

Curriculum Vitæ

Arnaud Legrand

Laboratoire d'Informatique de Grenoble
UMR 5217 (CNRS, INRIA, INPG, UJF, UPMF)

January, 2018

Personal Information

Last, First, and Middle Names: Legrand, Arnaud, Régis

Date of Birth: 02-20-1977

Place of Birth: Créteil (Val de Marne 94), France

Nationality: French

Marital Status: Married, five children

Business Address: LIG - Bâtiment IMAG
791 rue des Résidences, Domaine Universitaire
38401 St Martin d'Hères, France

Business Phone Number: (+33) 4 57 42 15 22

E-Mail Address: Arnaud.Legrand@imag.fr

Web Page: <https://team.inria.fr/polaris/members/arnaud-legrand/>



Contents

Personal Information	1
Education and Professional Experience	2
Research Interests	2
Administrative Responsibilities	4
Scientific Evaluation Responsibilities	5
Research Supervision	6
Software Development	8
Scientific Collaborations and Grants	9
Stays In Foreign Countries	11
Invited Talks	11
Teaching	13
Publications	14

Education and Professional Experience

Jan. 2016 : Leader of the POLARIS research team

Nov. 2015 : Habilitation à diriger des recherches, Université Grenoble Alpes

Thesis: *Scheduling for Large Scale Distributed Computing Systems: Approaches and Performance Evaluation Issues*.

Oct. 2004-... : Tenured Researcher for the CNRS (*Chargé de Recherche*) at Laboratoire d'Informatique de Grenoble (formerly known as "Informatique et Distribution" laboratory).

2003-2004: Post-Doctoral Research Associate, École Normale Supérieure de Lyon (France).

2000-Dec. 2003: Ph.D. Computer Science, École Normale Supérieure de Lyon. Laboratoire de l'Informatique du Parallélisme.

Thesis: *Heterogeneous parallel algorithms and scheduling : static and dynamic approaches*

Advisors: Prof. Olivier BEAUMONT and Prof. Yves ROBERT.

1999-2000; Diplôme d'études approfondies (DEA/M.Sc.) in Fundamental Computer Science (*Operating Systems, Networks and Parallel Algorithms* section), École Normale Supérieure de Lyon, France.

Thesis: *Parallel linear algebra kernels: heterogeneous and non-dedicated environments*.

Advisors: Prof. Olivier BEAUMONT and Prof. Yves ROBERT.

June-July 1999: Two month internship at University of California, San Diego (USA).

Thesis: *Scheduling Heuristics for Parameter-Sweep Applications on a Grid computing Platform*.

Advisors: Prof. Fran BERMAN and Prof. Henri CASANOVA.

1997-2000 Magistère d'Informatique et de Modélisation (B.Sc. in Computer Science and Modeling); École Normale Supérieure de Lyon, France.

Research Interests

My research targets the management (mostly from an algorithmic point of view, i.e., scheduling, load balancing, fairness, game theory...) and performance evaluation (in particular through simulation, visualization, statistical analysis, ...) of large scale distributed computing infrastructures such as clusters, grids, desktop grids, volunteer computing platforms, clouds,... when used for scientific computing.

Although my motivations are quite practical, my work is mostly theoretical but done in connection with practitioners whenever possible in order to validate my modeling assumptions. In this context, I consider computer science to have a strong experimental component just like natural sciences and I contribute accordingly to computer research methodology.

Scheduling for Distributed Platforms

I study **scheduling** problems arising on distributed platforms (like computing grids) with a particular emphasis on **heterogeneity** and **multi-user** issues, hence some background in **game theory**.

During my PhD thesis, I initially worked on scheduling and parallel algorithms for dense linear algebra kernels on heterogeneous platforms [22, 19] but my main results have been obtained in the context of **steady-state scheduling**, i.e., throughput optimization instead of more classical makespan minimization [18, 15, 13, 14] and **divisible load scheduling** [17, 12]. These

two models are relaxed versions of more classical scheduling frameworks and allow to easily account for key platform characteristics such as heterogeneity or complex topology while providing efficient practical solutions.

Since then, I have been particularly interested in trying to incorporate a notion of user in such scheduling problems, in particular using game theory notions:

- Centralized response time optimization [46, 50, 10],
- Max-min fair [45, 9] and distributed proportionally fair throughput optimization [40, 5],
- Non-cooperative throughput [44, 43, 7] and conjoint response time optimization [35].

Some of this work has been part of the [ANR ALPAGE](#) and the [ANR DOCCA](#). It is also the result of my participation to the [CloudShare](#) and [CloudComputing@home](#) associated teams with Berkeley (see Collaboration section on page 9 for more details).

Simulation and Performance Evaluation of Distributed Platforms

Since 2000, I am one of the main architects of the [SimGrid](#) project. SimGrid is a **simulation** toolkit for building simulators of distributed applications (originally designed for scheduling algorithm evaluation purposes). This software is developed in collaboration with [Henri Casanova](#), [Martin Quinson](#) and [Frédéric Suter](#) and has been the basis of at least 190 published articles in the last 12 years. Beside our long term investment and support, we try to provide high-quality software with **predictive** capabilities and to study the **validity** of our simulations. Predictive simulations are common in physics, chemistry or engineering and have revolutionized these fields but are rarely seen in parallel systems. Faithful simulations can yet be used to solve capacity planning problems, for co-design, to study resource usage, etc. The motivation behind the SimGrid research effort is to improve the experimental methodology and practices in the high performance computing field.

This software, which was initially mostly an incidental tool serving scheduling research, has finally turned under my influence into a real research vehicle around 2009. My most notable scientific contributions in the last 8 years in this software are the following:

- Deep assessment of the validity of fluid network models [6, 39],
- Fast and scalable implementation of fluid network models [37, 34],
- Reliable performance prediction capabilities in the context of complex HPC applications leveraging MPI [1, 24, 23, 82] or StarPU [3, 28]

A general overview of some of these contributions is given in [4].

SimGrid has been used at CERN for capacity planning and optimization of file replication strategies, in the biomed organization of the EGI grid to optimize task scheduling and check-pointing, by colleagues from Inria Bordeaux to perform daily non regression testing of their software, in Japan (NII) to design HPC optical networks, by Brazilian researchers to improve map-reduce scheduling strategies, etc.

All this work has been supported by the [INRIA](#) through 3 two year Technological Development Actions in 2006, 2010 (, which I coordinated), and 2017 and by the [ANR USS-SimGrid project](#) and [ANR SONGS project](#) (see Collaboration section on page 9 for more details). We used this framework in the context of the European [FP7-ICT Mont-Blanc 1 & 2](#) and [H2020 HPC4E](#) projects to study the performance and the energy consumption of real HPC applications. Such work has been regularly presented at the [Joint Laboratory on Extreme-Scale Computing](#) (Inria, Argonne National Labs, Barcelona Supercomputer Center, Julich, ...).

Scientific Methodology: Reproducible Research

In the beginning of my career, I have been commonly struggling with lack of information when trying to build on previous research results. Most of the time, too few information were provided (e.g., algorithmic details, experimental workload, or statistical procedures) in articles, which prevented me to know precisely what the authors had done and therefore to obtain similar results I could safely compare to. Yet as I felt my own work was also far from perfect from this point of view, I have thus decided to invest a significant part of my time on improving research methodology, in particular with the help of the **ANR SONGS project** where I played a major scientific coordination role.

