tilings of the whole plane. In *Proceedings of FPSAC 2007*, 2007.

[434] Alin Bostan, Claude-Pierre Jeannerod, and Éric Schost. Solving Toeplitz- and Vandermonde-like linear systems with large displacement rank. In *Proc. International Symposium on Symbolic and Algebraic Computation (ISSAC 2007), Waterloo, Canada*, pages 33–40. ACM Press, August 2007.

[435] Florent Bouchez, Alain Darte, and Fabrice Rastello. On the complexity of register coalescing. In *International Symposium on Code Generation and Optimization (CGO'07)*, pages 102–114. IEEE Computer Society Press, March 2007. **Best paper award**.

[436] Florent Bouchez, Alain Darte, and Fabrice Rastello. On the complexity of spill everywhere under SSA form. In *ACM SIGPLAN/SIGBED Conference on Languages, Compilers, and Tools for Embedded Systems (LCTES'07)*, pages 103–112, San Diego, USA, June 2007.

[437] A. Bouillard, B. Gaujal, S. Lagrange, and E. Thierry. Optimal routing for end-to-end guarantees: the price of multiplexing. In *Proceedings of ValueTools'2007*, 2007. Best paper award.

[438] Nicolas Brisebarre and Sylvain Chevillard. Efficient polynomial L^∞-approximations. In *Proceedings of the 18th IEEE Symposium on Computer Arithmetic (ARITH-18)*, pages 169–176. IEEE, 2007.

[439] Nicolas Brisebarre and Guillaume Hanrot. Floating-point L^2 approximations to functions. In *Proceedings of the 18th IEEE Symposium on Computer Arithmetic (ARITH-18)*, pages 177–184. IEEE, 2007.

[440] Yves Caniou, Eddy Caron, Hélène Courtois, Benjamin Depardon, and Romain Teyssier. Cosmological simulations using grid middleware. In *Fourth High-Performance Grid Computing Workshop (HPGC'07)*, Long Beach, California, USA, March 2007. IEEE.

[441] Eddy Caron, Frédéric Desprez, Franck Petit, and Cédric Tedeschi. Snap-stabilizing Prefix Tree for Peer-to-peer Systems. In *9th International Symposium on Stabilization, Safety, and Security of Distributed Systems*, volume 4838 of *Lecture Notes in Computer Science*, pages 82–96, Paris, France, November 2007. Springer Verlag Berlin Heidelberg.

[442] Hugo Carrao, Paulo Gonçalves, and Mário Caetano. Land cover characterization through parametric modeling of intra-annual reflectance time series: a comparative study with meris data. In *SPIE Europe Symposium on Remote Sensing*, Firenze (Italy), Sept. 2007.

[443] Guillaume Chelius, Antoine Fraboulet, and Eric Fleury. Worldsens: A fast and accurate development framework for sensor network applications. In *The 22nd Annual ACM Symposium on Applied Computing (SAC 2007)*, Seoul, Korea, March 2007. ACM.

[444] Bin Bin Chen and Pascale Vicat-Blanc Primet. Supporting bulk data transfers of high-end applications with guaranteed completion time. In IEEE, editor, *IEEE ICC2007 International Conference on Computer Communication*, 2007.

[445] Binbin Chen and Pascale Vicat-Blanc Primet. Scheduling bulk data transfers in grid networks. In IEEE, editor, *IEEE CCGRID2007*, 2007.

[446] Yu Chen and Eric Fleury. A distributed policy scheduling for wireless sensor networks. In *26th Annual IEEE Conference on Computer Communications IEEE INFOCOM 2007*, Anchorage États-Unis d'Amérique, 2007. External publication.

[447] Hadda Cherroun and Paul Feautrier. An exact resource constrained-scheduler using graph coloring technique. In *AICCSA'07: The 5th ACS/IEEE International Conference on Computer Systems and Applications*, pages 554–561. IEEE Computer Society, May 2007. **Best paper award**.

[448] Sylvain Chevillard and Christoph Quirin Lauter. A certified infinite norm for the implementation of elementary functions. In *Seventh International Conference on Quality Software (QSIC 2007)*, pages 153–160. IEEE, 2007.

[449] Alain Darte and Clément Quinson. Scheduling register-allocated codes in user-guided high-level synthesis. In *ASAP'07: The 18th IEEE International Conference on Application-specific Systems, Architectures and Processors*, pages 554–561. IEEE Computer Society, July 2007.

[450] Frédéric Desprez and Antoine Vernois. Semi-Static Algorithms for Data Replication and Scheduling Over the Grid. In *IEEE 3rd International Conference on Intelligent Computer Communication and Processing, workshop on Grid Computing*, Cluj-Napoca, Romania, September 2007.

[451] Jérémie Detrey and Florent de Dinechin. Floating-point trigonometric functions for FPGAs. In *International Conference on Field-Programmable Logic and Applications*, pages 29–34. IEEE, August 2007.

[452] Jérémie Detrey, Florent de Dinechin, and Xavier Pujol. Return of the hardware floating-point elementary function. In *18th Symposium on Computer Arithmetic (ARITH-18)*, pages 161–168. IEEE, June 2007.

[453] Jack DiGiovanna, Loris Marchal, Prapaporn Rattanatamrong, Ming Zhao, Shalom Darmanjian, Babak Mahmoudi, Justin Sanchez, José Príncipe, Linda Hermer-Vazquez, Renato Figueiredo, and José Fortes. Towards real-time distributed signal modeling for brain machine interfaces. In *Proceedings of Dynamic Data Driven Application Systems (workshop of ICCS)*, volume 4487 of *LNCS*, pages 964–971. Springer Verlag, 2007.

[454] Dinil Mon Divakaran and Pascale Vicat-Blanc Primet. Channel provisioning and flow scheduling in grid overlay networks. In *Web Proceedings of the HiPC 2007 poster*, December 2007.

[455] Dinil Mon Divakaran and Pascale Vicat-Blanc Primet. Channel provisioning in grid overlay networks (short paper). In *Workshop on IP QoS and Traffic Control*, December 2007.

[456] Wayne Eberly, Mark Giesbrecht, Pascal Giorgi, Arne Storjohann, and Gilles Villard. Faster inversion and other black box matrix computation using efficient block projections. In *Proc. International Symposium on Symbolic and Algebraic Computation, Waterloo, Canada*, pages 143–150. ACM Press, August 2007.

[457] Thomas Ehrhard and Olivier Laurent. Interpreting a finitary pi-calculus in differential interaction nets. In Luis Caires and Vasco T. Vasconcelos, editors, *Concurrency Theory (CONCUR '07)*, volume 4703 of *Lecture Notes in Computer Science*, pages 333–348. Springer, September 2007. External publication.

[458] Milos Ercegovac and Jean-Michel Muller. Complex multiply-add and other related operators. In *Proceedings of SPIE Conf. Advanced Signal Processing Algorithms, Architectures and Implementation XVII*, August 2007.

[459] Milos Ercegovac and Jean-Michel Muller. A hardware-oriented method for evaluating complex polynomials. In *Proceedings of 18th IEEE Conference on Application-Specific Systems, Architectures and Processors (ASAP 2007)*. IEEE, July 2007.

[460] Lionel Eyraud-Dubois, Arnaud Legrand, Martin Quinson, and Frédéric Vivien. A first step towards automatically building network representations. In *Proceedings of Euro-Par 2007*, volume 4641 of *LNCS*, pages 160–169, 2007.

[461] O. Finkel and D. Lecomte. There exist some ω-powers of any Borel rank. In *Proceedings of CSL 2007*, Lecture Notes in Computer Science 4646. Springer Verlag, 2007.

[462] U. Flarup, P. Koiran, and L. Lyaudet. On the expressive power of planar perfect matching and permanents of bounded treewidth matrices. In *Proceedings of ISAAC 2007*, Lecture Notes in Computer Science 4835. Springer, 2007.

[463] Nicolas Fournel, Antoine Fraboulet, Guillaume Chelius, Eric Fleury, Bruno Allard, and Olivier Brevet. Worldsens: Embedded sensor network application development and deployment. In *26th Annual IEEE Conference on Computer Communications (Infocom)*, Anchorage, Alaska, USA, May 2007. IEEE.

[464] Nicolas Fournel, Antoine Fraboulet, and Paul Feautrier. eSimu : a fast and accurate energy consumption simulator for embedded system. In *IEEE International Workshop: From Theory to Practice in Wireless Sensor Networks*, Helsinki, Finland, June 2007.

[465] Nicolas Fournel, Antoine Fraboulet, and Paul Feautrier. Fast and instruction accurate embedded systems energy characterization using non-intrusive measurements. In *PATMOS Workshop - International Workshop on Power And Timing Modeling, Optimization and Simulation*, Göteborg, Sweden, September 2007.

[466] Nicolas Fournel, Marine Minier, and Stéphane Ubéda. Survey and benchmark of stream ciphers for wireless sensor networks. In *Workshop in Information Security Theory and Practices (WISTP 2007)*, Heraklion, Crete, Greece, May 2007.

[467] Antoine Fraboulet, Guillaume Chelius, and Eric Fleury. Worldsens: Development and prototyping tools for application specific wireless sensors networks. In *IPSN'07 Track on Sensor Platforms, Tools and Design Methods (SPOTS)*, Cambridge, Massachusetts, USA., April 2007. ACM.

[468] Antoine Fraboulet, Guillaume Chelius, and Eric Fleury. Worldsens: development and prototyping tools for application specific wireless sensors networks. In ACM, editor, *6th ACM/IEEE international conference on Information Processing in Sensor Networks (IPSN 2007)*, pages 176–185, Cambridge États-Unis d'Amérique, 2007. ACM/IEEE. C.: Computer Systems Organization/C.2: COMPUTER-COMMUNICATION NETWORKS/C.2.2: Network Protocols, D.: Software/D.2: SOFTWARE ENGINEERING/D.2.5: Testing and Debugging/D.2.5.1: Debugging aids, D.: Software/D.2: SOFTWARE ENGINEERING/D.2.5: Testing and Debugging/D.2.5.6: Monitors, D.: Software/D.2: SOFTWARE ENGINEERING/D.2.6: Programming Environments/D.2.6.1: Integrated environments. External publication.

[469] Matthieu Gallet, Yves Robert, and Frédéric Vivien. Scheduling communication requests traversing a switch: complexity and algorithms. In *PDP'2007, 15th Euromicro Workshop on Parallel, Distributed and Network-based Processing*, pages 39–46. IEEE Computer Society Press, 2007.

[470] Matthieu Gallet, Yves Robert, and Frédéric Vivien. Scheduling multiple divisible loads on a linear processor network. In *ICPADS'2007, the 13th International Conference on Parallel and Distributed Systems*, 2007.

[471] Paulo Gonçalves, Patrice Abry, Gabriel Rilling, and Patrick Flandrin. Fractal dimension estimation: empirical mode decomposition versus wavelets. In *IEEE Int. Conf. on Acoust. Speech and Sig. Proc.*, Honolulu, Hawaii (US), April 2007.

[472] Daniel Grund and Sebastian Hack. A fast cutting-plane algorithm for optimal coalescing. In Shriram Krishnamurthi and Martin Odersky, editors, *Compiler Construction 2007*, Lecture Notes In Computer Science, Braga, Portugal, March 2007. Springer. **Best paper award**.

[473] Romaric Guillier, Ludovic Hablot, Yuetsu Kodama, Tomohiro Kudoh, Fumihiro Okazaki, Ryousei Takano, Pascale Vicat-Blanc Primet, and Sebastien Soudan. A study of large flow interactions in high-speed shared networks with grid5000 and gtrcnet-10 instruments. In *PFLDnet 2007*, Feb. 2007.

[474] Romaric Guillier, Sebastien Soudan, and Pascale Vicat-Blanc Primet. Tcp variants and transfer time predictability in very high speed networks. In *Infocom 2007 High Speed Networks Workshop*, May 2007.

[475] Ludovic Hablot, Olivier Gluck, Jean-Christophe Mignot, Stéphane Genaud, and Pascale Vicat-Blanc Primet. Comparison and tuning of mpi implementation in a grid context. In *In Proceedings of 2007 IEEE International Conference on Cluster Computing (CLUSTER)*, pages 458–463, September 2007.

[476] Guillaume Hanrot, Vincent Lefèvre, Damien Stehlé, and Paul Zimmermann. Worst cases of a periodic function for large arguments. In *Proceedings of the 18th IEEE Symposium on Computer Arithmetic (ARITH-18)*, pages 133–140. IEEE, 2007.

[477] Guillaume Hanrot and Damien Stehlé. Improved analysis of Kannan's shortest lattice vector algorithm (extended abstract). In *Proceedings of Crypto 2007*, volume 4622 of *LNCS*, pages 170–186. Springer-Verlag, 2007.

[478] Daniel Hirschkoff and Damien Pous. A Distribution Law for CCS and a New Congruence Result for the Pi-Calculus. In *Foundations of Software Science and Computational Structures, 10th International Conference, FOSSACS 2007, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2007, Braga, Portugal, March 24-April 1, 2007, Proceedings*, volume 4423 of *Lecture Notes in Computer Science*, pages 228–242. Springer, 2007.

[479] Michel Hirschowitz, André Hirschowitz, and Tom Hirschowitz. A theory for game theories. In *FSTTCS 2007: Foundations of Software Technology and Theoretical Computer Science, 27th International Conference, New Delhi, India, December 12-14, 2007, Proceedings*, volume 4855 of *Lecture Notes in Computer Science*, pages 192–203. Springer, 2007.

[480] Claude-Pierre Jeannerod, Hervé Knochel, Christophe Monat, and Guillaume Revy. Faster floating-point square root for integer processors. In *IEEE Symposium on Industrial Embedded Systems (SIES'07)*, 2007.

[481] S. Khalfallah, C. Sarr, and I. Guérin Lassous. Dynamic bandwidth management for multihop wireless ad hoc networks. In *VTC-Spring*, Dublin, Ireland, April 2007.

[482] P. Koiran and S. Perifel. VPSPACE and a transfer theorem over the complex field. In *Proceedings of MFCS 2007*, Lecture Notes in Computer Science 4708. Springer Verlag, 2007.

[483] P. Koiran and S. Perifel. VPSPACE and a transfer theorem over the reals. In *Proceedings of STACS 2007*, 2007.

[484] Peter Kornerup, Vincent Lefèvre, and Jean-Michel Muller. Computing integer powers in floating-point arithmetic. In *Proceedings of 41th Conference on signals, systems and computers*. IEEE, November 2007.

[485] Patrick Loiseau, Paulo Gonçalves, and Pascale Vicat-Blanc Primet. A comparative study of different heavy tail index estimators of the flow size from sampled data. In *MetroGrid Workshop, GridNets*, Lyon, France, October 2007. ACM Press.

[486] Dino M. Lopez Pacheco, Laurent Lefèvre, and Cong-Duc Pham. Fairness issues when transferring large volume of data on high speed networks with router-assisted transport protocols. In *High Speed Networks Workshop 2007, in conjunction with IEEE INFOCOM 2007*, Anchorage, Alaska, USA, May 2007.

[487] Edgar Magana, Laurent Lefèvre, Masum Hasan, and Joan Serrat. Snmp-based monitoring agents and heuristic scheduling for large scale grids. In *Grid computing, high-performAnce and Distributed Applications (GADA'07)*, Vilamoura, Algarve, Portugal, November 2007.

[488] Edgar Magana, Laurent Lefèvre, and Joan Serrat. Autonomic management architecture for flexible grid services deployment based on policies. In *Architecture of Computing Systems - ARCS 2007*, volume 4415, pages 157–170, ETH, Zurich, Switzerland, March 2007. Springer Berlin / Heidelberg.

[489] Loris Marchal, Veronika Rehn, Yves Robert, and Frédéric Vivien. Scheduling and data redistribution strategies on star platforms. In *PDP'2007, 15th Euromicro Workshop on Parallel, Distributed and Network-based Processing*, pages 288–295. IEEE Computer Society Press, 2007.

[490] S. Perifel. Symmetry of information and nonuniform lower bounds. In *Proceedings of CSR 2007*, Lecture Notes in Computer Science 4649. Springer Verlag, 2007.

