

RESEARCH CENTRE

**Inria Center  
at Université Grenoble Alpes**

IN PARTNERSHIP WITH:

Université de Grenoble Alpes, CNRS

2022

ACTIVITY REPORT

Project-Team

POLARIS

**Performance analysis and Optimization of  
LARge Infrastructures and Systems**

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

**DOMAIN**

**Networks, Systems and Services,  
Distributed Computing**

**THEME**

**Distributed and High Performance  
Computing**

*Inria*

# Contents

<b>Project-Team POLARIS</b>	<b>1</b>
<b>1 Team members, visitors, external collaborators</b>	<b>2</b>
<b>2 Overall objectives</b>	<b>3</b>
2.1 Context	3
2.2 Objectives	4
2.3 Contribution to AI/Learning	4
<b>3 Research program</b>	<b>5</b>
3.1 Performance Evaluation	5
3.2 Asymptotic Methods	7
3.3 Distributed Online Optimization and Learning in Games	8
3.4 Responsible Computer Science	10
<b>4 Application domains</b>	<b>12</b>
4.1 Large Computing Infrastructures	12
4.2 Next-Generation Wireless Networks	12
4.3 Energy and Transportation	13
4.4 Social Computing Systems	13
<b>5 Social and environmental responsibility</b>	<b>13</b>
5.1 Footprint of research activities	13
5.2 Raising awareness on the climate crisis	13
5.3 Impact of research results	13
<b>6 Highlights of the year</b>	<b>14</b>
6.1 Awards	14
<b>7 New software and platforms</b>	<b>14</b>
7.1 New software	14
7.1.1 SimGrid	14
7.1.2 PSI	15
7.1.3 marmoteCore	15
7.1.4 MarTO	16
7.1.5 GameSeer	16
7.1.6 rmf_tool	16
<b>8 New results</b>	<b>17</b>
8.1 Performance evaluation of large systems	17
8.2 Energy optimization	18
8.3 Large system performance optimization through learning techniques	18
8.4 Exploiting specific structures in reinforcement learning	19
8.5 Distributed learning and optimization	21
8.6 Learning in games	22
8.7 Advanced learning and optimization methods	23
8.8 Random matrix analysis and Machine Learning	24
8.9 Fairness and equity in digital (recommendation, advertising, persistent storage) systems	25
<b>9 Bilateral contracts and grants with industry</b>	<b>27</b>

<b>10 Partnerships and cooperations</b>	<b>27</b>
10.1 International initiatives	27
10.1.1 Inria associate team not involved in an IIL or an international program	27
10.2 International research visitors	28
10.2.1 Visits of international scientists	28
10.2.2 Visits to international teams	28
10.3 European initiatives	29
10.3.1 Other european programs/initiatives	29
10.4 National initiatives	29
<b>11 Dissemination</b>	<b>30</b>
11.1 Promoting scientific activities	30
11.1.1 Scientific events: organisation	30
11.1.2 Invited talks	31
11.1.3 Scientific expertise	32
11.1.4 Research administration	32
11.2 Teaching - Supervision - Juries	32
11.2.1 Teaching	32
11.2.2 Supervision	33
11.2.3 Juries	33
11.3 Popularization	34
11.3.1 Articles and contents	34
11.3.2 Education	34
11.3.3 Interventions	34
<b>12 Scientific production</b>	<b>35</b>
12.1 Publications of the year	35
12.2 Other	38
12.3 Cited publications	38

# Project-Team POLARIS

*Creation of the Project-Team: 2018 January 01*

## Keywords

### Computer sciences and digital sciences

- A1.2. – Networks
- A1.3.5. – Cloud
- A1.3.6. – Fog, Edge
- 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
- A9.2. – Machine learning
- A9.9. – Distributed AI, Multi-agent

### 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 members, visitors, external collaborators

## Research Scientists

- Arnaud Legrand [Team leader, CNRS, Researcher, until Jun 2022, HDR]
- Jonatha Anselmi [INRIA, Researcher]
- Nicolas Gast [INRIA, Researcher, HDR]
- Bruno Gaujal [INRIA, Senior Researcher, HDR]
- Panayotis Mertikopoulos [CNRS, Researcher, HDR]
- Bary Pradelski [CNRS, Researcher]

## Faculty Members

- Romain Couillet [GRENOBLE INP, Professor, HDR]
- Vincent Danjean [UGA, Associate Professor]
- Guillaume Huard [UGA, Associate Professor]
- Florence Perronnin [UGA, Associate Professor]
- Jean-Marc Vincent [UGA, Associate Professor]
- Philippe Waille [UGA, Associate Professor]

## Post-Doctoral Fellows

- Henry-Joseph Audéoud [INPG-SA Entreprise, until Jun 2022]
- Simon Finster [CNRS, from Nov 2022]
- Sacha Hodencq [FLORALIS, from Oct 2022]
- Dong Quan Vu [CNRS, until Mar 2022]

## PhD Students

- Sebastian Allmeier [INRIA]
- Achille Baucher [FLORALIS, from Oct 2022]
- Victor Boone [ENS DE LYON]
- Rémi Castera [UGA]
- Romain Cravic [INRIA]
- Yu-Guan Hsieh [UGA]
- Simon Jantschgi [UNIV ZURICH, from Aug 2022]
- Simon Jantschgi [CNRS, until Jul 2022]
- Kimang Khun [INRIA]
- Till Kletti [NAVER LABS]
- Lucas Leandro Nesi [GOUV BRESIL]

- Hugo Lebeau [UGA]
- Davide Legacci [UGA, from Nov 2022]
- Victor Leger [UGA]
- Mathieu Molina [INRIA]
- Louis-Sebastien Rebuffi [UGA]
- Chen Yan [UGA]

### Technical Staff

- Achille Baucher [FLORALIS, Engineer, from Apr 2022 until Jun 2022]

### Interns and Apprentices

- Achille Baucher [FLORALIS, until Mar 2022]
- Ran Cheng [UGA, Intern, from May 2022 until Jul 2022]
- Armand Grenier [UGA, Intern, from Jun 2022 until Aug 2022]

### Administrative Assistant

- Annie Simon [INRIA]

## 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 propose 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.

## 2.3 Contribution to AI/Learning

AI and Learning is everywhere now. Let us clarify how our research activities are positioned with respect to this trend.