In the last six years, I have been actively promoting better experimental practices and scientific methodology through numerous tutorials and keynotes in conferences and summer schools. To facilitate the dissemination of ideas, I have organized several webinars (all the slides and links to the videos are available on github¹) on reproducible research. I have conducted this promotion of **laboratory notebooks** and **reproducible research** along with the development of lightweight provenance tracking techniques suited to the parallel computing community [2, 25], which I have applied to all my recent work. As a consequence all the articles I have written in the last five years have been published along with a perennial repository comprising all our experimental data and meta-data as well as a clear explanation of the analysis workflow². As a consequence, every figure can be regenerated and it is possible to track down the description of the experimental conditions behind each point. To popularize this approach, I am currently designing a MOOC on reproducible research.

In parallel, since about 2009, I have also been working on **visualization** and **trace analysis** with a particular emphasis on the ability to build sound representations of very large systems through meaningful aggregations [36, 75, 32, 8]. More recently, building on the experience obtained in the context of reproducible research, I have proposed radically more agile data analysis workflows that allow to better understand the behavior of highly dynamic applications for which standard visualization tools are inappropriate [26, 83].

Administrative Responsibilities

- Contributor ("réfèrent") for the HRS4R labeling of Inria (2018). In charge of the animation and the improvement of tracability and reproducibility of research works.
- Coordinator of the "Reproducible Research" challenge (1 out of 22) of the 2018-2022 Inria Strategic Plan (2017).
- Leader of the Inria Project Laboratory HAC SPECIS (2016-...)
- **Leader of the Inria POLARIS research team** (2016-...). The team comprises 25 members: 11 permanent members and between 10 and 14 PhD students and postdocs.
- Coordination of the evaluation of the Parallel and Distributed Computing theme (9 research teams) at Inria (2016).
- Chargé de Mission by the CNRS of the "Distributed Systems, Parallel Computing and Networking" theme of the LIG (supervision of 2nd year PhD students, scientific animation) (2012-2015)
- Adjunct responsible of the Parallel, Distributed and Embedded Systems option of the Master Of Science in Informatics at Grenoble (2011-2016)

¹https://github.com/alegrand/RR_webinars/blob/master/README.org

²See for example https://github.com/Ezibenroc/simulating_mpi_applications_at_scale/blob/master/README.md for the laboratory notebook underlying one of my most recent work [82].

- Member of the CUMI (commission des moyens informatiques) of the LIG (2006-2013)
- Member of the webmasters working group for the LIG laboratory (2007-2013)
- Responsible of the seminars in the ID/LIG laboratory for the MESCAL/MOAIS/NANOSIM teams (2005 - 2015)

Scientific Evaluation Responsibilities

I participated several times to the evaluation of projects for the ANR or for the CAPES/COFECUB. I regularly do numerous reviews (around 25 to 35 reviews a year) for various journals (JPDC, ParCo, TPDS, JoGC, TKDE) and conferences (IPDPS, CCGrid, Grid, EuroPar, BDA, ISPA, ICCS, HiPC, ...).

I served in the program committee of the following conferences:

- 2017: CCgrid, Cluster, PPAM
- 2016: HIPC, EuroPar (head of the performance evaluation track), ICPP, COMPAS, PPAM
- 2014: ICPP, HIPC, COMPAS
- 2013: ICPP, IPDPS, PPAM, COMPAS
- 2012: IPDPS, Renpar
- 2011: CCgrid
- 2010: PCGrid
- 2009: HPDC, Renpar, PCGrid (Workshop on Desktop Grids and Volunteer Computing Systems)
- 2008: Grid, IPDPS, Renpar, PCGrid
- 2004: Grid

I also organized the following workshops:

- 2014-2017: International Workshop on Reproducibility in Parallel Computing (RepPar) in conjunction with Euro-Par
- 2007: Workshop on Programming Models for Grid Computing, held in conjunction with CCGrid'07 (7th IEEE International Symposium on Cluster Computing and the Grid)

Thesis Committees

I have been member of the jury for the following PhD thesis.

- Tien-Dat Phan (reviewer): Doctorat informatique de l'École Normale Supérieure de Rennes. Nov. 2017. *Energy-efficient Straggler Mitigation for Big Data Applications on the Clouds.*
- Rafife Nheili: Doctorat informatique de l'Université de Perpignan Via Domitia. Dec. 2016. *How to improve the numerical reproducibility of hydrodynamics simulations: analysis and solutions for one open-source HPC software*
- Harald Servat, Doctorat informatique de l'Universitat Politècnica de Catalunya, Apr. 2015. *Towards instantaneous performance analysis using coarse-grain sampled and instrumented data.*

- Sorina Camarasu Pop, Doctorat informatique de l'INSA Lyon, novembre 2013. *Exploitation d'infrastructures hétérogènes de calcul distribué pour la simulation Monte-Carlo dans le domaine médical.*
- Javier Celaya, Doctorat informatique de l'Université de Zaragoza, Dec. 2013. *STaRS: A Scalable Task Routing Approach to Distributed Scheduling.*
- Matthieu Pérotin, Doctorat informatique de l'Université François Rabelais, Tours, Dec. 2008. *Calcul haute performance sur matériel générique.*
- Maxime Martinasso, Doctorat informatique système et communication, Université Joseph Fourier, Grenoble (LIG), 2007. *Analyse et modélisation des communications concurrentes dans les réseaux haute-performance.*
- Feryal Moulai, Doctorat informatique système et communication, Laboratoire d'Informatique de Grenoble (LIG), 2007. *Nouvelles approches pour l'ordonnancement d'applications parallèles sous contraintes de déploiement d'environnements sur grappes.*

Hiring Committees (for Assistant Professor positions)

- 2012: Mcf ENSEIRB 0136, Bordeaux University.
- 2010: Chaire CNRS/G-SCOP 1163, Grenoble University.

Research Supervision

PhD. Students

- Pedro Bruel (co-tutelle with USP 2017-...): Design of experiments and autotuning of HPC computation kernels (co-advised with Alfredo Goldman and Brice Videau, funded by the Brazilian Government).
- Tom Cornebize (2017-...): Capacity planning and performance evaluation of supercomputers (funded by the French Ministry for Research).
- Bruno Luis de Moura Donassolo (CIFRE Orange 2017-...): Decentralized management of applications in Fog computing environments (co-advised with Panayotis Mertikopoulos and Ilhem Fajari, funded by Orange).
- Vinicius Garcia Pinto (co-tutelle with UFRGS 2013-...): Performance analysis and visualization of dynamic task-based applications (co-advised with Lucas Schnorr and Nicolas Maillard, funded by the Brazilian government).
- Rafael Tesser (co-tutelle with UFRGS 2013-...): Simulation and performance evaluation of dynamical load balancing of an over-decomposed Geophysics application (co-advised with Lucas Schnorr and Philippe Navaux, funded by the Brazilian government).
- Christian Heinrich (2015-...): Modeling of performance and energy consumption of HPC systems (funded by Inria).
- Luka Stanisic (2012-2015): Performance evaluation, modeling and simulation of HPC systems; Experimental methodology and reproducible research (co-advised with Jean-François Méhaut, funded by the French Ministry for Research). Current position: Post-doctoral Researcher in HPC at Max Planck Institute.

- Rémi Bertin (2007-2009): Collaboration Mechanisms in Peer-to-Peer and Collaborative Computing Systems (co-advised with Corinne Touati, funded by ANR DOCCA). Did not defend the PhD. Current position: R&D Engineer at Allegro DVT.
- Pedro Velho (2006-2011): Modeling and Simulation of Large Scale Distributed Platforms (co-advised with Jean-François Méhaut, funded by a Brazilian grant). Current position: R&D Senior Software Engineer at ActiveEon.

