



## 5<sup>th</sup> Workshop of

LICIA

<http://licia-lab.org>

October 22-23, 2015

Universidade Federal do Rio Grande do Sul  
Instituto de Informática  
Auditorium 0 and 1 – Building 43412

Porto Alegre<sup>1</sup>

---

### Scientific Program

---



---

<sup>1</sup>Photo by Fernando Stankuns (Flickr).

# 1 Scientific Program

---

Thursday, October 22nd, 2015		
09h00	09h30	Welcome & Reception
09h30	10h30	<b>Keynote:</b> DataMove: Data Aware Large Scale Computing • Bruno Raffin (INRIA)
10h30	11h00	Coffee Break
11h00	11h30	<b>IIHM Team (LIG)</b> Plasticity, Creativity, Persuasion Presentation of the IIHM Team • Gaëlle Calvary (LIG)
11h30	12h30	Parallel tracks : Associated Teams, Working groups ExaSe/AimWest/ Computer Graphics IHM Distributed Computing
12h30	14h00	Lunch
14h00	15h00	<b>Keynote:</b> Games vs Distributed Control: The Case of Virus Propagation in SIRV Models • Bruno Gaujal (INRIA)
15h00	15h30	Coffee break
15h30	16h30	<b>Collaborations UFRGS/France</b> Situation and funding opportunities • Nicolas Maillard (UFRGS) CAPES/Brafitec Cooperation • Lucas M. Schnorr (UFRGS)
16h30	17h30	Parallel tracks Associated Teams, working groups ExaSe/AimWest/. . . Computer Graphics IHM Distributed Computing
17h30	18h30	LICIA Scientific Committee Meeting
20h00		Dinner (Location to be announced)

---

---

Friday, October 23rd, 2015		
09h00	09h30	Welcome & Reception
09h30	10h15	<b>Opening session</b> • Vladimir P. do Nascimento (UFRGS-PROPG) • Nicolas B. Maillard (Relinter) • Luis Lamb (INF/UFRGS) • Abilio Baeta Neves (President FAPERGS) • Bruno Raffin (LICIA) • Gaëlle Calvary (LIG) • Jean-Marc Vincent (LICIA)
10h15	10h30	Coffee Break
10h30	11h30	The LICIA and the FAPERGS Assoc. Teams • LICIA Presentation • ExaSe (FAPERGS/INRIA) • AimWest (FAPERGS/CNRS)
11h30	12h30	<b>Keynote:</b> Research on automation, robotics and informatics for health • Carlos Eduardo Pereira (UFRGS)
12h30	14h00	Lunch
14h00	14h30	<b>STEAMER Team (LIG)</b> Spatio-Temporal data and knowledge processing in the Geoweb context : research agenda at the Steamer Team • Paule-Annick Davoine (LIG)
14h30	16h00	<b>Student Session</b> • Stefano Mor • Alan Velasques • David Beniamine • Francieli Boito
16h00	16h30	Coffee Break
16h30	17h30	<b>Emerging Trends</b> Interactive verification of outsourced computations • Jean-Louis Roch Behaviour and Profile Based Opinion simplification • Vinicius Woloszyn and Leandro Wives
17h30	18h00	Closing
20h00		Dinner (Location to be announced)

---

## 2 Detailed program

### 2.1 Keynotes

- **Bruno Raffin** (INRIA)

- **Title:** DataMove: Data Aware Large Scale Computing

- **Abstract:** Researchers from LIG/INRIA are gathering into a new team called DataMove. In this talk we will present the research axis of the DataMove with a specific focus on in situ analytics. DataMove propose to gather will focus on optimizing resource allocation for large scale computing with a strong focus on data movements. More specifically we will develop multi- objective allocation policies based on extended and refined models of applications and machine resources. We will take into account data movements and data dependencies internal to each application but also having a broader view of data lifespan from their origin down to their analysis, visualization or storage. Developing methods

for forecasting the application behavior and resource availability will further enable to improve resource allocation. For instance we will revisit the classical resource management methods taking into account the resource heterogeneity (network topology and location of I/O nodes) for optimizing the length of data routing paths. We will also consider new processes motivated by the ever growing amount of generated data. The emerging trend of in-situ analytics proposes to more tightly couple simulations and result analysis. The goal is to reduce the data traffic and to speed-up result analysis, by performing result processing (compression, indexation, analysis, visualization, etc.) as closely as possible to the locus and time of data generation. This paradigm requires to revisit the traditional workflow (batch processing + postmortem analysis). The application becomes a whole including the simulation, in-situ processing and I/Os. This motivates the development of adapted resource allocation strategies (cores, memory, network) to efficiently interleave the different components of the application and globally improve the performance.