A first line of research in POLARIS is devoted to the use of statistical learning techniques (Bayesian inference) to model the expected performance of distributed systems, to build aggregated performance views, to feed simulators of such systems, or to detect anomalous behaviours.

In a distributed context it is also essential to design systems that can seamlessly adapt to the workload and to the evolving behaviour of its components (users, resources, network). Obtaining faithful information on the dynamic of the system can be particularly difficult, which is why it is generally more efficient to design systems that dynamically learn the best actions to play through trial and errors. A key characteristic of the work in the POLARIS project is to leverage regularly game-theoretic modeling to handle situations where the resources or the decision is distributed among several agents or even situations where a centralised decision maker has to adapt to strategic users.

An important research direction in POLARIS is thus centered on reinforcement learning (Multi-armed bandits, Q-learning, online learning) and active learning in environments with one or several of the following features:

- Feedback is limited (e.g., gradient or even stochastic gradients are not available, which requires for example to resort to stochastic approximations);
- Multi-agent setting where each agent learns, possibly not in a synchronised way (i.e., decisions may be taken asynchronously, which raises convergence issues);
- Delayed feedback (avoid oscillations and quantify convergence degradation);
- Non stochastic (e.g., adversarial) or non stationary workloads (e.g., in presence of shocks);
- Systems composed of a very large number of entities, that we study through mean field approximation (mean-field games and mean field control).

As a side effect, many of the gained insights can often be used to dramatically improve the scalability and the performance of the implementation of more standard machine or deep learning techniques over supercomputers.

The POLARIS members are thus particularly interested in the design and analysis of adaptive learning algorithms for multi-agent systems, i.e. agents that seek to progressively improve their performance on a specific task. The resulting algorithms should not only learn an efficient (Nash) equilibrium but they should also be able of doing so quickly (low regret), even when facing the difficulties associated to a distributed context (lack of coordination, uncertain world, information delay, limited feedback, ...)

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

## 3 Research program

### 3.1 Performance Evaluation

**Participants:** Jonatha Anselmi, Vincent Danjean, Nicolas Gast, Guillaume Huard, Arnaud Legrand, Florence Perronnin, Jean-Marc Vincent.

#### Project-team positioning

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, etc.). Furthermore, such experiments are costly, difficult to control and therefore difficult to reproduce. Although many of these digital systems have been built by human, 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, in particular (1) Predictive Simulation, (2) Trace Analysis and Visualization, and (3) the Design of Experiments.

#### Scientific achievements

Large computing systems are particularly complex to understand because of the interplay between their discrete nature (originating from deterministic computer programs) and their stochastic nature (emerging from the physical world, long distance interactions, and complex hardware and software stacks). A first line of research in POLARIS is devoted to the design of relatively simple statistical models of key components of distributed systems and their exploitation to feed simulators of such systems, to build aggregated performance views, and to detect anomalous behaviors.

**Predictive Simulation** 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. 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.

Many simulators of MPI applications have been developed by renowned HPC groups (e.g., at SDSC [116], BSC [62], UIUC [124], Sandia Nat. Lab. [122], ORNL [63] or ETH Zürich [92]) but most of them build on restrictive network and application modeling assumptions that generally prevent to faithfully predict execution times, which limits the use of simulation to indication of gross trends at best.

The *SimGrid* simulation toolkit, whose development started more than 20 years ago in UCSD, is a renowned project which gathers more than 1,700 citations and has supported the research of at least 550 articles. The most important contribution of POLARIS to this project in the last years has been to improve the quality of SimGrid to the point where it can be used effectively on a daily basis by practitioners to accurately reproduce the dynamic of real HPC systems. In particular, *SMPI* [70], a simulator based on SimGrid that simulates unmodified MPI applications written in C/C++ or FORTRAN, has now become a very unique tool allowing to faithfully study particularly complex scenario such as legacy Geophysics application that suffers from spatial and temporal load balancing problem [95, 94] or the HPL benchmark [68][3]. We have shown that the performance (both for time and energy consumption [91]) predicted through our simulations was systematically within a few percents of real experiments, which allows to reliably tune the applications at very low cost. This capacity has also been leveraged to study (through *StarPU-SimGrid*) complex and modern task-based applications running on heterogeneous sets of hybrid (CPUs + GPUs) nodes [55]. The phenomenon studied through this approach would be particularly difficult to study through real experiments but yet allow to address real problems of these applications. Finally, SimGrid is also heavily used through *BatSim*, a batch simulator developed in the DATAMOVE team and which leverages SimGrid, to investigate the performance of machine learning strategies in a batch scheduling context [98, 16].