PostDoc, Engineers

- Lucas Schnorr (Invited Professor 2016-2017): Tracing, observation and visualization of large scale distributed systems.
- Augustin Degomme (Eng. 2012-2015): Simulation/performance prediction of MPI applications. Current position: Research Engineer at Basel University, Switzerland.
- Sascha Hunold (Post-doc 2011-2012): Design of Experiments, Reproducible Research, Fair Scheduling of Bag-of-Tasks Applications Using Distributed Lagrangian Optimization. Current position: Professor at TU Vienna, Austria.
- Lucas Schnorr (Post-doc 2009-2012): Tracing, observation and visualization of large scale distributed systems. Current position: Professor at UFRGS, Brazil.
- Pierre Navarro (Eng. 2010-2012): Improvement of the SimGrid Framework (scalability, robustness, new features, ...)
- Lionel Eyraud (Post-doc 2007): Automatically Building Sound Network Representations. Current position: Inria researcher.

Engineer/Master Students

- Tom Cornebize (MSc 2017): Capacity Planning of Supercomputers, Simulating MPI Applications at Scale.
- Steven Quinto Masnada (MSc 2016): Semi-Automatic Performance Optimization of HPC Kernels. Current position: Engineer at Inria.
- Wagner Kolberg (MSc 2012): Faithful Modeling of MapReduce Applications.
- Luka Stanistic (MSc 2012): Performance evaluation of cache performances on modern CPUs.
- Émile Morel (stage de fin d'étude IUT, 4 months, 2008-2009): Design and implementation of a collaborative bibliography entry management server.
- Bruno Luis de Moura Donassolo (Engineer, Msc, 2007-2009): Design and Implementation of a Scalable Scheduler for the SimGrid Project; Study of Non-Cooperative Optimization in Volunteer Computing Systems.
- Darina Dimitrova (Master Research, 4 months, 2006): Application-level Network Topology Discovery in Grid Computing Platforms.
- Rémi Vannier (Master Research, 4 months, 2006): Proportionally Fair and Distributed Scheduling of Multiple Bag-Of-Task Applications.

Software Development

SimGrid

SimGrid is a toolkit that provides core functionalities for the simulation of distributed applications in heterogeneous distributed environments. SimGrid is a 17 years-old free software project whose specific goal is to facilitate research in the area of distributed and parallel application scheduling on distributed computing platforms ranging from simple network of workstations to Computational Grids. It is highly scalable and can simulate hundreds of thousands up to millions of nodes on a single machine while using realistic network models, whose soundness was assessed through thorough (in)validation studies. SimGrid can also be used as an MPI Simulator to realistically simulate unmodified MPI programs. I am one of the main architects (together with Henri Casanova, Martin Quinson, and Frédéric Suter) of this project. The SimGrid project is open-source and hosted on the INRIA gforge³ and on github⁴.

My most notable scientific contributions in this software are related to the efficient implementation and to the assessment of the validity of fluid network and CPU models. I have used this software for my own research on scheduling in grid/volunteer computing and I am particularly active in making it evolve so that it can be used to predict the performance of complex HPC applications.

SimGrid is recognized in the HPC community as one of the most prominent simulation environments as shown by its large community of users and the number of publications that use it: in the past ten years, SimGrid has been the basis of at least 210 articles. Besides, the main four articles on SimGrid have been cited more than 1,550 times according to Google Scholar. There are currently about 220 members in the `simgrid-user` mailing list.

Among other success stories, SimGrid has been used at CERN for capacity planning and optimization of file replication strategies, in the biomed organization of the EGI grid to optimize task scheduling and checkpointing, by colleagues from Inria Bordeaux to perform daily non regression testing of their software, in Japan (NII) to design HPC optical networks, by Brazilian researchers to improve map-reduce scheduling strategies, etc.

The core of SimGrid represents 90,000 single lines of code (in C and C++) and more than 100,000 single lines of codes for examples and regression tests. Here is a rough evaluation of SimGrid's maturity following CNRS/Inria standards:

- Audience: Wide audience, large user's community (5/5)
- Originality: Original software implementing a fair number of original ideas (4/4)
- Maturity: Major software project, strong software engineering (4/5)
- Evolution and Maintenance: Well-defined and implemented plans for future maintenance and evolution (4/4)
- Software distribution and licensing: External packaging and distribution, as part of a popular open source distribution (5/5)

Although I was the main developer in 2002-2009, I have become now the fourth contributor in terms of lines of code as I have gradually handed over development to the community, focusing on experimental validation, on coordination of scientific developments and on potential technological transfer. Here is a rough evaluation of my own contribution to SimGrid following CNRS/Inria standards:

- Design and Architecture: A main contributor (4/4)
- Coding and Debugging: A regular contributor (3/4)
- Maintenance and Support: An occasional contributor (2/4).
- Team/Project Management: A main contributor (4/4)

³<http://simgrid.gforge.inria.fr>

⁴<https://github.com/simgrid/simgrid>

Pistou: a Bibliography Management Tool

I was asked by the direction of the LIG in 2006 to design a tool that would allow the LIG to collect and collaboratively manage the bibliography data required for the quadrennial AERES evaluation of the laboratory. By this time, HAL was not suited at all for such purpose. I designed and helped implementing such a tool⁵ with Émile Morel and Isabelle Guillet.

This open-source software (about 16K lines of ruby on rails and javascript) allowed us to collect and organize within a few months a relatively good quality bibliography database containing more than 10,000 entries. This tool no more used today as all the LIG teams have now been strongly encouraged to migrate their bibliography management to HAL.

Scientific Collaborations and Grants

In my career I have been heavily involved in the following national or international projects, which I briefly describe here:

Projects and Grants

IPL HAC SPECIS (2016-2020) High-performance Application and Computers, Studying Performance and Correctness In Simulation. I am the leader of this Inria Project Laboratory, which is a 4 year funding from Inria. This is a joint project between 8 Inria teams and whose goal is to answer methodological needs of HPC application and runtime developers and to allow to study real HPC systems both from the correctness and performance point of view. To this end, we gather experts from the HPC, formal verification and performance evaluation community. All the resulting research developments will be integrated in the open source SimGrid framework so that they can benefit as quickly as possible to the greatest number.

European HPC4E (2015-2017) The HPC4E involves European and Brazilian institutions. The European coordinator is the Barcelona Supercomputing Center. This project aims at applying the new exascale HPC techniques to energy industry simulations, customizing them, and going beyond the state-of-the-art in the required HPC exascale simulations for different energy sources: wind energy production and design, efficient combustion systems for biomass-derived fuels (biogas), and exploration geophysics for hydrocarbon reservoirs. I was in charge of Task 2.4 on performance analysis.

ANR SONGS (2012-2015) Simulation Of Next Generation Systems, a 4 year ANR grant (IN-FRA, Programme Systèmes embarqués et grandes infrastructures; 2012-2015). This is a joint project with Martin Quinson, Lionel Eyraud-Dubois, Frédéric Vivien, Frédéric Suter, Stéphane Genaud, Adrien Lèbre, Olivier Dalle, Abdou Guermouche and their respective teams. The goal of the SONGS project is to extend the applicability of the SimGrid simulation framework from Grids and Peer-to-Peer systems to Clouds and High Performance Computation systems. Each type of large-scale computing system will be addressed through a set of use cases and lead by researchers recognized as experts in this area. I was leading the work package on Analysis and Visualization and I was very involved in all the work packages (experimental methodology, HPC, data grid, peer-to-peer and volunteer computing, cloud computing) where I played a transverse scientific coordination role.

