

The logo for Inria, featuring the word "Inria" in a stylized, cursive red font.

IN PARTNERSHIP WITH:  
**CNRS**

**Université de Grenoble Alpes**

Activity Report 2019

## **Project-Team POLARIS**

Performance analysis and Optimization of  
LARge Infrastructures and Systems

IN COLLABORATION WITH: Laboratoire d'Informatique de Grenoble (LIG)

RESEARCH CENTER  
**Grenoble - Rhône-Alpes**

THEME  
**Distributed and High Performance  
Computing**



## Table of contents

<b>1. Team, Visitors, External Collaborators</b> .....	<b>1</b>
<b>2. Overall Objectives</b> .....	<b>3</b>
2.1. Context	3
2.2. Objectives	3
<b>3. Research Program</b> .....	<b>4</b>
3.1. Sound and Reproducible Experimental Methodology	4
3.2. Multi-Scale Analysis and Visualization	4
3.3. Fast and Faithful Performance Prediction of Very Large Systems	5
3.4. Local Interactions and Transient Analysis in Adaptive Dynamic Systems	5
3.5. Distributed Learning in Games and Online Optimization	6
<b>4. Application Domains</b> .....	<b>7</b>
4.1. Large Computing Infrastructures	7
4.2. Next-Generation Wireless Networks	7
4.3. Energy and Transportation	7
4.4. Social Computing Systems	8
<b>5. Highlights of the Year</b> .....	<b>8</b>
<b>6. New Software and Platforms</b> .....	<b>8</b>
6.1. Framesoc	8
6.2. GameSeer	8
6.3. marmoteCore	9
6.4. Moca	9
6.5. Ocelotl	9
6.6. PSI	9
6.7. SimGrid	9
6.8. Tabarnac	10
<b>7. New Results</b> .....	<b>10</b>
7.1. Design of Experiments	10
7.2. Predictive Simulation of HPC Applications	11
7.3. Simulation of Smart Grids	11
7.4. Batch Scheduling	12
7.5. Load Balancing	12
7.6. FoG Computing	13
7.7. Research Management: Research Reproducibility and Credit	14
7.8. Mean Field Games and Control	15
7.9. Energy and Network Optimization	15
7.10. Privacy, Fairness, and Transparency in Online Social Medias	16
7.11. Optimization Methods	18
7.12. Learning	19
<b>8. Bilateral Contracts and Grants with Industry</b> .....	<b>20</b>
<b>9. Partnerships and Cooperations</b> .....	<b>20</b>
9.1. Regional Initiatives	20
9.2. National Initiatives	20
9.2.1. Inria Project Labs	20
9.2.2. Grenoble INP grant	21
9.2.3. DGA Grants	21
9.2.4. PGMO Projects	21
9.2.5. PEPS	21
9.2.6. Fondation Blaise Pascal	21
9.2.7. MIAI @ Grenoble Alpes	21

---

9.2.8. ANR	21
9.3. International Initiatives	21
9.3.1.1. ReDaS	21
9.3.1.2. International Initiatives	22
9.4. International Research Visitors	22
9.4.1. Visits of International Scientists	22
9.4.2. Visits to International Teams	22
<b>10. Dissemination</b> .....	<b>23</b>
10.1. Promoting Scientific Activities	23
10.1.1. Scientific Events: Organisation	23
10.1.1.1. General Chair, Scientific Chair	23
10.1.1.2. Member of the Organizing Committees	23
10.1.2. Scientific Events: Selection	23
10.1.3. Journal	23
10.1.3.1. Member of the Editorial Boards	23
10.1.3.2. Reviewer - Reviewing Activities	23
10.1.4. Invited Talks	24
10.1.5. Research Administration	24
10.2. Teaching - Supervision - Juries	24
10.2.1. Teaching	24
10.2.2. Supervision	25
10.2.3. Juries	26
10.3. Popularization	26
10.3.1. Internal or external Inria responsibilities	26
10.3.2. Articles and contents	26
10.3.3. Education	27
10.3.4. Interventions	27
<b>11. Bibliography</b> .....	<b>27</b>

# Project-Team POLARIS

*Creation of the Team: 2016 January 01, updated into Project-Team: 2018 January 01*

## Keywords:

### Computer Science and Digital Science:

- A1.1.1. - Multicore, Manycore
- A1.1.2. - Hardware accelerators (GPGPU, FPGA, etc.)
- A1.1.4. - High performance computing
- A1.1.5. - Exascale
- A1.2. - Networks
- A1.2.3. - Routing
- A1.2.5. - Internet of things
- A1.6. - Green Computing
- A3.4. - Machine learning and statistics
- A3.5.2. - Recommendation systems
- A5.2. - Data visualization
- A6. - Modeling, simulation and control
- A6.2.3. - Probabilistic methods
- A6.2.4. - Statistical methods
- A6.2.6. - Optimization
- A6.2.7. - High performance computing
- A8.2. - Optimization
- A8.9. - Performance evaluation
- A8.11. - Game Theory

### Other Research Topics and Application Domains:

- B4.4. - Energy delivery
- B4.4.1. - Smart grids
- B4.5.1. - Green computing
- B6.2. - Network technologies
- B6.2.1. - Wired technologies
- B6.2.2. - Radio technology
- B6.4. - Internet of things
- B8.3. - Urbanism and urban planning
- B9.6.7. - Geography
- B9.7.2. - Open data
- B9.8. - Reproducibility

## 1. Team, Visitors, External Collaborators

### Research Scientists

Arnaud Legrand [Team leader, CNRS, Researcher, HDR]  
Jonatha Anselmi [Inria, Researcher, from Apr 2019]

Nicolas Gast [Inria, Researcher]  
Bruno Gaujal [Inria, Senior Researcher, HDR]  
Patrick Loiseau [Inria, Researcher, HDR]  
Panayotis Mertikopoulos [CNRS, Researcher, HDR]  
Bary Pradelski [CNRS, Researcher]

**Faculty Members**

Vincent Danjean [Univ Grenoble Alpes, Associate Professor]  
Guillaume Huard [Univ Grenoble Alpes, Associate Professor]  
Florence Perronnin [Univ Grenoble Alpes, Associate Professor, HDR]  
Jean-Marc Vincent [Univ Grenoble Alpes, Associate Professor]  
Philippe Waille [Univ Grenoble Alpes, Associate Professor]

**Post-Doctoral Fellows**

Olivier Bilenne [CNRS, Post-Doctoral Fellow]  
Takai Eddine Kennouche [Univ Grenoble Alpes, Post-Doctoral Fellow]  
Mouhcine Mendil [INP SA, Post-Doctoral Fellow, from Jul 2019]

**PhD Students**

Kimion Antonakopoulos [Inria, PhD Student]  
Thomas Barzola [Univ Grenoble Alpes, PhD Student, from Oct 2019]  
Tom Cornebize [Univ Grenoble Alpes, PhD Student]  
Bruno de Moura Donassolo [Orange Labs, PhD Student, granted by CIFRE]  
Vitalii Emelianov [Inria, PhD Student]  
Franz Christian Heinrich [Inria, PhD Student, until Feb 2019]  
Alexis Janon [Univ Grenoble Alpes, PhD Student]  
Baptiste Jonglez [Univ Grenoble Alpes, PhD Student]  
Kimang Khun [Inria, PhD Student, from Oct 2019]  
Alexandre Marcastel [Univ de Cergy Pontoise, until Feb 2019]  
Dimitrios Moustakas [Inria, PhD Student, from Dec 2019]  
Umar Ozeer [Orange Labs, PhD Student, until Oct 2019, granted by CIFRE]  
Stephan Plassart [Univ Grenoble Alpes, PhD Student]  
Pedro Rocha Bruel [UFRGS, PhD Student]  
Benjamin Roussillon [Univ Grenoble Alpes, PhD Student]  
Vera Sosnovik [Univ Grenoble Alpes, PhD Student, from Oct 2019]  
Dong Quan Vu [Nokia Bell Labs, PhD Student, granted by CIFRE]  
Chen Yan [Univ Grenoble Alpes, PhD Student, from Oct 2019]

**Technical staff**

Eleni Gkiouzepi [CNRS, Engineer, from Dec 2019]  
Mouhcine Mendil [INP SA, Engineer, until Jul 2019]

**Interns and Apprentices**

Manal Benaïssa [Inria, until Jul 2019]  
Nils Defauw [Inria, until Jul 2019]  
Leo Gayral [Univ Grenoble Alpes, until Jun 2019]  
Alice Jacquier [Inria, until Jul 2019]  
Hoel Jalmin [Inria, from Apr 2019 until Jul 2019]  
Kimang Khun [Inria, from Apr 2019 until Aug 2019]  
Dimitrios Moustakas [Univ Grenoble Alpes, from Feb 2019 until Aug 2019]  
Nicolas Rocher [Univ Grenoble Alpes, from Apr 2019]  
Chen Yan [Univ Grenoble Alpes, from Feb 2019 until Jul 2019]

**Administrative Assistant**

Annie Simon [Inria, Administrative Assistant]

## 2. Overall Objectives

### 2.1. Context

Large distributed infrastructures are rampant in our society. Numerical simulations form the basis of computational sciences and high performance computing infrastructures have become scientific instruments with similar roles as those of test tubes or telescopes. Cloud infrastructures are used by companies in such an intense way that even the shortest outage quickly incurs the loss of several millions of dollars. But every citizen also relies on (and interacts with) such infrastructures via complex wireless mobile embedded devices whose nature is constantly evolving. In this way, the advent of digital miniaturization and interconnection has enabled our homes, power stations, cars and bikes to evolve into smart grids and smart transportation systems that should be optimized to fulfill societal expectations.

Our dependence and intense usage of such gigantic systems obviously leads to very high expectations in terms of performance. Indeed, we strive for low-cost and energy-efficient systems that seamlessly adapt to changing environments that can only be accessed through uncertain measurements. Such digital systems also have to take into account both the users' profile and expectations to efficiently and fairly share resources in an online way. Analyzing, designing and provisioning such systems has thus become a real challenge.

Such systems are characterized by their **ever-growing size**, intrinsic **heterogeneity** and **distributedness**, **user-driven** requirements, and an unpredictable variability that renders them essentially **stochastic**. In such contexts, many of the former design and analysis hypotheses (homogeneity, limited hierarchy, omniscient view, optimization carried out by a single entity, open-loop optimization, user outside of the picture) have become obsolete, which calls for radically new approaches. Properly studying such systems requires a drastic rethinking of fundamental aspects regarding the system's **observation** (measure, trace, methodology, design of experiments), **analysis** (modeling, simulation, trace analysis and visualization), and **optimization** (distributed, online, stochastic).

### 2.2. Objectives

The goal of the POLARIS project is to **contribute to the understanding of the performance of very large scale distributed systems** by applying ideas from diverse research fields and application domains. We believe that studying all these different aspects at once without restricting to specific systems is the key to push forward our understanding of such challenges and to proposing innovative solutions. This is why we intend to investigate problems arising from application domains as varied as large computing systems, wireless networks, smart grids and transportation systems.

The members of the POLARIS project cover a very wide spectrum of expertise in performance evaluation and models, distributed optimization, and analysis of HPC middleware. Specifically, POLARIS' members have worked extensively on:

Experiment design: Experimental methodology, measuring/monitoring/tracing tools, experiment control, design of experiments, and reproducible research, especially in the context of large computing infrastructures (such as computing grids, HPC, volunteer computing and embedded systems).

Trace Analysis: Parallel application visualization (paje, triva/viva, framesoc/ocelotl, ...), characterization of failures in large distributed systems, visualization and analysis for geographical information systems, spatio-temporal analysis of media events in RSS flows from newspapers, and others.

Modeling and Simulation: Emulation, discrete event simulation, perfect sampling, Markov chains, Monte Carlo methods, and others.

Optimization: Stochastic approximation, mean field limits, game theory, discrete and continuous optimization, learning and information theory.

In the rest of this document, we describe in detail our new results in the above areas.

## 3. Research Program

### 3.1. Sound and Reproducible Experimental Methodology

**Participants:** Vincent Danjean, Nicolas Gast, Guillaume Huard, Arnaud Legrand, Patrick Loiseau, Jean-Marc Vincent.

Experiments in large scale distributed systems are costly, difficult to control and therefore difficult to reproduce. Although many of these digital systems have been built by men, they have reached such a complexity level that we are no longer able to study them like artificial systems and have to deal with the same kind of experimental issues as natural sciences. The development of a sound experimental methodology for the evaluation of resource management solutions is among the most important ways to cope with the growing complexity of computing environments. Although computing environments come with their own specific challenges, we believe such general observation problems should be addressed by borrowing good practices and techniques developed in many other domains of science.

This research theme builds on a transverse activity on *Open science and reproducible research* and is organized into the following two directions: (1) *Experimental design* (2) *Smart monitoring and tracing*. As we will explain in more detail hereafter, these transverse activity and research directions span several research areas and our goal within the POLARIS project is foremost to transfer original ideas from other domains of science to the distributed and high performance computing community.

### 3.2. Multi-Scale Analysis and Visualization

**Participants:** Vincent Danjean, Guillaume Huard, Arnaud Legrand, Jean-Marc Vincent, Panayotis Mertikopoulos.

As explained in the previous section, the first difficulty encountered when modeling large scale computer systems is to observe these systems and extract information on the behavior of both the architecture, the middleware, the applications, and the users. The second difficulty is to *visualize* and *analyze* such *multi-level traces to understand how the performance of the application can be improved*. While a lot of efforts are put into visualizing scientific data, in comparison little effort have gone into developing techniques specifically tailored for understanding the behavior of distributed systems. Many visualization tools have been developed by renowned HPC groups since decades (e.g., BSC [91], Jülich and TU Dresden [90], [61], UIUC [79], [94], [82] and ANL [107], Inria Bordeaux [67] and Grenoble [109], ...) but most of these tools build on the classical information visualization mantra [99] that consists in always first presenting an overview of the data, possibly by plotting everything if computing power allows, and then to allow users to zoom and filter, providing details on demand. However in our context, the amount of data comprised in such traces is several orders of magnitude larger than the number of pixels on a screen and displaying even a small fraction of the trace leads to harmful visualization artifacts [86]. Such traces are typically made of events that occur at very different time and space scales, which unfortunately hinders classical approaches. Such visualization tools have focused on easing interaction and navigation in the trace (through gantcharts, intuitive filters, pie charts and kiviats) but they are very difficult to maintain and evolve and they require some significant experience to identify performance bottlenecks.