[491] Jean-François Pineau, Yves Robert, and Frédéric Vivien. Revisiting matrix product on master-worker platforms. In *9th Workshop on Advances in Parallel and Distributed Computational Models APDCM 2007*. IEEE Computer Society Press, 2007.

[492] Damien Pous. Complete lattices and up-to techniques. In *Programming Languages and Systems, 5th Asian Symposium, APLAS 2007, Singapore, November 29-December 1, 2007, Proceedings*, volume 4807 of *Lecture Notes in Computer Science*, pages 351–366. Springer, 2007.

[493] Damien Regnault, Nicolas Schabanel, and Éric Thierry. Progresses in the analysis of stochastic 2D cellular automata: a study of asynchronous 2D minority. In *Proc. of 32nd Conf. on Mathematical Foundation of Computer Science (MFCS)*, volume LNCS 4708, pages 320–332, Aug. 2007.

[494] Veronika Rehn-Sonigo. Optimal Closest Policy with QoS and Bandwidth Constraints for Placing Replicas in Tree Networks. In *CoreGRID'2007, Core GRID Symposium 2007*. Springer Verlag, 2007.

[495] C. Riba. On the Stability by Union of Reducibility Candidates. In *Proceedings of FoSSaCS'07*, volume 4423 of *LNCS*, 2007. External publication.

[496] C. Riba. Strong Normalization as Safe Interaction. In *Proceedings of LiCS'07*, pages 13–22. IEEE Computer Society, 2007. External publication.

[497] Gabriel Rilling, Patrick Flandrin, and Paulo Gonçalves. Une extension bivariée pour la décomposition modale empirique: Application à des bruits blancs complexes. In *Proceedings of the 21th Colloquium GRETSI*, Troyes (France), September 2007.

[498] Julien Robert and Nicolas Schabanel. Non-clairvoyant batch set scheduling: Fairness is fair enough. In *Proc. of 15th European Symposium on Algorithms (ESA)*, volume LNCS 4698, pages 742–753, Oct. 2007.

[499] Julien Robert and Nicolas Schabanel. Pull-based data broadcast with dependencies: Be fair to users, not to items. In *Proc. of the 18th ACM/SIAM Symp. on Discrete Algorithms (SODA)*, pages 238–247, Jan. 2007.

[500] C. Sarr, C. Chaudet, G. Chelius, and I. Guérin Lassous. Amélioration de la précision pour l'estimation de la bande passante résiduelle dans les réseaux ad hoc basés sur ieee 802.11. In *8es Journées Doctorales Informatique et Réseau (JDIR)*, Marne-la-Vallée, France, January 2007.

[501] Sebastien Soudan, Romaric Guillier, Ludovic Hablot, Yuetsu Kodama, Tomohiro Kudoh, Fumihiro Okazaki, Ryousei Takano, and Pascale Vicat-Blanc Primet. Investigation of ethernet switches behavior in presence of contending flows at very high-speed. In *PFLDnet 2007*, Feb. 2007.

[502] Sebastien Soudan, Romaric Guillier, and Pascale Vicat-Blanc Primet. End-host based mechanisms for implementing flow scheduling in gridnetworks. In *GridNets 2007*, Oct. 2007.

[503] Gilles Villard. Certification of the QR factor R, and of lattice basis reducedness. In *Proc. International Symposium on Symbolic and Algebraic Computation, Waterloo, Canada*, pages 361–368. ACM Press, August 2007.

2008

[504] Kunal Agrawal, Anne Benoit, and Yves Robert. Mapping linear workflows with computation/communication overlap. In *ICPADS'2008, the 14th IEEE International Conference on Parallel and Distributed Systems*, pages 195–202. IEEE Computer Society Press, 2008.

[505] Emmanuel Agullo, Abdou Guermouche, and Jean-Yves L'Excellent. On the I/O volume in out-of-core multifrontal methods with a flexible allocation scheme. In *VECPAR'08 International Meeting on High Performance Computing for Computational Science*, volume 5336 of *LNCS*, pages 328–335. Springer-Verlag Berlin Heidelberg, June 2008.

[506] Ali Akhavi and Damien Stehlé. Speeding-up lattice reduction with random projections. In *Proc. 8th Latin American Theoretical Informatics (LATIN'08)*, volume 4957 of *Lecture Notes in Computer Science*, pages 293–305. Springer, 2008.

[507] P. Albert, E. Mayordoma, P. Moser, and S. Perifel. Pushdown compression. In *Proceedings of STACS 2008*, pages 39–48, 2008.

[508] Marco Aldinucci, Hinde Bouziane, Marco Danelutto, and Christian Pérez. Towards software component assembly language enhanced with workflows and skeletons. In *Joint Workshop on Component-Based High Performance Computing and Component-Based Software Engineering and Software Architecture (CBHPC/COMPARCH 2008)*. ACM, 14-17 October 2008. ISBN: 978-1-60558-311-2.

[509] Lachlan Andrew, Cesar Marcondes, Sally Floyd, Lawrence Dunn, Romaric Guillier, Wang Gang, Lars Eggert, Sangtae Ha, and Injong Rhee. Towards a common tcp evaluation suite. In *PFLDnet 2008*, March 2008.

[510] Hrachya Astsatryan, Vladimir Sahakyan, Yuri Shoukouryan, Michel Daydé, Aurélie Hurault, Marc Pantel, and Eddy Caron. A grid-aware web interface with advanced service trading for linear algebra calculations. In *8th International Meeting High Performance Computing for Computational Science (VECPAR'08)*, pages 106–113, Toulouse, June 2008.

[511] Narjess Ayari, Denis Barbaron, Laurent Lefèvre, and Pascale Vicat-Blanc Primet. On improving the reliability of cluster based voice over ip services. In *FastAbstract : DSN 2008 : The 38th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, Anchorage, Alaska, USA, June 2008.

[512] Narjess Ayari, Denis Barbaron, Laurent Lefèvre, and Pascale Vicat-Blanc Primet. A session aware admission control scheme for next generation ip services. In *Fifth annual IEEE Consumer Communications and Networking Conference (IEEE CCNC 2008)*, Las Vegas, USA, January 2008.

[513] Narjess Ayari, Laurent Lefèvre, and Denis Barbaron. On improving the reliability of internet services through active replication. In *PDCAT 2008 : The Ninth International Conference on Parallel and Distributed Computing, Applications and Technologies*, pages 259–262, Dunedin, New Zealand, December 2008.

[514] Narjess Ayari, Pablo Neira Ayuso, Laurent Lefèvre, and Denis Barbaron. Towards a dependable architecture for highly available internet services. In *ARES'08 : The Third International Conference on Availability, Reliability and Security*, pages 422–427, Barcelona, Spain, March 2008.

[515] Jean-Claude Bajard, Philippe Langlois, Dominique Michelucci, Géraldine Morin, and Nathalie Revol. Towards guaranteed geometric computations with approximate arithmetics (12 pages). In *Advanced Signal Processing Algorithms, Architectures, and Implementations XVIII, part of the SPIE Optics & Photonics 2008 Symposium*, volume 7074, August 2008.

[516] R. Baron, S. Béal, E. Rémila, and P. Solal. Average tree solution for graph games. In *Proceedings of the 15th International Colloquium on Structural Information and Communication Complexity (SIROCCO) 2008*, 2008.

[517] F. Becker, S. Rajsbaum, I. Rapaport, and E. Rémila. Average binary long-lived consensus: quantifying the stabilization role played by memory. In *Proceedings of the 15th International Colloquium on Structural Information and Communication Complexity (SIROCCO) 2008*, Lecture Notes in Computer Science 5058. Springer Verlag, 2008.

[518] F. Becker, E. Rémila, and N. Schabanel. Time optimal self-assembly of 2d and 3d shapes: the case of squares and cubes. In *proceedings of the 14th International Conference on DNA Computing (DNA), (2008)*, Lecture Notes in Computer Science. Springer Verlag, 2008.

[519] Elyes Ben Hamida, Guillaume Chelius, and Jean-Marie Gorce. Scalable versus Accurate Physical Layer Modeling in Wireless Network Simulations. In IEEE Computer Society, editor, *22nd ACM/IEEE/SCS workshop on Principles of Advanced and Distributed Simulation (PADS 2008)*, pages 127–134, Roma Italie, 2008. ACM/IEEE/SCS. External publication.

[520] Anne Benoit, Mourad Hakem, and Yves Robert. Fault tolerant scheduling of precedence task graphs on heterogeneous platforms. In *10th Workshop on Advances in Parallel and Distributed Computational Models APDCM 2008*. IEEE Computer Society Press, 2008.

[521] Anne Benoit, Mourad Hakem, and Yves Robert. Realistic models and efficient algorithms for fault tolerant scheduling on heterogeneous platforms. In *ICPP'2008, the 37th International Conference on Parallel Processing*. IEEE Computer Society Press, 2008.

[522] Anne Benoit, Harald Kosch, Veronika Rehn-Sonigo, and Yves Robert. Bi-criteria pipeline mappings for parallel image processing. In *ICCS'2008, the 8th International Conference on Computational Science*, volume 5101 of *LNCS*, pages 215–225. Springer Verlag, 2008.

[523] Anne Benoit, Loris Marchal, Jean-François Pineau, Yves Robert, and Frédéric Vivien. Offline and online scheduling of concurrent bags-of-tasks on heterogeneous platforms. In *10th Workshop on Advances in Parallel and Distributed Computational Models APDCM 2008*. IEEE Computer Society Press, 2008.

[524] Anne Benoit, Veronika Rehn-Sonigo, and Yves Robert. Optimizing latency and reliability of pipeline workflow applications. In *HCW'2008, the 17th Heterogeneous Computing Workshop*. IEEE Computer Society Press, 2008.

[525] Viktor Bertis, Raphaël Bolze, Frédéric Desprez, and Kevin Reed. Large scale execution of a bioinformatic application on a volunteer grid. In *Workshop on Parallel and Distributed Scientific and Engineering Computing (PDSEC08)*, April 2008.

[526] Jean-Luc Beuchat, Nicolas Brisebarre, Jérémie Detrey, Eiji Okamoto, and Francisco Rodríguez-Henríquez. A comparison between hardware accelerators for the modified Tate pairing over $\mathbb{F}_{2^m}$ and $\mathbb{F}_{3^m}$. In *Second International Conference on Pairing-based Cryptography (Pairing'08)*, volume 5209 of *Lecture Notes in Computer Science*, pages 297–315. Springer Verlag, 2008.

[527] Julien Bigot, Hinde Bouziane, Christian Pérez, and Thierry Priol. On abstractions of software component models for scientific applications. In *Proc. of the Abstractions for Distributed Systems workshop*, Las Palmas, Gran Canaria, Spain, 2008.

[528] Benoit Boissinot, Sebastian Hack, Daniel Grund, Benoît Dupont de Dinechin, and Fabrice Rastello. Fast liveness checking for SSA-form programs. In *Sixth Annual IEEE/ACM International Symposium on Code Generation and Optimization (CGO'08)*, pages 35–44, Boston, MA, USA, April 2008. ACM. **Best paper award**.

[529] Florent Bouchez, Alain Darte, and Fabrice Rastello. Advanced conservative and optimistic register coalescing. In *International Conference on Compilers, Architectures and Synthesis for Embedded Systems (CASES'08)*, pages 147–156, Atlanta, GA, USA, October 2008. ACM.

[530] A. Bouillard, L. Jouhet, and E. Thierry. Computation of a (min,+) multi-dimnesional convolution for end-to-end performance analysis. In *Proceedings of ValueTools'2008*, 2008.

[531] Hinde Bouziane, Christian Pérez, and Thierry Priol. A software component model with spatial and temporal compositions for grid infrastructures. In *Proc. 14th Intl. Euro-Par Conference (Euro-Par 08)*, volume 5168, pages 698 – 708, Las Palmas de Gran Canaria, Spain, 26-29 August 2008. Springer Berlin / Heidelberg. External publication.

[532] Nicolas Brisebarre, Sylvain Chevillard, Milos Ercegovac, Jean-Michel Muller, and Serge Torres. An efficient method for evaluating polynomial and rational function approximations. In *Application-specific Systems, Architectures and Processors (ASAP 2008)*, pages 245–250. IEEE, 2008.

[533] Nicolas Brisebarre, Florent de Dinechin, and Jean-Michel Muller. Integer and floating-point constant multipliers for FPGAs. In *Application-specific Systems, Architectures and Processors (ASAP 2008)*, pages 239–244. IEEE, 2008.

[534] Nicolas Brisebarre, Florent de Dinechin, and Jean-Michel Muller. Multiplieurs et diviseurs constants en virgule flottante avec arrondi correct. In *RenPar'18, SympA'2008, CFSE'6*, 2008.

[535] Yves Caniou, Eddy Caron, Ghislain Charrier, Andréea Chis, Frédéric Desprez, and Eric Maisonnave. Ocean-atmosphere modelization over the grid. In Wu-chi Feng and Yuanyuan Yang, editors, *The 37th International Conference on Parallel Processing (ICPP 2008)*, pages 206–213, Portland, Oregon. USA, September 2008. IEEE.

[536] Yves Caniou, Eddy Caron, Ghislain Charrier, Frédéric Desprez, Eric Maisonnave, and Vincent Pichon. Ocean-atmosphere application scheduling within DIET. In *APDCT-08 Symposium. International Symposium on Advanced in Parallel and Distributed Computing Techniques*, pages 675–680, Sydney, Australia, December 2008. In conjunction with ISPA'2008, IEEE Computer Society. Invited paper from the reviewed process of ISPA'08.

[537] Yves Caniou, Jean-Sébastien Gay, and Pierre Ramet. Tunable parallel experiments in a GridRPC framework: application to linear solvers. In *VECPAR'08 International Meeting on High Performance Computing for Computational Science*, volume 5336 of *LNCS*, pages 430–436, 2008.

[538] Yves Caniou, Noriyuki Kushida, and Naoya Teshima. Implementing interoperability between the AEGIS and DIET GridRPC middleware to build an international sparse linear algebra expert system. In *The Second IEEE International Conference on Advanced Engineering Computing and Applications in Sciences (ADVCOMP 2008)*, pages 205–210, 2008.

[539] Eddy Caron, Pushpinder Kaur Chouhan, and Frédéric Desprez. Automatic middleware deployment planning on heterogeneous platfoms. In *The 17th Heterogeneous Computing Workshop (HCW'08)*, Miami, Florida, April 2008. In conjunction with IPDPS 2008.

[540] Eddy Caron, Ajoy Datta, Franck Petit, and Cédric Tedeschi. Self-stabilization in tree-structured p2p service discovery systems. In *27th International Symposium on Reliable Distributed Systems (SRDS 2008)*, pages 207–216, Napoli, Italy, October 2008. IEEE.

[541] Eddy Caron, Frédéric Desprez, and David Loureiro. All-in-one graphical tool for the management of diet a GridRPC middleware. In Norbert Meyer, Domenico Talia, and Ramin Yahyapour, editors, *Grid and Services Evolution*, pages 169–187, Barcelona, Spain, June 2008. CoreGRID Workshop on Grid Middleware (in conjunction with OGF'23), Springer.

[542] Eddy Caron, Frédéric Desprez, and Cédric Tedeschi. Efficiency of tree-structured peer-to-peer service discovery systems. In *Fifth International Workshop on Hot Topics in Peer-to-Peer Systems (Hot-P2P)*, Miami, Florida, April 2008. In conjunction with IPDPS 2008.

[543] Lawrence Chang, Alex Galis, Bertrand Mathieu, Kerry Jean, Roel Ocampo, Lefteris Mamatas, Javier R. Loyola, Joan Serrat, Andreas Berl, Hermann de Meer, Steven Davy, Zeinab Movahedi, and Laurent Lefèvre. Self-organising management overlays for future internet services. In *MACE 2008 : 3rd IEEE International Workshop on Modelling Autonomic Communications Environments during the ManWeek 2008 : 4th International Week on Management of Networks and Services*, pages 74–89, Samos Island, Greece, September 2008.

[544] Sylvain Chevillard and Nathalie Revol. Computation of the error function erf in arbitrary precision with correct rounding. In *RNC 8 Proceedings, 8th Conference on Real Numbers and Computers*, pages 27–36. Javier D. Bruguera and Marc Daumas, July 2008.