European Mont-Blanc Projects (2011-2016) Energy efficiency is a primary concern for the design of any computer system and it is clear that designing the envisioned Exascale systems within a reasonable power envelope will require to fully redesign software and

⁵<http://pistou.imag.fr>

architecture. Since October 2011, the aim of the European project called Mont-Blanc has been to design a new type of computer architecture capable of setting future global HPC standards, built from energy efficient solutions used in embedded and mobile devices such as ARMv8 64-bit processors. This project was coordinated by the Barcelona Supercomputing Center (BSC) and funded by the European Commission. Two years later, the European Commission granted additional 8 million Euro funds to extend the Mont-Blanc project activities until September 2016.

My role in these two projects was to improve performance evaluation and simulation techniques to conduct network and resource provisioning studies raised by such context as well as possibly improve the parallel software development process.

ANR USS-SimGrid (2009-2011) Ultra-Scalable Simulations with SimGrid, a 3 year ANR grant (ARPEGE, Programme Systèmes embarqués et grandes infrastructures; 2009-2011). This was a joint project with Martin Quinson, Fabrice Le Fessant, Lionel Eyraud-Dubois, Frédéric Vivien, Frédéric Suter, Olivier Flauzac and their respective teams. This project aims at extending SimGrid more specifically for the very large scale distributed computing (peer-to-peer, desktop grids, ...), at improving the quality of its results and at providing tools for experiments management and steering. In the long term, we aim at federating the various efforts at the national scale in the area of simulation for large scale distributed computing. I was leading the work package on design and implementation of scalable models and I was highly involved in the work packages on visualization and analysis, experiment campaign management.

ANR DOCCA (2007-2010) Design and Optimization of Collaborative Computing Architecture, a 3 year ANR grant for young researchers. This was a joint project with Corinne Touati and Florence Perronnin.

ANR ALPAGE (2005-2008) ALgorithmique des Plates-formes À Grande Échelle, a 3 year ANR grant. This was a joint project with Yves Robert, Anne-Marie Kermarrec, Pierre Freignaud, Olivier Beaumont, and their respective teams.

Collaborations and Joint Laboratories

Inria/Orange laboratory (2017-...) The joint laboratory between Orange and Inria aims at reinforcing the research effort on network virtualization and on the convergence between communication networks and cloud computing (Fog, IoT, ...). I participate to this joint laboratory since the beginning of 2017 and coadvise a PhD thesis in this context.

Joint laboratory on *petascale* and *extreme-scale* computing (2011-2015) The joint laboratory between University of Illinois at Urbana-Champaign, Inria, the CNRS, Argonne National Laboratory, Barcelona Supercomputing Center and Jülich Supercomputing Center targets software and hardware issuers raised by the design and exploitation of supercomputers. I participate since 2011 to this joint laboratory in particular regarding modeling and performance evaluation topics.

In this context, I organized the summer school on *Performance Metrics, Modeling and Simulation of Large HPC Systems* funded by the Partner University Fund and the joint laboratory in June 2014 in Sophia Antipolis⁶. The audience was international with attendees working in the US, Germany, Spain, and France.

Action d'Envergure Inria HEMERA (2010-2014) HEMERA is an Inria research action, which started in 2010 and whose goal is to federate the research efforts linked to large-scale

⁶http://mescal.imag.fr/membres/arnaud.legrand/research/events/puf_jlpc_workshop_14.php

experimentation, particularly in the context of the Grid'5000 infrastructure. One of the goals of this action was to animate the high performance/distributed computing French research community. This project was lead by Christian Pérez and I was responsible with Martin Quinson of the *Modeling Large Scale Systems and Validating their Simulators* theme.

Associated Team Inria MESCAL/Berkeley (2008-2014) The MESCAL team has been associated to several researchers from the Bay area and in particular David Anderson, the leader of the BOINC project but also Walfredo Cirne from Google Inc. This collaboration was initially lead by Derrick Kondo and then by myself and focused on many performance evaluation aspects of very large scale computing systems such as volunteer computing systems and cloud computing systems. Derrick Kondo and Jean-Marc Vincent worked on statistically characterizing the availability and unavailability of resources in such systems based on traces collected by our American colleagues. Such models can then be used to predict load or improve replication strategies and collective availability. On my side I worked on how scheduling and game theory could be applied to possibly better understand such systems as well as on how to simulate them efficiently. This collaboration was initially lead by Derrick Kondo. I became the coordinator in 2012 and I organized the BOINC workshop in 2013.

Grenoble - Porto Alegre Associated Team and Joint Laboratory Grenoble and Porto Alegre Universities (in particular the Universidade Federal do Rio Grande do Sul) have a long standing collaboration that dates back from the end of the 1970s. There has been several associated teams and research/student exchanges (through Inria, CNRS, CAPES, CNPq, FAPERGS, ...) and these collaborations have recently evolved into the LICIA (*Laboratoire International en Calcul Intensif et Informatique Ambiante*), a joint laboratory between the computer science department of UFRGS and the LIG. I have thus visited regularly Porto Alegre in the last decade to give lectures and collaborate with Philippe Navaux, Nicolas Maillard, Claudio Geyer, Alexandre Carrissimi. This gave me the opportunity to advise many Brazilian students during the Msc, PhD or postdoc among which Pedro Velho, Lucas Schnorr, Bruno Donassolo, Wagner Kolberg, Rafael Tesser, ...

Stays In Foreign Countries

From November 2004 to October 2005, I worked at University of California, San Diego with Henri Casanova, Jeanne Ferrante and Larry Carter.

Besides, I stayed many times 2-3 weeks mostly in Porto Alegre (Brazil) and in the US (San Diego, Berkeley, Argonne, ...).

Invited Talks

I regularly participate to conferences and summer schools to give keynotes and tutorials. Here is a non exhaustive list of such interventions:

- On reproducible research:
 - Keynote (1 hour) at the LIRIS, Lyon, November 2017.
 - Keynote (1 hour) at the PRECIS (Précision, Reproductibilité en Calcul Scientifique) Spring school, Fréjus, May 2017.
 - Keynote (2× 1 hour) at the Grenoble Data Science Institute, April 2017.
 - Keynote (1 hour) at the LIG laboratory, Grenoble, March 2017.
 - Keynote (1 hour) at the ENS Rennes, February 2017.

- Talk (20 minutes) at Inria Scientific Days, Rennes, June 2016.
 - Keynote (1 hour) at the R⁴ conference, Orléans, Dec 2015.
 - Keynote (1 hour) in the PLAFRIM working group in Bordeaux, December 2014.
 - Keynote (1 hour) International Workshop on Reproducibility in Parallel Computing, August 2014.
 - Lecture (3 hours) in the «Summer school on Performance Metrics, Modeling and Simulation of Large HPC Systems», June 2014.
 - Lecture (3 hours) for the European project «sync-free» in Paris May 2014.
 - Tutorial (2 hours) Conférence d'informatique en Parallélisme, Architecture et Système (COMPAS) in Neuchâtel, en mars 2013.
- On Simulation and Performance Evaluation:
 - Keynote (1 hour) «Simulation of Large-Scale Distributed Computing Research» at the workshop du Laboratoire International en Calcul Intensif et Informatique Ambiante à Grenoble, September 2014.
 - Keynote (1 hour) «Simulation of Large-Scale Distributed Computing Research: The SimGrid Project» at the SimuTools conference, in Cannes, March 2013.
 - Tutorial (2 hours) «Simulation for Large-Scale Distributed Computing Research» given with Martin Quinson at the Conférence d'informatique en Parallélisme, Architecture et Système (COMPAS) in Grenoble, in January 2013.
 - Tutorial (2 hours) «Simulation for Large-Scale Distributed Computing Research». CLCAR (Latin American conference on HPC), Brazil, in 2010.
 - Keynote (1 hour) «Simulation for Large-Scale Distributed Computing Research». Summer school ERAD, Brazil, in 2009.
 - Tutorial (3 hours) «Simulation for Large-Scale Distributed Computing Research» given with Martin Quinson at the CCGrid conference, Lyon, 2008.
 - Keynote (1 hour) «Simulation, Emulation and Experiments» at the GRID'5000 spring school, 2006.
 - On Scheduling:
 - Talk on "the Use of Lagrangian Optimization For Designing Distributed Self-Stabilizing Protocols" at ROADEF (15ème congrès annuel de la Société française de recherche opérationnelle et d'aide à la décision), in 2014, Bordeaux.
 - Talk on "Non-Cooperative Scheduling of Multiple Bag-of-Tasks Applications", TU Wien Seminar in 2013 - Talk on "Non-Cooperative Scheduling of Multiple Bag-of-Tasks Applications", TU Wien Seminar in 2013.
 - «On the Impact of Platform Models» at the "École de Printemps d'Informatique Théorique" (EPIT) in 2007.
 - "Toward a Fully Decentralized Algorithm for Multiple Bag-of-tasks Application Scheduling on Grids" in CIRM workshop on "New Challenges on Scheduling Theory" in 2008.
 - "Scheduling Competing Regular Applications on a Heterogeneous Master-Worker Platforms" in the NSF/INRIA Workshop on "Scheduling for Large-Scale Distributed Platforms", La Jolla, California, in 2005