Therefore many groups have more recently proposed in combination to these tools some techniques to help identifying the structure of the application or regions (applicative, spatial or temporal) of interest. For example, researchers from the SDSC [89] propose some segment matching techniques based on clustering (Euclidean or Manhattan distance) of start and end dates of the segments that enables to reduce the amount of information to display. Researchers from the BSC use clustering, linear regression and Kriging techniques [98], [85], [78] to identify and characterize (in term of performance and resource usage) application phases and present aggregated representations of the trace [97]. Researchers from Jülich and TU Darmstadt have proposed techniques to identify specific communication patterns that incur wait states [104], [54]



### 3.3. Fast and Faithful Performance Prediction of Very Large Systems

**Participants:** Jonatha Anselmi, Vincent Danjean, Bruno Gaujal, Arnaud Legrand, Florence Perronnin, Jean-Marc Vincent.

Evaluating the scalability, robustness, energy consumption and performance of large infrastructures such as exascale platforms and clouds raises severe methodological challenges. The complexity of such platforms mandates empirical evaluation but direct experimentation via an application deployment on a real-world testbed is often limited by the few platforms available at hand and is even sometimes impossible (cost, access, early stages of the infrastructure design, ...). Unlike direct experimentation via an application deployment on a real-world testbed, simulation enables fully repeatable and configurable experiments that can often be conducted quickly for arbitrary hypothetical scenarios. In spite of these promises, current simulation practice is often not conducive to obtaining scientifically sound results. To date, most simulation results in the parallel and distributed computing literature are obtained with simulators that are ad hoc, unavailable, undocumented, and/or no longer maintained. For instance, Naicken et al. [53] point out that out of 125 recent papers they surveyed that study peer-to-peer systems, 52% use simulation and mention a simulator, but 72% of them use a custom simulator. As a result, most published simulation results build on throw-away (short-lived and non validated) simulators that are specifically designed for a particular study, which prevents other researchers from building upon it. There is thus a strong need for recognized simulation frameworks by which simulation results can be reproduced, further analyzed and improved.

The *SimGrid* simulation toolkit [65], whose development is partially supported by POLARIS, is specifically designed for studying large scale distributed computing systems. It has already been successfully used for simulation of grid, volunteer computing, HPC, cloud infrastructures and we have constantly invested on the software quality, the scalability [57] and the validity of the underlying network models [55], [102]. Many simulators of MPI applications have been developed by renowned HPC groups (e.g., at SDSC [100], BSC [51], UIUC [108], Sandia Nat. Lab. [103], ORNL [64] or ETH Zürich [80] for the most prominent ones). Yet, to scale most of them build on restrictive network and application modeling assumptions that make them difficult to extend to more complex architectures and to applications that do not solely build on the MPI API. Furthermore, simplistic modeling assumptions generally prevent to faithfully predict execution times, which limits the use of simulation to indication of gross trends at best. Our goal is to improve the quality of SimGrid to the point where it can be used effectively on a daily basis by practitioners to *reproduce the dynamic of real HPC systems*.

We also develop another simulation software, *PSI* (Perfect SIMulator) [69], [62], dedicated to the simulation of very large systems that can be modeled as Markov chains. PSI provides a set of simulation kernels for Markov chains specified by events. It allows one to sample stationary distributions through the Perfect Sampling method (pioneered by Propp and Wilson [92]) or simply to generate trajectories with a forward Monte-Carlo simulation leveraging time parallel simulation (pioneered by Fujimoto [73], Lin and Lazowska [84]). One of the strength of the PSI framework is its expressiveness that allows us to easily study networks with finite and infinite capacity queues [63]. Although PSI already allows to simulate very large and complex systems, our main objective is to push its scalability even further and *improve its capabilities by one or several orders of magnitude*.

### 3.4. Local Interactions and Transient Analysis in Adaptive Dynamic Systems

**Participants:** Jonatha Anselmi, Nicolas Gast, Bruno Gaujal, Florence Perronnin, Jean-Marc Vincent, Panayotis Mertikopoulos.

Many systems can be effectively described by stochastic population models. These systems are composed of a set of  $n$  entities interacting together and the resulting stochastic process can be seen as a continuous-time Markov chain with a finite state space. Many numerical techniques exist to study the behavior of Markov chains, to solve stochastic optimal control problems [93] or to perform model-checking [52]. These techniques, however, are limited in their applicability, as they suffer from the *curse of dimensionality*: the state-space grows exponentially with  $n$ .

This results in the need for approximation techniques. Mean field analysis offers a viable, and often very accurate, solution for large  $n$ . The basic idea of the mean field approximation is to count the number of entities that are in a given state. Hence, the fluctuations due to stochasticity become negligible as the number of entities grows. For large  $n$ , the system becomes essentially deterministic. This approximation has been originally developed in statistical mechanics for vary large systems composed of more than  $10^{20}$  particles (called entities here). More recently, it has been claimed that, under some conditions, this approximation can be successfully used for stochastic systems composed of a few tens of entities. The claim is supported by various convergence results [74], [83], [106], and has been successfully applied in various domains: wireless networks [56], computer-based systems [77], [88], [101], epidemic or rumour propagation [66], [81] and bike-sharing systems [70]. It is also used to develop distributed control strategies [105], [87] or to construct approximate solutions of stochastic model checking problems [58], [59], [60].

Within the POLARIS project, we will continue developing both the theory behind these approximation techniques and their applications. Typically, these techniques require a homogeneous population of objects where the dynamics of the entities depend only on their state (the state space of each object must not scale with  $n$  the number of objects) but neither on their identity nor on their spatial location. Continuing our work in [74], we would like to be able to handle heterogeneous or uncertain dynamics. Typical applications are caching mechanisms [77] or bike-sharing systems [71]. A second point of interest is the use of mean field or large deviation asymptotics to compute the time between two regimes [96] or to reach an equilibrium state. Last, mean-field methods are mostly descriptive and are used to analyse the performance of a given system. We wish to extend their use to solve optimal control problems. In particular, we would like to implement numerical algorithms that use the framework that we developed in [75] to build distributed control algorithms [68] and optimal pricing mechanisms [76].

### 3.5. Distributed Learning in Games and Online Optimization

**Participants:** Nicolas Gast, Bruno Gaujal, Arnaud Legrand, Patrick Loiseau, Panayotis Mertikopoulos, Bary Pradelski.

Game theory is a thriving interdisciplinary field that studies the interactions between competing optimizing agents, be they humans, firms, bacteria, or computers. As such, game-theoretic models have met with remarkable success when applied to complex systems consisting of interdependent components with vastly different (and often conflicting) objectives – ranging from latency minimization in packet-switched networks to throughput maximization and power control in mobile wireless networks.

In the context of large-scale, decentralized systems (the core focus of the POLARIS project), it is more relevant to take an inductive, “bottom-up” approach to game theory, because the components of a large system cannot be assumed to perform the numerical calculations required to solve a very-large-scale optimization problem. In view of this, POLARIS’ overarching objective in this area is to *develop novel algorithmic frameworks that offer robust performance guarantees when employed by all interacting decision-makers.*

A key challenge here is that most of the literature on learning in games has focused on *static* games with a *finite number of actions* per player [72], [95]. While relatively tractable, such games are ill-suited to practical applications where players pick an action from a continuous space or when their payoff functions evolve over time – this being typically the case in our target applications (e.g., routing in packet-switched networks or energy-efficient throughput maximization in wireless). On the other hand, the framework of online convex optimization typically provides worst-case performance bounds on the learner’s *regret* that the agents can attain irrespectively of how their environment varies over time. However, if the agents’ environment is determined chiefly by their interactions these bounds are fairly loose, so more sophisticated convergence criteria should be applied.

From an algorithmic standpoint, a further challenge occurs when players can only observe their own payoffs (or a perturbed version thereof). In this bandit-like setting regret-matching or trial-and-error procedures guarantee convergence to an equilibrium in a weak sense in certain classes of games. However, these results apply exclusively to static, finite games: learning in games with continuous action spaces and/or nonlinear

payoff functions cannot be studied within this framework. Furthermore, even in the case of finite games, the complexity of the algorithms described above is not known, so it is impossible to decide a priori which algorithmic scheme can be applied to which application.

## 4. Application Domains

### 4.1. Large Computing Infrastructures

Supercomputers typically comprise thousands to millions of multi-core CPUs with GPU accelerators interconnected by complex interconnection networks that are typically structured as an intricate hierarchy of network switches. Capacity planning and management of such systems not only raises challenges in term of computing efficiency but also in term of energy consumption. Most legacy (SPMD) applications struggle to benefit from such infrastructure since the slightest failure or load imbalance immediately causes the whole program to stop or at best to waste resources. To scale and handle the stochastic nature of resources, these applications have to rely on dynamic runtimes that schedule computations and communications in an opportunistic way. Such evolution raises challenges not only in terms of programming but also in terms of observation (complexity and dynamicity prevents experiment reproducibility, intrusiveness hinders large scale data collection, ...) and analysis (dynamic and flexible application structures make classical visualization and simulation techniques totally ineffective and require to build on *ad hoc* information on the application structure).

### 4.2. Next-Generation Wireless Networks

Considerable interest has arisen from the seminal prediction that the use of multiple-input, multiple-output (MIMO) technologies can lead to substantial gains in information throughput in wireless communications, especially when used at a massive level. In particular, by employing multiple inexpensive service antennas, it is possible to exploit spatial multiplexing in the transmission and reception of radio signals, the only physical limit being the number of antennas that can be deployed on a portable device. As a result, the wireless medium can accommodate greater volumes of data traffic without requiring the reallocation (and subsequent re-regulation) of additional frequency bands. In this context, throughput maximization in the presence of interference by neighboring transmitters leads to games with convex action sets (covariance matrices with trace constraints) and individually concave utility functions (each user's Shannon throughput); developing efficient and distributed optimization protocols for such systems is one of the core objectives of Theme 5.

Another major challenge that occurs here is due to the fact that the efficient physical layer optimization of wireless networks relies on perfect (or close to perfect) channel state information (CSI), on both the uplink and the downlink. Due to the vastly increased computational overhead of this feedback – especially in decentralized, small-cell environments – the ongoing transition to fifth generation (5G) wireless networks is expected to go hand-in-hand with distributed learning and optimization methods that can operate reliably in feedback-starved environments. Accordingly, one of POLARIS' application-driven goals will be to leverage the algorithmic output of Theme 5 into a highly adaptive resource allocation framework for next-generation wireless systems that can effectively "learn in the dark", without requiring crippling amounts of feedback.

### 4.3. Energy and Transportation

Smart urban transport systems and smart grids are two examples of collective adaptive systems. They consist of a large number of heterogeneous entities with decentralised control and varying degrees of complex autonomous behaviour. We develop an analysis tools to help to reason about such systems. Our work relies on tools from fluid and mean-field approximation to build decentralized algorithms that solve complex optimization problems. We focus on two problems: decentralized control of electric grids and capacity planning in vehicle-sharing systems to improve load balancing.

## 4.4. Social Computing Systems

Social computing systems are online digital systems that use personal data of their users at their core to deliver personalized services directly to the users. They are omnipresent and include for instance recommendation systems, social networks, online medias, daily apps, etc. Despite their interest and utility for users, these systems pose critical challenges of privacy, security, transparency, and respect of certain ethical constraints such as fairness. Solving these challenges involves a mix of measurement and/or audit to understand and assess issues, and modeling and optimization to propose and calibrate solutions.

# 5. Highlights of the Year

## 5.1. Highlights of the Year

- N. Gast received an ANR JCJC grant.
- The team was highly involved in the 3IA institute MIAI @ Grenoble Alpes: P. Loiseau is co-holder of the chair “Explainable and Responsible AI” of which N. Gast and B. Pradelski are members; and P. Mertikopoulos is a member of the chair “Optimization & Learning”.
- Arnaud Legrand participated in the writing of a book [39] on Reproducible Research, which aims at helping students and engineers and researchers to find efficient and accessible ways leading them to improve their reproducible research practices.

### 5.1.1. Awards

- The paper “Privacy Risks with Facebook’s PII-based Targeting: Auditing a Data Broker’s Advertising Interface” by P. Loiseau and co-authors (IEEE S&P ’18) was runner up for the 2019 Caspar Bowden Award for Outstanding Research in Privacy Enhancing Technologies.
- The paper “Investigating ad transparency mechanisms in social media: A case study of Facebook’s explanations” by P. Loiseau and co-authors (NDSS ’18) was runner up for the 2019 CNIL-Inria Award for Privacy Protection.

# 6. New Software and Platforms

## 6.1. Framesoc

**FUNCTIONAL DESCRIPTION:** Framesoc is the core software infrastructure of the SoC-Trace project. It provides a graphical user environment for execution-trace analysis, featuring interactive analysis views as Gantt charts or statistics views. It provides also a software library to store generic trace data, play with them, and build other analysis tools (e.g., Ocelotl).

- Participants: Arnaud Legrand and Jean-Marc Vincent
- Contact: Guillaume Huard
- URL: <http://soctrace-inria.github.io/framesoc/>

## 6.2. GameSeer

**FUNCTIONAL DESCRIPTION:** GameSeer is a tool for students and researchers in game theory that uses Mathematica to generate phase portraits for normal form games under a variety of (user-customizable) evolutionary dynamics. The whole point behind GameSeer is to provide a dynamic graphical interface that allows the user to employ Mathematica’s vast numerical capabilities from a simple and intuitive front-end. So, even if you’ve never used Mathematica before, you should be able to generate fully editable and customizable portraits quickly and painlessly.

- Contact: Panayotis Mertikopoulos
- URL: <http://mescal.imag.fr/membres/panayotis.mertikopoulos/>

### 6.3. marmoteCore

*Markov Modeling Tools and Environments - the Core*

KEYWORDS: Modeling - Stochastic models - Markov model

FUNCTIONAL DESCRIPTION: marmoteCore is a C++ environment for modeling with Markov chains. It consists in a reduced set of high-level abstractions for constructing state spaces, transition structures and Markov chains (discrete-time and continuous-time). It provides the ability of constructing hierarchies of Markov models, from the most general to the particular, and equip each level with specifically optimized solution methods.

This software is developed within the ANR MARMOTE project: ANR-12-MONU-00019.

- Participants: Alain Jean-Marie, Hlib Mykhailenko, Benjamin Briot, Franck Quessette, Issam Rabhi, Jean-Marc Vincent and Jean-Michel Fourneau
- Partner: UVSQ
- Contact: Alain Jean-Marie
- Publications: [marmoteCore: a Markov Modeling Platform](#) - [marmoteCore: a software platform for Markov modeling](#)
- URL: <http://marmotecore.gforge.inria.fr/>

### 6.4. Moca

*Memory Organisation Cartography and Analysis*

KEYWORDS: High-Performance Computing - Performance analysis

- Contact: David Beniamine
- URL: <https://github.com/dbeniamine/MOCA>

### 6.5. Ocelotl

*Multidimensional Overviews for Huge Trace Analysis*

FUNCTIONAL DESCRIPTION: Ocelotl is an innovative visualization tool, which provides overviews for execution trace analysis by using a data aggregation technique. This technique enables to find anomalies in huge traces containing up to several billions of events, while keeping a fast computation time and providing a simple representation that does not overload the user.

- Participants: Arnaud Legrand and Jean-Marc Vincent
- Contact: Jean-Marc Vincent
- URL: <http://soctrace-inria.github.io/ocelotl/>

### 6.6. PSI

*Perfect Simulator*

FUNCTIONAL DESCRIPTION: Perfect simulator is a simulation software of markovian models. It is able to simulate discrete and continuous time models to provide a perfect sampling of the stationary distribution or directly a sampling of functional of this distribution by using coupling from the past. The simulation kernel is based on the CFTP algorithm, and the internal simulation of transitions on the Aliasing method.