[545] Philippe Combes, Eddy Caron, Frédéric Desprez, Bastien Chopard, and Julien Zory. Relaxing synchronization in a parallel systemc kernel. In *ISPA 2008. International Symposium on Parallel and Distributed Processing with Applications*, pages 180–187, Sydney, Australia, December 2008. IEEE Computer Society.

[546] Ugo Dal Lago and Olivier Laurent. Quantitative game semantics for linear logic. In Michael Kaminski and Simone Martini, editors, *Computer Science Logic '08*, volume 5213 of *Lecture Notes in Computer Science*, pages 230–245. Springer, September 2008. External publication.

[547] Florent de Dinechin, Bogdan Pasca, Octavian Creţ, and Radu Tudoran. An FPGA-specific approach to floating-point accumulation and sum-of-products. In *Field-Programmable Technologies*, pages 33–40. IEEE, 2008.

[548] Carole Delporte-Gallet, Stéphane Devismes, Hugues Fauconnier, Franck Petit, and Sam Toueg. With finite memory consensus is easier than reliable broadcast. In *12th International Conference On Principles of Distributed Systems (OPODIS 2008)*, volume 5401 of *Lecture Notes in Computer Science, Springer*, pages 41–57, Luxor, Egypt, 2008.

[549] Romain Demangeon, Daniel Hirschkoff, Naoki Kobayashi, and Davide Sangiorgi. On the complexity of termination inference for processes. In *Trustworthy Global Computing, Third Symposium, TGC 2007, Sophia-Antipolis, France, November 5-6, 2007, Revised Selected Papers*, volume 4912 of *Lecture Notes in Computer Science*, pages 140–155. Springer, 2008.

[550] Romain Demangeon, Daniel Hirschkoff, and Davide Sangiorgi. Static and dynamic typing for the termination of mobile processes. In *Fifth IFIP International Conference On Theoretical Computer Science - TCS 2008, IFIP 20th World Computer Congress, TC 1, Foundations of Computer Science, September 7-10, 2008, Milano, Italy*, volume 273 of *IFIP*, pages 413–427. Springer, 2008.

[551] J. Demongeot, M. Morvan, and S. Sené. Impact of fixed boundary conditions on the basins of attraction in the flower's morphogenesis of arabidopsis thaliana. In *Proceedings of the 22nd International Conference on Advanced Information Networking and Applications*, pages 782–789. IEEE Computer Society, 2008.

[552] J. Demongeot, M. Morvan, and S. Sené. Robustness of dynamical systems attraction basins against state perturbations: Theoretical protocol and application in systems biology. In *CISIS archive Proceedings of the 2008 International Conference on Complex, Intelligent and Software Intensive Systems*, pages 675–681. IEEE Computer Society, 2008.

[553] Benjamin Depardon. Déploiement générique d'applications sur plates-formes hétérogènes distribuées. In *RenPar'18*, Fribourg, Switzerland, feb 2008.

[554] Frédéric Desprez, Eddy Caron, and Gaël Le Mahec. Dagda: Data arrangement for the grid and distributed applications. In *AHEMA 2008. International Workshop on Advances in High-Performance E-Science Middleware and Applications. In conjunction with eScience 2008*, pages 680–687, Indianapolis, Indiana, USA, December 2008.

[555] Jack Dongarra, Jean-François Pineau, Yves Robert, and Frédéric Vivien. Matrix product on heterogeneous master-worker platforms. In *13th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 53–62, Salt Lake City, Utah, February 2008.

[556] Jean Duprat. Une axiomatique de la géométrie plane en Coq. In *Actes du Congrès Journées Francophones des Langages Applicatifs*, 2008. Etretat, january 26-29.

[557] C. Dürr, M. Chrobak, M. Hurand, and J. Robert. Algorithms for temperature-aware task scheduling in microprocessor systems. In *Proc. of the 4th International Conference on Algorithmic Aspects in Information and Management (AAIM)*, 2008.

[558] U. Flarup and L. Lyaudet. On the expressive power of permanents and perfect matchings of matrices of bounded pathwidth/cliquewidth. In *Proceedings of CSR 2008 (3rd International Computer Science Symposium in Russia)*, 2008.

[559] Eric Fleury, Antoine Scherrer, Pierre Borgnat, Jean-Loup Guillaume, and Céline Robardet. Description and simulation of dynamic mobility networks. In *Dynamics on and of Complex Networks II*, Jerusalem Israël, 2008.

[560] Matthieu Gallet, Loris Marchal, and Frédéric Vivien. Allocating series of workflows on computing grids. In *Proceedings of the 14th IEEE International Conference on Parallel and Distributed Systems (ICPADS'08)*, pages 48–55. IEEE Computer Society Press, 2008.

[561] Jean-Patrick Gelas, Olivier Mornard, and Pascale Vicat-Blanc Primet. évaluation des performances réseau dans le contexte de la virtualisation xen. In *CFIP 2008 : Colloque Francophone sur l'Ingénierie des Protocoles*, Les Arcs, France, March 2008.

[562] Romaric Guillier and Pascale Vicat-Blanc Primet. Methodologies and tools for exploring transport protocols in the context of high-speed networks. In *IEEE TCSC Doctoral Symposium*, May 2008.

[563] Ludovic Hablot, Olivier Gluck, Jean-Christophe Mignot, and Pascale Vicat-Blanc Primet. Etude d'implémentations mpi dans une grille de calcul. In *Actes de Renpar'08*, Février 2008.

[564] Syed Hasan, Laurent Lefèvre, Zhiyi Huang, and Paul Werstein. Cross layer protocol support for live streaming media. In *AINA-08 : The IEEE 22nd International Conference on Advanced Information Networking and Applications*, pages 319–326, Okinawa, Japan, March 2008.

[565] Syed Hasan, Laurent Lefèvre, Zhiyi Huang, and Paul Werstein. Supporting large scale eresearch infrastructures with adapted live streaming capabilities. In *6th Australasian Symposium on Grid Computing and e-Research*, volume 82, pages 55–63, Wollongong, Australia, January 2008. Conferences in Research and Practice in Information Technology - Australian Computer Society.

[566] Daniel Hirschkoff, Tom Hirschowitz, Samuel Hym, Aurélien Pardon, and Damien Pous. Encapsulation and dynamic modularity in the Pi-calculus. In V.T. Vasconcelos and N. Yoshida, editors, *Proc. of the PLACES'08 workshop*, volume 241 of *ENTCS*, pages 85–100, 2008.

[567] P. Kacsuk, Z. Farkas, and G. Fedak. Towards Making BOINC and EGEE Interoperable. In *Proceedings of 4th IEEE International Conference on e-Science (e-Science 2008), International Grid Interoperability and Interoperation Workshop 2008 (IGIIW 2008)*, pages 478–484, Indianapolis, USA, December 2008.

[568] Daniel Kahn, Clément Rezvoy, and Frédéric Vivien. Parallel large scale inference of protein domain families. In *ICPADS'2008, the 14th IEEE International Conference on Parallel and Distributed Systems*, pages 72–79. IEEE Computer Society Press, 2008.

[569] E. Kaltofen and P. Koiran. Expressing a fraction of two determinants as a determinant. In *Proceedings of ISSAC 2008 (International Symposium on Symbolic and Algebraic Computation)*. ACM Press, 2008.

[570] P. Koiran, J. Landes, N. Portier, and P. Yao. Adversary lower bounds for nonadaptive quantum algorithms. In *Proceedings of Wollic'08 (15th Workshop on Logic, Language, Information and Computation)*, Lecture Notes in Computer Science 5110. Springer, 2008. long version to appear in *Journal of Computer and System Sciences*.

[571] P. Koiran and K. Meer. On the expressive power of CNF formulas of bounded tree- and clique-width. In *Proceedings of WG'08 (34th Workshop on Graph-Theoretic Concepts in Computer Science)*, Lecture Notes in Computer Science 5344. Springer, 2008.

[572] Guilherme Koslovski, Pascale Vicat-Blanc Primet, and Andrea Schwertner Charao. Vxdl: Virtual resources and interconnection networks description language. In *GridNets 2008*, Oct. 2008.

[573] Noriyuki Kushida, Yoshio Suzuki, Naoya Teshima, Norihiro Nakajima, Yves Caniou, Michel Daydé, and Pierre Ramet. Toward an international sparse linear algebra expert system by interconnecting the ITBL computational grid with the GridTLSE platform. In *VECPAR'08 International Meeting on High Performance Computing for Computational Science*, volume 5336 of *LNCS*, pages 424–429, 2008.

[574] C. La Rota, F. Tarissan, and L. Liberti. Inferring parameters of genetic regulatory networks. In Jaime Amador, Carlos Paternina-Arboleda, and Jesús Velásquez, editors, *Proceedings of the XIV Latino Ibero American Congress of Operational Research (CLAIO).*, page 127, September 2008. ISBN: 978 958 825283-4.

[575] Philippe Langlois and Nicolas Louvet. Compensated Horner algorithm in K times the working precision. In *RNC 8 Proceedings, 8th Conference on Real Numbers and Computers*, pages 157–166. Javier D. Bruguera and Marc Daumas, July 2008.

[576] Christoph Quirin Lauter and Florent de Dinechin. Optimising polynomials for floating-point implementation. In *RNC 8 Proceedings, 8th Conference on Real Numbers and Computers*, pages 7–16, July 2008.

[577] Vincent Lefèvre, Damien Stehlé, and Paul Zimmermann. Worst cases for the exponential function in the IEEE 754r decimal64 format. In *Reliable Implementation of Real Number Algorithms: Theory and Practice*, volume 5045 of *Lecture Notes in Computer Science*, pages 114–126, Dagstuhl, Germany, 2008. Springer.

[578] Dino M. Lopez Pacheco, Laurent Lefèvre, and Cong-Duc Pham. Lightweight fairness solutions for xcp and tcp cohabitation. In *IFIP/TC6 Networking 2008*, pages 715–726, Singapore, May 2008.

[579] Yannis Mazzer and Bernard Tourancheau. MPI in wireless sensor networks. In *PVM/MPI Recent Advances in Parallel Virtual Machine and Message Passing Interface*, volume 5205 of *Lecture Notes in Computer Science*, pages 334–339. Springer, 2008.

[580] A. Muchnik and A. Romashchenko. A random oracle does not help extract the mutual information. In *Proceedings of MFCS 2008 (33rd International Symposium on Mathematical Foundations of Computer Science)*, Lecture Notes in Computer Science 5162. Springer Verlag, 2008.

[581] Pablo Neira Ayuso, Laurent Lefèvre, and Rafael M. Gasca. hft-fw : Hybrid fault-tolerance for cluste-based stateful firewalls. In *ICPADS 2008 : The 14th IEEE International Conference on Parallel and Distributed Systems*, Melbourne, Australia, December 2008.

[582] Pablo Neira Ayuso, Rafael M. Gasca, and Laurent Lefèvre. Ft-fw: Efficient connection failover in cluster-based stateful firewall. In *PDP2008 : 16th Euromicro International Conference on Parallel, Distributed and network-based Processing*, pages 573–580, Toulouse, France, February 2008.

[583] Pablo Neira Ayuso, Rafael M. Gasca, and Laurent Lefèvre. Multiprimary support for the availability of cluster-based stateful firewalls using ft-fw. In *ESORICS 2008 : 13th European Symposium on Research in Computer Security*, pages 1–17, Malaga, Spain, October 2008.

[584] Pablo Neira Ayuso, Leonardo Maccari, Laurent Lefèvre, and Rafael M. Gasca. Stateful firewalling for wireless mesh networks. In *NTMS 2008 : The second IFIP International Conference on New Technologies, Mobility and Security*, Tangier, Morocco, November 2008.

[585] Anne-Cécile Orgerie, Laurent Lefèvre, and Jean-Patrick Gelas. Chasing gaps between bursts : Towards energy efficient large scale experimental grids. In *PDCAT 2008 : The Ninth International Conference on Parallel and Distributed Computing, Applications and Technologies*, pages 381–389, Dunedin, New Zealand, December 2008.

[586] Anne-Cécile Orgerie, Laurent Lefèvre, and Jean-Patrick Gelas. Save watts in your grid: Green strategies for energy-aware framework in large scale distributed systems. In *ICPADS 2008 : The 14th IEEE International Conference on Parallel and Distributed Systems*, pages 171–178, Melbourne, Australia, December 2008.

[587] Andreea Picu, Eric Fleury, and Antoine Fraboulet. On Frequency Optimisation for Power Saving in WSNs. In *14th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, St Louis États-Unis d'Amérique, 2008.

[588] Alexandru Plesco and Tanguy Risset. Coupling loop transformations and high-level synthesis. In *SYMPosium en Architectures nouvelles de machines (SYMPA'08)*, February 2008.

[589] Xavier Pujol and Damien Stehlé. Rigorous and efficient short lattice vectors enumeration. In *Proceedings of ASIACRYPT'08*, volume 5350 of *Lecture Notes in Computer Science*, pages 390–405. Springer, 2008.

[590] R. Razafindralambo and I. Guérin Lassous. Sba: un algorithme simple de backoff pour les réseaux ad hoc. In *Algotel*, St Malo, France, May 2008.

[591] D. Regnault. Directed percolation arising in stochastic cellular automata anaslysis. In *Proceedings of the 33rd International Symposium on Mathematical Foundations of Computer Science*, volume 5162 of *LNCS*, pages 563–574, 2008.

[592] D. Regnault. Quick energy drop in stochastic 2d minority. In *Proceedings of 8th International Conference on Cellular Automata for Research and Industry (ACRI2008)*, volume LNCS 5191, 2008. to appear.

[593] D. Regnault, N. Schabanel, and E. Thierry. On the analysis of "simple" 2d stochastic cellular automata. In *Proceedings of the 2nd International Conference on Language and Automata Theory and Applications*, LNCS, pages 463–474. Springer, 2008. to appear.

[594] C. Riba. Stability by Union of Reducibility Candidates for Orthogonal Constructor Rewriting. In *Proceedings of CiE'08*, volume 5028 of *LNCS*, pages 498–510. Springer, 2008. Invited Paper for the Special Session "Higher-type Recursion and Applications". External publication.

[595] J. Robert and N. Schabanel. Non-clairvoyant scheduling with precedence constraints. In *Proc. of the 19th ACM/SIAM Symp on Discrete Algorithms (SODA)*, 2008.

[596] Stéphane Le Roux. Graphs and path equilibria. In Rudolf Fleischer and Jinhui Xu, editors, *4th International Conference on Algorithmic Aspects in Information and Management*, volume 5034 of *Lecture Notes in Computer Science*, pages 247–258. Springer, 2008.

[597] Antoine Scherrer, Pierre Borgnat, Eric Fleury, Jean-Loup Guillaume, and Céline Robardet. A Methodology to Identify Characteristics of the Dynamic of Mobile Networks. In *ACM ASIAN INTERNET ENGINEERING CONFERENCE (AINTEC)*, Bangkok Thaïlande, 2008.

[598] Fabien Tarissan, Leo Liberti, and Camilo La Rota. Biological regulatory network reconstruction: a mathematical programming approach. In *Proceedings of the European Conference on Complex Systems (ECCS08)*, 2008.

[599] Cédric Tedeschi. Arbre de préfixes auto-stable pour les systèmes pair-à-pair. In *Fribourg'2008 - Conférences conjointes RenPar'18 / SympA'2008 / CFSE'6*, Fribourg, Suisse, February 2008.

[600] Bernard Tourancheau, Gérard Krauss, and Richard Blanchard. Parametric sensitivity study and optimization of the SDHW and PV subsystems in an energy positive house. In *IBPSA-France*, 2008.

[601] Bernard Tourancheau, Yannis Mazzer, Gérard Krauss, Valentin Marvan, and Frederic Kuznick. Calibration of sensors embedded on a micro-controler based wireless network. In *IBPSA-France*, 2008.

[602] Ionuţ Trestian, Octavian Creţ, Laura Creţ, Lucia Văcariu, Radu Tudoran, and Florent de Dinechin. FPGA-based computation of the inductance of coils used for the magnetic stimulation of the nervous system. In *Biomedical Electronics and Devices*, volume 1, pages 151–155, 2008.