**Trace Analysis and Visualization** Many monolithic visualization tools have been developed by renowned HPC groups since decades (e.g., BSC [109], Jülich and TU Dresden [106, 65], UIUC [90, 112, 93] and ANL [123]) but most of these tools build on the classical information visualization [113] 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. Such traces are typically made of events that occur at very different time and space scales and originate from different sources, which hinders classical approaches, especially when the application structure departs from classical MPI programs with a BSP/SPMD structure. In particular, modern HPC applications that build on a task-based runtime and run on hybrid nodes are particularly challenging to analyze. Indeed, the underlying task-graph is dynamically scheduled to avoid spurious synchronizations, which prevents classical visualizations to exploit and reveal the application structure.

In [77], we explain how modern data analytics tools can be used to build, from heterogeneous information sources, custom, reproducible and insightful visualizations of task-based HPC applications at a very low development cost in the *StarVZ* framework. By specifying and validating statistical models of the performance of HPC applications/systems, we manage to identify when their behavior departs from what is expected and detect performance anomalies. This approach has first been applied to state-of-the-art linear algebra libraries in [77] and more recently to a sparse direct solver [13]. In both cases, we have been able to identify and fix several non-trivial anomalies that had not been noticed even by the application and runtime developers. Finally, these models not only allow to reveal when applications depart from what is expected but also to summarize the execution by focusing on the most important features, which is particularly useful when comparing two executions.

**Design of Experiments and Reproducibility** Part of our work is devoted to the control of experiments on both classical (HPC) and novel (IoT/Fog in a smart home context) infrastructures. To this end, we heavily rely on experimental testbeds such as Grid5000 and FIT-IoT Lab that can be well-controlled but real experiments are nonetheless quite resource-consuming. *Design of experiments* has been successfully applied in many fields (e.g., agriculture, chemistry, industrial processes) where experiments are considered expensive. Building on concrete use cases, we explore how *Design of Experiments* and *Reproducible Research* techniques can be used to (1) design transparent auto-tuning strategies of scientific computation kernels [64, 56] (2) set up systematic performance non regression tests on Grid5000 (450 nodes for 1.5

year) and detect many abnormal events (related to bios and system upgrades, cooling, faulty memory and power instability) that had a significant effect on the nodes, from subtle performance changes of 1% to much more severe degradation of more than 10%, and had yet been unnoticed by both Grid'5000 technical team and Grid'5000 users (3) design and evaluate the performance of service provisioning strategies [4][71] in Fog infrastructures.

## 3.2 Asymptotic Methods

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

### Project-team positioning

Stochastic models often suffer from the curse of dimensionality: their complexity grows exponentially with the number of dimensions of the system. At the same time, very large stochastic systems are sometimes easier to analyze: it can be shown that some classes of stochastic systems simplify as their dimension goes to infinity because of averaging effects such as the law of large numbers, or the central limit theorem. This forms the basis of what is called an *asymptotic method*, which consists in studying what happens when a system gets large in order to build an approximation that is easier to study or to simulate.

Within the team, the research that we conduct in this axis is to foster the applicability of these asymptotic methods to new application areas. This leads us to work on the application of classical methods to new problems, but also to develop new approximation methods that take into account special features of the systems we study (i.e., moderate number of dimensions, transient behavior, random matrices). Typical applications are mean field method for performance evaluation, application to distributed optimization, and more recently statistical learning. One originality of our work is to quantify precisely what is the error made by such approximations. This allows us to define refinement terms that lead to more accurate approximations.

### Scientific achievements

**Refined mean field approximation** Mean field approximation is a well-known technique in statistical physics, that was originally introduced to study systems composed of a very large number of particles (say  $n > 10^{20}$ ). The idea of this approximation is to assume that objects are independent and only interact between them through an average environment (the *mean* field). Nowadays, variants of this technique are widely applied in many domains: in game theory for instance (with the example of mean field games), but also to quantify the performance of distributed algorithms. Mean field approximation is often justified by showing that a system of  $n$  well-mixed interacting objects converges to its deterministic mean field approximation as  $n$  goes to infinity. Yet, this does not explain why mean field approximation provides a very accurate approximation of the behavior of systems composed by a few hundreds of objects or less. Until recently, this was essentially an open question.

In [78], we give a partial answer to this question. We show that, for most of the mean field models used for performance evaluation, the error made when using a mean field approximation is a  $\Theta(1/n)$ . This results greatly improved compared to previous work that showed that the error made by mean field approximation was smaller than  $O(1/\sqrt{n})$ . On the contrary, we obtain the *exact* rate of accuracy. This result came from the use of Stein's method that allows one to quantify precisely the distance between two stochastic processes. Subsequently, in [81], we show that the constant in the  $\Theta(1/n)$  can be computed numerically by a very efficient algorithm. By using this, we define the notion of refined approximation which consists in adding the  $1/n$ -correction term. This methods can also be generalize to higher order extensions or [83, 79].

**Design and analysis of distributed control algorithms** Mean field approximation is widely used in the performance evaluation community to analyze and design distributed control algorithms. Our

contribution in this domain has covered mainly two applications: cache replacement algorithms and load balancing algorithms.

Cache replacement algorithms are widely used in content delivery networks. In [66, 85, 84], we show how mean field and refined mean field approximation can be used to evaluate the performance of list-based cache replacement algorithms. In particular, we show that such policies can outperform the classically used LRU algorithm. A methodological contribution of our work is that, when evaluating precisely the behavior of such a policy, the refined mean field approximation is both faster and more accurate than what could be obtained with a stochastic simulator.