- Contact: Jean-Marc Vincent
- URL: <http://psi.gforge.inria.fr/>

### 6.7. SimGrid

KEYWORDS: Large-scale Emulators - Grid Computing - Distributed Applications

**SCIENTIFIC DESCRIPTION:** SimGrid is a toolkit that provides core functionalities for the simulation of distributed applications in heterogeneous distributed environments. The simulation engine uses algorithmic and implementation techniques toward the fast simulation of large systems on a single machine. The models are theoretically grounded and experimentally validated. The results are reproducible, enabling better scientific practices.

Its models of networks, cpus and disks are adapted to (Data)Grids, P2P, Clouds, Clusters and HPC, allowing multi-domain studies. It can be used either to simulate algorithms and prototypes of applications, or to emulate real MPI applications through the virtualization of their communication, or to formally assess algorithms and applications that can run in the framework.

The formal verification module explores all possible message interleavings in the application, searching for states violating the provided properties. We recently added the ability to assess liveness properties over arbitrary and legacy codes, thanks to a system-level introspection tool that provides a finely detailed view of the running application to the model checker. This can for example be leveraged to verify both safety or liveness properties, on arbitrary MPI code written in C/C++/Fortran.

**NEWS OF THE YEAR:** There were 3 major releases in 2019: Python bindings were introduced, SMPI now partially supports some of the MPI/IO functions, a new model for Wifi networks was proposed, and the API for the simulation of storage resources was completely revisited. We also pursued our efforts to improve the documentation of the software, simplified the web site, and made a lot of bug fixing and code refactoring.

- Participants: Adrien Lèbre, Arnaud Legrand, Augustin Degomme, Florence Perronnin, Frédéric Suter, Jean-Marc Vincent, Jonathan Pastor, Luka Stanisic and Martin Quinson
- Partners: CNRS - ENS Rennes
- Contact: Martin Quinson
- URL: <https://simgrid.org/>

## 6.8. Tabarnac

*Tool for Analyzing the Behavior of Applications Running on NUMA Architecture*

**KEYWORDS:** High-Performance Computing - Performance analysis - NUMA

- Contact: David Beniamine
- URL: <https://dbeniamine.github.io/Tabarnac/>

# 7. New Results

## 7.1. Design of Experiments

Performance engineering of scientific HPC applications requires to measure repeatedly the performance of applications or of computation kernels, which consume a large amount of time and resources. It is essential to design experiments so as to reduce this cost as much as possible. Our contribution along this axis is twofold: (1) the investigation sound exploration techniques and (2) the control of experiments to ensure the measurements are as representative as possible of real workload.

Writing, porting, and optimizing scientific applications makes autotuning techniques fundamental to lower the cost of leveraging the improvements on execution time and power consumption provided by the latest software and hardware platforms. Despite the need for economy, most autotuning techniques still require large budgets of costly experimental measurements to provide good results, while rarely providing exploitable knowledge after optimization. In [16], we investigate the use of *Design of Experiments* to propose a user-transparent autotuning technique that operates under tight budget constraints by significantly reducing the measurements needed to find good optimizations. Our approach enables users to make informed decisions on which optimizations to pursue and when to stop. We present an experimental evaluation of our approach and show it is capable of leveraging user decisions to find the best global configuration of a GPU Laplacian kernel using half of the measurement budget used by other common autotuning techniques. We show that our approach is also capable of finding speedups of up to  $50\times$ , compared to gcc's -O3, for some kernels from the SPAPT benchmark suite, using up to  $10\times$  fewer measurements than random sampling. Although the results are very encouraging, our approach relies on assumptions on the geometry of the search space that are difficult to test in very large dimension. We are thus currently pursuing this line of research using non parametric approaches based on gaussian process regression, space filling designs and iteratively selecting configurations that yield the best expected improvement.

Our second contribution is related to the control of measurements. In [40], we relate a surprising observation on the performance of the highly optimized and regular DGEMM function on modern processors. The DGEMM function is a widely used implementation of the matrix product. While the asymptotic complexity of the algorithm only depends on the sizes of the matrices, we show that the performance is significantly impacted by the matrices content. Although it would be expected that special values like 1 or 0 may yield to specific behavior, we show that arbitrary constant values are no different and that random values incur a significant performance drop. Our experiments show that this may be due to bit flips in the CPU causing an energy consumption overhead. Such phenomenon reminds the importance of thoroughly randomizing every single parameter of experiments to avoid bias toward specific behavior.

## 7.2. Predictive Simulation of HPC Applications

Finely tuning MPI applications (number of processes, granularity, collective operation algorithms, topology and process placement) is critical to obtain good performance on supercomputers. With a rising cost of modern supercomputers, running parallel applications at scale solely to optimize their performance is extremely expensive. Using SimGrid, we work toward providing a methodology allowing to provide inexpensive but faithful predictions of expected performance.

The methodology we propose relies on SimGrid/SMPI and captures the complexity of adaptive applications by emulating the MPI code while skipping insignificant parts. In [18] we demonstrate its capability with High Performance Linpack (HPL), the benchmark used to rank supercomputers in the TOP500 and which requires a careful tuning. We explain (1) how we both extended the SimGrid's SMPI simulator and slightly modified the open-source version of HPL to allow a fast emulation on a single commodity server at the scale of a supercomputer and (2) how to model the different components (network, BLAS, ...) of the system. We show that a careful modeling of both spatial and temporal node variability allows us to obtain predictions within a few percents of real experiments. The modeling of BLAS operations is particularly important and we have thus started investigating in the context of simulating a sparse direct solver how to automatically performance models for commonly used BLAS kernels [33]. A key difficulty remains the acquisition of faithful performance measurements as modern processors are often quite unstable. This effort is therefore particularly related to the aforementioned "Design of Experiments" line of research.

## 7.3. Simulation of Smart Grids

In [35], we present ASGriDS, an asynchronous Smart Grid simulation framework. ASGriDS is multi-domain, it simultaneously models the power network along with its physical loads/generators, controllers, and communication infrastructure. ASGriDS provides a unified workflow in a pythonic environment, to describe, run and control complex SmartGrid deployment scenarios. ASGriDS is an event-driven simulator that can run

in either real-time or accelerated real-time. As it is modular and its components interact asynchronously, it can run either locally on a distributed infrastructure, also in hardware-in-the-loop setups, and on top of emulated/physical communication links. In this paper, we present the design of our simulator and we demonstrate its use with a generation control problem on a low voltage network. We use ASGridS to deploy a real-time controller based on optimal power flow, on top of TCP and UDP based communication network, under various packet loss conditions.

## 7.4. Batch Scheduling

Despite the impressive growth and size of super-computers, the computational power they provide still cannot match the demand. Efficient and fair resource allocation is a critical task. Super-computers use Resource and Job Management Systems to schedule applications, which is generally done by relying on generic index policies such as First Come First Served and Shortest Processing time First in combination with Backfilling strategies. Unfortunately, such generic policies often fail to exploit specific characteristics of real workloads.

In [36], we focus on improving the performance of online schedulers by studying mixed policies, which are created by combining multiple job characteristics in a weighted linear expression, as opposed to classical pure policies which use only a single characteristic. This larger class of scheduling policies aims at providing more flexibility and adaptability. We use space coverage and black-box optimization techniques to explore this new space of mixed policies and we study how can they adapt to the changes in the workload. We perform an extensive experimental campaign through which we show that (1) the best pure policy is far from optimal and that (2) using a carefully tuned mixed policy would allow to significantly improve the performance of the system. (3) We also provide empirical evidence that there is no one size fits all policy, by showing that the rapid workload evolution seems to prevent classical online learning algorithms from being effective.

A careful investigation of why such mixed strategy fail to globally exploit weekly workload features reveal that some users sometimes provide widely inaccurate information, which dramatically fools the batch scheduling heuristic. Indeed, users typically provide a loose upper bound estimate for job execution times that are hardly useful. Previous studies attempted to improve these estimates using regression techniques. Although these attempts provide reasonable predictions, they require a long period of training data. Furthermore, aiming for perfect prediction may be of limited use for scheduling purposes. In [50], we propose a simpler approach by classifying jobs as small or large and prioritizing the execution of small jobs over large ones. Indeed, small jobs are the most impacted by queuing delays but they typically represent a light load and incur a small burden on the other jobs. The classifier operates online and learns by using data collected over the previous weeks, facilitating its deployment and enabling fast adaptations to changes in workload characteristics. We evaluate our approach using four scheduling policies on six HPC platform workload traces. We show that: (i) incorporating such classification significantly reduces the average bounded slowdown of jobs in all scenarios, and (ii) the obtained improvements are comparable, in most scenarios, to the ideal hypothetical situation where the scheduler would know the exact running time of jobs in advance.

## 7.5. Load Balancing

In distributed systems, load balancing is a powerful concept to improve the distribution of jobs across multiple computing resources and to control performance metrics such as delays and throughputs while avoiding the overload of any single resource. This section describes three contributions:

- In multi-server distributed queueing systems, the access of stochastically arriving jobs to resources is often regulated by a dispatcher, also known as load balancer. A fundamental problem consists in designing a load balancing algorithm that minimizes the delays experienced by jobs. During the last two decades, the power-of- $d$ -choice algorithm, based on the idea of dispatching each job to the least loaded server out of  $d$  servers randomly sampled at the arrival of the job itself, has emerged as a breakthrough in the foundations of this area due to its versatility and appealing asymptotic properties. In [8], we consider the power-of- $d$ -choice algorithm with the addition of a local memory that keeps track of the latest observations collected over time on the sampled servers. Then, each job is sent to a server with the lowest observation. We show that this algorithm is asymptotically optimal



in the sense that the load balancer can always assign each job to an idle server in the large-system limit. This holds true if and only if the system load  $\lambda$  is less than  $1 - 1/d$ . If this condition is not satisfied, we show that queue lengths are tightly bounded by  $\lceil \frac{-\log(1-\lambda)}{\log(\lambda d+1)} \rceil$ . This is in contrast with the classic version of the power-of- $d$ -choice algorithm, where at the fluid scale a strictly positive proportion of servers containing  $i$  jobs exists for all  $i \geq 0$ , in equilibrium. Our results quantify and highlight the importance of using memory as a means to enhance performance in randomized load balancing.

- When dispatching jobs to parallel servers, or queues, the highly scalable round-robin (RR) scheme reduces the variance of interarrival times at all queues to a great extent but has no impact on the variances of service processes. Contrariwise, size-interval task assignment (SITA) routing has little impact on the variances of interarrival times but makes the service processes as deterministic as possible. In [6], we unify both 'static' approaches to design a scalable load balancing framework able to control the variances of the arrival and service processes jointly. It turns out that the resulting combination significantly improves performance and is able to drive the mean job delay to zero in the large-system limit; it is known that this property is not achieved when both approaches are considered separately. Within realistic parameters, we show that the optimal number of size intervals that partition the support of the job size distribution is small with respect to the system size. This enhances the applicability of the proposed load balancing scheme at a large scale. In fact, we find that adding a little bit of information about job sizes to a dispatcher operating under RR improves performance a lot. Under the optimal scaling of size intervals and assuming highly variable job sizes, numerical simulations indicate that the proposed algorithm is competitive with the (less scalable) join-the-shortest-workload algorithm even when the system size grows large.
- Size-based routing provides robust strategies to improve the performance of computer and communication systems with highly variable workloads because it is able to isolate small jobs from large ones in a static manner. The basic idea is that each server is assigned all jobs whose sizes belong to a distinct and continuous interval. In the literature, dispatching rules of this type are referred to as SITA (Size Interval Task Assignment) policies. Though their evident benefits, the problem of finding a SITA policy that minimizes the overall mean (steady-state) waiting time is known to be intractable. In particular it is not clear when it is preferable to balance or unbalance server loads and, in the latter case, how. In [7], we provide an answer to these questions in the celebrated limiting regime where the system capacity grows linearly with the system demand to infinity. Within this framework, we prove that the minimum mean waiting time achievable by a SITA policy necessarily converges to the mean waiting time achieved by SITA-E, the SITA policy that equalizes server loads, provided that servers are homogeneous. However, within the set of SITA policies we also show that SITA-E can perform arbitrarily bad if servers are heterogeneous. In this case we prove that there exist exactly  $C!$  asymptotically optimal policies, where  $C$  denotes the number of server types, and all of them are linked to the solution of a single strictly convex optimization problem. It turns out that the mean waiting time achieved by any of such asymptotically optimal policies does not depend on how job-size intervals are mapped to servers. Our theoretical results are validated by numerical simulations with respect to realistic parameters and suggest that the above insights are also accurate in small systems composed of a few servers, i.e., ten.

## 7.6. FoG Computing

To this day, the Internet of Things (IoT) continues its explosive growth. Nevertheless, with the exceptional evolution of traffic demand, existing infrastructures are struggling to resist. In this context, Fog computing is shaping the future of IoT applications. It offers nearby computational, networking and storage resources to respond to the stringent requirements of these applications. However, despite its several advantages, Fog computing raises new challenges which slow its adoption down. Hence, there is a lack of practical solutions to enable the exploitation of this novel concept.

In [19], we propose FITOR, an orchestration system for IoT applications in the Fog environment. This solution builds a realistic Fog environment while offering efficient orchestration mechanisms. In order to optimize the provisioning of Fog-Enabled IoT applications, FITOR relies on O-FSP, an optimized fog service provisioning strategy which aims to minimize the provisioning cost of IoT applications, while meeting their requirements. Based on extensive experiments, the results obtained show that O-FSP optimizes the placement of IoT applications and outperforms the related strategies in terms of i) provisioning cost ii) resource usage and iii) acceptance rate. In [46], we propose a novel strategy, which we call GO-FSP and which optimizes the placement of IoT application components while coping with their strict performance requirements. To do so, we first propose an Integer Linear Programming (ILP) formulation for the IoT application provisioning problem. The latter targets to minimize the deployment cost while ensuring a load balancing between heterogeneous devices. Then, a GRASP-based approach is proposed to achieve the aforementioned objectives. Finally, we make use of the FITOR orchestration system to evaluate the performance of our solution under real conditions. Obtained results show that our scheme outperforms the related strategies. We are currently comparing such strategy with other strategies based on online learning mechanisms under various information scenarios (delayed and noisy feedback, inaccurate application load information, etc.).

Last, fog computing also extends the capacities of the cloud to the edge of the network, near the physical world, so that Internet of Things (IoT) applications can benefit from properties such as short delays, real-time and privacy. Unfortunately, devices in the Fog-IoT environment are usually unstable and prone to failures. In this context, the consequences of failures may impact the physical world and can, therefore, be critical. In [28], we present a framework for end-to-end resilience of Fog-IoT applications. The framework was implemented and experimented on a smart home testbed.