[603] Bora Uçar. Heuristics for a matrix symmetrization problem. In R. Wyrzykowski, J. Dongarra, K. Karczewski, and J. Wasniewski, editors, *Proceedings of Parallel Processing and Applied Mathematics (PPAM'07)*, volume 4967 of *Lecture Notes in Computer Science*, pages 718–727, 2008. External publication.

[604] R. Vannier and I. Guérin Lassous. Partage équitable de la bande passante dans les réseaux ad hoc. In *CFIP*, Les Arcs, France, march 2008.

[605] R. Vannier and I. Guérin Lassous. Towards a practical and fair rate allocation for multihop wireless networks based on a simple node model. In *11th ACM/IEEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, Vancouver, Canada, October 2008.

[606] Dominique Verchere, Olivier Audouin, Bela Berde, Agostino Chiosi, Richard Douville, Helia Pouyllau, Pascale Vicat-Blanc Primet, Marcelo Pasin, Sebastien Soudan, Thierry Marcot, Veronique Piperaud, Remi Theillaud, Dohy Hong, Dominique Barth, Christian Cadéré, V. Reinhart, and Joanna Tomasik. Automatic network services aligned with grid application requirements in carriocas project. In *GridNets 2008*, Oct. 2008.

[607] Gilles Villard. Differentiation of Kaltofen's division-free determinant algorithm. In *MICA'2008 : Milestones in Computer Algebra, Stonehaven Bay, Trinidad and Tobago*, May 2008.

2009

[608] Andrei Agapi, Sebastien Soudan, Marcelo Pasin, Pascale Pascale Vicat-Blanc Primet, and Thilo Kielmann. Optimizing deadline-driven bulk data transfers in overlay networks. In *ICCCN 2009 Track on Pervasive Computing and Grid Networking (PCGN)*, San Francisco, USA, 2009.

[609] Kunal Agrawal, Anne Benoit, Fanny Dufossé, and Yves Robert. Mapping filtering streaming applications with communication costs. In *21st ACM Symposium on Parallelism in Algorithms and Architectures SPAA 2009*. ACM Press, 2009.

[610] Marco Aldinucci, Hinde Lilia Bouziane, Marco Danelutto, and Christian Pérez. STKM on SCA: a unified framework with components, workflows and algorthmic skeletons. In *15th International European Conference on Parallel and Distributed Computing (Euro-Par 2009)*, volume 5704 of *LNCS*, pages 678 – 690, Delft, Netherlands, August 2009. Springer.

[611] Françoise André, Guillaume Gauvrit, and Christian Pérez. Dynamic adaptation of the master-worker paradigm. In *Proc. of the IEEE 9th International Conference on Computer and Information Technology*, Xiamen, China, oct 2009. IEEE Computer Society. to appear.

[612] Fabienne Anhalt, Guilherme Koslovski, Marcelo Pasin, Jean-Patrick Gelas, and Pascale Vicat-Blanc Primet. Les infrastructures virtuelles à la demande pour un usage flexible de l'internet. In *JDIR 09: Journees Doctorales en Informatique et Reseaux*, Belfort, France, February 2009.

[613] Fabienne Anhalt and Pascale Vicat-Blanc Primet. Analysis and experimental evaluation of data plane virtualization with xen. In *ICNS 09 : International Conference on Networking and Services*, Valencia, Spain, April 2009.

[614] Narjess Ayari, Denis Barbaron, and Laurent Lefèvre. Evaluating session aware admission control strategies for improving the profitability of service providers. In *The 3rd IEEE Workshop on Enabling the Future Service-Oriented Internet: Towards Socially-Aware Networks - Held in conjunction with IEEE GLOBECOM 2009*, Hawai, USA, December 2009.

[615] Suleyman Baykut, Paulo Gonçalves, Pierre-Hervé Luppi, Patrice Abry, Edmundo Pereira de Souza Neto, and Damien Gervasoni. Emd-based analysis of rat eeg data for sleep state classification. In *Biosignals*. Springer, January 2009.

[616] Leila Ben Saad and Bernard Tourancheau. Multiple mobile sinks positioning in wireless sensor networks for buildings. In *SensorComm*. IEEE-IARIA, 2009.

[617] Anne Benoit. Comparison of Access Policies for Replica Placement in Tree Networks. In *15th International Euro-Par Conference*, LNCS, Delft, The Netherlands, August 2009. Springer Verlag.

[618] Anne Benoit, Henri Casanova, Veronika Rehn-Sonigo, and Yves Robert. Resource allocation for multiple concurrent in-network stream-processing applications. In *HeteroPar'2009: Seventh Int. Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms, jointly held with Euro-Par 2009*, LNCS. Springer Verlag, 2009. To appear. Received the Best Paper Award.

[619] Anne Benoit, Henri Casanova, Veronika Rehn-Sonigo, and Yves Robert. Resource allocation strategies for in-network stream processing. In *11th Workshop on Advances in Parallel and Distributed Computational Models APDCM 2009*. IEEE Computer Society Press, 2009.

[620] Anne Benoit, Alexandru Dobrila, Jean-Marc Nicod, and Laurent Philippe. Throughput optimization for micro-factories subject to failures. In *ISPDC'2009, 8th International Symposium on Parallel and Distributed Computing*, pages 11–18, Lisbon, Portugal, July 2009.

[621] Anne Benoit, Fanny Dufossé, and Yves Robert. Filter placement on a pipelined architecture. In *11th Workshop on Advances in Parallel and Distributed Computational Models APDCM 2009*. IEEE Computer Society Press, 2009.

[622] Anne Benoit, Fanny Dufossé, and Yves Robert. On the complexity of mapping pipelined filtering services on heterogeneous platforms. In *IPDPS'2009, the 23rd IEEE International Parallel and Distributed Processing Symposium*. IEEE Computer Society Press, 2009.

[623] Anne Benoit, Bruno Gaujal, Matthieu Gallet, and Yves Robert. Computing the throughput of replicated workflows on heterogeneous platforms. In *ICPP'2009, the 38th International Conference on Parallel Processing*. IEEE Computer Society Press, 2009. To appear.

[624] Anne Benoit, Mourad Hakem, and Yves Robert. Optimizing the latency of streaming applications under throughput and reliability constraints. In *ICPP'2009, the 38th International Conference on Parallel Processing*. IEEE Computer Society Press, 2009. To appear.

[625] Anne Benoit, Loris Marchal, Jean-Francois Pineau, Yves Robert, and Frédéric Vivien. Resource-aware allocation strategies for divisible loads on large-scale systems. In *18th International Heterogeneity in Computing Workshop HCW 2009*. IEEE Computer Society Press, 2009.

[626] Anne Benoit and Yves Robert. Multi-criteria mapping techniques for pipeline workflows on heterogeneous platforms. In *Recent Developments in Grid Technology and Applications*. Nova Science Publishers, 2009. To appear.

[627] Anne Benoit, Yves Robert, Arnold Rosenberg, and Frédéric Vivien. Static strategies for worksharing with unrecoverable interruptions. In *IPDPS'2009, the 23rd IEEE International Parallel and Distributed Processing Symposium*. IEEE Computer Society Press, 2009.

[628] Anne Benoit, Yves Robert, Arnold Rosenberg, and Frédéric Vivien. Static worksharing strategies for heterogeneous computers with unrecoverable failures. In *HeteroPar'2009: Seventh Int. Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms, jointly held with Euro-Par 2009*, LNCS. Springer Verlag, 2009. To appear.

[629] Julien Bigot and Christian Pérez. Increasing reuse in component models through genericity. In *Proc of the 11th International Conference on Software Reuse*, LNCS, Falls Church, Virginia, USA, oct 2009. Springer Verlag. To appear.

[630] O. Bodini, T. Fernique, and E. Rémila. Distances on lozenge tilings. In *Proceedings of the 15th International Conference on Discrete Geometry for Computer Imagery (DGCI) 2009*, 2009.

[631] Benoit Boissinot, Alain Darte, Benoît Dupont de Dinechin, Christophe Guillon, and Fabrice Rastello. Revisiting out-of-SSA translation for correctness, code quality, and efficiency. In *International Symposium on Code Generation and Optimization (CGO'09)*, pages 114–125, Seattle, WA, USA, March 2009. IEEE Computer Society Press. **Best paper award**.

[632] Yves Caniou, Eddy Caron, Ghislain Charrier, and Frédéric Desprez. Meta-scheduling and task reallocation in a grid environment. In *The Third IEEE International Conference on Advanced Engineering Computing and Applications in Sciences (ADVCOMP 2009)*, page 6p., Sliema, Malta, October 2009.

[633] Yves Caniou and Jean-Sébastien Gay. Simbatch: an API for simulating and predicting the performance of parallel resources managed by batch systems. In *Workshop on Secure, Trusted, Manageable and Controllable Grid Services (SGS), held in conjunction with EuroPar'08*, volume 5415 of *LNCS*, pages 223–234, 2009.

[634] E. Caron, F. Desprez, D. Loureiro, and A. Muresan. Cloud computing resource management through a grid middleware: A case study with diet and eucalyptus. In IEEE, editor, *IEEE International Conference on Cloud Computing (CLOUD 2009)*, Bangalore, India, September 2009. To appear in the Work-in-Progress Track from the CLOUD-II 2009 Research Track.

[635] Eddy Caron, Ajoy Datta, Benjamin Depardon, and Lawrence Larmore. A self-stabilizing k-clustering algorithm using an arbitrary metric. In *Euro-Par 2009*, volume LNCS 5704, pages 602–614, Delft, The Netherlands, August 25-28 2009. TU Delft.

[636] Eddy Caron, Frédéric Desprez, David Loureiro, and Adrian Muresan. Cloud computing resource management through a grid middleware: A case study with diet and eucalyptus. In *IEEE International Conference on Cloud Computing (CLOUD 2009)*, Bangalore, India, Septembre 2009. IEEE. To appear in the Work-in-Progress Track from the CLOUD-II 2009 Research Track.

[637] Eddy Caron, Cristian Klein, and Christian Perez. Efficient grid resource selection for a cem application. In *RenPar'19. 19ème Rencontres Francophones du Parallélisme*, Toulouse, France, September 2009.

[638] Fabienne Carrier, Stéphane Devismes, Franck Petit, and Yvan Rivierre. Space-optimal deterministic rendezvous. In *Second International Workshop on Reliability, Availability, and Security (WRAS 2009)*, Hiroshima, Japan, 2009. IEEE Computer Society. To appear.

[639] Ghislain Charrier and Yves Caniou. Ordonnancement et réallocation de tâches sur une grille de calcul. In *19e Rencontres Francophones du Parallélisme, des Architectures et des Systèmes*, Toulouse, September 2009.

[640] Guillaume Chelius, Eric Fleury, Antoine Fraboulet, and Jean-Christophe Lucet. A wireless sensor network to measure the health care workers exposure to tuberculosis. In *International Workshop and Conference on Complex Networks and their Applications (NetSci 09)*, Venice Italie, 2009.

[641] Sylvain Chevillard, Mioara Joldes, and Christoph Lauter. Certified and fast computation of supremum norms of approximation errors. In *19th IEEE Symposium on Computer Arithmetic (ARITH-19)*, Portland, Oregon, U.S.A., 2009.

[642] Luca Chiarabini and Philippe Audebaud. New development in extracting tail recursive programs from proofs. In *Logic-Based Program Synthesis and Transformation, 19th International Symposium, LOPSTR 2009, Coimbra, Portugal, Portugal, Sep 9-11, 2009*, LNCS. Springer Verlag, 2009. to appear.

[643] Georges Da-Costa, Jean-Patrick Gelas, Yiannis Georgiou, Laurent Lefèvre, Anne-Cécile Orgerie, Jean-Marc Pierson, Olivier Richard, and Kamal Sharma. The green-net framework: Energy efficiency in large scale distributed systems. In *HPPAC 2009 : High Performance Power Aware Computing Workshop in conjunction with IPDPS 2009*, Rome, Italy, May 2009.

[644] Florent de Dinechin, Mioara Joldes, Bogdan Pasca, and Guillaume Revy. Racines carrées multiplicatives sur FPGA. In *13^e SYMPosium en Architectures nouvelles de machines (SYMPA)*, Toulouse, September 2009.

[645] Florent de Dinechin, Cristian Klein, and Bogdan Pasca. Generating high-performance custom floating-point pipelines. In *Proceedings of the 19th International Conference on Field Programmable Logic and Applications*. IEEE, August 2009.

[646] Florent de Dinechin and Bogdan Pasca. Large multipliers with less DSP blocks. In *Proceedings of the 19th International Conference on Field Programmable Logic and Applications*. IEEE, August 2009.

[647] Romain Demangeon, Daniel Hirschkoff, and Davide Sangiorgi. Termination in higher-order concurrent calculi. In *Proc. of FSEN'09*, LNCS. Springer Verlag, 2009. to appear.

[648] Benjamin Depardon. Un algorithme auto-stabilisant pour le problème du k-partitionnement sur graphe pondéré. In *RenPar'19. 19ème Rencontres Francophones du Parallélisme*, Toulouse, 2009. To appear.

[649] Stéphane Devismes, Carole Delporte-Gallet, Hugues Fauconnier, Franck Petit, and Sam Toueg. Quand le consensus est plus simple que la diffusion fiable. In *11^{e} rencontres francophones sur les aspects algorithmiques des télécommunications (Algotel 2009)*, Carry-Le-Rouet, France, 2009. To appear.

[650] Stéphane Devismes, Franck Petit, and Sébastien Tixeuil. Exploration optimale probabiliste d'un anneau par des robots asynchrones et amnésiques. In *11^{e} rencontres francophones sur les aspects algorithmiques des télécommunications (Algotel 2009)*, Carry-Le-Rouet, France, 2009. To appear.

[651] Stéphane Devismes, Franck Petit, and Sébastien Tixeuil. Optimal probabilistic ring exploration by asynchronous oblivious robots. In *16th International Colloquium on Structural Information and Communication Complexity (SIROCCO 2009)*, volume 5804 of *Lecture Notes in Computer Science, Springer*, pages 230–241, Piran, Slovenia, 2009.

[652] Sékou Diakité, Loris Marchal, Jean-Marc Nicod, and Laurent Philippe. Steady-state for batches of identical task graphs. In *Euro-Par 2009*, volume 5704 of *LNCS*, pages 203–215, Delft University of Technology, Delft, the Netherlands, August 2009.

[653] Yoann Dieudonné, Shlomi Dolev, Franck Petit, and Michael Segal. Brief announcement: deaf, dumb, and chatting robots. In *28th Annual ACM Symposium on Principles of Distributed Computing (PODC 2009)*, pages 308–309, Calgary, Canada, 2009. ACM.

[654] Yoann Dieudonné and Franck Petit. Squaring the circle with weak mobile robots. In *11^{e} rencontres francophones sur les aspects algorithmiques des télécommunications (Algotel 2009)*, Carry-Le-Rouet, France, 2009. To appear.

[655] Dinil Mon Divakaran, Eitan Altman, Georg Post, Ludovic Noirie, and Pascale Vicat-Blanc Primet. Analysis of the effects of xlframes in a network. In *IFIP/TC6 NETWORKING 2009*, pages 364–377, May 2009.

[656] Dinil Mon Divakaran, Eitan Altman, Georg Post, Ludovic Noirie, and Pascale Vicat-Blanc Primet. From packets to xlframes: Sand and rocks for transfer of mice and elephants. In *IEEE High-Speed Networks Workshop*, Rio di Janerio, Brazil, April 2009.

[657] Pouya Dormiani, Milos Ercegovac, and Jean-Michel Muller. Design and implementation of a radix-4 complex division unit with prescaling. In *Proceedings of the 20th IEEE International Conference on Application-specific Systems, Architectures and Processors (ASAP 2009)*, Boston, U.S.A., 2009. IEEE Computer Society.

[658] E. Dramitinos and I. Guérin Lassous. A bandwidth allocation mechanism for 4g. In *European Wireless Technology Conference*, Roma, Italy, September 2009.

[659] Eric Fleury, Kevin Huguenin, and Anne-Marie Kermarrec. Route in Mobile WSN and Get Self-Deployment for Free. In *5th IEEE/ACM International Conference on Distributed Computing in Sensor Systems (DCOSS)*, Marina del Rey États-Unis d'Amérique, 2009.