Teaching

I teach in the 2nd year of Master Research since I arrived in Grenoble in 2006 and I have been co-responsible of the "Parallel, Distributed and Embedded Systems" option of the MOSIG (International Master of Science of Grenoble) from 2011 to 2016. I am responsible of the Parallel Systems lecture in this option since 2008.

In the last years, I have also been regularly invited to give a series of lectures (about 15-18 hours) on Scientific Methodology and Performance Evaluation at the master level (in August 2015 at Federal University of Rio Grande do Sul, in December 2016 at ENS Lyon, in September 2017 at ENS Rennes...). The content of these lectures is freely available on [github](https://github.com/alegrand/SMPE)⁷. I am currently designing a MOOC on Reproducible Research with Konrad Hinsén (CNRS/Centre de Biophysique Moléculaire) and Christophe Pouzat (CNRS/ Mathématiques Appliquées à Paris 5) with the support of the Inria MOOC-lab. The diffusion of this MOOC is planned for 2018 on FUN-MOOC.

I also worked several years for the Écoles Normales Supérieures to design and examine the competitive exam for admission. I designed the algorithm subject of the "second concours" in 2005-2006. I designed and conducted exams for the "Épreuve Pratique d'Algorithmique et Programmation" of the "premier concours" from 2006 to 2009. I designed the computer science subject of the "first concours" in 2015.

Finally, in June 2014, I have organized the summer school on *Performance Metrics, Modeling and Simulation of Large HPC Systems* funded by the Partner University Fund and the Joint Laboratory on Extreme Scale Computing in Sophia. The audience was international with attendees working in the US, Germany, Spain, and France.

Année	Intitulé	Public	Lieu	Durée
2015,2016,2017	Scientific Methodology and Performance Evaluation	Master/PhD	UFRGS (2015) ENS Lyon (2016) ENS Rennes (2017)	15-18 hours each time
2012-...	Probability and Simulation Performance Evaluation	RICM4 (engineers M1)	UGA Polytech	20+20 hours per year
2011-...	Scientific Methodology and Performance Evaluation	International Master of Science in Informatics at Grenoble (MOSIG) M2R	Université Grenoble Alpes	15 hours per year
2008-...	Parallel Systems (coordinateur + enseignant)	MOSIG M2	Université Grenoble Alpes	8-18 hours per year
2010-2011	Operating Systems	MOSIG M1	Université Joseph Fourier	32 hours per year
2008-2009	Scheduling for Large Distributed Computing Systems	Master 2	UFRGS (Brésil)	9 hours per year
2007-2010	Systèmes d'exploitation	Master 1	Université Joseph Fourier	6-9 hours per year
2006-2008	Mesure et analyse de donnée pour l'évaluation de performance	Master 2 Recherche	Université Joseph Fourier	6 hours per year
2006-2007	Architectures hautes performances	Master 2 Recherche	Université Joseph Fourier	8 hours per year
2005-2006	Introduction à l'ordonnancement	Graduate Students	UCSD	2 hours
1 ^{er} semestre 2003-2004 2 nd semestre 2002-2003 1 ^{er} semestre 2002-2003	Introduction à l'algorithmique	DEUG SM/STPI/SV 2 nd year	UCBL	16 hours
	courses			16 hours
	courses, tutorials			32 hours
2003-2004 2002-2003 2001-2002	Algorithmique parallèle	Maîtrise	ENS Lyon	32 hours 32 hours 32 hours
2001-2002	Langage C	DEUG MIAS 2 nd year	UCBL	15 hours
2000-2001	Graphes, automates et langages formels	DEUG MIAS 2 nd year	UCBL	18 hours
1999-2000	CAML	Spé MP et MP*	Lycée du Parc	2 hours per week over 8 months
1998-1999 2000-2001	Colles de Mathématique	Sup PCSI	Lycée du Parc	2 hours per week over a year

⁷<https://github.com/alegrand/SMPE>

Publications

My publication strategy depends on the domain and whether PhD students and postdocs are involved or not. I particularly value high quality (e.g., JPDC, TPDS, CCPE) journal publications as it allows to publish results with deeper content and better interactions with reviewers. A few recent journals have a B ranking as it was the best way to target a particular audience (e.g., through a special issue on reproducible research for *Operating Systems Review* from ACM SIGOPS) or reviewers with a particular expertise (e.g., for ACM TOMACS). Regarding conferences, I mostly target rank A conferences (e.g., CCgrid, Cluster, Euro-Par) but also a few specialized workshops (e.g., PMBS or VPA at the SuperComputing conference, or LSAP) that allow a fast and efficient dissemination of ideas toward a particular audience.

Since it may be difficult to know the different venue in computer science, I indicate between parenthesis the ERA⁸ and QUALIS⁹ (for conferences) ranking.

Regarding authorship, the author order depends on when and with who the article was written. When contributions are fairly shared, we favor alphabetical order. When the article involves PhD students or postdoc who need to build a reputation, I generally insist on having their name to appear first and my name rather to appear in the end as a supervisor.

Journal articles

- [1] A. Degomme, A. Legrand, G. Markomanolis, M. Quinson, M. L. Stillwell, and F. Suter. Simulating MPI applications: the SMPI approach. *IEEE Transactions on Parallel and Distributed Systems*, page 14, February 2017. (**A***).
- [2] L. Stanisic, A. Legrand, and V. Danjean. An Effective Git And Org-Mode Based Workflow For Reproducible Research. *Operating Systems Review*, 49:61 – 70, 2015.
- [3] L. Stanisic, S. Thibault, A. Legrand, B. Videau, and J.-F. Méhaut. Faithful Performance Prediction of a Dynamic Task-Based Runtime System for Heterogeneous Multi-Core Architectures. *Concurrency and Computation: Practice and Experience*, page 16, May 2015. (**A**).
- [4] H. Casanova, A. Giersch, A. Legrand, M. Quinson, and F. Suter. Versatile, Scalable, and Accurate Simulation of Distributed Applications and Platforms. *Journal of Parallel and Distributed Computing*, 74(10):2899–2917, June 2014. (**A***).
- [5] R. Bertin, S. Hunold, A. Legrand, and C. Touati. Fair scheduling of bag-of-tasks applications using distributed Lagrangian optimization. *Journal of Parallel and Distributed Computing*, August 2013. (**A***).
- [6] P. Velho, L. Schnorr, H. Casanova, and A. Legrand. On the Validity of Flow-level TCP Network Models for Grid and Cloud Simulations. *ACM Transactions on Modeling and Computer Simulation*, 23(4), October 2013. (**B**).
- [7] H. Kameda, E. Altman, C. Touati, and A. Legrand. Nash Equilibrium Based Fairness. *Mathematical Methods of Operations Research*, 76(1), 2012. (**B**).
- [8] L. M. Schnorr, A. Legrand, and J.-M. Vincent. Detection and analysis of resource usage anomalies in large distributed systems through multi-scale visualization. *Concurrency and Computation: Practice and Experience*, 24:1792–1816, 2011. (**A**).
- [9] O. Beaumont, L. Carter, J. Ferrante, A. Legrand, L. Marchal, and Y. Robert. Centralized Versus Distributed Schedulers for Multiple Bag-of-Tasks Applications. *IEEE Transactions on Parallel and Distributed Systems*, 19:698–709, 2008. (**A***).