## 7.7. Research Management: Research Reproducibility and Credit

We are actively promoting better research practices, in particular in term of research reproducibility and contribution recognition. Our contribution this year is threefold

First, we have participated to the writing of a book introducing reproducible research [39]. For a researcher, there is nothing more frustrating than the failure to reproduce major results obtained a few months back. The causes of such disappointments can be multiple and insidious. This phenomenon plays an important role in the so-called "research reproducibility crisis". This book takes a current perspective onto a number of potentially dangerous situations and practices, to exemplify and highlight the symptoms of non-reproducibility in research. Each time, it provides efficient solutions ranging from good-practices that are easily and immediately implementable to more technical tools, all of which are free and have been put to the test by the authors themselves. Students and engineers and researchers should find efficient and accessible ways leading them to improve their reproducible research practices.

Second, to allow students and engineers and researchers to receive proper training in reproducible research, we have run the second session of the [Mooc "Reproducible research: Methodological principles for a transparent science"](#) on the FUN platform from April, 1 to June, 13 2019. This MOOC allows scientists to learn modern and reliable tools such as Markdown for taking structured notes, Desktop search applications, GitLab for version control and collaborative working, and Computational notebooks (Jupyter, RStudio, and Org-Mode) for efficiently combining the computation, presentation, and analysis of data. More than 2,100 persons registered to this session and we are currently working on a third session which is expected to start in the beginning of the year 2020.

Third, software is a fundamental pillar of modern scientific research, not only in computer science, but actually across all fields and disciplines. However, there is a lack of adequate means to cite and reference software, for many reasons. An obvious first reason is software authorship, which can range from a single developer to a whole team, and can even vary in time. The panorama is even more complex than that, because many roles can be involved in software development: software architect, coder, debugger, tester, team manager, and so on. Arguably, the researchers who have invented the key algorithms underlying the software can also claim a part of the authorship. And there are many other reasons that make this issue complex. We provide in [5] a contribution to the ongoing efforts to develop proper guidelines and recommendations for software

citation, building upon the internal experience of Inria, the French research institute for digital sciences. As a central contribution, we make three key recommendations. (1) We propose a richer taxonomy for software contributions with a qualitative scale. (2) We claim that it is essential to put the human at the heart of the evaluation. And (3) we propose to distinguish citation from reference which is particularly important in the context of reproducible research.

## 7.8. Mean Field Games and Control

In [10], we consider mean field games with discrete state spaces (called discrete mean field games in the following) and we analyze these games in continuous and discrete time, over finite as well as infinite time horizons. We prove the existence of a mean field equilibrium assuming continuity of the cost and of the drift. These conditions are more general than the existing papers studying finite state space mean field games. Besides, we also study the convergence of the equilibria of  $N$ -player games to mean field equilibria in our four settings. On the one hand, we define a class of strategies in which any sequence of equilibria of the finite games converges weakly to a mean field equilibrium when the number of players goes to infinity. On the other hand, we exhibit equilibria outside this class that do not converge to mean field equilibria and for which the value of the game does not converge. In discrete time this non-convergence phenomenon implies that the Folk theorem does not scale to the mean field limit.

In [20], we consider a class of nonlinear systems of differential equations with uncertainties, i.e., with lack of knowledge in some of the parameters that is represented by a time-varying unknown bounded functions. An under-approximation of such systems consists of a subset of its reachable set, for any value of the unknown parameters. By relying on optimal control theory through Pontryagin's principle, we provide an algorithm for the under-approximation of a linear combination of the state variables in terms of a fully automated tool-chain named UTOPIA. This allows to establish tight under-approximations of common benchmarks models with dimensions as large as sixty-five.

## 7.9. Energy and Network Optimization

This section describes four contributions on energy and network optimization.

- One of the key challenges in Internet of Things (IoT) networks is to connect many different types of autonomous devices while reducing their individual power consumption. This problem is exacerbated by two main factors: first, the fact that these devices operate in and give rise to a highly dynamic and unpredictable environment where existing solutions (e.g., water-filling algorithms) are no longer relevant; and second, the lack of sufficient information at the device end. To address these issues, we propose a regret-based formulation that accounts for arbitrary network dynamics: this allows us to derive an online power control scheme that is provably capable of adapting to such changes, while relying solely on strictly causal feedback. In so doing, we identify an important tradeoff between the amount of feedback available at the transmitter side and the resulting system performance: if the device has access to unbiased gradient observations, the algorithm's regret after  $T$  stages is  $O(T^{-1/2})$  (up to logarithmic factors); on the other hand, if the device only has access to scalar, utility-based information, this decay rate drops to  $O(T^{-1/4})$ . The above is validated by an extensive suite of numerical simulations in realistic channel conditions, which clearly exhibit the gains of the proposed online approach over traditional water-filling methods. This contribution appeared in [11].
- Many businesses possess a small infrastructure that they can use for their computing tasks, but also often buy extra computing resources from clouds. Cloud vendors such as Amazon EC2 offer two types of purchase options: on-demand and spot instances. As tenants have limited budgets to satisfy their computing needs, it is crucial for them to determine how to purchase different options and utilize them (in addition to possible self-owned instances) in a cost-effective manner while respecting their response-time targets. In this paper, we propose a framework to design policies to allocate self-owned, on-demand and spot instances to arriving jobs. In particular, we propose a near-optimal policy to determine the number of self-owned instances and an optimal policy to determine the

number of on-demand instances to buy and the number of spot instances to bid for at each time unit. Our policies rely on a small number of parameters and we use an online learning technique to infer their optimal values. Through numerical simulations, we show the effectiveness of our proposed policies, in particular that they achieve a cost reduction of up to 64.51% when spot and on-demand instances are considered and of up to 43.74% when self-owned instances are considered, compared to previously proposed or intuitive policies. This contribution appeared in [13].

- In [22], we consider the classical problem of minimizing offline the total energy consumption required to execute a set of  $n$  real-time jobs on a single processor with varying speed. Each real-time job is defined by its release time, size, and deadline (all integers). The goal is to find a sequence of processor speeds, chosen among a finite set of available speeds, such that no job misses its deadline and the energy consumption is minimal. Such a sequence is called an optimal speed schedule. We propose a linear time algorithm that checks the schedulability of the given set of  $n$  jobs and computes an optimal speed schedule. The time complexity of our algorithm is in  $O(n)$ , to be compared with  $O(n \log(n))$  for the best known solutions. Besides the complexity gain, the main interest of our algorithm is that it is based on a completely different idea: instead of computing the critical intervals, it sweeps the set of jobs and uses a dynamic programming approach to compute an optimal speed schedule. Our linear time algorithm is still valid (with some changes) with an arbitrary power function (not necessarily convex) and arbitrary switching times
- Network utility maximization (NUM) is an iconic problem in network traffic management which is at the core of many current and emerging network design paradigms - and, in particular, software-defined networks (SDNs). Thus, given the exponential growth of modern-day networks (in both size and complexity), it is crucial to develop scalable algorithmic tools that are capable of providing efficient solutions in time which is dimension-free, i.e., independent-or nearly-independent-on the size of the system. To do so, we leverage a suite of modified gradient methods known as “mirror descent” and we derive a scalable and efficient algorithm for the NUM problem based on gradient exponentiation. We show that the convergence speed of the proposed algorithm only carries a logarithmic dependence on the size of the network, so it can be implemented reliably and efficiently in massively large networks where traditional gradient methods are prohibitively slow. These theoretical results are sub-sequently validated by extensive numerical simulations showing an improvement of several order of magnitudes over standard gradient methods in large-scale networks. This contribution appeared in [31].
- In the DNS resolution process, packet losses and ensuing retransmission timeouts induce marked latencies: the current UDP-based resolution process takes up to 5 seconds to detect a loss event. In [24], [24], we find that persistent DNS connections based on TCP or TLS can provide an elegant solution to this problem. With controlled experiments on a testbed, we show that persistent DNS connections significantly reduces worst-case latency. We then leverage a large-scale platform to study the performance impact of TCP/TLS on recursive resolvers. We find that off-the-shelf software and reasonably powerful hardware can effectively provide recursive DNS service over TCP and TLS, with a manageable performance hit compared to UDP.

## 7.10. Privacy, Fairness, and Transparency in Online Social Medias

This section describes four contributions on privacy, fairness and transparency in online social medias

- The Facebook advertising platform has been subject to a number of controversies in the past years regarding privacy violations, lack of transparency, as well as its capacity to be used by dishonest actors for discrimination or propaganda. In this study, we aim to provide a better understanding of the Facebook advertising ecosystem, focusing on how it is being used by advertisers. We first analyze the set of advertisers and then investigate how those advertisers are targeting users and customizing ads via the platform. Our analysis is based on the data we collected from over 600 real-world users via a browser extension that collects the ads our users receive when they browse their Facebook timeline, as well as the explanations for why users received these ads. Our results reveal that users are targeted

by a wide range of advertisers (e.g., from popular to niche advertisers); that a non-negligible fraction of advertisers are part of potentially sensitive categories such as news and politics, health or religion; that a significant number of advertisers employ targeting strategies that could be either invasive or opaque; and that many advertisers use a variety of targeting parameters and ad texts. Overall, our work emphasizes the need for better mechanisms to audit ads and advertisers in social media and provides an overview of the platform usage that can help move towards such mechanisms.

This contribution appeared in [14].

- To help their users to discover important items at a particular time, major websites like Twitter, Yelp, TripAdvisor or NYTimes provide Top-K recommendations (e.g., 10 Trending Topics, Top 5 Hotels in Paris or 10 Most Viewed News Stories), which rely on crowd-sourced popularity signals to select the items. However, different sections of a crowd may have different preferences, and there is a large silent majority who do not explicitly express their opinion. Also, the crowd often consists of actors like bots, spammers, or people running orchestrated campaigns. Recommendation algorithms today largely do not consider such nuances, hence are vulnerable to strategic manipulation by small but hyper-active user groups. To fairly aggregate the preferences of all users while recommending top-K items, we borrow ideas from prior research on social choice theory, and identify a voting mechanism called Single Transferable Vote (STV) as having many of the fairness properties we desire in top-K item (s)elections. We develop an innovative mechanism to attribute preferences of silent majority which also make STV completely operational. We show the generalizability of our approach by implementing it on two different real-world datasets. Through extensive experimentation and comparison with state-of-the-art techniques, we show that our proposed approach provides maximum user satisfaction, and cuts down drastically on items disliked by most but hyper-actively promoted by a few users.

This contribution appeared in [17].

- The rise of algorithmic decision making led to active researches on how to define and guarantee fairness, mostly focusing on one-shot decision making. In several important applications such as hiring, however, decisions are made in multiple stage with additional information at each stage. In such cases, fairness issues remain poorly understood. In this paper we study fairness in k-stage selection problems where additional features are observed at every stage. We first introduce two fairness notions, local (per stage) and global (final stage) fairness, that extend the classical fairness notions to the k-stage setting. We propose a simple model based on a probabilistic formulation and show that the locally and globally fair selections that maximize precision can be computed via a linear program. We then define the price of local fairness to measure the loss of precision induced by local constraints; and investigate theoretically and empirically this quantity. In particular, our experiments show that the price of local fairness is generally smaller when the sensitive attribute is observed at the first stage; but globally fair selections are more locally fair when the sensitive attribute is observed at the second stage—hence in both cases it is often possible to have a selection that has a small price of local fairness and is close to locally fair.

This contribution appeared in [21].

- Most social platforms offer mechanisms allowing users to delete their posts, and a significant fraction of users exercise this right to be forgotten. However, ironically, users' attempt to reduce attention to sensitive posts via deletion, in practice, attracts unwanted attention from stalkers specifically to those (deleted) posts. Thus, deletions may leave users more vulnerable to attacks on their privacy in general. Users hoping to make their posts forgotten face a "damned if I do, damned if I don't" dilemma. Many are shifting towards ephemeral social platform like Snapchat, which will deprive us of important user-data archival. In the form of intermittent withdrawals, we present, Lethe, a novel solution to this problem of (really) forgetting the forgotten. If the next-generation social platforms are willing to give up the uninterrupted availability of non-deleted posts by a very small fraction, Lethe provides privacy to the deleted posts over long durations. In presence of Lethe, an adversarial observer becomes unsure if some posts are permanently deleted or just temporarily withdrawn by

Lethe; at the same time, the adversarial observer is overwhelmed by a large number of falsely flagged un-deleted posts. To demonstrate the feasibility and performance of Lethe, we analyze large-scale real data about users' deletion over Twitter and thoroughly investigate how to choose time duration distributions for alternating between temporary withdrawals and resurrections of non-deleted posts. We find a favorable trade-off between privacy, availability and adversarial overhead in different settings for users exercising their right to delete. We show that, even against an ultimate adversary with an uninterrupted access to the entire platform, Lethe offers deletion privacy for up to 3 months from the time of deletion, while maintaining content availability as high as 95% and keeping the adversarial precision to 20%.

This contribution appeared in [27],

## 7.11. Optimization Methods

This section describes six contributions on optimization.

- In [9], we propose an interior-point method for linearly constrained – and possibly nonconvex – optimization problems. The proposed method – which we call the Hessian barrier algorithm (HBA) – combines a forward Euler discretization of Hessian Riemannian gradient flows with an Armijo backtracking step-size policy. In this way, HBA can be seen as an alternative to mirror descent (MD), and contains as special cases the affine scaling algorithm, regularized Newton processes, and several other iterative solution methods. Our main result is that, modulo a non-degeneracy condition, the algorithm converges to the problem's critical set; hence, in the convex case, the algorithm converges globally to the problem's minimum set. In the case of linearly constrained quadratic programs (not necessarily convex), we also show that the method's convergence rate is  $O(1/k^\rho)$  for some  $\rho \in (0, 1]$  that depends only on the choice of kernel function (i.e., not on the problem's primitives). These theoretical results are validated by numerical experiments in standard non-convex test functions and large-scale traffic assignment problems.
- In [15], Lipschitz continuity is a central requirement for achieving the optimal  $O(1/T)$  rate of convergence in monotone, deterministic variational inequalities (a setting that includes convex minimization, convex-concave optimization, nonatomic games, and many other problems). However, in many cases of practical interest, the operator defining the variational inequality may exhibit singularities at the boundary of the feasible region, precluding in this way the use of fast gradient methods that attain this optimal rate (such as Nemirovski's mirror-prox algorithm and its variants). To address this issue, we propose a novel regularity condition which we call Bregman continuity, and which relates the variation of the operator to that of a suitably chosen Bregman function. Leveraging this condition, we derive an adaptive mirror-prox algorithm which attains the optimal  $O(1/T)$  rate of convergence in problems with possibly singular operators, without any prior knowledge of the degree of smoothness (the Bregman analogue of the Lipschitz constant). We also show that, under Bregman continuity, the mirror-prox algorithm achieves a  $O(1/\sqrt{T})$  convergence rate in stochastic variational inequalities.
- In [23] Variational inequalities have recently attracted considerable interest in machine learning as a flexible paradigm for models that go beyond ordinary loss function minimization (such as generative adversarial networks and related deep learning systems). In this setting, the optimal  $O(1/t)$  convergence rate for solving smooth monotone variational inequalities is achieved by the Extra-Gradient (EG) algorithm and its variants. Aiming to alleviate the cost of an extra gradient step per iteration (which can become quite substantial in deep learning applications), several algorithms have been proposed as surrogates to Extra-Gradient with a *single* oracle call per iteration. In this paper, we develop a synthetic view of such algorithms, and we complement the existing literature by showing that they retain a  $O(1/t)$  ergodic convergence rate in smooth, deterministic problems. Subsequently, beyond the monotone deterministic case, we also show that the last iterate of single-call, *stochastic* extra-gradient methods still enjoys a  $O(1/t)$  local convergence rate to solutions of non-monotone variational inequalities that satisfy a second-order sufficient condition.