[660] Alex Galis, Spyros Denazis, Alessandro Bassi, Pierpaolo Giacomin, Andreas Berl, Andreas Fischer, Herman de Meer, John Srassner, Steven Davy, Daniel Macedo, Guy Pujolle, Javier R. Loyola, Joan Serrat, Laurent Lefèvre, and Abderhaman Cheniour. Management architecture and systems for future internet networks. In *FIA Book : Towards the Future Internet - A European Research Perspective*, pages 112–122, Prague, May 2009. IOS Press. ISBN 978-1-60750-007-0.

[661] Matthieu Gallet, Loris Marchal, and Frédéric Vivien. Efficient scheduling of task graph collections on heterogeneous resources. In *International Parallel and Distributed Processing Symposium IPDPS'2009*. IEEE Computer Society Press, 2009.

[662] Richard H. G. Garner, Tom Hirschowitz, and Aurélien Pardon. Variable binding, symmetric monoidal closed theories, and bigraphs. In *Proc. of CONCUR'09*, volume 5710 of *LNCS*, pages 321–337. Springer, 2009.

[663] Laure Gonnord and Jean-Philippe Babau. Quantity of resource properties expression and runtime assurance for embedded systems. In *ACS/IEEE International Conference on Computer Systems and Applications (AICCSA'09)*, pages 428–435, Rabbat, Morocco, May 2009.

[664] Stef Graillat, Jean-Luc Lamotte, and Hong Diep Nguyen. Error-free transformation in rounding mode toward zero. In *Dagstuhl Seminar on Numerical Validation in Current Hardware Architectures*, volume 5492 of *Lecture Notes in Computer Science*, pages 217–229, 2009.

[665] Yi Gu, Qishi Wu, Anne Benoit, and Yves Robert. Optimizing end-to-end performance of distributed applications with linear computing pipelines. In *ICPADS'2009, the 15th International Conference on Parallel and Distributed Systems*. IEEE Computer Society Press, 2009.

[666] Romaric Guillier and Pascale Vicat-Blanc Primet. A user-oriented test suite for transport protocols comparison in datagrid context. In *ICOIN 2009*, January 2009.

[667] Haiwu He, Gilles Fedak, Bing Tran, and Franck Cappello. BLAST Application with Data-aware Desktop Grid Middleware. In *Proceedings of 9th IEEE International Symposium on Cluster Computing and the Grid CCGRID'09*, pages 284–291, Shanghai, China, May 2009.

[668] Mathias Jacquelin, Loris Marchal, and Yves Robert. Complexity analysis and performance evaluation of matrix product on multicore architectures. In *ICPP'2009, the 38th International Conference on Parallel Processing*. IEEE Computer Society Press, 2009. To appear.

[669] Claude-Pierre Jeannerod, Hervé Knochel, Christophe Monat, Guillaume Revy, and Gilles Villard. A new binary floating-point division algorithm and its software implementation on the ST231 processor. In *Proceedings of the 19th IEEE Symposium on Computer Arithmetic (ARITH-19)*, Portland, OR, June 2009.

[670] R. Baker Kearfott, John D. Pryce, and Nathalie Revol. Discussions on an interval arithmetic standard at Dagstuhl seminar 08021. In *Dagstuhl Seminar on Numerical Validation in Current Hardware Architectures*, volume 5492 of *Lecture Notes in Computer Science*, pages 1–6, 2009.

[671] P. Koiran and I. Briquel. A dichotomy theorem for polynomial evaluation. In *Proceedings of MFCS 2009 (34th International Symposium on Mathematical Foundations of Computer Science)*. Springer, 2009.

[672] P. Koiran and S. Perifel. A superpolynomial lower bound on the size of uniform non-constant-depth threshold circuits for the permanent. In *Proceedings of CCC 2009 (24th IEEE Conference on Computational Complexity)*. IEEE, 2009.

[673] Peter Kornerup, Vincent Lefèvre, Nicolas Louvet, and Jean-Michel Muller. On the computation of correctly-rounded sums. In *19th IEEE Symposium on Computer Arithmetic (ARITH-19)*, Portland, Oregon, U.S.A., 2009.

[674] Ouassila Labbani, Paul Feautrier, Eric Lenormand, and Michel Barreteau. Elementary transformation analysis for Array-OL. In *ACS/IEEE International Conference on Computer Systems and Applications, AICCSA'09*, pages 362–367, Rabat, Morocco, May 2009.

[675] Laurent Lefèvre and Anne-Cécile Orgerie. When clouds become green. In *Parco2009 : International Conference on Parallel Computing*, Lyon, France, September 2009.

[676] Patrick Loiseau, Paulo Gonçalves, Stéphane Girard, Florence Forbes, and Pascale Vicat-Blanc Primet. Maximum likelihood estimation of the flow size distribution tail index from sampled packet data. In *ACM Sigmetrics*, June 2009.

[677] Patrick Loiseau, Paulo Gonçalves, Romaric Guillier, Matthieu Imbert, Yuetsu Kodama, and Pascale Vicat-Blanc Primet. Metroflux: A high performance system for analyzing flow at very fine-grain. In *TridentCom*, April 2009.

[678] Qingda Lu, Christophe Alias, Uday Bondhugula, Sriram Krishnamoorthy, J. Ramanujam, Atanas Rountev, P. Sadayappan, Yongjian Chen, Haibo Lin, and Tin fook Ngai. Data layout transformation for enhancing locality on NUCA chip multiprocessors. In *International ACM/IEEE Conference on Parallel Architectures and Compilation Techniques*, September 2009.

[679] Qingda Lu, Christophe Alias, Uday Bondhugula, Sriram Krishnamoorthy, J. Ramanujam, Atanas Rountev, P. Sadayappan, Yongjian Chen, Haibo Lin, and Tin fook Ngai. Data layout transformation for enhancing locality on NUCA chip multiprocessors. In *International ACM/IEEE Conference on Parallel Architectures and Compilation Techniques (PACT'09)*, September 2009. External publication.

[680] Yannis Mazzer and Bernard Tourancheau. Comparisons for 6LoWPAN implementation on wireless sensor networks. In *SensorComm*. IEEE-IARIA, 2009.

[681] Alexandre Miquel. Relating classical realizability and negative translation for existential witness extraction. In P.-L. Curien, editor, *Proc. of TLCA'09*, volume 5608 of *LNCS*, pages 188–202. Springer Verlag, 2009.

[682] Ivan Morel, Damien Stehlé, and Gilles Villard. H-LLL: Using Householder inside LLL. In *Proc. International Symposium on Symbolic and Algebraic Computation (ISSAC 2009), Seould, Korea*, pages 271–278. ACM Press, August 2009.

[683] Kashif Munir, Pascale Vicat-Blanc Primet, and Michael Welzl. Grid network dimensioning by modeling the deadline constrained bulk data transfers. In *11th IEEE International Conference on High Performance Computing and Communications (HPCC-09)*, Seoul, Korea, June 2009.

[684] Lucas Nussbaum. Rebuilding debian using distributed computing. In *CLADE 2009*, June 2009.

[685] Lucas Nussbaum, Fabienne Anhalt, Olivier Mornard, and Jean-Patrick Gelas. Linux-based virtualization for hpc clusters. In *Linux Symposium 2009*, July 2009.

[686] Anne-Cécile Orgerie, Laurent Lefèvre, and Jean-Patrick Gelas. Economies d'energie dans les systèmes distribués à grande echelle : l'approche eari. In *JDIR 09 : Journèes Doctorales en Informatique et Réseaux*, Belfort, France, February 2009.

[687] Barbara Petit. A polymorphic type system for the lambda-calculus with constructors. In P.-L. Curien, editor, *Proc. of TLCA'09*, volume 5608 of *LNCS*, pages 234–248. Springer Verlag, 2009.

[688] Jean-François Pineau, Yves Robert, and Frédéric Vivien. Energy-aware scheduling of flow applications on master-worker platforms. In *Euro-Par 2009 - Parallel Processing*, volume 5704 of *LNCS*, pages 281–292. Springer Verlag, 2009.

[689] T. Razafindralambo and I. Guérin Lassous. Sba: a simple backoff algorithm for wireless ad hoc networks. In *IFIP Networking*, Aachen, Germany, May 2009.

[690] Colin Riba. On the values of reducibility candidates. In P.-L. Curien, editor, *Proc. of TLCA'09*, volume 5608 of *LNCS*, pages 264–278. Springer Verlag, 2009.

[691] Stéphane Le Roux. Acyclic preferences and existence of sequential Nash equilibria: A formal and constructive equivalence. In Stefan Berghofer, Tobias Nipkow, Christian Urban, and Makarius Wenzel, editors, *22nd International Conference on Theorem Proving in Higher Order Logics*, volume 5674 of *Lecture Notes in Computer Science*, pages 293–309. Springer, 2009.

[692] C. Sarr, S. Khalfallah, and I. Guérin Lassous. Gestion dynamique de la bande passante dans les réseaux ad hoc multi-sauts. In *9es Journées Doctorales Informatique et Réseau (JDIR)*, Belfort, France, January 2009.

[693] Sebastien Soudan, Dinil Mon Divakaran, Eitan Altman, and Pascale Vicat-Blanc Primet. Equilibrium in size-based scheduling systems. In *16th International Conference on Analytical and Stochastic Modelling Techniques and Applications*, Madrid, Spain, June 2009.

[694] Sebastien Soudan and Pascale Vicat-Blanc Primet. Mixing malleable and rigid bandwidth requests for optimizing network provisioning. In *21st International Teletraffic Congress*, Paris, France, sept 2009.

[695] Mark Stillwell, David Schanzenbach, Frédéric Vivien, and Henri Casanova. Resource allocation using virtual clusters. In *9th IEEE International Symposium on Cluster Computing and the Grid (CCGrid 09)*, 2009.

[696] Bernard Tourancheau, Yannis Mazzer, Gérard Krauss, Valentin Marvan, and Frederic Kuznick. Software calibration of wirelessly networked sensors. In *SensorComm*. IEEE-IARIA, 2009.

[697] Pascale Vicat-Blanc Primet, Fabienne Anhalt, and Guilherme Koslovski. Exploring the virtual infrastructure service concept in grid'5000. In *20th ITC Specialist Seminar on Network Virtualization*, Hoi An, Vietnam, May 2009.

[698] Pascale Vicat-Blanc Primet, Jean-Patrick Gelas, Olivier Mornard, Guilherme Koslovski, Vincent Roca, Lionel Giraud, Johan Montagnat, and Tram Truong Huu. A scalable security model for enabling dynamic virtual private execution infrastructures on the internet. In *IEEE International Conference on Cluster Computing and the Grid CCGrid2009*, Shanghai, May 2009.

Short communications [COM] and posters [AFF] in conferences and workshops

2005

[699] Martine Chaudier, Jean-Patrick Gelas, and Laurent Lefèvre. IAN2 : Industrial autonomic network node. Poster INRIA Booth, Supercomputing 2005, Seattle, USA, November 2005.

[700] Claude-Pierre Jeannerod, Saurabh Kumar Raina, and Arnaud Tisserand. High-radix floating-point division algorithms for embedded VLIW integer processors (extended abstract). In *Proc. 17th World Congress on Scientific Computation, Applied Mathematics and Simulation IMACS*, Paris, France, July 2005.

[701] Nathalie Revol. Bounding roundoff errors in Taylor models arithmetic. In *17th IMACS Conf. on Scientific Computation, Applied Math. and Simulation*, Paris, France, July 2005.

2006

[702] Sylvie Boldo, Marc Daumas, William Kahan, and Guillaume Melquiond. Proof and certification for an accurate discriminant. In *SCAN 2006 - 12th GAMM - IMACS International Symposium on Scientific Computing, Computer Arithmetic and Validated Numerics*, Duisburg, Germany, 2006.

[703] Alin Bostan, Claude-Pierre Jeannerod, and Éric Schost. Solving structured linear systems of large displacement rank. Poster at ISSAC 2006, July 2006.

[704] Hervé Brönnimann, Guillaume Melquiond, and Sylvain Pion. Proposing interval arithmetic for the C++ standard. In *SCAN 2006 - 12th GAMM - IMACS International Symposium on Scientific Computing, Computer Arithmetic and Validated Numerics*, Duisburg, Germany, 2006.

[705] Sylvain Chevillard and Christoph Quirin Lauter. Certified infinite norm using interval arithmetic. In *SCAN 2006 - 12th GAMM - IMACS International Symposium on Scientific Computing, Computer Arithmetic and Validated Numerics*, Duisburg, Germany, 2006.

[706] Pushpinder Kaur Chouhan, Holly Dail, Eddy Caron, and Frédéric Vivien. How should you structure your hierarchical scheduler? In *HPDC-15. 15th IEEE International Symposium on High Performance Distributed Computing*, pages 339–340 (Poster), Paris, France, June 2006. IEEE.

[707] Florent de Dinechin. Elementary functions for double-precision interval arithmetic. In *SCAN 2006 - 12th GAMM - IMACS International Symposium on Scientific Computing, Computer Arithmetic and Validated Numerics*, Duisburg, Germany, 2006.

[708] Romaric Guillier, Ludovic Hablot, Sebastien Soudan, and Pascale Vicat-Blanc Primet. Evaluating transport protocols in very high speed grid networks: How to benefit from huge available bandwidth? poster, SuperComputing 2006, October 2006.

[709] Christoph Quirin Lauter. A survey of multiple precision computation using floating-point arithmetic. In *Taylor Models 2006*, Florida, 2006.

[710] Laurent Lefèvre and Jean-Patrick Gelas. High performance autonomic gateways for large scale grids and distributed environments. Poster INRIA Booth, Supercomputing 2006, Tampa Bay USA, November 2006.

[711] Nathalie Revol. Design choices and their limits for an interval arithmetic library. The example of MPFI. In *SCAN 2006 - 12th GAMM - IMACS International Symposium on Scientific Computing, Computer Arithmetic and Validated Numerics*, Duisburg, Germany, 2006.

[712] Sebastien Soudan and Pascale Vicat-Blanc Primet. Passerelle à base de processeurs réseaux pour le controle des flux de grille. poster, CFIP 2006, October 2006.

2007

[713] Narjess Ayari, Denis Barbaron, Laurent Lefèvre, and Pascale Vicat-Blanc Primet. Implementation of an active replication based framework for highly available services. NetFilter Workshop 2007, Karlsruhe, Germany, September 2007.

[714] Francisco José Cháves Alonso, Marc Daumas, Cesar Muñoz, and Nathalie Revol. Automatic strategies to evaluate formulas on Taylor models and generate proofs in PVS. In *6th International Congress on Industrial and Applied Mathematics (ICIAM'07)*, July 2007.

[715] Nicolas Fournel, Antoine Fraboulet, Guillaume Chelius, Eric Fleury, Bruno Allard, and Olivier Brevet. Worldsens: From lab to sensor network application development and deployment. In *International Conference on Information Processing in Sensor Networks (IPSN), demo session*, Cambridge, MA, USA, April 2007. ACM.

[716] Jean-Patrick Gelas and Laurent Lefèvre. MoonGrid: Bring processing power to the moon. ISU Annual International Symposium: Why the Moon? , Strasbourg, France, February 2007.

[717] Jean-Patrick Gelas, Laurent Lefèvre, and Eric Rohmer. Network support for long distance telerobotic platform. Poster INRIA Booth in collaboration with Tohoku University (Japan), Supercomputing 2007, Reno, USA, November 2007.

[718] Romaric Guillier and Pascale Vicat-Blanc Primet. High speed transport protocol test suite. poster, SuperComputing 2007, November 2007.

[719] Romaric Guillier and Pascale Vicat-Blanc Primet. TCP variants and transfer time predictability in very high speed networks. poster, Ecole d'été RESCOM 2007, session doctorant, June 2007.

[720] Sebastian Hack. Register Allocation for Programs in SSA Form. Date 2007 PhD Forum Poster, April 2007.

[721] André Hirschowitz, Michel Hirschowitz, and Tom Hirschowitz. Playing for truth on mutable graphs. Talk at the Colloquium in honor of Gérard Huet's 60th birthday, ENS, Paris, 2007.