- **Short bio:** Bruno Raffin (PhD, HDR) is associate researcher at INRIA Grenoble since 2001. He has a PhD from the Université d'Orléans on parallel programming language design (1997). After a 2 year postdoc at Iowa State University he refocused his research activity on 3D high performance interactive computing. He led the development of the FlowVR middleware dedicated to parallel interactive applications, initiated and steered the multi-camera Grimage platform (<http://grimage.inrialpes.fr>) used to develop real-time full-body 3D interactions and 3D telepresence. He also works on parallel algorithms and cache-efficient parallel 3D data structures (cache oblivious mesh layouts), targeting multi-CPU and multi-GPU machines with applications to real-time physics simulations and scientific visualization. Bruno Raffin accounts for more than 50 international publications, 11 advised PhD students. He is a long time committee member of IEEE VR, the major international virtual reality conference, and steering committee member of the Eurographics Symposium on Parallel Graphics and Visualization.

- **Carlos Eduardo Pereira** (UFRGS)

- **Title:** Research on automation, robotics and informatics for health
- **Abstract:**
- **Short bio:** Prof. Carlos Eduardo Pereira is full professor at the Engineering department of the Federal University of Rio Grande do Sul and vice-director of the science and technology park at UFRGS – Porto Alegre, Brazil. He earned his PhD. at the University of Stuttgart in Electrical Engineering. He is also member of the Board of the International Federation of Automatic Control - IFAC, was chairman of the Board of the Brazilian Control Automation (SBA). He is researcher level 1 of the CNPq, and member of the director board of the Secretariat of the CAPES International Affairs Committee. He is author of over 400 technical publications, having served as a member of the International Program Committees of several conferences in the areas of industrial automation, manufacturing, industrial protocols and distributed object computing in real time. In 2012, he received the Friedrich Wilhelm Bessel prize for research, from the Alexander von Humboldt Foundation in Germany, and in 2010 he was awarded the Santander Prize University, the first placed in the Technology area of information, Communication and Education. His research focuses on methodologies and support tools for the development of distributed embedded systems for real-time computing, with particular emphasis on automation applications.

- **Bruno Gaujal** (INRIA)

- **Title:** Games vs Distributed Control: The Case of Virus Propagation in SIRV Models
- **Abstract:** We consider a simple SIRV (Susceptible, Infected, Recovered, Vaccinated) model where each individual chooses its vaccination strategy.  
This simple problem can be seen as a mean field game and/or a distributed control problem. We discuss both approaches and show their similarities and differences. We also show how a pricing mechanism can be used to make all the individuals adopt the social optimal vaccination strategy given by the optimal control approach.
- **Short bio:** Bruno Gaujal is an Inria Research Director. He is the head of the large-scale computing group, MESCAL, in the research center of Inria Grenoble Rhône-Alpes. He has held several positions in AT&T Bell Labs, INRIA Sophia-Antipolis, Loria and Ecole Normale Supérieure in Lyon. He obtained his PhD from University of Nice, under the supervision of François Baccelli, in 1994. He got his « Habilitation à diriger des recherches » in 2001 from the University of Nancy. He is the author of more than 100 scientific publications in journals and international conferences. He is a founding partner and a scientific advisor of a start-up company, RTaW, since 2007. His main interests are in performance evaluation, optimization and control of large discrete event dynamic systems with applications to telecommunications networks and large computing infrastructures.

## 2.2 Presentations

- **Nicolas Maillard** (UFRGS)

- **Title:** Collaborations UFRGS/France: situation and funding opportunities
- **Abstract:** Nicolas Maillard will present a panorama of the scientific collaborations between UFRGS and France in terms of student exchange, scientific missions and joint research. This context will highlight the unique situation of Grenoble,

where all means of international partnership have been exploited for more than three decades. A sketch of every modality will be provided (CAPES/COFECUB, joint or double degrees, BRAFITEC/BRAFAGRI, Erasmus \* programs, STIC-AMSUD, CNPq/INRIA, FAP/INRIA, INRIA Équipe Associée, CNPq/CNRS, Science without Border,...) and possible new trends will be discussed.

- **Short bio:** Nicolas Maillard is currently head of UFRGS International Office, where he started to work in Oct. 2012. Of French nationality, he graduated in 1996 at the French “Grande École d’Ingénieur” ENSIMAG – INPG and obtained a PhD in Information Sciences and Technologies at the Université Joseph Fourier, in 2001. He has started his work as “Professor Adjunto” at the Federal University of Rio Grande do Sul, Porto Alegre, Brazil, in 2004. He has taught Operating System, Compilers and Parallel Programming from 2004 to 2012. In 2012, He spent 6 months as invited scholar at the University of Pittsburgh, USA. Nicolas’ research field is High-Performance Computing. He has advised 14 Master’s and 4 PhD students and is currently advising 1 PhD student. He has published 19 papers in international conferences with editorial review and 8 in international journals.