- In [25], we study a class of online convex optimization problems with long-term budget constraints that arise naturally as reliability guarantees or total consumption constraints. In this general setting, prior work by Mannor et al. (2009) has shown that achieving no regret is impossible if the functions defining the agent's budget are chosen by an adversary. To overcome this obstacle, we refine the agent's regret metric by introducing the notion of a " $K$ -benchmark", i.e., a comparator which meets the problem's allotted budget over any window of length  $K$ . The impossibility analysis of Mannor et al. (2009) is recovered when  $K = T$ ; however, for  $K = o(T)$ , we show that it is possible to minimize regret while still meeting the problem's long-term budget constraints. We achieve this via an online learning policy based on Cautious Online Lagrangian Descent (COLD) for which we derive explicit bounds, in terms of both the incurred regret and the residual budget violations.
- In [26], owing to their connection with generative adversarial networks (GANs), saddle-point problems have recently attracted considerable interest in machine learning and beyond. By necessity, most theoretical guarantees revolve around convex-concave (or even linear) problems; however, making theoretical inroads towards efficient GAN training depends crucially on moving beyond this classic framework. To make piecemeal progress along these lines, we analyze the behavior of mirror descent (MD) in a class of non-monotone problems whose solutions coincide with those of a naturally associated variational inequality - a property which we call coherence. We first show that ordinary, "vanilla" MD converges under a strict version of this condition, but not otherwise; in particular, it may fail to converge even in bilinear models with a unique solution. We then show that this deficiency is mitigated by optimism: by taking an "extra-gradient" step, optimistic mirror descent (OMD) converges in all coherent problems. Our analysis generalizes and extends the results of Daskalakis et al. (2018) for optimistic gradient descent (OGD) in bilinear problems, and makes concrete headway for establishing convergence beyond convex-concave games. We also provide stochastic analogues of these results, and we validate our analysis by numerical experiments in a wide array of GAN models (including Gaussian mixture models, as well as the CelebA and CIFAR-10 datasets).
- In [30], we develop a new stochastic algorithm with variance reduction for solving pseudo-monotone stochastic variational inequalities. Our method builds on Tseng's forward-backward-forward algorithm, which is known in the deterministic literature to be a valuable alternative to Korpelevich's extragradient method when solving variational inequalities over a convex and closed set governed with pseudo-monotone and Lipschitz continuous operators. The main computational advantage of Tseng's algorithm is that it relies only on a single projection step, and two independent queries of a stochastic oracle. Our algorithm incorporates a variance reduction mechanism, and leads to a.s. convergence to solutions of a merely pseudo-monotone stochastic variational inequality problem. To the best of our knowledge, this is the first stochastic algorithm achieving this by using only a single projection at each iteration.

## 7.12. Learning

This section describes three contributions on machine learning.

- In [12], we examine the convergence of no-regret learning in games with continuous action sets. For concreteness, we focus on learning via "dual averaging", a widely used class of no-regret learning schemes where players take small steps along their individual payoff gradients and then "mirror" the output back to their action sets. In terms of feedback, we assume that players can only estimate their payoff gradients up to a zero-mean error with bounded variance. To study the convergence of the induced sequence of play, we introduce the notion of variational stability, and we show that stable equilibria are locally attracting with high probability whereas globally stable equilibria are globally attracting with probability 1. We also discuss some applications to mixed-strategy learning in finite games, and we provide explicit estimates of the method's convergence speed.
- Resource allocation games such as the famous Colonel Blotto (CB) and Hide-and-Seek (HS) games are often used to model a large variety of practical problems, but only in their one-shot versions.

Indeed, due to their extremely large strategy space, it remains an open question how one can efficiently learn in these games. In this work, we show that the online CB and HS games can be cast as path planning problems with side-observations (SOPPP): at each stage, a learner chooses a path on a directed acyclic graph and suffers the sum of losses that are adversarially assigned to the corresponding edges; and she then receives semi-bandit feedback with side-observations (i.e., she observes the losses on the chosen edges plus some others). We propose a novel algorithm, EXP3-OE, the first-of-its-kind with guaranteed efficient running time for SOPPP without requiring any auxiliary oracle. We provide an expected-regret bound of EXP3-OE in SOPPP matching the order of the best benchmark in the literature. Moreover, we introduce additional assumptions on the observability model under which we can further improve the regret bounds of EXP3-OE. We illustrate the benefit of using EXP3-OE in SOPPP by applying it to the online CB and HS games.

This contribution appeared in [29], [49]. In an earlier article [38], we also looked at the sequential Colonel Blotto game under bandit feedback and we proposed a blackbox optimization based method to optimize the exploration distribution of the classical COMBAND algorithm.

- In [32], we study nonzero-sum hypothesis testing games that arise in the context of adversarial classification, in both the Bayesian as well as the Neyman-Pearson frameworks. We first show that these games admit mixed strategy Nash equilibria, and then we examine some interesting concentration phenomena of these equilibria. Our main results are on the exponential rates of convergence of classification errors at equilibrium, which are analogous to the well-known Chernoff-Stein lemma and Chernoff information that describe the error exponents in the classical binary hypothesis testing problem, but with parameters derived from the adversarial model. The results are validated through numerical experiments.

## 8. Bilateral Contracts and Grants with Industry

### 8.1. Bilateral Grants with Industry

Nicolas Gast obtained funding Enedis for a study on the PLC-G3 protocol ( $\approx 50k$  euros).

## 9. Partnerships and Cooperations

### 9.1. Regional Initiatives

#### 9.1.1. IDEX UGA

- Nicolas Gast received a grant from the IDEX UGA that funds a two-years post-doctoral researcher (Takai Kennouche) for two years (2018 and 2019) to work on the smart-grid project that focus on distributed optimization in electrical distribution networks.
- Patrick Loiseau and Panayotis Mertikopoulos received a grant from the IDEX UGA that partly funds a PhD student (Benjamin Roussillon) to work on game theoretic models for adversarial classification.

### 9.2. National Initiatives

#### 9.2.1. Inria Project Labs

Arnaud Legrand is the leader of the HAC SPECIS project. The goal of the HAC SPECIS (High-performance Application and Computers: Studying PErformance and Correctness In Simulation) project 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. Inria Teams: AVALON, POLARIS, MYRIADS, SUMO, HIEPACS, STORM, MEXICO, VERIDIS.



### 9.2.2. Grenoble INP grant

Patrick Loiseau and Bary Pradelski received a grant from the Presidency of Grenoble INP that covers half of the funding of PhD student Dimitrios Moustakas to work on dynamic matching. This PhD is done in collaboration with Univ. Zurich (Heinrich Nax), which covers the rest.

### 9.2.3. DGA Grants

Patrick Loiseau and Panayotis Mertikopoulos received a grant from DGA that complements the funding of PhD student (Benjamin Roussillon) to work on game theoretic models for adversarial classification.

### 9.2.4. PGMO Projects

PGMO projects are supported by the Jacques Hadamard Mathematical Foundation (FMJH). Our project (HEAVY.NET) is focused on congested networks and their asymptotic properties.

### 9.2.5. PEPS

Panayotis Mertikopoulos is co-PI of a PEPS I3A project: MixedGAN ("Mixed-strategy generative adversarial networks") (PI: R. Laraki, U. Dauphine).

### 9.2.6. Fondation Blaise Pascal

Project IAM (Informatique à la Main) funded by fondation Blaise Pascal (Jean-Marc Vincent).

### 9.2.7. MIAI @ Grenoble Alpes

MIAI @ Grenoble Alpes (Multidisciplinary Institute in Artificial Intelligence) is the 3IA institute of Grenoble that was selected by the government in 2019. With the MIAI institute, Patrick Loiseau is the co-holder of a chair on "Explainable and Responsible AI" of which Nicolas Gast and Bary Pradelski are also members; and Panayotis Mertikopoulos is a member of the "Optimization and Learning" chair.

### 9.2.8. ANR

- Nicolas Gast obtained funding from the ANR JCJC for the project REFINO. 250k euros. Duration: 4 years
- Bary Pradelski (PI), P. Mertikopoulos and P. Loiseau obtained funding from the ANR for the project ALIAS (Adaptive Learning for Interactive Agents and Systems). This is a bilateral PRCI (collaboration internationale) project joint with Singapore University of Technology and Design (SUTD). The Singapore team consists of G. Piliouras and G. Panageas.
- *ORACLESS (2016–2021)*  
ORACLESS is an ANR starting grant (JCJC) coordinated by Panayotis Mertikopoulos. The goal of the project is to develop highly adaptive resource allocation methods for wireless communication networks that are provably capable of adapting to unpredictable changes in the network. In particular, the project will focus on the application of online optimization and online learning methodologies to multi-antenna systems and cognitive radio networks.
- *CONNECTED (2016–2019)*  
CONNECTED is an ANR Tremplin-ERC (T-ERC) grant coordinated by Patrick Loiseau. The goal of the project is to work on several game-theoretic models involving learning agents and data revealed by strategic agents in response to the learning algorithms, so as to derive better learning algorithms for such special data.

## 9.3. International Initiatives

### 9.3.1. Inria International Labs

#### 9.3.1.1. ReDaS

Title: Reproducible Data Science

International Partner (Institution - Laboratory - Researcher):

Universidade Federal do Rio Grande do Sul (Brazil) - Industrial Engineering and Operations Research Departments - Lucas Mello Schnorr

Start year: 2019

See also: <https://associatedteam.gitlabpages.inria.fr/redas>

Data science builds on a variety of technique and tools that makes analysis often difficult to follow and reproduce. The goal of this project is to develop interactive, reproducible and scalable analysis workflows that provide uncertainty and quality estimators about the analysis.

### 9.3.1.2. International Initiatives

#### **GENE**

Title: Stochastic dynamics of large games and networks

International Partners (Institution - Laboratory - Researcher):

Universidad de Buenos Aires (Argentina) - Matthieu Jonckheere

Universidad de la Republica Uruguay (Uruguay) - Federico La Rocca

CNRS (France) - Balakrishna Prabhu

Universidad ORT Uruguay (Uruguay) - Andrés Ferragut

Duration: 2018 - 2019

Start year: 2018

Through the creation and consolidation of strong research and formation exchanges between Argentina, France and Uruguay, the GENE project will contribute to the fields of performance evaluation and control of communication networks, using tools of game theory, probability theory and control theory. Some of the challenges this project will address are: - Mean-field games and their application to load balancing and resource allocations, - Scaling limits for centralized and decentralized load balancing strategies and implementation of practical policies for web servers farms, - Information diffusion and communication protocols in large and distributed wireless networks.

## 9.4. International Research Visitors

### 9.4.1. Visits of International Scientists

We have hosted multiple international scientists for short (typically one-week) visits: Jonathan Newton, Paul Duetting, Jason Marden, Bruno Ziliotto

### 9.4.2. Visits to International Teams

- V. Danjean spent one week at Porto Alegre (Brasil) at UFRGS, hosted by Lucas M. Schnorr to work on the research subject: Tracing of multi-tasked OpenMP Application.
- A. Legrand spent 10 days at Porto Alegre (Brasil) at UFRGS, hosted by Lucas M. Schnorr to teach scientific methodology and Performance Evaluation and to work on the visual performance analysis of dynamic task-based applications.
- G. Huard visited UFRGS (Porto Alegre, Brasil) in the context of the ReDaS Inria associated team from Nov. 27th to Dec 16th along with Alexis Janon. During this visit we worked with Lucas Schnorr on several application trace analysis cases using our own custom analysis framework and leveraging UFRGS expertise on the design and conduct of practical data analysis.
- B. Pradelski was invited for seminars at several places: IHP Game Theory Seminar, Bar-Ilan University Economic Theory seminar, University of Oxford Game Theory seminar. He is also an associate member of the Oxford Man Institute.

#### 9.4.2.1. Research Stays Abroad

P. Mertikopoulos was invited to spend a three-month research visit at the Ecole Polytechnique Fédérale de Lausanne (EPFL). He was hosted by the LIONS lab (headed by V. Cevher).

## 10. Dissemination

### 10.1. Promoting Scientific Activities

#### 10.1.1. Scientific Events: Organisation

##### 10.1.1.1. General Chair, Scientific Chair

- P. Mertikopoulos was a technical program co-chair of the 10th International Conference on NETWORK Games, CONTROL and OPTimisation (NetGCoop 2020).
- B. Pradelski was a technical program co-chair of the 14th Workshop on the Economics of Networks, Systems and Computation (NetEcon), colocated with ACM SIGMETRICS and EC.

##### 10.1.1.2. Member of the Organizing Committees

- B. Gaujal organized a special day on Potential games at Gamenets ( Paris)
- P. Mertikopoulos co-organized the workshop “20PoA: Twenty years of the Price of Anarchy” (Chania, Greece, July 2019).

#### 10.1.2. Scientific Events: Selection

##### 10.1.2.1. Member of the Conference Program Committees

- B. Gaujal: WiOpt, NeuIPS
- J. Anselmi: Valuetools, ASMTA
- P. Mertikopoulos: The 2020 French Days on Optimization and Decision Science (SMAI MODE 2020)
- A. Legrand: ISC HIGH PERFORMANCE
- N. Gast: SIGMETRICS, NeurIPS, ValueTools
- J-M. Vincent: Epew, Valuetools, Simultech
- P. Loiseau: NeurIPS, ICML, SIGMETRICS, PETS, NetEcon
- P. Mertikopoulos: ICML, NeurIPS (area chair).

#### 10.1.3. Journal

##### 10.1.3.1. Member of the Editorial Boards

- P. Mertikopoulos serves as an associate editor for JDG (Journal of Dynamics and Games) and MCAP (Methodology and Computing in Applied Probability).
- N. Gast serves as an associate editor for Performance Evaluation and Stochastic Models.
- P. Loiseau is an associate editor for ACM Transactions on Internet Technology and IEEE Transactions on Big Data.

##### 10.1.3.2. Reviewer - Reviewing Activities

All members of the team are active reviewers for several international journals and conferences.