[722] Nathalie Revol. Implementing Taylor models arithmetic with floating-point arithmetic. In *6th International Congress on Industrial and Applied Mathematics (ICIAM'07)*, July 2007.

2008

[723] Florent de Dinechin, Jérémie Detrey, Ionuţ Trestian, Octavian Creţ, and Radu Tudoran. When FPGAs are better at floating-point than microprocessors. Poster at FPGA'2008, 2008.

[724] Romain Demangeon. Type systems for the termination of mobile processes. talk given at the 20th Nordic Workshop on Programming Theory, Tallinn, Estonia, 2008.

[725] Nicolas Farrugia, Michel Paindavoine, and Clément Quinson. On the need for semi-automated source to source transformations in the user-guided high-level synthesis tool. In *High-level Synthesis: Back to the Future (DAC workshop)*, June 2008. Poster.

[726] Romaric Guillier and Pascale Vicat-Blanc Primet. Congestion collapse in Grid5000. demo, Stanford Congestion Collapse workshop, April 2008.

[727] Claude-Pierre Jeannerod, Nicolas Louvet, Nathalie Revol, and Gilles Villard. Computing condition numbers with automatic differentiation. In *SCAN 2008 - 13th GAMM - IMACS International Symposium on Scientific Computing, Computer Arithmetic and Validated Numerics*, El Paso, Texas, 2008.

[728] Claude-Pierre Jeannerod, Nicolas Louvet, Nathalie Revol, and Gilles Villard. On the computation of some componentwise condition numbers. In *SNSC'08 - 4th International Conference on Symbolic and Numerical Scientific Computing*, Hagenberg, Austria, 2008.

[729] Philippe Langlois and Nicolas Louvet. Accurate solution of triangular linear system. In *SCAN 2008 - 13th GAMM - IMACS International Symposium on Scientific Computing, Computer Arithmetic and Validated Numerics*, El Paso, Texas, 2008.

[730] Laurent Lefèvre, Jean-Patrick Gelas, and Anne-Cécile Orgerie. How an experimental grid is used : the Grid5000 case and its impact on energy usage. Poster CCGrid2008 : 8th IEEE International Symposium on Cluster Computing and the Grid, Lyon, France, May 2008.

[731] Laurent Lefèvre and Anne-Cécile Orgerie. EARI : Energy aware resource infrastructure for large scale distributed systems. Poster INRIA Booth, Supercomputing 2008, Austin, USA, November 2008.

[732] Laurent Lefèvre and Anne-Cécile Orgerie. GREEN-NET : Power aware software frameworks for high performance data transport and computing in large scale distributed systems. Poster National ARC Days - INRIA Sophia Antipolis, France, October 2008.

[733] Patrick Loiseau, Paulo Gonçalves, Guillaume Dewaele, Pierre Borgnat, Patrice Abry, and Pascale Vicat-Blanc Primet. Vérification du lien entre auto-similarité et distributions à queues lourdes sur un dispositif grande échelle, June 2008. 9 ième Atelier en Evaluation de Performances, Aussois, France.

[734] Patrick Loiseau, Paulo Gonçalves, Yuetsu Kodama, and Pascale Vicat-Blanc Primet. Metroflux: A fully operational high speed metrology platform, September 2008. Euro-NF workshop: New trends in modeling, quantitative methods and measurements, in cooperation with Net-coop, THOMSON Paris Research Labs, France.

[735] Patrick Loiseau, Paulo Gonçalves, and Pascale Vicat-Blanc Primet. How TCP can kill self-similarity, September 2008. Euro-NF workshop: Traffic Engineering and Dependability in the Network of the Future, VTT, Finland.

[736] Ivan Morel. From an LLL-reduced basis to another. Poster at ISSAC 2008, July 2008.

[737] Hong-Diep Nguyen and Nathalie Revol. Solving and certifying the solution of a linear system. In *SCAN 2008 - 13th GAMM - IMACS International Symposium on Scientific Computing, Computer Arithmetic and Validated Numerics*, El Paso, Texas, 2008.

[738] Nathalie Revol. Automatic adaptation of the computing precision. In *V Taylor Models Workshop*, May 2008.

[739] Nathalie Revol. Survey of proposals for the standardization of interval arithmetic. In *SWIM 08: Small Workshop on Interval Methods*, June 2008.

2009

[740] Gabriel Caillat, Oleg Lodygensky, Gilles Fedak, Haiwu He, Zoltan Balaton, Zoltan Farkas, Gabor Gombas, Peter Kacsuk, Robert Lovas, Attila Csaba Maros, Ian Kelley, Ian Taylor, Gabor Terstyanszky, Tamas Kiss, Miguel Cardenas-Montes, Ad Emmen, and Filipe Araujo. EDGeS: The art of bridging EGEE to BOINC and XtremWeb. In *Proceedings of Computing in High Energy and Nuclear Physics (CHEP'09) (Abstract)*, Prague, Czech Republic, March 2009.

[741] Romain Demangeon. Type systems for the termination of mobile processes. talk given at the 10th International Workshop on Termination, Leipzig, Germany, 2009.

[742] Jean Duprat. Plane geometry in Coq. talk at the TYPES'09 meeting, Aussois, France, 2009.

[743] Yi Gu, Qishi Wu, Anne Benoit, and Yves Robert. Complexity analysis and algorithmic development for pipeline mappings in heterogeneous networks. In *28th ACM Symposium on Principles of Distributed Computing PODC 2009*. ACM Press, 2009. Short communication.

[744] I. Guérin Lassous. Standard pour les réseaux sans fil : IEEE 802.11. Journal Techniques de l'Ingénieur, 2009.

[745] Claude-Pierre Jeannerod, Christophe Mouilleron, and Gilles Villard. Extending Cardinal's algorithm to a broader class of structured matrices. Poster at ISSAC 2009, July 2009.

[746] Patrick Loiseau, Romaric Guillier, Oana Goga, Matthieu Imbert, Paulo Goncalves, and Pascale Vicat-Blanc Primet. Automated traffic measurements and analysis in grid'5000. Best Demonstration award at ACM SIGMETRICS/PERFORMANCE (Seattle, US), June 2009.

[747] Anne-Cécile Orgerie and Laurent Lefèvre. Greening the clouds ! Poster during Rescom 2009 Summer School, La Palmyre, France, June 2009.

[748] Anne-Cécile Orgerie and Laurent Lefèvre. Towards a green Grid5000. Best presentation award of the Grid5000 school, April 2009.

[749] Barbara Petit. A polymorphic type system for the lambda-calculus with constructors. talk at the TYPES'09 meeting, Aussois, France, 2009.

Scientific books and book chapters [OS]

2005

[750] C. Pham and M. Maimour. *Le controle de congestion dans les communications multicast*, chapter 7. TRAITE IC2, Multicast Multimédia sur l'Internet. Hermes-Lavoiser, March 2005.

2006

[751] Gabriel Antoniu, Marin Bertier, Luc Bougé, Eddy Caron, Frédéric Desprez, Mathieu Jan, Sébastien Monnet, and Pierre Sens. GDS: An architecture proposal for a grid data-sharing service. In Vladimir Getov, Domenico Laforenza, and Alexander Reinefeld, editors, *Future Generation Grids*, volume XVIII, CoreGrid Series of *Proceedings of the Workshop on Future Generation Grids November 1-5, 2004, Dagstuhl, Germany*. Springer Verlag, 2006.

[752] E. Caron, F. Desprez, J.-Y. L'Excellent, C. Hamerling, M. Pantel, and C. Puglisi-Amestoy. Use of a network enabled server system for a sparse linear algebra application. In Vladimir Getov, Domenico Laforenza, and Alexander Reinefeld, editors, *Future Generation Grids*, volume XVIII, CoreGrid Series of *Proceedings of the Workshop on Future Generation Grids November 1-5, 2004, Dagstuhl, Germany*. Springer Verlag, 2006.

[753] Jean-Michel Muller. *Elementary Functions, Algorithms and Implementation*. Birkhäuser Boston, 2nd Edition, 2006.

2007

[754] Paul Feautrier. Les compilateurs. In Jean-Éric Pin, editor, *Encyclopédie de l'Informatique*. Vuibert, 2007.

2008

[755] G. Barthe, B. Grégoire, and C. Riba. A Tutorial on Type-Based Termination. In *LERNET 2008 Summer School*, volume 5520 of *LNCS Tutorial Series*, page ? Springer, 2008. To Appear. External publication.

[756] H. Casanova, A. Legrand, and Y. Robert. *Parallel Algorithms*. Chapman & Hall/CRC Press, 2008.

[757] Florent de Dinechin, Milos Ercegovac, Jean-Michel Muller, and Nathalie Revol. *Encyclopedia of Computer Science and Engineering*, chapter Digital Arithmetic. Wiley, 2008.

[758] Paulo Gonçalves, Jean-Philippe Ovarlez, and Richard Baraniuk. *Quadratic Time-Frequency Analysis III: The Affine Class and Other Covariant Classes*, pages 193–226. ISTE – John Wiley & Sons, Inc., 2008.

2009

[759] Olivier Beaumont and Loris Marchal. Steady-state scheduling. In *Introduction to Scheduling*. Chapman and Hall/CRC Press, 2009. To appear.

[760] Vincent Breton, Eddy Caron, Frédéric Desprez, and Gaël Le Mahec. *Handbook of Research on Computational Grid Technologies for Life Sciences, Biomedicine and Healthcare*, chapter High Performance BLAST Over the Grid. IGI Global, 2009.

[761] Yves Caniou, Eddy Caron, Frédéric Desprez, Hidemoto Nakada, Keith Seymour, and Yoshio Tanaka. High performance GridRPC middleware. In G.A. Gravvanis, J.P Morrison, H.R. Arabnia, and D.A. Power, editors, *Grid Technology and Applications: Recent Developments*. Nova Science Publishers, 2009. At Prepress. Pub. Date: 2009, 2nd quarter. ISBN 978-1-60692-768-7.

[762] Franck Cappello, Gilles Fedak, Derrick Kondo, Paul Malécot, and Ala Rezmerita. *Handbook of Research on Scalable Computing Technologies*, chapter Desktop Grids: From Volunteer Distributed Computing to High Throughput Computing Production Platforms. IGI Global, 2009. to appear.

[763] Eddy Caron, Frédéric Desprez, Franck Petit, and Cédric Tedeschi. DLPT: A P2P tool for Service Discovery in Grid Computing. In Nick Antonopoulos, Georgios Exarchakos, Maozhen Li, and Antonio Liotta, editors, *Handbook of Research on P2P and Grid Systems for Service-Oriented Computing: Models, Methodologies and Applications*. IGI Global, December 2009. Released: December 2009. ISBN-13: 978-1615206865.

[764] C. Chaudet and I. Guérin Lassous. *Réseaux de capteurs*, chapter Principes et protocoles d'accès au médium. ISTE/Wiley, 2009.

[765] Patrick Flandrin, Paulo Gonçalves, and Patrice Abry. *Scale Invariance and Wavelets*, pages 71–102. ISTE – John Wiley & Sons, Inc., 2009.

[766] Eric Fleury and Yu Chen. Scheduling Activities in Wireless Sensor Networks. In Sudip Misra, Isaac Woungang, and Subhas Chandra Misra, editors, *Guide to Wireless Sensor Networks Computer Communications and Networks*, Computer Communications and Networks, page 707. Springer, 2009.

[767] Eric Fleury and David Simplot Ryl. *Réseaux de capteurs : théorie et modélisation*. Architecture, Applications, Service. Hermes, 2009.

[768] Matthieu Gallet, Yves Robert, and Frédéric Vivien. Divisible load scheduling. In *Introduction to Scheduling*. Chapman and Hall/CRC Press, 2009. To appear.

[769] Jean-Michel Muller, Nicolas Brisebarre, Florent de Dinechin, Claude-Pierre Jeannerod, Vincent Lefèvre, Guillaume Melquiond, Nathalie Revol, Damien Stehlé, and Serge Torres. *Handbook of Floating-Point Arithmetic*. Birkhauser Boston, 2009. ACM G.1.0; G.1.2; G.4; B.2.0; B.2.4; F.2.1. To appear.

[770] Yves Robert and Frédéric Vivien. Algorithmic Issues in Grid Computing. In *Algorithms and Theory of Computation Handbook*. Chapman and Hall/CRC Press, second edition, 2009. To appear.

[771] Damien Stehlé. *LLL+25: 25th Anniversary of the LLL Algorithm Conference*, chapter Floating-point LLL: theoretical and practical aspects. Springer, 2009. To appear.

Scientific popularization [OV]

2005

[772] Pierre Lescanne. Entretien avec Michel Morvan. *Bulletin de SPECIF*, 54, December 2005.

2006

[773] Pierre Lescanne. Interview de Jacques Cohen. *Bulletin de SPECIF*, 55, March 2006.

[774] Pierre Lescanne. Review of "Alfred Tarski: Life and Logic" by Anita Burdman Feferman and Solomon Feferman, Cambridge University Press 2004. *SIGACT News*, 37(1):27–28, 2006.

2007

[775] Michel Cosnard and Yves Robert. Algorithmique parallèle. In *Encyclopédie de l'Informatique et des Systèmes d'Information*, pages 955–965. Vuibert, 2007.

[776] M. Delorme and J. Mazoyer. La riche zoologie des automates cellulaires. *Interstices*, 2007. http://interstices.info/.

[777] Pierre Lescanne. Entretien avec Gilles Dowek. *Bulletin SPECIF*, 58, December 2007.

[778] Pierre Lescanne. In memoriam John Backus. *Bulletin SPECIF*, 57, April 2007.

[779] Pierre Lescanne. Les propositions vraies sont démontrables. *La Recherche*, 412, oct 2007. special issue *Le dictionnaire des idées reçues en science*.

[780] Jean-Michel Muller. Ordinateurs en quête d'arithmétique. In *Histoire des Nombres*. Taillendier, 2007.

2008

[781] Pierre Lescanne. Interview d'Alexandre Miquel. *Bulletin SPECIF*, 59, April 2008.

[782] Pierre Lescanne. La pensée informatique, traduction d'un article de Jeannette Wing. *Bulletin SPECIF*, December 2008.

2009

[783] I. Guérin Lassous. Standard pour les réseaux sans fil : IEEE 802.11. Journal Techniques de l'Ingénieur, 2009.

[784] Laurent Lefèvre. Towards green computing platforms - vers des plateformes de calcul vertes. *INEDIT : the Newsletter of INRIA*, May 2009.

[785] Vincent Lefèvre and Jean-Michel Muller. Erreurs en arithmétique des ordinateurs. *Images des mathématiques*, June 2009. http://images.math.cnrs.fr/Erreurs-en-arithmetique-des.html.

[786] Pierre Lescanne. Interview de Gérard Berry. *Bulletin SPECIF*, 61:41–44, April 2009.

Book or Proceedings editing [DO]

2005

[787] David Hutchison, Spiros Denazis, Laurent Lefèvre, and Gary Minden. *Proceedings of IWAN2005 : Seventh Annual International Working Conference on Active and Programmable Networks*, volume 4388 of *LNCS*. Nice, France, November 2005.

[788] Craig A. Lee, Thilo Kielmann, Laurent Lefèvre, and Joao Gabriel Silva. *Topic 6 Grid and Cluster Computing: Models, Middleware and Architectures : EuroPar 2005 Parallel Processing*, volume 3648 of *Lecture Notes in Computer Science*. Springer Berlin, Lisboa, Portugal, August 2005. ISBN : 978-3-540-28700-1.

[789] C. Pham and B. Tourancheau. *Grid Infrastructures: practice and perspectives - Future Generation Computer System*, volume 21. January 2005. Editorial.

[790] Pascale Vicat-Blanc Primet, Richard Hughes-Jones, and Cheng Jin. *Proceedings of the 3rd International Workshop on Protocols for Very Long Distance networks*. February 2005.

2006

[791] L. Bao and I. Guérin Lassous, editors. *Third ACM International Workshop on Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks (PE-WASUN)*, ACM, Torremolinos, Spain, October 2006. ACM.

[792] Lionel Brunie, Salim Hariri, Laurent Lefèvre, and Jean-Marc Pierson. *Proceedings of ICPS2006 : International Conference on Pervasive Services*. IEEE, Lyon, France, June 2006.