## 2.3 Emerging Trends

- **Jean-Louis Roch** (Université de Grenoble, LIG-Inria)

- **Title:** Interactive verification of outsourced computations
- **Abstract:** When a client outsources most of the computation to an external server, typically a HPC platform or a cloud, how to efficiently check the correctness of the obtained result? Thanks to randomization, interaction between different computing models appears successful. We present in this talk some examples of efficient verifications, either checking or interacting with the server for proof or for algorithm-based fault tolerance.

- **Vinicius Woloszyn (vinicius@open.inf.br) Leandro Wives (wives@inf.ufrgs.br) - advisor (UFRGS)**

- **Title:** Behaviour and Profile Based Opinion simplification
- **Abstract:** Opinion mining and summary generation is usually performed without considering the user vocabulary neither their interests. A simplified version of a document may not be based on the use of more common (and simple) words but one that uses words that are known to the reader. We also believe that their interest and behaviour is important. For instance, previous user experiences either positive and negative may affect their judgement. In this talk we will show the current state of the research.  
Keywords: opinion mining, text simplification, behaviour analysis.

## 2.4 Student Session

- **Stéfano Drimon Kurz Mor** (UFRGS)

- **Title:** Analysis of Synchronizations in Greedy-Scheduled Executions and Applications to Efficient Generation of Pseudorandom Numbers in Parallel
- **Abstract:** We present two contributions to the field of parallel programming. The first contribution is theoretical: we introduce SIPS analysis, a novel approach to estimate the number of synchronizations performed during the execution of a parallel algorithm. It generalizes the concept of logical clocks to account the number of synchronizations performed by a parallel algorithm and is able to deliver worst-case bounds on it even in the presence of non-deterministic parallel executions, which are generally not covered by state-of-the-art analysis. It also allows us to estimate new worst-case bounds for computations scheduled by the popular work stealing algorithm and to design more efficient adaptable parallel programs. The second contribution is pragmatic: we present a parallelization of generic algorithms for pseudorandom number generation in current hardware. As an alternative to fixed implementations of random generator components we introduce a generic API called PAR-R, which is designed and analyzed using SIPS. Its main insight is the use of a sequential generator R’s capability of performing an efficient “jump-ahead” in the generated stream; we bound it to steal operations performed by a work stealing scheduler, which our SIPS-based analysis demonstrates to occur rarely in expectation. PAR-R is compared to the parallel pseudorandom number generator DotMix, written for the Cilk Plus dynamic multithreading platform. The polylogarithm overhead of PAR-R compares favorably to the linear overhead of DotMix’s re-seedings, while not requiring a fixed generator underneath.

- **David Beniamine** (LIG/INRIA)

- **Title:** Analyzing an application memory behavior
- **Abstract:** In modern parallel architectures, memory accesses represent a common bottleneck. Thus, optimizing the way applications access the memory is an important way to improve performance and energy consumption. Memory accesses are even more important with NUMA machines, as the access time to data depends on its location in the memory. Most efforts on memory optimization, focus on mappings tools that tries to dynamically move thread and/or data on the machine to insure locality. Only a few profilers were designed to help the user understanding its application memory

behavior. Furthermore as collection memory trace is hard and costly, most of these tools relies on indirect information such a performance counters.

We present two tools to collect memory trace and analyze an application memory behavior. The first focus on data sharing on NUMA nodes while the second aims at providing a global overview of the memory pattern over time.

- **Francieli Zanon Boito** (UFRGS)

- **Title:** I/O Scheduling Algorithm Selection for Parallel File Systems

- **Abstract:** High Performance Computing applications rely on Parallel File Systems (PFS) to achieve good performance even when handling large amounts of data. It is usual for HPC systems to provide a shared storage infrastructure for applications. In this situation, when multiple applications are concurrently accessing the shared PFS, their accesses will affect each other in a phenomenon called “interference”, which compromises I/O optimization techniques’ efficacy. In this talk, we focus on I/O scheduling as a tool to alleviate interference’s effects.

We have conducted an extensive performance evaluation of five scheduling algorithms at a parallel file system’s data servers. Experiments were executed on different platforms and under different access patterns. Results indicate that schedulers’ results are deeply affected by applications’ access patterns and by the underlying I/O system characteristics - especially by storage devices. Our results have shown that there is no scheduling algorithm able to improve performance for all situations, and the best choice depends on applications’ and storage devices’ characteristics.

For these reasons, we will discuss our approach to provide I/O scheduling with adaptivity to applications and devices. We use information about these two aspects to automatically select the best fit in scheduling algorithm to each situation. Our approach has provided better results than using the same algorithm to all situations - without adaptability - due to successfully applying I/O scheduling techniques to improve performance while avoiding situations where it would lead to performance impairment.