#### 10.1.4. Invited Talks

- B. Gaujal:
  - 23/01: ENS Lyon seminars (Le Pleyret) “Evolutionary games and bounded rationality”
  - 7/02: Eva Tardos seminar (Grenoble): “price of anarchy in routing games”
  - 3/05: Workshop for Eitan Altman’s 60th Birthday (Avignon) “Sturmian words at work in optimal routing”
- P. Mertikopoulos:
  - Invited instructor at the CONNECT Summer School on Machine Learning for Communications “Online learning and optimization for wireless systems”, Trinity College, Dublin
  - Invited talk at ICCOPT 2019 (2019 International Conference on Continuous Optimization), Berlin, August 2019
  - Invited talk at NPCG 2019 (Workshop on Network, Population and Congestion Games), Paris, April 2019
  - Invited talk at GDO 2019 (Workshop on Games, Dynamics and Optimization), Cluj-Napoca, April 2019
  - Invited talk at OSL 2019 (Workshop on Optimization and Statistical Learning), Les Houches, March 2019
  - Invited talk at EPFL Machine Learning Seminar, March 2019
  - Invited talk at the Criteo AI Lab, February 2019
- A. Legrand:
  - Simulation of HPC applications and predictions, Scheduling workshop, Bordeaux (27/6/19)
  - Series of talks about reproducible research: TILECS workshop, Grenoble (3/7/19); UFRGS keynote, Porto Alegre (9/10/19); SBAC-PAD conference, Campo Grande (17/10/19); Formidex, UGA (6/11/19); Doctoral school, Neuchatel (7/11/19); Inria Alumni, Paris (12/11/19)
- Nicolas Gast was invited to give a tutorial about “Mean field and refined mean field approximation” at the conference ITC.

#### 10.1.5. Research Administration

- B. Gaujal is a member of the scientific committee of GDR-IM and a member of the council of ‘pole MSTIC’ Grenoble
- P. Mertikopoulos is a member of the steering committee (comité de liaison) of the optimization and decision theory group of the French Society for Industrial and Applied Mathematics (SMAI)
- P. Mertikopoulos is the working group coordinator, core group member and management committee (MC) representative for France in the European Network for Game Theory (GAMENET).
- P. Loiseau is the chair of the steering committee of NetEcon.

## 10.2. Teaching - Supervision - Juries

### 10.2.1. Teaching

We only list the master level teaching.

- B. Gaujal was involved in multiple courses:
  - M2 course in ENS Lyon with Panayotis Mertikopoulos : Online Optimization
  - M2 course in MPRI (Paris) with Ana Busic: performance evaluation in communication networks

- M2 course (Ensimag) on network performance models
- M1 exercise session (Ensimag) applied probability
- P. Mertikopoulos gave an invited PhD level course at EPFL on “Min-max optimization and variational inequalities”.
- V. Danjean was involved in INFO3 and INFO4 at Polytech Grenoble (System Architecture, Internship supervising, ...) and in M1 Info (Operating systems and Parallel Programming course, Operating System project)
- A. Legrand was involved in multiple courses:
  - Scientific Methodology and Performance Evaluation (M2 MOSIG, UGA)
  - Scientific Methodology and Performance Evaluation (M2 Univ. Federale do Rio Grande do Sul, Porto Alegre)
  - Parallel Systems (M2 MOSIG, UGA)
  - Probability and Simulation (M1, Polytech/UGA)
  - Performance Evaluation (M1, Polytech/UGA)
  - Reproducible Research (Doctoral School MSTII, UGA)
- J. Anselmi taught in the course Probability and Simulation (M1, Polytech/UGA).
- P. Loiseau taught in the courses Probability and Simulation (M1, Polytech/UGA) and “Algorithms for data processing” (M1 INFO, UGA).
- N. Gast is responsible of the master course “Optimization under Uncertainties” (Master 2 ORCO in Grenoble).
- J.-M. Vincent teaches Probability for Informatics and Performance Evaluation at Ensimag, and Mathematics for Computer Science (1st year) and Scientific Methodology and Performance Evaluation (2nd year) at the Master of Computer Science.
- G. Huard taught the course Object Oriented Design class for the M1 INFO, UGA.

### 10.2.2. Supervision

Supervision of PhD students and postdocs:

- B. Jonglez (Bruno Gaujal and Martin Heusse)
- S. Plassart (Bruno Gaujal and Alain Girault)
- K. Khun (Bruno Gaujal and Nicolas Gast)
- C. Yan (Bruno Gaujal and Nicolas Gast)
- K. Antonakopoulos (P. Mertikopoulos and E. V. Belmega, ETIS/ENSEA)
- B. Roussillon (P. Mertikopoulos and P. Loiseau)
- B. Donassolo (P. Mertikopoulos and A. Legrand):
- P. Rocha Bruel (A. Legrand and Alfredo Goldman)
- T. Cornebize (A. Legrand)
- C. Heinrich (A. Legrand)
- S. Zrigui (A. Legrand and D. Trystram)
- A. Janon (G. Huard and A. Legrand)
- V. Emelianov (N. Gast and P. Loiseau)
- T. Barzolla (N. Gast with Vincent Jost and Van-Dat Cung from G-SCOP laboratory)
- M. Mendil (N. Gast)
- T. Kennouche (N. Gast)
- U. Ozeer (J.-M. Vincent)
- Dong Quan Vu (P. Loiseau)
- Vera Sosnovik (O. Goga and P. Loiseau)
- Eleni Gkiouzepi (P. Loiseau)
- Lucas Leandro Nesi (A. Legrand and Lucas Mello Schnorr)
- Dimitrios Moustakas (B. Pradelski and P. Loiseau, with H. Nax from UZH)
- Simon Jantscheg (B. Pradelski and P. Loiseau, with H. Nax from UZH)

Supervision of M2 Students:

- Manal Benaissa (V. Danjean)
- Leo Gayral (Bruno Gaujal and Federica Garin)
- Kimang Khun (Bruno Gaujal and Nicolas Gast)
- Nicolas Rocher (Patrick Loiseau and Panayotis Mertikopoulos)
- Chen Yan (Nicolas Gast)
- Dimitrios Moustakas (B. Pradelski)

### 10.2.3. *Juries*

- B. Gaujal was a reviewer of the PhD Thesis of Paulin Jacquot (Ecole Polytechnique).
- V. Danjean was involved in several teaching juries: INFO3 at Polytech Grenoble, L3 M&I, M1 Info, DU ISN and DIU EIL at UGA.
- A. Legrand was a reviewer of the PhD Thesis of Mohamad El Sayah (Univ. Franche Comté, Besançon)
- N. Gast was member of the PhD Jury of Celine Comte and Eyal Castiel.

## 10.3. Popularization

Patrick Loiseau wrote with Oana Goga an article “Publicité en Ligne : reprenons la main !”, that was co-published co-published by the blog Binaire (Le Monde) and The Conversation France, June 3, 2019.

### 10.3.1. *Internal or external Inria responsibilities*

- B. Gaujal is a member of the CR2 hiring committee in Grenoble.
- J.-M. Vincent is in charge of the relation Rectorat / Inria-Grenoble for the organization of scientific events (Festival of Science, Schools Visits, organization of Conference Cycles on research in CS and Applied Mathematics for teachers in Colleges)
- J.-M. Vincent is
  - Member of the national coordination of the Diplôme Inter-Universitaire “Enseigner l’Informatique au Lycée” (50 universities involved).
  - Local Head of DIU EIL in Academy of Grenoble
  - Member of the organization of the teaching sessions for all the teachers in CS coming from abroad
  - Member of the national Commission Inter-Irem in Informatics
  - Member of the first national jury for the competitive recruitment of teachers in computer science (Capes NSI 2019-20)

### 10.3.2. *Articles and contents*

Arnaud Legrand participated in the writing of a book [39] on Reproducible Research, which aims at helping students and engineers and researchers to find efficient and accessible ways leading them to improve their reproducible research practices.

### 10.3.3. Education

- V. Danjean is the head of the DU ISN formation (Diplôme Universitaire Informatique et Sciences du Numérique)
- V. Danjean co-organized the new DIU EIL formation (Diplôme Inter-Universitaire Enseigner l'Informatique au Lycée). He is involved both at the national level (for the coordination and the definition of the content of this formation provided in more than 30 universities in France), and at the local level (coordination of the local teams, courses scheduling, conference organization, ...)

### 10.3.4. Interventions

- V. Danjean participated in “La Fête de la Science”, animating several sessions of “unplugged computer science”
- P. Loiseau co-organized and animated a workshop “IA, éthique et société”, Forum Ecobiz Grenoble, October 2019.
- P. Loiseau participated in a debate “Éthique et numérique : quels enjeux sociétaux ?”. Festival Transfo, Grenoble, France, January 2019.

## 11. Bibliography

### Publications of the year

#### Doctoral Dissertations and Habilitation Theses

- [1] F. HEINRICH. *Modeling, Prediction and Optimization of Energy Consumption of MPI Applications using SimGrid*, Université Grenoble Alpes, May 2019, <https://tel.archives-ouvertes.fr/tel-02269894>
- [2] A. MARCASTEL. *Optimisation en ligne et apprentissage adaptatif pour les réseaux dans les bandes ISM*, Université de Cergy Pontoise, February 2019
- [3] P. MERTIKOPOULOS. *Online optimization and learning in games: Theory and applications*, Grenoble 1 UGA - Université Grenoble Alpes, December 2019, Habilitation à diriger des recherches, <https://hal.inria.fr/tel-02428077>
- [4] U. OZEER. *Autonomic Resilience of Distributed IoT Applications in the Fog*, UGA - Université Grenoble Alpes ; MSTII, December 2019

#### Articles in International Peer-Reviewed Journals

- [5] P. ALLIEZ, R. DI COSMO, B. GUEDJ, A. GIRAULT, M.-S. HACID, A. LEGRAND, N. P. ROUGIER. *Attributing and Referencing (Research) Software: Best Practices and Outlook from Inria*, in "Computing in Science & Engineering", 2019, pp. 1-14, <https://arxiv.org/abs/1905.11123> [DOI : 10.1109/MCSE.2019.2949413], <https://hal.archives-ouvertes.fr/hal-02135891>
- [6] J. ANSELMI. *Combining Size-Based Load Balancing with Round-Robin for Scalable Low Latency*, in "IEEE Transactions on Parallel and Distributed Systems", 2019, pp. 1-3, forthcoming [DOI : 10.1109/TPDS.2019.2950621], <https://hal.archives-ouvertes.fr/hal-02276789>
- [7] J. ANSELMI, J. DONCEL. *Asymptotically Optimal Size-Interval Task Assignments*, in "IEEE Transactions on Parallel and Distributed Systems", 2019, vol. 30, n<sup>o</sup> 11, pp. 2422-2433 [DOI : 10.1109/TPDS.2019.2920121], <https://hal.archives-ouvertes.fr/hal-02318576>

- [8] J. ANSELMINI, F. DUFOUR. *Power-of-d-Choices with Memory: Fluid Limit and Optimality*, in "Mathematics of Operations Research", 2019, pp. 1-31, forthcoming, <https://hal.archives-ouvertes.fr/hal-02394147>
- [9] I. M. BOMZE, P. MERTIKOPOULOS, W. SCHACHINGER, M. STAUDIGL. *Hessian barrier algorithms for linearly constrained optimization problems*, in "SIAM Journal on Optimization", 2019, vol. 29, pp. 2100 - 2127 [DOI : 10.1137/18M1215682], <https://hal.inria.fr/hal-02403531>
- [10] J. DONCEL, N. GAST, B. GAUJAL. *Discrete Mean Field Games: Existence of Equilibria and Convergence*, in "Journal of Dynamics and Games", 2019, vol. 6, n<sup>o</sup> 3, pp. 1-19, <https://arxiv.org/abs/1909.01209> [DOI : 10.3934/JDG.2019016], <https://hal.inria.fr/hal-01277098>
- [11] A. MARCASTEL, E.-V. BELMEGA, P. MERTIKOPOULOS, I. FIJALKOW. *Online Power Optimization in Feedback-Limited, Dynamic and Unpredictable IoT Networks*, in "IEEE Transactions on Signal Processing", 2019, vol. 67, n<sup>o</sup> 11, pp. 2987 - 3000, forthcoming [DOI : 10.1109/TSP.2019.2910479], <https://hal.archives-ouvertes.fr/hal-02189523>
- [12] P. MERTIKOPOULOS, Z. ZHOU. *Learning in games with continuous action spaces and unknown payoff functions*, in "Mathematical Programming, Series A", 2019, vol. 173, n<sup>o</sup> 1-2, pp. 465-507, <https://arxiv.org/abs/1608.07310> [DOI : 10.1007/s10107-018-1254-8], <https://hal.archives-ouvertes.fr/hal-01382282>
- [13] X. WU, P. LOISEAU, E. HYYTIÄ. *Towards Designing Cost-Optimal Policies to Utilize IaaS Clouds with Online Learning*, in "IEEE Transactions on Parallel and Distributed Systems", 2019, vol. 14, forthcoming [DOI : 10.1109/TPDS.2019.2935199], <https://hal.inria.fr/hal-02303480>

### International Conferences with Proceedings

- [14] A. ANDREOU, M. SILVA, F. BENEVENUTO, O. GOGA, P. LOISEAU, A. MISLOVE. *Measuring the Facebook Advertising Ecosystem*, in "NDSS 2019 - Proceedings of the Network and Distributed System Security Symposium", San Diego, United States, February 2019, pp. 1-15 [DOI : 10.14722/NDSS.2019.23280], <https://hal.archives-ouvertes.fr/hal-01959145>
- [15] K. ANTONAKOPOULOS, E.-V. BELMEGA, P. MERTIKOPOULOS. *An adaptive mirror-prox algorithm for variational inequalities with singular operators*, in "NeurIPS 2019", Vancouver, Canada, 2019, <https://hal.inria.fr/hal-02403562>
- [16] P. BRUEL, S. QUINTO MASNADA, B. VIDEAU, A. LEGRAND, J.-M. VINCENT, A. GOLDMAN. *Auto-tuning under Tight Budget Constraints: A Transparent Design of Experiments Approach*, in "CCGrid 2019 - International Symposium in Cluster, Cloud, and Grid Computing", Larcana, Cyprus, May 2019, pp. 1-10 [DOI : 10.1109/CCGRID.2019.00026], <https://hal.inria.fr/hal-02110868>
- [17] A. CHAKRABORTY, G. K. PATRO, N. GANGULY, K. P. GUMMADI, P. LOISEAU. *Equality of Voice: Towards Fair Representation in Crowdsourced Top-K Recommendations*, in "FAT\* 2019 - ACM Conference on Fairness, Accountability, and Transparency", Atlanta, United States, Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAT\*), ACM, January 2019, pp. 129-138 [DOI : 10.1145/3287560.3287570], <https://hal.archives-ouvertes.fr/hal-01959135>
- [18] T. CORNEBIZE, A. LEGRAND, F. C. HEINRICH. *Fast and Faithful Performance Prediction of MPI Applications: the HPL Case Study*, in "2019 IEEE International Conference on Cluster Computing (CLUSTER)", Al-