[793] L. Carter, H. Casanova, F. Desprez, J. Ferrante, and Y. Robert, editors. *Special issue on* Scheduling techniques for large-scale distributed platforms. Int. J. High Performance Computing Applications 20(1), 2006.

[794] L. Carter, H. Casanova, F. Desprez, J. Ferrante, and Y. Robert, editors. *Special issue on* Scheduling techniques for large-scale distributed platforms. Int. J. High Performance Computing Applications 20(4), 2006.

[795] H. Casanova, Y. Robert, and H.J. Siegel, editors. *Special issue on* Algorithm design and scheduling techniques (realistic platform models) for heterogeneous clusters. IEEE Trans. Parallel Distributed Systems 17(2), 2006.

[796] Marc Daumas and Nathalie Revol, editors. *Special issue on Real Numbers and Computers*, volume 351 of *Theoretical Computer Science*, 2006.

[797] René David, Danièle Gardy, Pierre Lescanne, and Marek Zaiuonc, editors. *Computational Logic and Applications, CLA '05*, 2006. http://www.dmtcs.org/dmtcs-ojs/index.php/proceedings/issue/view/83.

[798] Jack Dongarra and Bernard Tourancheau, editors. *Cluster and Computational Grids for Scientific Computing, Special Issue on Tools*. International Journal of High Performance Computing Applications 20(3), 2006.

[799] Y. Robert, M. Parashar, R. Badrinath, and V. Prasanna, editors. *High Performance Computing (HiPC) 2006*, LNCS 4297. Springer Verlag, 2006.

2007

[800] O. Akan and I. Guérin Lassous, editors. *Fourth ACM International Workshop on Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks (PE-WASUN)*, ACM, Chania, Crete Islands, Greece, October 2007. ACM.

[801] Jack Dongarra and Bernard Tourancheau, editors. *Clusters and Computational Grids for Scientific Computing*. Parallel Processing Letters 17(1), March 2007.

[802] Peter Kornerup and Jean-Michel Muller, editors. *Proceedings of the 18th IEEE Symposium on Computer Arithmetic*. IEEE Conference Publishing Services, June 2007.

[803] Laurent Lefèvre and Jean-Marc Pierson. *Special issue from International Conference on Pervasive Services - Journal of System and Software*, volume 80. Elsevier, December 2007.

[804] Joe Touch, Katsushi Kobayashi, and Pascale Vicat-Blanc Primet. *Special issue Hot topics in Transport Protocols for Very Long distance networks - International Journal of Computer Networks (COMNET)*. january 2007.

2008

[805] Philippe Audebaud and Christine Paulin-Mohring, editors. *Mathematics of Program Construction, 9th International Conference, MPC 2008, Marseille, France, July 15-18, 2008. Proceedings*, volume 5133 of *Lecture Notes in Computer Science*. Springer, 2008.

[806] L. Bononi and I. Guérin Lassous, editors. *Fifth ACM International Symposium on Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks (PE-WASUN)*, ACM, Vancouver, Canada, October 2008. ACM.

[807] Jack Dongarra and Bernard Tourancheau, editors. *Special section: Cluster and computational grids for scientific computing*. Future Generation Comp. Syst. 24 (1), 2008.

[808] Peter Hertling, Christoph M. Hoffmann, Wolfram Luther, and Nathalie Revol, editors. *Special issue on Reliable Implementation of Real Number Algorithms: Theory and Practice*, volume 5045 of *Lecture Notes in Computer Science*, 2008.

[809] Zhiyi Huang, Zhiwei Xu, Nathan Rountree, Laurent Lefèvre, Hong Shen, John Hine, and Yi Pan. *Proceedings of PDCAT 2008 : The Ninth International Conference on Parallel and Distributed Computing, Applications and Technologies*. IEEE Computer Society, Dunedin, New Zealand, December 2008.

[810] Thierry Priol, Laurent Lefèvre, and Rajkumar Buyya. *Proceedings of CCGrid2008 : Eight IEEE International Symposium on Cluster Computing and Grid*. IEEE, Lyon, France, May 2008.

[811] Pascale Vicat-Blanc Primet, Joe Mambretti, and Tomohiro Kudoh. *Proceedings of the 2nd International ACM/ICST GRIDNETS conference on Networks for Grid Applications*. October 2008.

2009

[812] Patrice Abry, Paulo Gonçalves, and Jaques Lévy Véhel. *Scaling, Fractals and wavelets*. Digital signal and image processing series. ISTE – John Wiley & Sons, Inc., London (UK) and Hoboken (NJ, USA), 2009.

[813] Philippe Audebaud and Christine Paulin-Mohring, editors. *Special issue dedicated to MPC'08*, Science of Computer Programming. Elsevier, 2009. In preparation.

[814] Patrick Baillot, Jean-Yves Marion, and Simona Ronchi Della Rocca, editors. *Special Issue on Implicit Computational Complexity*, volume 10 of *ACM Transactions on Computational Logic*. ACM, 2009.

[815] A. Boukerche and I. Guérin Lassous, editors. *Sixth ACM International Symposium on Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks (PE-WASUN)*, ACM, Canary Islands, Spain, October 2009. ACM.

[816] Jack Dongarra and Bernard Tourancheau, editors. *Cluster and Computational Grids for Scientific Computing*. Parallel Processing Letters, 2009.

[817] Thomas Ehrhard, Claudia Faggian, and Olivier Laurent, editors. *Jean-Yves Girard's Festschrift*, Special issue of Theoretical Computer Science. Elsevier, 2009. In preparation.

[818] Zhiyi Huang, Zhiwei Xu, Laurent Lefèvre, Hong Shen, John Hine, and Yi Pan. *Special Issue on Emerging Research in Parallel and Distributed Computing - Journal of Supercomputing*. December 2009. To Appear.

[819] Peter Kornerup, Paolo Montuschi, Jean-Michel Muller, and Eric Schwarz, editors. *Special Section on Computer Arithmetic*, volume 58 of *IEEE Transactions on Computers*, February 2009.

[820] Y. Robert, editor. *Special issue on* IPDPS'2008. J. Parallel and Distributed Computing 69, 9, 2009.

[821] Yves Robert and Frédéric Vivien, editors. *Introduction to Scheduling*. Chapman and Hall/CRC Press, 2009. To appear.

[822] Piero Spinnato, Pascale Vicat-Blanc Primet, Chris Edwards, and Michael Welzl, editors. *Special Section on Networks for Grid Applications*. Elsevier, april 2009. International Journal on Future Generation Computer Systems.

[823] Pascale Vicat-Blanc Primet, Tomohiro Kudoh, and Joe Mambretti, editors. *Networks for Grid Applications*. Springer, Beijing, China, june 2009.

Other Publications [AP]

2005

[824] Sylvie Boldo and Cesar Muñoz. A formalization of floating-point numbers in PVS. Technical report, National Institute for Aerospace, 2005. A much shorter version of this paper has been published in *Proceedings of the 21st Annual ACM Symposium on Applied Computing* (2006).

[825] Christoph Quirin Lauter. Basic building blocks for a triple-double intermediate format. Technical Report RR2005-38, LIP, September 2005.

[826] Clément Pernet, Aude Rondepierre, and Gilles Villard. Computing the Kalman form. Research report ccsd-00009558, arXiv cs.SC/0510014, IMAG, Grenoble, October 2005.

2006

[827] Sylvie Boldo and Guillaume Melquiond. Emulation of a FMA and correctly-rounded sums: proved algorithms using rounding to odd. Technical report, INRIA, 2006. Available at http://hal.inria.fr/inria-00080427.

[828] Guillaume Chelius, Antoine Fraboulet, and Eric Fleury. WSNet: A modular event-driven wireless network simulator, 2006. IDDN 06-370013-000.

[829] Sylvain Chevillard. Polynômes de meilleure approximation à coefficients flottants. Master thesis, Laboratoire de l'Informatique du Parallélisme - ENS Lyon, 2006.

[830] Nicolas Fournel. ARM Linux support for ARM integrator CM922T-XA10 in standalone mode. available online, March 2006.

[831] Nicolas Fournel. Mutek OS support for ARM integrator CM922T-XA10 in standalone mode. available online, April 2006.

[832] Antoine Fraboulet, Guillaume Chelius, and Eric Fleury. WSim: A hardware platform simulator, 2006. IDDN 06-370012-000.

[833] Claude-Pierre Jeannerod. LSP matrix decomposition revisited. Research report RR2006-28, Laboratoire de l'Informatique du Parallélisme, ENS Lyon, September 2006.

[834] Christoph Quirin Lauter. Exact and mid-point rounding cases of power(x,y). Technical Report 2006-46, Laboratoire de l'Informatique du Parallélisme, 2006.

[835] Guillaume Revy. Analyse et implantation d'algorithmes rapides pour l'évaluation polynomiale sur les nombres flottants. Master thesis, Laboratoire de l'Informatique du Parallélisme - ENS Lyon, 2006.

[836] V. Vazirani. *Algorithmes d'approximation (2ème édition)*. Springer, collection IRIS, 2006. Traduction de N. Schabanel.

[837] Francois Vigouroux, Bernard Tourancheau, and Cedric Koch-Hofer. Optimal communication path routing in a system employing interconnected integrated circuit technology. US Patent Application 0018261, 2006. External publication.

2007

[838] Narjess Ayari, Denis Barbaron, and Laurent Lefèvre. Procédés de gestion de sessions multi-flux. france telecom r&d patent. June 2007.

[839] Florent de Dinechin, Jérémie Detrey, Ionuţ Trestian, Octavian Creţ, and Radu Tudoran. When FPGAs are better at floating-point than microprocessors. Technical Report ensl-00174627, École Normale Supérieure de Lyon, 2007.

[840] Florent de Dinechin, Christoph Quirin Lauter, and Guillaume Melquiond. Certifying floating-point implementations using Gappa. Technical Report ensl-00200830, arXiv:0801.0523, École Normale Supérieure de Lyon, 2007.

[841] P. Koiran. Decision versus evaluation in algebraic complexity. In *Proceedings of MCU 2007 (Modèles de Calcul Universels)*, Lecture Notes in Computer Science 4664. Springer, 2007. Invited paper.

[842] Yusuke Tanimura, Keith Seymour, Eddy Caron, Abelkader Amar, Hidemoto Nakada, Yoshio Tanaka, and Frédéric Desprez. *Interoperability Testing for The GridRPC API Specification*. Open Grid Forum, May 2007. OGF Reference: GFD.102.

[843] Bernard Tourancheau, Francois Vigouroux, and Cedric Koch-Hofer. System and method for determining a carrier layout using cornered chip-to-chip input/output. US Patent 7191422, 2007. External publication.

[844] Gilles Villard. Certification of the QR Factor R, and of Lattice Basis Reducedness (extended version of [503]). Research report hal-00127059, arXiv:0701183, CNRS, Université de Lyon, INRIA, 2007.

2008

[845] Patrick R. Amestoy, Alfredo Buttari, Philippe Combes, Abdou Guermouche, Jean-Yves L'Excellent, Tzvetomila Slavova, and Bora Uçar. Overview of MUMPS (a multifrontal massively parallel solver), June 2008. Presentation at the 2nd French-Japanese workshop on Petascale Applications, Algorithms and Programming (PAAP'08).

[846] Patrick R. Amestoy, Alfredo Buttari, and Jean-Yves L'Excellent. Towards a parallel analysis phase for a multifrontal sparse solver, June 2008. Presentation at the 5th International workshop on Parallel Matrix Algorithms and Applications (PMAA'08).

[847] Alain Darte. Quelques propriétés mathématiques et algorithmiques des ensembles convexes. énoncé et corrigé de l'épreuve de mathématiques et informatique, concours d'entrée aux ENS de Cachan, Lyon et Ulm, session 2008. *Revue de Mathématiques Spéciales*, 119(1), 2008.

[848] Alain Darte and Rob Schreiber. System and method of optimizing memory usage with data lifetimes. US patent number 7363459, April 2008.

[849] Guillaume Hanrot and Damien Stehlé. Worst-case Hermite-Korkine-Zolotarev reduced lattice bases. Technical report, Maths Arxiv, 2008.

[850] Claude-Pierre Jeannerod, Hervé Knochel, Christophe Monat, and Guillaume Revy. Computing floating-point square roots via bivariate polynomial evaluation. Technical Report 2008-38, ensl-00335792, LIP, École Normale Supérieure de Lyon, 2008.

[851] Mioara Joldes. Computation of various infinite norms, in one or several variables, designed for the development of mathematical libraries. Master thesis, Laboratoire de l'Informatique du Parallélisme - ENS Lyon, 2008.

[852] Bogdan Pasca. Customizing floating-point operators for linear algebra acceleration on FPGAs. Master thesis, Laboratoire de l'Informatique du Parallélisme - ENS Lyon, 2008.

[853] Bernard Tourancheau. Method and apparatus for routing data across an n-dimensional grid network. US Patent 7444424, 2008. External publication.

[854] Bernard Tourancheau, Ronald Ho, and Robert Drost. Method and apparatus for performing error-detection and error-correction. US Patent 7395483, 2008. External publication.

[855] Gilles Villard. Kaltofen's division-free determinant algorithm differentiated for matrix adjoint computation (journal submission). Research report ensl-00335918, arXiv:0810.5647, CNRS, Université de Lyon, INRIA, 2008.

2009

[856] Patrick R. Amestoy, Iain S. Duff, Daniel Ruiz, and Bora Uçar. Towards parallel bipartite matching algorithms. Presentation at Scheduling for large-scale systems, Knoxville, Tennessee, USA, May 2009.

[857] Patrick R. Amestoy, Iain S. Duff, Tzvetomila Slavova, and Bora Uçar. Out-of-core solution for singleton rhs vectors. Presentation at SIAM Conference on Computational Science and Engineering (CSE09), Miami, Florida, USA, March 2009.

[858] Patrick R. Amestoy, Iain S. Duff, Tzvetomila Slavova, and Bora Uçar. Out-of-core solution for singleton right hand-side vectors. Poster at Dagstuhl Seminar on Combinatorial Scientific Computing, February 2009.

[859] Sylvain Chevillard. The functions erf and erfc computed with arbitrary precision. Technical Report RR2009-04, Laboratoire de l'Informatique du Parallélisme (LIP), 46, allée d'Italie, 69 364 Lyon Cedex 07, January 2009.

[860] Romaric Guillier and Pascale Vicat-Blanc Primet. PATHNIF. INRIA patent. 2009.

[861] Claude-Pierre Jeannerod, Nicolas Louvet, Jean-Michel Muller, and Adrien Panhaleux. Midpoints and exact points of some algebraic functions in floating-point arithmetic. Technical Report ensl-00409366, LIP, École Normale Supérieure de Lyon, 2009.

[862] Claude-Pierre Jeannerod and Guillaume Revy. Optimizing correctly-rounded reciprocal square roots for embedded VLIW cores. Technical Report 2009-21, ensl-00391185, LIP, École Normale Supérieure de Lyon, 2009.

[863] Ümit V. Çatalyürek and Bora Uçar. Partitioning sparse matrices. Presentation at SIAM Conference on Computational Science and Engineering (CSE09), Miami, Florida, USA delivered by Ümit V. Çatalyürek, March 2009.

Doctoral Dissertations and Habilitation Theses [TH]

2005

[864] F. Chavanon. *Aspects combinatoires des pavages*. PhD thesis, Ecole Normale Supérieure de Lyon, 2005.

[865] Brice Goglin. *R seaux rapides et stockage distribu dans les grappes de calculateurs : propositions pour une interaction efficace*. PhD thesis, cole normale sup rieure de Lyon, 46, all e d'Italie, 69364 Lyon cedex 07, France, October 2005. 194 pages.

[866] E. Jeandel. *Techniques algébriques en calcul quantique*. PhD thesis, Ecole Normale Supérieure de Lyon, 2005.

[867] Silvia Likavec. *Types for object oriented and functional programming languages*. PhD thesis, ENS Lyon and Università di Torino, 2005.

2006

[868] Pushpinder Kaur Chouhan. *Automatic Deployment for Application Service Provider Environments*. PhD thesis, École normale supérieure de Lyon, France, September 2006.