Computing resources are often spread across many machines. An efficient use of such resources requires the design of a good load balancing strategy, to distribute the load among the available machines. In [60, 61, 59], we study two paradigms that we use to design asymptotically optimal load balancing policies where a central broker sends tasks to a set of parallel servers. We show in [60, 59] that combining the classical round-robin allocation plus an evaluation of the tasks sizes can yield a policy that has a zero delay in the large system limit. This policy is interesting because the broker does not need any feedback from the servers. At the same time, this policy needs to estimate or know job durations, which is not always possible. A different approach is used in [61] where we consider a policy that does not need to estimate job durations but that uses some feedback from the servers plus a memory of where jobs where send. We show that this paradigm can also be used to design zero-delay load balancing policies as the system size grows to infinity.

**Mean field games** Various notions of mean field games have been introduced in the years 2000-2010 in theoretical economics, engineering or game theory. A mean field game is a game in which an individual tries to maximize its utility while evolving in a population of other individuals whose behavior are not directly affected by the individual. An equilibrium is a population dynamics for which a selfish individual would behave as the population. In [73], we develop the notion of discrete space mean field games, that is more amenable to study than the previously introduced notions of mean field games. This leads to two interesting contributions: mean field games are not always the limits of stochastic games as the number of players grow [72], mean field games can be used to study how much vaccination should be subsidized to encourage people to adapt a socially optimal behaviour [52].

### 3.3 Distributed Online Optimization and Learning in Games

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

#### Project-team positioning

Online learning concerns the study of *repeated decision-making in changing environments*. Of course, depending on the context, the words “learning” and “decision-making” may refer to very different things: in economics, this could mean predicting how rational agents react to market drifts; in data networks, it could mean adapting the way packets are routed based on changing traffic conditions; in machine learning and AI applications, it could mean training a neural network or the guidance system of a self-driving car; etc. In particular, the changes in the learner’s environment could be either *exogenous* (that is, independent of the learner’s decisions, such as the weather affecting the time of travel), or *endogenous* (i.e., they could depend on the learner’s decisions, as in a game of poker), or any combination thereof. However, the goal for the learner(s) is always the same: *to make more informed decisions that lead to better rewards over time*.

The study of online learning models and algorithms dates back to the seminal work of Robbins, Nash and Bellman in the 50’s, and it has since given rise to a vigorous research field at the interface of game theory, control and optimization, with numerous applications in operations research, machine learning, and data science. In this general context, our team focuses on the asymptotic behavior of online learning

and optimization algorithms, both single- and multi-agent: whether they converge, at what speed, and/or what type of non-stationary, off-equilibrium behaviors may arise when they do not.

The focus of POLARIS on game-theoretic and Markovian models of learning covers a set of specific challenges that dovetail in a highly synergistic manner with the work of other learning-oriented teams within Inria (like SCOOOL in Lille, SIERRA in Paris, and THOTH in Grenoble), and it is an important component of Inria’s activities and contributions in the field (which includes major industrial stakeholders like Google / DeepMind, Facebook, Microsoft, Amazon, and many others).

### Scientific achievements

Our team’s work on online learning covers both single- and multi-agent models; in the sequel, we present some highlights of our work structured along these basic axes.

In the single-agent setting, an important problem in the theory of Markov decision processes – i.e., discrete-time control processes with decision-dependent randomness – is the so-called “restless bandit” problem. Here, the learner chooses an action – or “arm” – from a finite set, and the mechanism determining the action’s reward changes depending on whether the action was chosen or not (in contrast to standard Markov problems where the activation of an arm does not have this effect). In this general setting, Whittle conjectured – and Weber and Weiss proved – that Whittle’s eponymous index policy is asymptotically optimal. However, the result of Weber and Weiss is purely asymptotic, and the rate of this convergence remained elusive for several decades. This gap was finally settled in a series of POLARIS papers [80][42], where we showed that Whittle indices (as well as other index policies) become optimal at a geometric rate under the same technical conditions used by Weber and Weiss to prove Whittle’s conjecture, plus a technical requirement on the non-singularity of the fixed point of the mean-field dynamics. We also propose the first sub-cubic algorithm to compute Whittle and Gittins indexes. As for reinforcement learning in Markovian bandits, we have shown that Bayesian and optimistic approaches do not use the structure of Markovian bandits similarly: While Bayesian learning has both a regret and a computational complexity that scales linearly with the number of arms, optimistic approaches all incur an exponential computation time, at least in their current versions [40].

In the multi-agent setting, our work has focused on the following fundamental question:

*Does the concurrent use of (possibly optimal) single-agent learning algorithms ensure convergence to Nash equilibrium in multi-agent, game-theoretic environments?*

Conventional wisdom might suggest a positive answer to this question because of the following “folk theorem”: under no-regret learning, the agents’ empirical frequency of play converges to the game’s set of coarse correlated equilibria. However, the actual implications of this result are quite weak: First, it concerns the empirical frequency of play and not the day-to-day sequence of actions employed by the players. Second, it concerns coarse correlated equilibria which may be supported on strictly dominated strategies – and are thus unacceptable in terms of rationalizability. These realizations prompted us to make a clean break with conventional wisdom on this topic, ultimately showing that the answer to the above question is, in general, “no”: specifically, [102, 100] showed that the (optimal) class of “follow-the-regularized-leader” (FTRL) learning algorithms leads to Poincaré recurrence even in simple,  $2 \times 2$  min-max games, thus precluding convergence to Nash equilibrium in this context.

This negative result generated significant interest in the literature as it contributed in shifting the focus towards identifying *which* Nash equilibria may arise as stable limit points of FTRL algorithms and dynamics. Earlier work by POLARIS on the topic [69, 103, 104] suggested that *strict* Nash equilibria play an important role in this question. This suspicion was recently confirmed in a series of papers [76, 53] where we established a sweeping negative result to the effect that *mixed Nash equilibria are incompatible with no-regret learning*. Specifically, we showed that any Nash equilibrium which is not strict cannot be stable and attracting under the dynamics of FTRL, especially in the presence of randomness and uncertainty. This result has significant implications for predicting the outcome of a multi-agent learning process because, combined with [103], it establishes the following far-reaching equivalence: *a state is asymptotically stable under no-regret learning if and only if it is a strict Nash equilibrium*.