⁸<http://www.core.edu.au/conference-portal>

⁹<http://www.conferenceranks.com/>

- [10] A. Legrand, A. Su, and F. Vivien. Minimizing the Stretch When Scheduling Flows of Divisible Requests. *Journal of Scheduling*, 2008. **(B)**.
- [11] O. Beaumont, A. Legrand, L. Marchal, and Y. Robert. Pipelining broadcasts on heterogeneous platforms. *IEEE Transactions on Parallel and Distributed Systems*, 2005. **(A*)**.
- [12] O. Beaumont, A. Legrand, L. Marchal, and Y. Robert. Steady-state scheduling on heterogeneous clusters. *International Journal of Foundations of Computer Science*, 2005. **(B)**.
- [13] A. Legrand, L. Marchal, and Y. Robert. Optimizing the Steady-State Throughput of Scatter and Reduce Operations on Heterogeneous Platforms. *Journal of Parallel and Distributed Computing*, 65:1497–1514, 2005. **(A*)**.
- [14] C. Banino, O. Beaumont, L. Carter, J. Ferrante, A. Legrand, and Y. Robert. Scheduling Strategies for Master-Slave Tasking on Heterogeneous Processor Platforms. *IEEE Trans. Parallel Distributed Systems*, 15:319–330, 2004. **(A*)**.
- [15] A. Legrand, H. Renard, Y. Robert, and F. Vivien. Mapping and Load-Balancing Iterative Computations on Heterogeneous Clusters with Shared Links. *IEEE Trans. Parallel Distributed Systems*, 15:546–558, 2004. **(A*)**.
- [16] O. Beaumont, A. Legrand, L. Marchal, and Y. Robert. Scheduling Strategies for Mixed Data and Task Parallelism on Heterogeneous Clusters. *Parallel Processing Letters*, 13:225–244, 2003. **(B)**.
- [17] O. Beaumont, A. Legrand, and Y. Robert. Scheduling Divisible Workloads on Heterogeneous Platforms. *Parallel Computing*, 29:1121–1152, 2003. **(A)**.
- [18] O. Beaumont, A. Legrand, and Y. Robert. The Master-Slave Paradigm with Heterogeneous Processors. *IEEE Trans. Parallel Distributed Systems*, 14:897–908, 2003. **(A*)**.
- [19] O. Beaumont, A. Legrand, F. Rastello, and Y. Robert. Dense Linear Algebra Kernels on Heterogeneous Platforms: Redistribution Issues. *Parallel Computing*, 28:155–185, 2002. **(A)**.
- [20] O. Beaumont, A. Legrand, and Y. Robert. Static Scheduling Strategies for Heterogeneous Systems. *Computing and Informatics*, 21:413–430, 2002.
- [21] A. Legrand. Équilibrage de Charge Statique Pour Noyaux Dalgèbre Linéaire Sur Plate-Forme Hétérogène. *Technique Et Science Informatique, Numéro Spécial RenPar13*, pages 711–734, 2002.
- [22] O. Beaumont, A. Legrand, F. Rastello, and Y. Robert. Static LU Decomposition on Heterogeneous Platforms. *Int. Journal of High Performance Computing Applications*, 15:310–323, 2001. **(B)**.

Conference articles

- [23] F. C. Heinrich, T. Cornebize, A. Degomme, A. Legrand, A. Carpen-Amarie, S. Hunold, A.-C. Orgerie, and M. Quinson. Predicting the Energy Consumption of MPI Applications at Scale Using a Single Node. In *Cluster 2017*, Hawaii, United States, September 2017. IEEE. **(A/A2)**.
- [24] R. Keller Tesser, L. Mello Schnorr, A. Legrand, F. Dupros, and P. O. A. Navaux. Using Simulation to Evaluate and Tune the Performance of Dynamic Load Balancing of an Over-decomposed Geophysics Application. In *Euro-Par 2017: 23rd International European Conference on Parallel and Distributed Computing*, page 15, Santiago de Compostela, Spain, August 2017. **(A/A2)**.

- [25] L. Stanisic, L. C. Mello Schnorr, A. Degomme, F. C. Heinrich, A. Legrand, and B. Videau. Characterizing the Performance of Modern Architectures Through Opaque Benchmarks: Pitfalls Learned the Hard Way. In *IPDPS 2017 - 31st IEEE International Parallel & Distributed Processing Symposium (RepPar workshop)*, Orlando, United States, June 2017. **(A/A1)**.
- [26] V. Garcia Pinto, L. Stanisic, A. Legrand, L. Mello Schnorr, S. Thibault, and V. Danjean. Analyzing Dynamic Task-Based Applications on Hybrid Platforms: An Agile Scripting Approach. In *3rd Workshop on Visual Performance Analysis (VPA)*, Salt Lake City, United States, November 2016. Held in conjunction with SC16.
- [27] A. Lebre, A. Legrand, F. Suter, and P. Veyre. Adding Storage Simulation Capacities to the SimGrid Toolkit: Concepts, Models, and API. In *CCGrid 2015 - Proceedings of the 15th IEEE/ACM Symposium on Cluster, Cloud and Grid Computing*, pages 251–260, Shenzhen, China, May 2015. IEEE/ACM. **(A/A1)**.
- [28] L. Stanisic, E. Agullo, A. Buttari, A. Guermouche, A. Legrand, F. Lopez, and B. Videau. Fast and Accurate Simulation of Multithreaded Sparse Linear Algebra Solvers. In *The 21st IEEE International Conference on Parallel and Distributed Systems*, Melbourne, Australia, December 2015. **(B/B1)**.
- [29] L. Stanisic and A. Legrand. Effective Reproducible Research with Org-Mode and Git. In *1st International Workshop on Reproducibility in Parallel Computing*, Porto, Portugal, August 2014.
- [30] L. Stanisic, S. Thibault, A. Legrand, B. Videau, and J.-F. Méhaut. Modeling and Simulation of a Dynamic Task-Based Runtime System for Heterogeneous Multi-Core Architectures. In *Euro-par - 20th International Conference on Parallel Processing*, Euro-Par 2014, LNCS 8632, pages 50–62, Porto, Portugal, August 2014. Springer International Publishing Switzerland. **(A/A2)**.
- [31] P. Bedaride, A. Degomme, S. Genaud, A. Legrand, G. Markomanolis, M. Quinson, M. L. Stillwell, F. Suter, and B. Videau. Toward Better Simulation of MPI Applications on Ethernet/TCP Networks. In *PMBS13 - 4th International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems*, Denver, United States, November 2013.
- [32] L. Mello Schnorr, A. Legrand, and J.-M. Vincent. Interactive Analysis of Large Distributed Systems with Scalable Topology-based Visualization. In *International Symposium on Performance Analysis of Systems and Software (ISPASS'13)*, Austin, Texas, United States, 2013. IEEE Computer Society Press. **(C/B1)**.
- [33] L. Stanisic, B. Videau, J. Cronsioe, A. Degomme, V. Marangozova-Martin, A. Legrand, and J.-F. Mehaut. Performance Analysis of HPC Applications on Low-Power Embedded Platforms. In *DATE - Design, Automation & Test in Europe*, pages 475–480, Grenoble, France, March 2013. **(B/A1)**.
- [34] L. Bobelin, A. Legrand, M. A. G. David, P. Navarro, M. Quinson, F. Suter, and C. Thiery. Scalable Multi-Purpose Network Representation for Large Scale Distributed System Simulation. In *CCGrid 2012 – The 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*, page 19, Ottawa, Canada, May 2012. **(A/A1)**.
- [35] B. De Moura Donassolo, A. Legrand, and C. Geyer. Non-Cooperative Scheduling Considered Harmful in Collaborative Volunteer Computing Environments. In *Proceedings of the 11th IEEE International Symposium on Cluster Computing and the Grid (CCGrid'11)*, pages 144–153, Newport Beach, United States, 2011. IEEE Computer Society Press. **(A/A1)**.