[869] Julien Laganier. *Architecture de sécurité décentralisée basé sur l'identification cryptographique*. PhD thesis, ENS-Lyon, Université de Lyon, 46, allée d'Italie, 69364 Lyon cedex 07, France, 2006.

[870] Loris Marchal. *Communications collectives et ordonnancement en régime permanent sur plates-formes hétérogènes*. PhD thesis, École Normale Supérieure de Lyon, France, October 2006.

[871] Guillaume Melquiond. *De l'arithmétique d'intervalles à la certification de programmes*. PhD thesis, École Normale Supérieure de Lyon, France, November 2006.

[872] B. Nouvel. *Rotations discrètes et automates cellulaires*. PhD thesis, Ecole Normale Supérieure de Lyon, 2006.

[873] V. Poupet. *Automates cellulaires : temps réel et voisinages*. PhD thesis, Ecole Normale Supérieure de Lyon, 2006.

[874] Saurabh Kumar Raina. *FLIP: a floating-point library for integer processors*. PhD thesis, École Normale Supérieure de Lyon, France, September 2006.

[875] Antoine Scherrer. *Analyses statistiques des communications sur puce*. PhD thesis, École normale supérieure de Lyon, December 2006.

[876] Antoine Vernois. *Ordonnancement et réplication de données bioinformatiques dans un contexte de grille de calcul*. PhD thesis, École normale supérieure de Lyon, France, October 2006.

2007

[877] Francisco José Cháves Alonso. *Utilisation et certification de l'arithmétique d'intervalles dans un assistant de preuves*. PhD thesis, École Normale Supérieure de Lyon, France, September 2007.

[878] Hadda Cherroun. *Scheduling for High-Level Synthesis*. PhD thesis, Université des Sciences et de la Technologie Houari Boumediene, Alger, December 2007.

[879] Florent de Dinechin. *Matériel et logiciel pour l'évaluation de fonctions numériques. Précision, performance et validation*. Habilitation à diriger des recherches, Université Claude Bernard - Lyon 1, June 2007.

[880] Jérémie Detrey. *Arithmétiques réelles sur FPGA : virgule fixe, virgule flottante et système logarithmique*. PhD thesis, École Normale Supérieure de Lyon, France, January 2007.

[881] Nicolas Fournel. *Estimation et optimisation de performances temporelles et énergétiques pour la conception de logiciels embarqués*. PhD thesis, École normale supérieure de Lyon, November 2007.

[882] Philippe Grosse. *Gestion dynamique des tâches dans une architecture microélectronique intégrée à des fins de basse consommation*. PhD thesis, École normale supérieure de Lyon, December 2007.

[883] Romain Kervarc. *Systèmes de types purs et substitutions explicites*. PhD thesis, Ecole Normale Supérieure de Lyon, 2007.

[884] L. Lyaudet. *Graphes et hypergraphes : complexités algorithmique et algébrique*. PhD thesis, Ecole Normale Supérieure de Lyon, 2007.

[885] V. Nesme. *Complexité en requêtes et symétries*. PhD thesis, Ecole Normale Supérieure de Lyon, 2007.

[886] Louvet Nicolas. *Algorithmes compensés en arithmétique flottante : précision, validation, performances.* PhD thesis, Université de Perpignan, France, November 2007.

[887] S. Perifel. *Problèmes de décision et d'évaluation en complexité algébrique*. PhD thesis, Ecole Normale Supérieure de Lyon, 2007.

[888] Nicolas Veyrat-Charvillon. *Opérateurs arithmétiques matériels pour des applications spécifiques*. PhD thesis, École Normale Supérieure de Lyon, France, June 2007.

[889] Dragisa Zunic. *Computing with sequents and diagrams in classical logic*. PhD thesis, Ecole Normale Supérieure de Lyon, 2007.

2008

[890] Emmanuel Agullo. *On the Out-of-core Factorization of Large Sparse Matrices*. PhD thesis, École Normale Supérieure de Lyon, November 2008.

[891] Narjess Ayari. *Contributions to Session Aware Frameworks for Next Generation Internet Services*. PhD thesis, École normale supérieure de Lyon, 46, allée d'Italie, 69364 Lyon cedex 07, France, October 2008. 153 pages.

[892] F. Becker. *Géométrie pour l'auto-assemblage*. PhD thesis, Ecole Normale Supérieure de Lyon, 2008.

[893] Raphaël Bolze. *Analyse et déploiement de solutions algorithmiques et logicielles pour des applications bioinformatiques à grande échelle sur la grille*. PhD thesis, École Normale Supérieure de Lyon, October 2008.

[894] Christoph Quirin Lauter. *Arrondi correct de fonctions mathématiques - fonctions univariées et bivariées, certification et automatisation*. PhD thesis, École Normale Supérieure de Lyon, France, October 2008.

[895] Romain Michard. *Opérateurs arithmétiques matériels optimisés*. PhD thesis, École Normale Supérieure de Lyon, France, June 2008.

[896] Dino M. Lopez P. *Propositions for a robust and inter-operable eXplicit Control Protocol on heterogeneous high speed networks*. PhD thesis, École normale supérieure de Lyon, 46, allée d'Italie, 69364 Lyon cedex 07, France, June 2008. 137 pages.

[897] Jean-François Pineau. *Communication-aware scheduling on heterogeneous master-worker platforms*. PhD thesis, École Normale Supérieure de Lyon, September 2008.

[898] Damien Pous. *Techniques modulo pour les bisimulations*. PhD thesis, Ecole Normale Supérieure de Lyon, 2008.

[899] D. Regnault. *Sur les automates cellulaires probabilistes : comportements asynchrones*. PhD thesis, Ecole Normale supérieure de Lyon, 2008.

[900] Jean-Baptiste Rouquier. *Robustesse et émergence dans les systémes complexes: le modéle des automates cellulaires*. PhD thesis, Ecole Normale Supérieure de Lyon, December 2008.

[901] Stéphane Le Roux. *Generalisation and formalisation in game theory*. PhD thesis, Ecole Normale Supérieure de Lyon, 2008.

[902] Cédric Tedeschi. *Peer-to-Peer Prefix Tree for Large Scale Service Discovery*. PhD thesis, École normale supérieure de Lyon, October 2008.

[903] Frédéric Vivien. *On scheduling for distributed heterogeneous platforms*. Habilitation à diriger des recherches, École normale supérieure de Lyon, May 2008.

2009

[904] Anne Benoit. *Scheduling pipelined applications: models, algorithms and complexity*. Habilitation à diriger des recherches, École normale supérieure de Lyon, July 2009.

[905] Florent Bouchez. *A Study of Spilling and Coalescing in Register Allocation as Two Separate Phases*. PhD thesis, École normale supérieure de Lyon, April 2009.

[906] Hugo Carrao. *The contribution of multitemporal information from multispectral satellite images for automatic land cover classification at the national scale*. PhD thesis, Université de Lisbonne (Portugal), 2009.

[907] Sylvain Chevillard. *Évaluation efficace de fonctions numériques - Outils et exemples*. PhD thesis, École Normale Supérieure de Lyon, France, July 2009.

[908] Paulo Gonçalves. *Some scale based tools and applications*. Habilitation à diriger des recherches, École doctorale de mathématiques et d'informatique fondamentale de Lyon, 2009. In preparation.

[909] Romaric Guillier. *Methodologies and Tools for the Evaluation of Transport Protocols in the Context of Highspeed Networks*. PhD thesis, ENS-Lyon, Université de Lyon, September 2009.

[910] Ludovic Hablot. *Exécution d'applications MPI dans les réseaux de grille*. PhD thesis, ENS Lyon, Université de Lyon, 2009.

[911] Olivier Laurent. *Investigations classiques, complexes et concurrentes à l'aide de la logique linéaire*. Habilitation à diriger des recherches, Université Paris Diderot – Paris 7, 2009. In preparation.

[912] Patrick Loiseau. *Contribution to the analysis of scaling behavior and quality of service in networks: experimental and theoretical aspects*. PhD thesis, École Normale Supérieure de Lyon, 2009.

[913] Veronika Rehn-Sonigo. *Multi-criteria Mapping and Scheduling of Workflow Applications onto Heterogeneous Platforms*. PhD thesis, École Normale Supérieure de Lyon, July 2009.

[914] Vannier Rémi. *Partage de la bande passante équitable dans les réseaux ad hoc*. PhD thesis, ENS Lyon, INRIA, 2009.

[915] Guillaume Revy. *"Mathematical functions in floating-point arithmetic for integer processors - Algorithms, implementation, polynomial evaluation and code generation*. PhD thesis, École Normale Supérieure de Lyon, France, 2009.

[916] Sebastien Soudan. *Bandwidth Sharing and Control in High-Speed Networks: Combining Packet- and Circuit-Switching Paradigms*. PhD thesis, ENS-Lyon, Université de Lyon, 46, allée d'Italie, 69364 Lyon cedex 07, France, 2009.

Other Productions: Software [AP]

[917] Adage: Automatic deployment of applications in a grid environment.
http://adage.gforge.inria.fr

[918] E. Boix, J. Beltran, G. Chiquillo, A. Grignard, M. Morvan, M. Malaterre, and R. Uribe. Plateforme de simulation des systemes complexes.

[919] David Cadé, Xavier Pujol, and Damien Stehlé. fplll, 2005–2009.
http://perso.ens-lyon.fr/damien.stehle

[920] Martine Chaudier, Jean-Patrick Gelas, and Laurent Lefèvre. Tamanoir embedded : an industrial active networking framework, 2006.

[921] Sylvain Chevillard, Mioara Joldes, Nicolas Jourdan, and Christoph Lauter. Sollya, a toolbox for manipulating numerical functions, 2006–2009.
http://sollya.gforge.inria.fr/

[922] Sylvain Collange, Jérémie Detrey, Florent de Dinechin, Bogdan Pasca, Cristian Klein, and Xavier Pujol. FloPoCo, a generator of arithmetic cores for FPGAs, 2008–2009.
http://www.ens-lyon.fr/LIP/Arenaire/Ware/FloPoCo/

[923] Sylvain Collange, Jérémie Detrey, Florent de Dinechin, and Xavier Pujol. FPLibrary, a library of floating-point operators for FPGAs, 2002–2008.
http://www.ens-lyon.fr/LIP/Arenaire/Ware/FPLibrary/

[924] Catherine Daramy-Loirat, David Defour, Florent de Dinechin, Matthieu Gallet, Nicolas Gast, Christoph Lauter, and Jean-Michel Muller. CRlibm, a library of correctly rounded elementary functions in double-precision, 2004–2009.
http://lipforge.ens-lyon.fr/www/crlibm/

[925] DIET: Distributed Interactive Engineering Toolbox.
http://graal.ens-lyon.fr/DIET

[926] Jean Duprat. Ruler and compass geometry axiomatization, 2007. Coq contribution.
http://coq.inria.fr/contribs-eng.html

[927] Laurent Fousse, Guillaume Hanrot, Vincent Lefèvre, Patrick Pélissier, Philippe Théveny, and Paul Zimmermann. GNU MPFR, March 2000.
http://www.mpfr.org/

[928] Jean-Patrick Gelas and Laurent Lefèvre. Tamanoir :a high performance active networking framework, 2006.

[929] Jean-Patrick Gelas, Laurent Lefèvre, and Anne-Cecile Orgerie. Showwatts: Real time energy consumption grapher, 2008.

[930] Jean-Patrick Gelas, Laurent Lefèvre, and Anne-Cecile Orgerie. Wattm : Monitoring framework for energy consumption of data centers, 2009.

[931] The LinBox Group. Linbox, 1998–2009.
http://www.linalg.org

[932] Romaric Guillier and Pascale Vicat-Blanc Primet. Nxe, 2009.

[933] Romaric Guillier and Pascale Vicat-Blanc Primet. Pathnif, 2009.

[934] Claude-Pierre Jeannerod, Nicolas Jourdan, Saurabh Kumar Raina, Guillaume Revy, and Arnaud Tisserand. FLIP, a floating-point library for integer processors, 2003–2009.
http://flip.gforge.inria.fr/

[935] Guilherme Koslovski and Pascale Vicat-Blanc Primet. Vxdl parser, 2009.

[936] S. Lagrange. COINC (Computational issues on Network Calculus). Contribution de MC2 (E. Thierry) pour la partie algorithmique. http://www.istia.univ-angers.fr/~lagrange/spip.php?article21.

[937] Christoph Lauter. Metalibm, a code generator for elementary functions, 2008.
http://lipforge.ens-lyon.fr/www/metalibm/

[938] Pierre Lescanne. Coq development of common knowledge logic. Corresponds to article [62], 2006. Coq contribution.
http://perso.ens-lyon.fr/pierre.lescanne/COQ/EpistemicLogic.v8

[939] Pierre Lescanne. Coq development of infinite extensive games and escalation. Corresponds to report http://prunel.ccsd.cnrs.fr/ensl-00376141/fr/, 2009. Coq contribution.
http://perso.ens-lyon.fr/pierre.lescanne/COQ/INFGAMES/infinite_extensive_games.html

[940] Dino Lopez Pacheco, Anne-Cecile Orgerie, and Laurent Lefèvre. Xcp-i : Interoperable explicit control protocol, 2006.

[941] Guillaume Melquiond. Gappa, a tool for certifying numerical programs, 2005–2009.
http://lipforge.ens-lyon.fr/www/gappa/

[942] Alexandre Miquel. The classical extraction module for Coq, 2009. Software: Coq module.
http://perso.ens-lyon.fr/alexandre.miquel/kextraction

[943] Alexandre Miquel. The Jivaro head reduction machine for the λ_c-calculus, 2009. Software.
http://perso.ens-lyon.fr/alexandre.miquel/jivaro

[944] Olivier Mornard and Pascale Vicat-Blanc Primet. Hipernet, 2009.
http://www.ens-lyon.fr/LIP/RESO/Software/hipernet/index.html

[945] MUMPS: a MUltifrontal Massively Parallel Solver.
http://graal.ens-lyon.fr/MUMPS

[946] Pablo Neira Ayuso and Laurent Lefèvre. Sne : Stateful network equipment, 2006.

[947] Marcelo Pasin, Sebastien Soudan, and Pascale Vicat-Blanc Primet. Buld data transfer scheduling service, 2008.
http://www.ens-lyon.fr/LIP/RESO/software/bdts

[948] Pablo Pazos Rey and Laurent Lefèvre. Lscan : Large scale deployment of autonomic networks, 2006.

[949] Damien Pous. GcPan: an Efficient Abstract Machine for Safe Ambients, 2005. Software.
http://sardes.inrialpes.fr/~pous

[950] Damien Pous. A coq formalisation of up-to techniques for weak bisimulation, 2007. Coq contribution.
http://sardes.inrialpes.fr/~pous

[951] J.-B. Rouquier. CimulA, a cellular automata analyser. Contribution de F. Becker. http://cimula.sourceforge.net/.

[952] Sebastien Soudan and Romaric Guillier. Tcp raw, 2009.

[953] Sebastien Soudan and Pascale Vicat-Blanc Primet. Floc : Flow control, 2009.

[954] Sebastien Soudan and Pascale Vicat-Blanc Primet. Vx calendar, 2009.

[955] Sebastien Soudan and Pascale Vicat-Blanc Primet. Vx scheduler, 2009.

[956] Sebastien Soudan and Pascale Vicat-Blanc Primet. Vx topology, 2009.

[957] Ulcmi: Unified lego component model implementation. APP in progress.

Interteam Publications

Anne Benoit, Eric Thierry, and Yves Robert. On the complexity of mapping linear chain applications onto heterogeneous platforms. *Parallel Processing Letters*, 19(3):383–397, 2009.

Alain Darte, Rob Schreiber, and Gilles Villard. Lattice based memory allocation. *IEEE Transactions on Computers*, 54(10):1242–1257, 2005.

Loris Marchal, Pascale Primet, Yves Robert, and Jingdi Zeng. Optimizing network resource sharing in grids. In *IEEE Global Telecommunications Conference GlobeCom'2005*. IEEE Communications Society Press, 2005.

Loris Marchal, Pascale Primet, Yves Robert, and Jingdi Zeng. Optimal bandwidth sharing in grid environment. In *HPDC'2006, the 15th International Symposium on High Performance Distributed Computing*. IEEE Computer Society Press, 2006.