- buquerque, United States, 2019 IEEE International Conference on Cluster Computing (CLUSTER), September 2019 [DOI : 10.1109/CLUSTER.2019.8891011], <https://hal.inria.fr/hal-02096571>
- [19] B. DONASSOLO, I. FAJJARI, A. LEGRAND, P. MERTIKOPOULOS. *Fog Based Framework for IoT Service Provisioning*, in "CCNC 2019 - IEEE Consumer Communications & Networking Conference", Las Vegas, United States, IEEE, January 2019, pp. 1-6 [DOI : 10.1109/CCNC.2019.8651835], <https://hal.inria.fr/hal-01859695>
- [20] J. DONCEL, N. GAST, M. TRIBASTONE, M. TSCHAIKOWSKI, A. VANDIN. *UTOPIC: Under-Approximation Through Optimal Control*, in "QEST 2019 - 16th International Conference on Quantitative Evaluation of SysTems", Glasgow, United Kingdom, Springer, September 2019, pp. 277-291 [DOI : 10.1007/978-3-030-30281-8\_16], <https://hal.inria.fr/hal-02283189>
- [21] V. EMELIANOV, G. ARVANITAKIS, N. GAST, K. P. GUMMADI, P. LOISEAU. *The Price of Local Fairness in Multistage Selection*, in "IJCAI-2019 - Twenty-Eighth International Joint Conference on Artificial Intelligence", Macao, France, International Joint Conferences on Artificial Intelligence Organization, May 2019, pp. 5836-5842, <https://arxiv.org/abs/1906.06613> [DOI : 10.24963/IJCAI.2019/809], <https://hal.inria.fr/hal-02145071>
- [22] B. GAUJAL, A. GIRAULT, S. PLASSART. *A Linear Time Algorithm for Computing Off-line Speed Schedules Minimizing Energy Consumption*, in "MSR 2019 - 12ème Colloque sur la Modélisation des Systèmes Réactifs", Angers, France, November 2019, pp. 1-14, <https://hal.archives-ouvertes.fr/hal-02372136>
- [23] Y.-G. HSIEH, F. IUTZELER, J. MALICK, P. MERTIKOPOULOS. *On the convergence of single-call stochastic extra-gradient methods*, in "NeurIPS 2019", Vancouver, Canada, 2019, <https://arxiv.org/abs/1908.08465> - 27 pages, 3 figures, <https://hal.inria.fr/hal-02403555>
- [24] B. JONGLEZ, S. BIRBALTA, M. HEUSSE. *Persistent DNS connections for improved performance*, in "NETWORKING 2019 - IFIP Networking 2019", Warsaw, Poland, May 2019, pp. 1-2, <https://hal.inria.fr/hal-02149978>
- [25] N. LIAKOPOULOS, A. S. DESTOUNIS, G. PASCHOS, T. SPYROPOULOS, P. MERTIKOPOULOS. *Cautious regret minimization: Online optimization with long-term budget constraints*, in "ICML 2019 - 36th International Conference on Machine Learning", Long Beach, United States, June 2019, pp. 1-9, <https://hal.inria.fr/hal-02405753>
- [26] P. MERTIKOPOULOS, B. LECOAT, H. ZENATI, C.-S. FOO, V. CHANDRASEKHAR, G. PILIOURAS. *Optimistic Mirror Descent in Saddle-Point Problems: Going the Extra (Gradient) Mile*, in "ICLR 2019 - 7th International Conference on Learning Representations", New Orleans, United States, May 2019, pp. 1-23, <https://hal.inria.fr/hal-02111937>
- [27] M. MINAEI, M. MONDAL, P. LOISEAU, K. P. GUMMADI, A. KATE. *Forgetting the Forgotten with Lethé: Conceal Content Deletion from Persistent Observers*, in "PETS 2019 - 19th Privacy Enhancing Technologies Symposium", Stockholm, Sweden, July 2019, pp. 1-21, <https://hal.archives-ouvertes.fr/hal-01959119>
- [28] U. OZEER, L. LETONDEUR, F.-G. OTTOGALLI, G. SALAÜN, J.-M. VINCENT. *Designing and Implementing Resilient IoT Applications in the Fog: A Smart Home Use Case*, in "ICIN 2019 - 22nd Conference on Innovation in Clouds, Internet and Networks", Paris, France, IEEE, February 2019, pp. 230-232 [DOI : 10.1109/ICIN.2019.8685909], <https://hal.archives-ouvertes.fr/hal-01979686>

- [29] D. QUAN VU, P. LOISEAU, A. SILVA, L. TRAN-THANH. *Path Planning Problems with Side Observations—When Colonels Play Hide-and-Seek*, in "AAAI 2020 - Thirty-Fourth AAAI Conference on Artificial Intelligence", New-York, United States, February 2020, pp. 1-15, <https://hal.inria.fr/hal-02375789>
- [30] M. STAUDIGL, P. MERTIKOPOULOS. *Convergent Noisy forward-backward-forward algorithms in non-monotone variational inequalities*, in "LSS 2019 - 15th IFAC Symposium on Large Scale Complex Systems: Theory and Applications", Delft, Pays-Bas, May 2019, pp. 120-125 [DOI : 10.1016/j.ifacol.2019.06.021], <https://hal.inria.fr/hal-02405750>
- [31] L. VIGNERI, G. PASCHOS, P. MERTIKOPOULOS. *Large-Scale Network Utility Maximization: Countering Exponential Growth with Exponentiated Gradients*, in "INFOCOM 2019 - IEEE International Conference on Computer Communications", Paris, France, IEEE, April 2019, pp. 1630-1638 [DOI : 10.1109/INFOCOM.2019.8737600], <https://hal.inria.fr/hal-02405759>
- [32] S. YASODHARAN, P. LOISEAU. *Nonzero-sum Adversarial Hypothesis Testing Games*, in "NeurIPS 2019 - Thirty-third Conference on Neural Information Processing Systems", Vancouver, Canada, 2019, pp. 1-23, <https://hal.inria.fr/hal-02299451>

### Conferences without Proceedings

- [33] E. AGULLO, A. BUTTARI, A. GUERMOUCHE, A. LEGRAND, I. MASLIAH, L. STANISIC. *Simulation of a Sparse Direct Solver on Heterogeneous Systems using Starpu and Simgrid*, in "CSE 2019 - SIAM Conference on Computational Science and Engineering", Spokane, United States, SIAM, February 2019, <https://hal.inria.fr/hal-02073725>
- [34] J. ASSUNÇÃO, J.-M. VINCENT, P. FERNANDES. *Piecewise Aggregation for HMM fitting. A pre-fitting model for seamless integration with time series data*, in "SEKE 2019 - 31st International Conference on Software Engineering and Knowledge Engineering", Lisbon, Portugal, July 2019, pp. 729-734 [DOI : 10.18293/SEKE2019-185], <https://hal.archives-ouvertes.fr/hal-02409589>
- [35] T.-E. KENNOUCHE, F. CADOUX, N. GAST, B. VINOT. *ASGridS: Asynchronous Smart-Grids Distributed Simulator*, in "APPEEC 2019 - 11th IEEE PES Asia-Pacific Power and Energy Engineering Conference", Macao, Macau SAR China, IEEE, December 2019, pp. 1-5, <https://hal.archives-ouvertes.fr/hal-02384051>
- [36] A. LEGRAND, D. TRYSTRAM, S. ZRIGUI. *Adapting Batch Scheduling to Workload Characteristics: What can we expect From Online Learning?*, in "IPDPS 2019 - 33rd IEEE International Parallel & Distributed Processing Symposium", Rio de Janeiro, Brazil, IEEE, May 2019, pp. 686-695 [DOI : 10.1109/IPDPS.2019.00077], <https://hal.archives-ouvertes.fr/hal-02044903>
- [37] A. MARCASTEL, E.-V. BELMEGA, P. MERTIKOPOULOS, I. FIJALKOW. *Gradient-free Online Resource Allocation Algorithms for Dynamic Wireless Networks*, in "SPAWC 2019 - 20th IEEE International Workshop on Signal Processing Advances in Wireless Communications", Cannes, France, IEEE, July 2019, pp. 1-4 [DOI : 10.1109/SPAWC.2019.8815409], <https://hal.archives-ouvertes.fr/hal-02189108>
- [38] D. QUAN VU, P. LOISEAU, A. SILVA. *Combinatorial Bandits for Sequential Learning in Colonel Blotto Games*, in "CDC 2019 - 58th IEEE Conference on Decision and Control", Nice, France, December 2019, <https://arxiv.org/abs/1909.04912> , <https://hal.archives-ouvertes.fr/hal-02283535>

### Scientific Books (or Scientific Book chapters)

- [39] L. DESQUILBET, S. GRANGER, B. HEJBLUM, A. LEGRAND, P. PERNOT, N. P. ROUGIER, E. DE CASTRO GUERRA, M. COURBIN-COULAUD, L. DUVAUX, P. GRAVIER, G. LE CAMPION, S. ROUX, F. SANTOS. *Towards reproducible research : Evolve your practices*, Unité régionale de formation à l'information scientifique et technique de Bordeaux, May 2019, pp. 1-161, <https://hal.archives-ouvertes.fr/hal-02144142>

### Research Reports

- [40] T. CORNEBIZE, A. LEGRAND. *DGEMM performance is data-dependent*, Université Grenoble Alpes ; Inria ; CNRS, December 2019, n<sup>o</sup> RR-9310, <https://arxiv.org/abs/1912.05381> , <https://hal.inria.fr/hal-02401760>
- [41] B. GAUJAL, A. GIRAULT, S. PLASSART. *A Discrete Time Markov Decision Process for Energy Minimization Under Deadline Constraints*, Grenoble Alpes ; Inria Grenoble Rhône-Alpes, Université de Grenoble, December 2019, n<sup>o</sup> RR-9309, 46 p. , <https://hal.inria.fr/hal-02391948>
- [42] B. GAUJAL, A. GIRAULT, S. PLASSART. *Exploiting Job Variability to Minimize Energy Consumption under Real-Time Constraints*, Inria Grenoble Rhône-Alpes, Université de Grenoble ; Université Grenoble - Alpes, November 2019, n<sup>o</sup> RR-9300, 23 p. , <https://hal.inria.fr/hal-02371742>
- [43] B. GAUJAL, A. GIRAULT, S. PLASSART. *Feasibility of on-line speed policies in real-time systems*, Inria Grenoble Rhône-Alpes, Université de Grenoble ; Univ. Grenoble Alpes, November 2019, n<sup>o</sup> RR-9301, 38 p. , <https://hal.inria.fr/hal-02371996>

### Software

- [44] S. ARCHIPOFF, C. AUGONNET, O. AUMAGE, G. BEAUCHAMP, B. BRAMAS, A. BUTTARI, A. CASSAGNE, J. CLET-ORTEGA, T. COJEAN, N. COLLIN, V. DANJEAN, A. DENIS, L. EYRAUD-DUBOIS, N. FURMENTO, S. HENRY, A. HUGO, M. JUHOOR, A. JUVEN, M. KERYELL-EVEN, Y. KHORSI, T. LAMBERT, E. LERIA, B. LIZÉ, M. MAKNI, S. NAKOV, R. NAMYST, L. NESI LUCAS, J. PABLO, D. PASQUALINOTTO, S. PITOISET, N. QUÔC-DINH, C. ROELANDT, C. SAKKA, C. SALINGUE, L. MELLO SCHNORR, M. SERGENT, A. SIMONET, L. STANISIC, S. BÉRANGÈRE, F. TESSIER, S. THIBAUT, V. BRICE, L. VILLEVEY-GOUX, P.-A. WACRENIER. *StarPU*, January 2020, Version : 1.3.3  
[SWH-ID : swh:1:dir:b6e19d99449a78805e7a55a341fbaba2bc431973], Software, <https://hal.inria.fr/hal-02443512>

### Other Publications

- [45] R. I. BOT, P. MERTIKOPOULOS, M. STAUDIGL, P. T. VUONG. *Forward-backward-forward methods with variance reduction for stochastic variational inequalities*, December 2019, working paper or preprint, <https://hal.inria.fr/hal-02405776>
- [46] B. DONASSOLO, I. FAJJARI, A. LEGRAND, P. MERTIKOPOULOS. *Load Aware Provisioning of IoT Services on Fog Computing Platform*, March 2019, working paper or preprint, <https://hal.inria.fr/hal-02056743>
- [47] B. JONGLEZ, S. BIRBALTA, M. HEUSSE. *Persistent DNS connections for improved performance*, May 2019, 1 p. , NETWORKING 2019 - IFIP Networking 2019, Poster, <https://hal.inria.fr/hal-02149975>
- [48] D. QUAN VU, P. LOISEAU, A. SILVA. *Approximate Equilibria in Non-constant-sum Colonel Blotto and Lottery Blotto Games with Large Numbers of Battlefields*, October 2019, <https://arxiv.org/abs/1910.06559> - working paper or preprint, <https://hal.archives-ouvertes.fr/hal-02315698>

- [49] D. QUAN VU, P. LOISEAU, A. SILVA, L. TRAN-THANH. *Colonel Blotto and Hide-and-Seek Games as Path Planning Problems with Side Observations*, May 2019, <https://arxiv.org/abs/1905.11151> - working paper or preprint, <https://hal.archives-ouvertes.fr/hal-02139519>
- [50] S. ZRIGUI, R. Y. DE CAMARGO, D. TRYSTRAM, A. LEGRAND. *Improving the Performance of Batch Schedulers Using Online Job Size Classification*, October 2019, working paper or preprint, <https://hal.archives-ouvertes.fr/hal-02334116>

## References in notes

- [51] R. M. BADIA, J. LABARTA, J. GIMÉNEZ, F. ESCALÉ. *Dimemas: Predicting MPI Applications Behaviour in Grid Environments*, in "Proc. of the Workshop on Grid Applications and Programming Tools", June 2003
- [52] C. BAIER, B. HAVERKORT, H. HERMANN, J.-P. KATOEN. *Model-checking algorithms for continuous-time Markov chains*, in "Software Engineering, IEEE Transactions on", 2003, vol. 29, n<sup>o</sup> 6, [http://ieeexplore.ieee.org/xpls/abs\\_all.jsp?arnumber=1205180](http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1205180)
- [53] A. BASU, S. FLEMING, J. STANIER, S. NAICKEN, I. WAKEMAN, V. K. GURBANI. *The State of Peer-to-peer Network Simulators*, in "ACM Computing Survey.", August 2013, vol. 45, n<sup>o</sup> 4
- [54] D. BECKER, F. WOLF, W. FRINGS, M. GEIMER, B. WYLIE, B. MOHR. *Automatic Trace-Based Performance Analysis of Metacomputing Applications*, in "Parallel and Distributed Processing Symposium, 2007. IPDPS 2007. IEEE International", March 2007, <http://dx.doi.org/10.1109/IPDPS.2007.370238>
- [55] P. BEDARIDE, A. DEGOMME, S. GENAUD, A. LEGRAND, G. MARKOMANOLIS, M. QUINSON, M. L. STILLWELL, F. SUTER, 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, <https://hal.inria.fr/hal-00919507>
- [56] G. BIANCHI. *Performance analysis of the IEEE 802.11 distributed coordination function*, in "Selected Areas in Communications, IEEE Journal on", 2000, vol. 18, n<sup>o</sup> 3
- [57] L. BOBELIN, A. LEGRAND, M. A. G. DAVID, P. NAVARRO, M. QUINSON, F. SUTER, 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", Ottawa, Canada, May 2012, 19 p. , <https://hal.inria.fr/hal-00650233>
- [58] L. BORTOLUSSI, J. HILLSTON. *Model checking single agent behaviours by fluid approximation*, in "Information and Computation", 2015, vol. 242, <http://dx.doi.org/10.1016/j.ic.2015.03.002>
- [59] L. BORTOLUSSI, R. LANCIANI. *Model Checking Markov Population Models by Central Limit Approximation*, in "Quantitative Evaluation of Systems", Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2013, n<sup>o</sup> 8054
- [60] L. BORTOLUSSI, R. LANCIANI. *Fluid Model Checking of Timed Properties*, in "Formal Modeling and Analysis of Timed Systems", Springer International Publishing, 2015