- [36] L. Mello Schnorr, A. Legrand, and J.-M. Vincent. Multi-scale analysis of large distributed computing systems. In *Proceedings of the third international workshop on Large-scale system and application performance*, pages 27–34, San Jose, CA, United States, June 2011. ACM.
- [37] B. Donassolo, H. Casanova, A. Legrand, and P. Velho. Fast and scalable simulation of volunteer computing systems using SimGrid. In ACM, editor, *LSAP 2010 : 2nd Workshop on Large-Scale System and Application Performance*, pages 605–612, Chicago, Illinois, United States, June 2010. ACM.
- [38] H. Kameda, E. Altman, C. Touati, and A. Legrand. Nash Equilibrium Based Fairness. In *International Conference on Game Theory for Networks (GameNets) - 2009, Game Theory for Networks, 2009. GameNets '09. International Conference on*, pages 533 – 539, Istanbul, Turkey, May 2009. IEEE.
- [39] P. Velho and A. Legrand. Accuracy Study and Improvement of Network Simulation in the SimGrid Framework. In *SIMUTools'09, 2nd International Conference on Simulation Tools and Techniques*, Rome, Italy, March 2009.
- [40] R. Bertin, A. Legrand, and C. Touati. Toward a Fully Decentralized Algorithm for Multiple Bag-of-tasks Application Scheduling on Grids. In *IEEE/ACM International Conference on Grid Computing (Grid)*, Tsukuba, Japan, 2008. **(A/A2)**.
- [41] H. Casanova, A. Legrand, and M. Quinson. SimGrid: a Generic Framework for Large-Scale Distributed Experiments. In *10th IEEE International Conference on Computer Modeling and Simulation - EUROSIM / UKSIM 2008*, Cambridge, United Kingdom, April 2008. IEEE.
- [42] L. Eyraud-Dubois, A. Legrand, M. Quinson, and F. Vivien. A First Step Towards Automatically Building Network Representations. In *13th International Euro-Par Conference - Euro-Par 2007*, volume 4641/2007 of *Lecture Notes in Computer Science*, pages 160–169, Rennes, France, August 2008. Springer Berlin / Heidelberg. **(A/A2)**.
- [43] A. Legrand and C. Touati. How to measure efficiency? In *Proceedings of the 1st International Workshop on Game theory for Communication networks (Game-Comm07)*, 2007.
- [44] A. Legrand and C. Touati. Non-Cooperative Scheduling of Multiple Bag-of-Task Applications. In *INFOCOM*, Anchorage, Alaska, May 2007. **(A*/A1)**.
- [45] O. Beaumont, L. Carter, J. Ferrante, A. Legrand, L. Marchal, and Y. Robert. Centralized Versus Distributed Schedulers Multiple Bag-of-Task Applications. In *International Parallel and Distributed Processing Symposium IPDPS'2006*, Rhodes Island, Greece, 2006. IEEE Computer Society Press. **(A/A1)**.
- [46] A. Legrand, A. Su, and F. Vivien. Minimizing the Stretch When Scheduling Flows of Biological Requests. In *Symposium on Parallelism in Algorithms and Architectures SPAA'2006*. ACM Press, 2006. **(A/A2)**.
- [47] M. Quinson, H. Casanova, A. Legrand, and K. Fujiwara. The SIMGRID Project: Simulation and Deployment of Distributed Applications. In IEEE, editor, *The 15th IEEE International Symposium on High Performance Distributed Computing (HPDC'06)*, Paris, France, June 2006. **(A/A1)**.
- [48] O. Beaumont, V. Boudet, P.-F. Dutot, A. Legrand, and Y. Robert. Fondements théoriques pour la conception d'algorithmes efficaces de gestion de ressources. In *Informatique répartie: architecture, parallélisme et systèmes*. Hermes Sciences, 2005.

- [49] O. Beaumont, A. Legrand, L. Marchal, and Y. Robert. Independent and Divisible Tasks Scheduling on Heterogeneous Star-Shaped Platforms with Limited Memory. In *PDP'2005, 13th Euromicro Workshop on Parallel, Distributed and Network-Based Processing*, pages 179–186. IEEE Computer Society Press, 2005. **(C/B1)**.
- [50] A. Legrand, A. Su, and F. Vivien. Off-Line Scheduling of Divisible Requests on an Heterogeneous Collection of Databanks. In *Proceedings of the 14th Heterogeneous Computing Workshop*, Denver, Colorado, USA, 2005. IEEE Computer Society Press. **(/B1)**.
- [51] O. Beaumont, A. Legrand, L. Marchal, and Y. Robert. Assessing the Impact and Limits of Steady-State Scheduling for Mixed Task and Data Parallelism on Heterogeneous Platforms. In *HeteroPar'2004: International Conference on Heterogeneous Computing, Jointly Published with ISPDC'2004: International Symposium on Parallel and Distributed Computing*, pages 296–302. IEEE Computer Society Press, 2004.
- [52] O. Beaumont, A. Legrand, L. Marchal, and Y. Robert. Complexity Results and Heuristics for Pipelined Multicast Operations on Heterogeneous Platforms. In *2004 International Conference on Parallel Processing (ICPP'2004)*, pages 267–274. IEEE Computer Society Press, 2004. **(A/A2)**.
- [53] O. Beaumont, A. Legrand, L. Marchal, and Y. Robert. Pipelining broadcasts on heterogeneous platforms. In *International Parallel and Distributed Processing Symposium IPDPS'2004*. IEEE Computer Society Press, 2004. **(A/A1)**.
- [54] O. Beaumont, A. Legrand, L. Marchal, and Y. Robert. Steady-State Scheduling on Heterogeneous Clusters: Why and How? In *6th Workshop on Advances in Parallel and Distributed Computational Models APDCM*, page 171a (8 pages). IEEE Computer Society Press, 2004.
- [55] E. Caron, P. K. Chouhan, and A. Legrand. Automatic Deployment for Hierarchical Network Enabled Server. In *Heterogeneous Computing Workshop*. IEEE Computer Society Press, 2004. **(/B1)**.
- [56] A. Legrand, L. Marchal, and Y. Robert. Optimizing the Steady-State Throughput of Scatter and Reduce Operations on Heterogeneous Platforms. In *6th Workshop on Advances in Parallel and Distributed Computational Models APDCM 2004*, page 176a (8 pages). IEEE Computer Society Press, 2004.
- [57] A. Legrand and M. Quinson. Automatic deployment of the Network Weather Service using the Effective Network View. In *High-Performance Grid Computing Workshop, associated to IPDPS'04*, Santa Fe, New Mexico, USA, April 2004. **(A/A1)**.
- [58] O. Beaumont, A. Legrand, and Y. Robert. Optimal Algorithms for Scheduling Divisible Workloads on Heterogeneous Systems. In *IPDS'2003, Parallel and Distributed Processing Symposium, 2003*, Nice, France, February 2003. IEEE Computer Society Press. **(A/A1)**.
- [59] O. Beaumont, A. Legrand, and Y. Robert. Scheduling Strategies for Mixed Data and Task Parallelism on Heterogeneous Clusters and Grids. In *PDP'2003, 11th Euromicro Workshop on Parallel, Distributed and Network-Based Processing*, pages 209–216, Gênes, Italy, February 2003. IEEE Computer Society Press. **(C/B1)**.
- [60] H. Casanova, A. Legrand, and L. Marchal. Scheduling Distributed Applications: the Sim-Grid Simulation Framework. In *Proceedings of the Third IEEE International Symposium on Cluster Computing and the Grid (CCGrid'03)*. IEEE Computer Society Press, 2003. **(A/A1)**.