Going beyond finite games, this further raised the question of what type of non-convergent behaviors can be observed in continuous games – such as the class of stochastic min-max problems that are typically

associated to generative adversarial networks (GANs) in machine learning. This question was one of our primary collaboration axes with EPFL, and led to a joint research project focused on the characterization of the convergence properties of zeroth-, first-, and (scalable) second-order methods in non-convex/non-concave problems. In particular, we showed in [54] that these state-of-the-art min-max optimization algorithms may converge with arbitrarily high probability to attractors that are in no way min-max optimal or even stationary – and, in fact, may not even contain a single stationary point (let alone a Nash equilibrium). Spurious convergence phenomena of this type can arise even in two-dimensional problems, a fact which corroborates the empirical evidence surrounding the formidable difficulty of training GANs.

### 3.4 Responsible Computer Science

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

#### Project-team positioning

The topics in this axis emerge from current social and economic questions rather than from a fixed set of mathematical methods. To this end we have identified large trends such as energy efficiency, fairness, privacy, and the growing number of new market places. In addition, COVID has posed new questions that opened new paths of research with strong links to policy making.

Throughout these works, the focus of the team is on modeling aspects of the aforementioned problems, and obtaining strong theoretical results that can give high-level guidelines on the design of markets or of decision-making procedures. Where relevant, we complement those works by measurement studies and audits of existing systems that allow identifying key issues. As this work is driven by topics, rather than methods, it allows for a wide range of collaborations, including with enterprises (e.g., Naverlabs), policy makers, and academics from various fields (economics, policy, epidemiology, etc.).

Other teams at Inria cover some of the societal challenges listed here (e.g., PRIVATICS, COMETE) but rather in isolation. The specificity of POLARIS resides in the breadth of societal topics covered and of the collaborations with non-CS researchers and non-research bodies; as well as in the application of methods such as game theory to those topics.

#### Scientific achievements

**Algorithmic fairness** As algorithmic decision-making became increasingly omnipresent in our daily lives (in domains ranging from credits to advertising, hiring, or medicine); it also became increasingly apparent that the outcome of algorithms can be discriminatory for various reasons. Since 2016, the scientific community working on the problem of algorithmic fairness has been exponentially increasing. In this context, in the early days, we worked on better understanding the extent of the problem through measurement in the case of social networks [115]. In particular, in this work, we showed that in advertising platforms, discrimination can occur from multiple different internal processes that cannot be controlled, and we advocate for measuring discrimination on the outcome directly. Then we worked on proposing solutions to guarantee fair representation in online public recommendations (aka trending topics on Twitter) [67]. This is an example of an application in which it was observed that recommendations are typically biased towards some demographic groups. In this work, our proposed solution draws an analogy between recommendation and voting and builds on existing works on fair representation in voting. Finally, in most recent times, we worked on better understanding the sources of discrimination, in the particular simple case of selection problems, and the consequences of fixing it. While most works attribute discrimination to implicit bias of the decision maker [97], we identified a fundamentally different source of discrimination: Even in the absence of implicit bias in a decision maker's estimate of candidates' quality, the estimates may differ between the different groups in their variance—that is, the decision maker's ability to precisely estimate a candidate's quality may depend on the candidate's group [75]. We show that this differential variance leads to discrimination for two reasonable baseline decision makers (group-oblivious and Bayesian optimal). Then we analyze the consequence on the selection utility of

imposing fairness mechanisms such as demographic parity or its generalization; in particular we identify some cases for which imposing fairness can improve utility. In [74], we also study similar questions in the two-stage setting, and derive the optimal selector and the “price of local fairness” one pays in utility by imposing that the interim stage be fair.

**Privacy and transparency in social computing system** Online services in general, and social networks in particular, collect massive amounts of data about their users (both online and offline). It is critical that (i) the users’ data is protected so that it cannot leak and (ii) users can know what data the service has about them and understand how it is used—this is the transparency requirement. In this context, we did two kinds of work. First, we studied social networks through measurement, in particular using the use case of Facebook. We showed that their advertising platform, through the PII<sup>1</sup>-based targeting option, allowed attackers to discover some personal data of users [117]. We also proposed an alternative design—valid for any system that proposed PII-based targeting—and proved that it fixes the problem. We then audited the transparency mechanisms of the Facebook ad platform, specifically the “Ad Preferences” page that shows what interests the platform inferred about a user, and the “Why am I seeing this” button that gives some reasons why the user saw a particular ad. In both cases, we laid the foundation for defining the quality of explanations and we showed that the explanations given were lacking key desirable properties (they were incomplete and misleading, they have since been changed) [58]. A follow-up work shed further light on the typical uses of the platform [57]. In another work, we proposed an innovative protocol based on randomized withdrawal to protect public posts deletion privacy [105]. Finally, in [82], we study an alternative data sharing ecosystem where users can choose the precision of the data they give. We model it as a game and show that, if users are motivated to reveal data by a public good component of the outcome’s precision, then certain basic statistical properties (the optimality of generalized least squares in particular) no longer hold.