- [61] H. BRUNST, D. HACKENBERG, G. JUCKELAND, H. ROHLING. *Comprehensive Performance Tracking with Vampir 7*, in "Tools for High Performance Computing 2009", M. S. MÜLLER, M. M. RESCH, A. SCHULZ, W. E. NAGEL (editors), Springer Berlin Heidelberg, 2010, [http://dx.doi.org/10.1007/978-3-642-11261-4\\_2](http://dx.doi.org/10.1007/978-3-642-11261-4_2)
- [62] A. BUSIC, B. GAUJAL, G. GORGO, J.-M. VINCENT. *PSI2 : Envelope Perfect Sampling of Non Monotone Systems*, in "QEST 2010 - International Conference on Quantitative Evaluation of Systems", Williamsburg, VA, United States, IEEE, September 2010, pp. 83-84, <https://hal.inria.fr/hal-00788884>
- [63] A. BUSIC, B. GAUJAL, F. PERRONNIN. *Perfect Sampling of Networks with Finite and Infinite Capacity Queues*, in "19th International Conference on Analytical and Stochastic Modelling Techniques and Applications (ASMTA) 2012", Grenoble, France, K. AL-BEGAIN, D. FIEMS, J.-M. VINCENT (editors), Lecture Notes in Computer Science, Springer, 2012, vol. 7314, pp. 136-149 [DOI : 10.1007/978-3-642-30782-9\_10], <https://hal.inria.fr/hal-00788003>
- [64] S. BÖHM, C. ENGELMANN. *xSim: The Extreme-Scale Simulator*, in "Proceedings of the International Conference on High Performance Computing and Simulation (HPCS) 2011", Istanbul, Turkey, IEEE Computer Society, Los Alamitos, CA, USA, July 2011
- [65] H. CASANOVA, A. GIERSCH, A. LEGRAND, M. QUINSON, F. SUTER. *Versatile, Scalable, and Accurate Simulation of Distributed Applications and Platforms*, in "Journal of Parallel and Distributed Computing", June 2014, vol. 74, n<sup>o</sup> 10, pp. 2899-2917 [DOI : 10.1016/J.JPDC.2014.06.008], <https://hal.inria.fr/hal-01017319>
- [66] A. CHAINTREAU, J.-Y. LE BOUDEC, N. RISTANOVIC. *The Age of Gossip: Spatial Mean Field Regime*, in "SIGMETRICS Perform. Eval. Rev.", June 2009, vol. 37, n<sup>o</sup> 1, <http://doi.acm.org/10.1145/2492101.1555363>
- [67] K. COULOMB, M. FAVERGE, J. JAZEIX, O. LAGRASSE, J. MARCOUEILLE, P. NOISSETTE, A. REDONDY, C. VUCHENER. *Visual trace explorer (ViTE)*, October, 2009
- [68] J. DONCEL, N. GAST, B. GAUJAL. *Mean-Field Games with Explicit Interactions*, February 2016, <https://hal.inria.fr/hal-01277098>
- [69] S. DURAND, B. GAUJAL, F. PERRONNIN, J.-M. VINCENT. *A perfect sampling algorithm of random walks with forbidden arcs*, in "QEST 2014 - 11th International Conference on Quantitative Evaluation of Systems", Florence, Italy, Springer, September 2014, vol. 8657, pp. 178-193 [DOI : 10.1007/978-3-319-10696-0\_15], <https://hal.inria.fr/hal-01069975>
- [70] C. FRICKER, N. GAST. *Incentives and redistribution in homogeneous bike-sharing systems with stations of finite capacity*, in "EURO Journal on Transportation and Logistics", June 2014, 31 p. [DOI : 10.1007/s13676-014-0053-5], <https://hal.inria.fr/hal-01086009>
- [71] C. FRICKER, N. GAST, H. MOHAMED. *Mean field analysis for inhomogeneous bike sharing systems*, in "AofA", Montreal, Canada, July 2012, <https://hal.inria.fr/hal-01086055>
- [72] D. FUDENBERG, D. K. LEVINE. *The Theory of Learning in Games*, Economic learning and social evolution, MIT Press, Cambridge, MA, 1998, vol. 2
- [73] R. M. FUJIMOTO. *Parallel Discrete Event Simulation*, in "Commun. ACM", October 1990, vol. 33, n<sup>o</sup> 10, <http://doi.acm.org/10.1145/84537.84545>

- [74] N. GAST, B. GAUJAL. *Markov chains with discontinuous drifts have differential inclusion limits*, in "Performance Evaluation", 2012, vol. 69, n<sup>o</sup> 12, pp. 623-642 [DOI : 10.1016/J.PEVA.2012.07.003], <https://hal.inria.fr/hal-00787999>
- [75] N. GAST, B. GAUJAL, J.-Y. LE BOUDEC. *Mean field for Markov Decision Processes: from Discrete to Continuous Optimization*, in "IEEE Transactions on Automatic Control", 2012, vol. 57, n<sup>o</sup> 9, pp. 2266–2280 [DOI : 10.1109/TAC.2012.2186176], <https://hal.inria.fr/hal-00787996>
- [76] N. GAST, J.-Y. LE BOUDEC, D.-C. TOMOZEI. *Impact of Demand-Response on the Efficiency and Prices in Real-Time Electricity Markets*, in "ACM e-Energy 2014", Cambridge, United Kingdom, June 2014 [DOI : 10.1145/2602044.2602052], <https://hal.inria.fr/hal-01086036>
- [77] N. GAST, B. VAN HOUTD. *Transient and Steady-state Regime of a Family of List-based Cache Replacement Algorithms*, in "ACM SIGMETRICS 2015", Portland, United States, June 2015 [DOI : 10.1145/2745844.2745850], <https://hal.inria.fr/hal-01143838>
- [78] J. GONZALEZ, J. GIMENEZ, J. LABARTA. *Automatic detection of parallel applications computation phases*, in "Parallel and Distributed Processing Symposium, International", 2009, vol. 0, <http://doi.ieeecomputersociety.org/10.1109/IPDPS.2009.5161027>
- [79] M. HEATH, J. ETHERIDGE. *Visualizing the performance of parallel programs*, in "IEEE software", 1991, vol. 8, n<sup>o</sup> 5
- [80] T. HOEFLER, T. SCHNEIDER, A. LUMSDAINE. *LogGOPSim - Simulating Large-Scale Applications in the LogGOPS Model*, in "Proc. of the ACM Workshop on Large-Scale System and Application Performance", June 2010
- [81] L. HU, J.-Y. LE BOUDEC, M. VOJNOVIĆ. *Optimal channel choice for collaborative ad-hoc dissemination*, in "INFOCOM, 2010 Proceedings IEEE", IEEE, 2010
- [82] L. V. KALÉ, G. ZHENG, C. W. LEE, S. KUMAR. *Scaling applications to massively parallel machines using Projections performance analysis tool*, in "Future Generation Comp. Syst.", 2006, vol. 22, n<sup>o</sup> 3
- [83] T. G. KURTZ. *Approximation of population processes*, SIAM, 1981, vol. 36
- [84] Y.-B. LIN, E. D. LAZOWSKA. *A Time-division Algorithm for Parallel Simulation*, in "ACM Trans. Model. Comput. Simul.", January 1991, vol. 1, n<sup>o</sup> 1, <http://doi.acm.org/10.1145/102810.214307>
- [85] G. LLORT, J. GONZÁLEZ, H. SERVAT, J. GIMÉNEZ, J. LABARTA. *On-line Detection of Large-scale Parallel Application's Structure*, in "24th IEEE International Parallel and Distributed Processing Symposium (IPDPS'2010)", 2010
- [86] L. MELLO SCHNORR, A. LEGRAND. *Visualizing More Performance Data Than What Fits on Your Screen*, in "Tools for High Performance Computing 2012", A. CHEPTSOV, S. BRINKMANN, J. GRACIA, M. M. RESCH, W. E. NAGEL (editors), Springer Berlin Heidelberg, 2013, pp. 149-162 [DOI : 10.1007/978-3-642-37349-7\_10], <https://hal.inria.fr/hal-00842761>

- [87] S. MEYN, P. BAROOAH, A. BUSIC, J. EHREN. *Ancillary service to the grid from deferrable loads: the case for intelligent pool pumps in Florida*, in "Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on", IEEE, 2013
- [88] M. MITZENMACHER. *The power of two choices in randomized load balancing*, in "Parallel and Distributed Systems, IEEE Transactions on", 2001, vol. 12, n<sup>o</sup> 10
- [89] K. MOHROR, K. KARAVANIC, A. SNAVELY. *Scalable Event Trace Visualization*, in "Euro-Par 2009 – Parallel Processing Workshops", H.-X. LIN, M. ALEXANDER, M. FORSELL, A. KNÜPFER, R. PRODAN, L. SOUSA, A. STREIT (editors), Lecture Notes in Computer Science, Springer Berlin / Heidelberg, 2010, vol. 6043, [http://dx.doi.org/10.1007/978-3-642-14122-5\\_27](http://dx.doi.org/10.1007/978-3-642-14122-5_27)
- [90] W. NAGEL, A. ARNOLD, M. WEBER, H. HOPPE, K. SOLCHENBACH. *VAMPIR: Visualization and Analysis of MPI Resources*, in "Supercomputer", 1996, vol. 12, n<sup>o</sup> 1
- [91] V. PILLET, J. LABARTA, T. CORTES, S. GIRONA. *PARAVER: A tool to visualise and analyze parallel code*, in "Proceedings of Transputer and occam Developments, WOTUG-18", Transputer and Occam Engineering, IOS Press, 1995, vol. 44
- [92] J. PROPP, D. WILSON. *Coupling from the past: a user's guide*, in "DIMACS Series on Discrete Mathematics and Theoretical Computer Science", 1998, vol. 41, Microsurveys in discrete probability
- [93] M. L. PUTERMAN. *Markov decision processes: discrete stochastic dynamic programming*, John Wiley & Sons, 2014
- [94] D. REED, P. ROTH, R. AYDT, K. SHIELDS, L. TAVERA, R. NOE, B. SCHWARTZ. *Scalable performance analysis: the Pablo performance analysis environment*, in "Scalable Parallel Libraries Conference, 1993., Proceedings of the", 1993
- [95] W. H. SANDHOLM. *Population Games and Evolutionary Dynamics*, Economic learning and social evolution, MIT Press, Cambridge, MA, 2010
- [96] W. H. SANDHOLM, M. STAUDIGL. *A Sample Path Large Deviation Principle for a Class of Population Processes*, in "arXiv preprint arXiv:1511.07897", 2015
- [97] H. SERVAT, G. LLORT, J. GIMÉNEZ, K. HUCK, J. LABARTA. *Folding: detailed analysis with coarse sampling*, in "Tools for High Performance Computing 2011", Springer Berlin Heidelberg, 2012
- [98] H. SERVAT, G. LLORT, J. GONZALEZ, J. GIMENEZ, J. LABARTA. *Identifying code phases using piece-wise linear regressions*, in "Parallel and Distributed Processing Symposium, 2014 IEEE 28th International", IEEE, 2014
- [99] B. SHNEIDERMAN. *The eyes have it: A task by data type taxonomy for information visualizations*, in "Visual Languages, 1996. Proceedings., IEEE Symposium on", IEEE, 1996
- [100] M. TIKIR, M. LAURENZANO, L. CARRINGTON, A. SNAVELY. *PSINS: An Open Source Event Tracer and Execution Simulator for MPI Applications*, in "Proc. of the 15th International Euro-Par Conference on Parallel Processing", LNCS, Springer, August 2009, n<sup>o</sup> 5704

- 
- [101] B. VAN HOUDT. *A Mean Field Model for a Class of Garbage Collection Algorithms in Flash-based Solid State Drives*, in "Proceedings of the ACM SIGMETRICS", New York, NY, USA, SIGMETRICS '13, ACM, 2013, <http://doi.acm.org/10.1145/2465529.2465543>
- [102] P. VELHO, L. MELLO SCHNORR, H. CASANOVA, A. LEGRAND. *On the Validity of Flow-level TCP Network Models for Grid and Cloud Simulations*, in "ACM Transactions on Modeling and Computer Simulation", October 2013, vol. 23, n<sup>o</sup> 4, <https://hal.inria.fr/hal-00872476>
- [103] J. J. WILKE, K. SARGSYAN, J. P. KENNY, B. DEBUSSCHERE, H. N. NAJM, G. HENDRY. *Validation and Uncertainty Assessment of Extreme-Scale HPC Simulation through Bayesian Inference*, in "Euro-Par 2013 Parallel Processing: 19th International Conference, Aachen, Germany, August 26-30, 2013. Proceedings", Springer Berlin Heidelberg, Berlin, Heidelberg, 2013
- [104] F. WOLF, B. MOHR. *Automatic performance analysis of hybrid MPI/OpenMP applications*, in "Journal of Systems Architecture", 2003, vol. 49, n<sup>o</sup> 10-11
- [105] T. YANG, P. G. MEHTA, S. P. MEYN. *A mean-field control-oriented approach to particle filtering*, in "American Control Conference (ACC), 2011", IEEE, 2011
- [106] L. YING. *On the Rate of Convergence of Mean-Field Models: Stein's Method Meets the Perturbation Theory*, in "arXiv preprint arXiv:1510.00761", 2015
- [107] O. ZAKI, E. LUSK, W. GROPP, D. SWIDER. *Toward Scalable Performance Visualization with Jumpshot*, in "International Journal of High Performance Computing Applications", 1999, vol. 13, n<sup>o</sup> 3, <http://dx.doi.org/10.1177/109434209901300310>
- [108] G. ZHENG, G. KAKULAPATI, L. KALÉ. *BigSim: A Parallel Simulator for Performance Prediction of Extremely Large Parallel Machines*, in "Proc. of the 18th International Parallel and Distributed Processing Symposium (IPDPS)", April 2004
- [109] J. C. DE KERGOMMEAUX, B. STEIN, P. BERNARD. *Paje, an interactive visualization tool for tuning multi-threaded parallel applications*, in "Parallel Computing", 2000, vol. 10, n<sup>o</sup> 26, pp. 1253–1274