- [61] A. Legrand, H. Renard, Y. Robert, and F. Vivien. Load-Balancing Iterative Computations on Heterogeneous Clusters with Shared Communication Links. In *PPAM-2003: Fifth International Conference on Parallel Processing and Applied Mathematics*, LNCS 3019, pages 930–937. Springer Verlag, 2003.
- [62] A. Legrand, H. Renard, Y. Robert, and F. Vivien. Mapping and Load-Balancing Iterative Computations on Heterogeneous Clusters. In *Euro-PVM-MPI-2003: Recent Advances in Parallel Virtual Machine and Message Passing Interface*, LNCS 2840, pages 586–594. Springer Verlag, 2003.
- [63] C. Banino, O. Beaumont, A. Legrand, and Y. Robert. Scheduling Strategies for Master-Slave Tasking on Heterogeneous Processor Grids. In *PARA02: International Conference on Applied Parallel Computing*, LNCS 2367, pages 423–432, Espoo, Finland, June 2002. Springer Verlag.
- [64] O. Beaumont, L. Carter, J. Ferrante, A. Legrand, and Y. Robert. Bandwidth-Centric Allocation of Independent Tasks on Heterogeneous Platforms. In *International Parallel and Distributed Processing Symposium IPDPS'2002*. IEEE Computer Society Press, 2002. **(A/A1)**.
- [65] O. Beaumont, A. Legrand, and Y. Robert. A Polynomial-Time Algorithm for Allocating Independent Tasks on Heterogeneous Fork-Graphs. In *ISCIS XVII, Seventeenth International Symposium On Computer and Information Sciences*, pages 115–119. CRC Press, 2002. **(C/B2)**.
- [66] O. Beaumont, A. Legrand, and Y. Robert. Mixed Task and Data Parallelism. In *Parallel Matrix Algorithms and Applications*. Université de Neuchâtel, 2002.
- [67] O. Beaumont, A. Legrand, and Y. Robert. Static Scheduling Strategies for Dense Linear Algebra Kernels on Heterogeneous Clusters. In *Parallel Matrix Algorithms and Applications*. Université de Neuchâtel, 2002.
- [68] O. Beaumont, A. Legrand, and Y. Robert. Static Scheduling Strategies for Heterogeneous Systems. In *ISCIS XVII, Seventeenth International Symposium On Computer and Information Sciences*, pages 18–22. CRC Press, 2002. **(C/B2)**.
- [69] O. Beaumont, V. Boudet, A. Legrand, F. Rastello, and Y. Robert. Heterogeneous Matrix-Matrix Multiplication, or Partitioning a Square into Rectangles: NP-Completeness and Approximation Algorithms. In *EuroMicro Workshop on Parallel and Distributed Computing (EuroMicro2001)*, pages 298–305. IEEE Computer Society Press, 2001. **(C/B1)**.
- [70] O. Beaumont, A. Legrand, and Y. Robert. Master-Slave Tasking with Heterogeneous Processors. In *2001 International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA2001)*, pages 857–863. CSREA Press, 2001. **(B/B2)**.
- [71] O. Beaumont, A. Legrand, and Y. Robert. The Master-Slave Paradigm with Heterogeneous Processors. In D. S. Katz, T. Sterling, M. Baker, L. Bergman, M. Paprzycki, and R. Buyya, editors, *Cluster2001*, pages 419–426. IEEE Computer Society Press, 2001. **(A/A2)**.
- [72] O. Beaumont, V. Boudet, A. Legrand, F. Rastello, and Y. Robert. Dense Linear Algebra Kernels on Heterogeneous Platforms. In *Parallel Matrix Algorithms and Applications*. Université de Neuchâtel, 2000.
- [73] O. Beaumont, V. Boudet, A. Legrand, F. Rastello, and Y. Robert. Heterogeneity Considered Harmful to Algorithm Designers. In *Cluster'2000*, pages 403–404. IEEE Computer Society Press, 2000. **(A/A2)**.
- [74] H. Casanova, A. Legrand, D. Zagorodnov, and F. Berman. Heuristics for Scheduling Parameter Sweep Applications in Grid Environments. In *Heterogeneous Computing Workshop*, pages 349–363, Cancun, Mexico, 2000. **(B1)**.

Book chapters

- [75] L. Mello Schnorr and A. Legrand. Visualizing More Performance Data Than What Fits on Your Screen. In A. Cheptsov, S. Brinkmann, J. Gracia, M. M. Resch, and W. E. Nagel, editors, *Tools for High Performance Computing 2012*, pages 149–162. Springer Berlin Heidelberg, 2013.
- [76] L. Eyraud-Dubois and A. Legrand. The impact of platform models on scheduling techniques. In Y. Robert and F. Vivien, editors, *Introduction to Scheduling*. Taylor and Francis, 2009.
- [77] O. Beaumont, V. Boudet, A. Legrand, F. Rastello, and Y. Robert. Static Data Allocation and Load Balancing Techniques for Heterogeneous Systems. In C. K. Yuen, editor, *Annual Review of Scalable Computing*, volume 4, pages 1–37. World Scientific Publishing, 2002.

PhD thesis and Accreditation to Supervise Research (H.D.R.)

- [78] A. Legrand. *Scheduling for Large Scale Distributed Computing Systems: Approaches and Performance Evaluation Issues*. Habilitation à diriger des recherches, Université Grenoble Alpes, November 2015.
- [79] A. Legrand. *Algorithmique Parallèle Hétérogène Et Techniques D’ordonnancement : Approches Statiques Et Dynamiques*. PhD thesis, École Normale Supérieure de Lyon, December 2003.

Books

- [80] H. Casanova, A. Legrand, and Y. Robert. *Parallel Algorithms*. Chapman & Hall, 2008.
- [81] A. Legrand and Y. Robert. *Algorithmique Parallèle – Cours Et Exercices Corrigés*. Dunod, 2003.

Pre-prints and other documents

- [82] T. Cornebize, F. C. Heinrich, A. Legrand, and J. Vienne. Emulating High Performance Linpack on a Commodity Server at the Scale of a Supercomputer. working paper or preprint, December 2017.
- [83] V. Garcia Pinto, L. M. Schnorr, L. Stanisic, A. Legrand, S. Thibault, and V. Danjean. A Visual Performance Analysis Framework for Task-based Parallel Applications running on Hybrid Clusters. working paper or preprint, October 2017.