**Online markets** Market design operates at the intersection of computer science and economics and has become increasingly important as many markets are redesigned on digital platforms. Studying markets for commodities, in an ongoing project we evaluate how different fee models alter strategic incentives for both buyers and sellers. We identify two general classes of fees: for one, strategic manipulation becomes infeasible as the market grows large and agents therefore have no incentive to misreport their true valuation. On the other hand, strategic manipulation is possible and we show that in this case agents aim to maximally shade their bids. This has immediate implications for the design of such markets. By contrast, [101] considers a matching market where buyers and sellers have heterogeneous preferences over each other. Traders arrive at random to the market and the market maker, having limited information, aims to optimize when to open the market for a clearing event to take place. There is a tradeoff between thickening the market (to achieve better matches) and matching quickly (to reduce waiting time of traders in the market). The tradeoff is made explicit for a wide range of underlying preferences. These works are adding to an ongoing effort to better understand and design markets [111][48].

**COVID** The COVID-19 pandemic has put humanity to one of the defining challenges of its generation and as such naturally trans-disciplinary efforts have been necessary to support decision making. In a series of articles [50][108] we proposed Green Zoning. ‘Green zones’—areas where the virus is under control based on a uniform set of conditions—can progressively return to normal economic and social activity levels, and mobility between them is permitted. By contrast, stricter public health measures are in place in ‘red zones’, and mobility between red and green zones is restricted. France and Spain were among the first countries to introduce green zoning in April 2020. The initial success of this proposal opened up the way to a large amount of follow-up work analyzing and proposing various tools to effectively deploy different tools to combat the pandemic (e.g., focus-mass testing [110] and a vaccination policy [107]). In a joint work with a group of leading economists, public health researchers and sociologists it was found that countries that opted to aim to eliminate the virus fared better not only for public health, but also for the economy and civil liberties [49]. Overall this work has been characterized by close interactions with

---

<sup>1</sup>Personally Identifiable Information

policy makers in France, Spain and the European Commission as well as substantial activity in public discourse (via TV, newspapers and radio).

**Energy efficiency** Our work on energy efficiency spanned multiple different areas and applications such as embedded systems and smart grids. Minimizing the energy consumption of embedded systems with real-time constraints is becoming more important for ecological as well as practical reasons since batteries are becoming standard power supplies. Dynamically changing the speed of the processor is the most common and efficient way to reduce energy consumption [114]. In fact, this is the reason why modern processors are equipped with Dynamic Voltage and Frequency Scaling (DVFS) technology [121]. In a stochastic environment, with random job sizes and arrival times, combining hard deadlines and energy minimization via DVFS-based techniques is difficult because forcing hard deadlines requires considering the worst cases, hardly compatible with random dynamics. Nevertheless, progress have been made to solve these types of problems in a series of papers using constrained Markov decision processes, both on the theoretical side (proving existence of optimal policies and showing their structure [88, 86, 87]) as well as on the experimental side (showing the gains of optimal policies over classical solutions [89]).

In the context of a collaboration with Enedis and Schneider Electric (via the Smart Grid chair of Grenoble-INP), we also study the problem of using smart meters to optimize the behavior of electrical distribution networks. We made three kinds of contributions on this subject: (1) how to design efficient control strategies in such a system [120, 119, 118], (2) how to co-simulate an electrical network and a communication network [96], and (3) what is the performance of the communication protocol (PLC G3) used by the Linky smart meters [99].

## 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 the research theme presented in Section 3.3.

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 continued 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 tool 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 Social and environmental responsibility

### 5.1 Footprint of research activities

The carbon footprint of the team has been quite minimal in 2021 since there has been no travel allowed with most of us working from home. Our team does not train heavy ML models requiring important processing power although some of us perform computer science experiments, mostly using the Grid5000 platforms. We keep this usage very reasonable and rely on cheaper alternatives (e.g., simulations) as much as possible.

### 5.2 Raising awareness on the climate crisis

Romain Couillet has organized several introductory seminars on the Anthropocene, which he has presented to students at the UGA and Grenoble-INSP, as well as to associations in Grenoble (FNE, AgriAlter-natif). He is also co-responsible of the *Digital Transformation DU*. He has published three articles on the issue of "usability" of artificial intelligence, and is the organizer of a special session on "Signal processing and resilience" for the GRETSI 2022 conference. He is also co-creator of the *sustainable AI* transversal axis of the MIAI project in Grenoble. He connects his professional activity with public action (Lowtechlab de Grenoble, Université Autogérée, Arche des Innovateurs, etc). Finally, he is a trainer for the "Fresque du Climat" and a member of Adrastia and FNE Isère. See section 11.2 and 11.3.3 for more details.

### 5.3 Impact of research results

The efforts of Barry Pradelski on COVID policy have received lots of media coverage and he was appointed to share his expertise to the *Centre européen de prévention et de contrôle des maladies* (ECDC).

Jean-Marc Vincent is heavily engaged since several years in the training of computer science teachers at the elementary/middle/high school levels [47]. Among one of his many activities, we can mention his

involvement in the design of the *Numérique et Sciences Informatiques, NSI : les fondamentaux* MOOC. See section 11.2 and 11.3.3 for more details.

## 6 Highlights of the year

### 6.1 Awards

- The paper "Energy optimal activation of processors for the execution of a single task with unknown size" [17], authored by J. Anselmi and B. Gaujal, won the best paper award at the international conference IEEE MASCOTS 2022.
- The paper "Robust power management via learning and game design" [51], by Z. Zhou, P. Mertikopoulos, A. L. Moustakas, N. Bambos, and P. W. Glynn, published in *Operations Research*, 69(1):331–345, in January 2021 won the INFORMS best paper award in the network analytics section.
- The paper "UnderGrad: A universal black-box optimization method with almost dimension-free convergence rate guarantees" [19], authored by K. Antonakopoulos, D. Q. Vu, V. Cevher, K. Levy, and P. Mertikopoulos, won the best paper award at the ICML'22 conference.
- Panayotis Mertikopoulos was recognized as an *Outstanding reviewer* at ICLR 2022.
- The book "Random matrix methods for machine learning" [35] by R. Couillet and Z. Liao has been published at Cambridge University Press in June 2022.

## 7 New software and platforms

### 7.1 New software

#### 7.1.1 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.

**Functional Description:** SimGrid is a simulation toolkit that provides core functionalities for the simulation of distributed applications in large scale heterogeneous distributed environments.

**News of the Year:** There were 3 major releases in 2022. The SimDag API for the simulation of the scheduling of Directed Acyclic Graphs has been dropped and replaced by the SimDag++ API which provides the different features of SimDag directly on top of the S4U API. We also dropped the old and clumsy Lua bindings to create platforms in a programmatic way. It can be done in C++ in a much cleaner

way now, which motivates this suppression. The C++ platform description has been improved to reject forbidden topologies, improve exporting for visualization, and allow users to dynamically change injected costs for MPI\_\* operations. The Python API to S4U has been extended. A new solver for parallel task (BMF) has been introduced and provides with more realistic sharing of heterogeneous resources compared to the fair bottleneck solver used by ptask\_L07. Although this is still ongoing work, this paves the way to efficient macroscopic modeling of streaming activities and parallel applications. The internals of the Model Checker have been heavily reworked and new test suites from the MPI Bugs Initiative (MBI) are now used. The documentation was thoroughly overhauled to ease the use of the framework. We also pursued our efforts to improve the overall framework, through bug fixes, code refactoring and other software quality improvements.

**URL:** <https://simgrid.org/>

**Contact:** Martin Quinson

**Participants:** Adrien Lebre, Anne-Cécile Orgerie, Arnaud Legrand, Augustin Degomme, Arnaud Giersch, Emmanuelle Saillard, Frédéric Suter, Jonathan Pastor, Martin Quinson, Samuel Thibault

**Partners:** CNRS, ENS Rennes

### 7.1.2 PSI

**Name:** Perfect Simulator

**Keywords:** Markov model, Simulation

**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 a 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.

**News of the Year:** No active development. Maintenance is ensured by the POLARIS team. The next generation of PSI lies in the MARTO project.

**URL:** <https://gitlab.inria.fr/PSI/psi3/>

**Contact:** Jean-Marc Vincent

### 7.1.3 marmoteCore

**Name:** 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 was started within the ANR MARMOTE project: ANR-12-MONU-00019.

**News of the Year:** No active development. Current development lies now in the MARTO project (next generations of PSI and marmoteCore) and in the forthcoming Marmote project.

**URL:** <https://gitlab.inria.fr/PSI/marmotecore/>

**Publications:** [hal-01651940](#), [hal-01276456](#)

**Contact:** Alain Jean-Marie

**Participants:** Alain Jean-Marie, Hlib Mykhailenko, Benjamin Briot, Franck Quessette, Issam Rabhi, Jean-Marc Vincent, Jean-Michel Fourneau

**Partners:** Université de Versailles St-Quentin-en-Yvelines, Université Paris Nanterre

#### 7.1.4 MarTO

**Name:** Markov Toolkit for Markov models simulation: perfect sampling and Monte Carlo simulation

**Keywords:** Perfect sampling, Markov model

**Functional Description:** MarTO 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. This software is a rewrite, more efficient and flexible, of PSI

**News of the Year:** No official release yet. The code development is in progress.

**URL:** <https://gitlab.inria.fr/MarTo/marto>

**Contact:** Vincent Danjean

#### 7.1.5 GameSeer

**Keyword:** Game theory

**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.

**News of the Year:** No new release but the development is still active.

**URL:** <http://polaris.imag.fr/panayotis.mertikopoulos/publications/s01-gameseer/>

**Contact:** Panayotis Mertikopoulos

#### 7.1.6 rmf\_tool

**Name:** A library to Compute (Refined) Mean Field Approximations

**Keyword:** Mean Field

**Functional Description:** The tool accepts three model types:

- homogeneous population processes (HomPP)
- density dependent population processes (DDPPs)
- heterogeneous population models (HetPP)

In particular, it provides a numerical algorithm to compute the constant of the refined mean field approximation provided in the paper "A Refined Mean Field Approximation" by N. Gast and B. Van Houdt, SIGMETRICS 2018, and a framework to compute heterogeneous mean field approximations as proposed in "Mean Field and Refined Mean Field Approximations for Heterogeneous Systems: It Works!" by N. Gast and S. Allmeier, SIGMETRICS 2022.

**URL:** [https://github.com/ngast/rmf\\_tool](https://github.com/ngast/rmf_tool)

**Publications:** hal-01622054, tel-02509756, hal-03485044

**Contact:** Nicolas Gast

## 8 New results

The new results produced by the team in 2022 can be grouped into the following categories.

### 8.1 Performance evaluation of large systems

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

Finely tuning applications and understanding the influence of key parameters (number of processes, granularity, collective operation algorithms, virtual topology, and process placement) is critical to obtain good performance on supercomputers. With the high consumption of running applications at scale, doing so solely to optimize their performance is particularly costly. We have shown in [3] that SimGrid and SMPI [simgrid.org](https://simgrid.org) could be used to obtain inexpensive but faithful predictions of expected performance. The methodology we propose decouples the complexity of the platform, which is captured through statistical models of the performance of its main components (MPI communications, BLAS operations), from the complexity of adaptive applications by emulating the application and skipping regular non-MPI parts of the code. We demonstrate the capability of our method with High-Performance Linpack (HPL), the benchmark used to rank supercomputers in the TOP500, which requires careful tuning. This work presents an extensive (in)validation study that compares simulation with real experiments and demonstrates our ability to predict the performance of HPL within a few percent consistently. This study allows us to identify the main modeling pitfalls (e.g., spatial and temporal node variability or network heterogeneity and irregular behavior) that need to be considered. Our “surrogate” also allows studying several subtle HPL parameter optimization problems while accounting for uncertainty on the platform.

We have also shown in [13] how the structure of complex applications such as a multifrontal sparse linear solvers could be exploited to detect and correct non-trivial performance problems. Efficiently exploiting computational resources in heterogeneous platforms is a real challenge which has motivated the adoption of the task-based programming paradigm where resource usage is dynamic and adaptive. Unfortunately, classical performance visualization techniques used in routine performance analysis often fail to provide any insight in this new context, especially when the application structure is irregular. We propose and implement in StarVZ several performance visualization techniques tailored for the analysis of task-based multifrontal sparse linear solvers and show that by building on both a performance model of irregular tasks and on structure of the application (in particular the elimination tree), we can detect and highlight anomalies and understand resource utilization from the application point-of-view in a very insightful way. We validate these novel performance analysis techniques with the QR\_mumps sparse parallel solver by describing a series of case studies where we identify and address non trivial performance issues thanks to our visualization methodology.

Large systems can be particularly difficult to analyze because of inherent state-space explosion and most computations become untractable. Mean field approximation is a powerful technique to study the performance of very large stochastic systems represented as systems of interacting objects. Applications include load balancing models, epidemic spreading, cache replacement policies, or large-scale data centers, for which mean field approximation gives very accurate estimates of the transient or steady-state behaviors. In a series of recent papers, a new and more accurate approximation, called the refined mean field approximation has been presented. A key strength of this technique lies in its applicability to not-so-large systems. Yet, computing this new approximation can be cumbersome, which is why develop a tool, called `rmf_tool` and available at [github.com/ngast/rmf\\_tool](https://github.com/ngast/rmf_tool), that takes the description of a mean field model, and can numerically compute its mean field approximations and refinement.

Mean field approximation is asymptotically exact for systems composed of  $n$  homogeneous objects under mild conditions. In [1], we study what happens when objects are heterogeneous. This can represent servers with different speeds or contents with different popularities. We define an interaction model that allows obtaining asymptotic convergence results for stochastic systems with heterogeneous object behavior, and show that the error of the mean field approximation is of order  $O(1/n)$ . More importantly, we show how to adapt the refined mean field approximation and show that the error of this approximation is reduced to  $O(1/n^2)$ . To illustrate the applicability of our result, we present two examples. The first

addresses a list-based cache replacement model,  $RANDOM(m)$ , which is an extension of the  $RANDOM$  policy. The second is a heterogeneous supermarket model. These examples show that the proposed approximations are computationally tractable and very accurate. They also show that for moderate system sizes (30) the refined mean field approximation tends to be more accurate than simulations for any reasonable simulation time.

## 8.2 Energy optimization

**Participants:** Jonatha Anselmi, Bruno Gaujal, Louis-Sébastien Rebuffi.

A key objective in the management of modern computer systems consists in minimizing the electrical energy consumed by processing resources while satisfying certain target performance criteria. In [17], we consider the execution of a single task with unknown size on top of a service system that offers a limited number of processing speeds, say  $N$ , and investigate the problem of finding a speed profile that minimizes the resulting energy consumption subject to a deadline constraint. Existing works mainly investigated this problem when speed profiles are continuous functions. In contrast, the novelty of our work is to consider discontinuous speed profiles, i.e., a case that arises naturally when the underlying computational platform offers a finite number of speeds. In our main result, we show that the computation of an optimal speed profile boils down to solving a convex optimization problem. Under mild assumptions, for such convex optimization we prove some structural results that yield the formulation of an extremely efficient solution algorithm. Specifically, we show that the optimal speed profile can be computed by solving  $O(\log N)$  one-dimensional equations. Our results hold when the task size follows a known probability distribution function and the set of available speeds, if listed in increasing order, forms a sublinear concave sequence.

More generally, energy optimization should be performed at a global scale and requires to revisit load balancing strategies. Techniques like replication and speculation are double-edged weapons that must be handled with caution as the resource overhead may be detrimental when used too aggressively. We have studied such strategies in previous work and presented an overview [2] in the special issue of Queueing Systems, *100 views on queues*.

## 8.3 Large system performance optimization through learning techniques

**Participants:** Arnaud Legrand, Lucas Leandro Nesi, Panayotis Mertikopoulos.

Large infrastructures and computing applications typically exhibit some form of regularity which should be exploited but their stochastic nature makes their optimization difficult. In this series of work, we demonstrate that simple machine and reinforcement learning techniques can be tailored to optimize these systems.

Parallel applications performance strongly depends on the number of resources. Although adding new nodes usually reduces execution time, excessive amounts are often detrimental as they incur substantial communication overhead, which is difficult to anticipate. Characteristics like network contention, data distribution methods, synchronizations, and how communications and computations overlap generally impact the performance. Finding the correct number of resources can thus be particularly tricky for multi-phase applications as each phase may have very different needs, and the popularization of hybrid (CPU+GPU) machines and heterogeneous partitions makes it even more difficult. In [32], we study and propose, in the context of a task-based GeoStatistic application, strategies for the application to actively learn and adapt to the best set of heterogeneous nodes it has access to. We propose strategies that use the Gaussian Process method with trends, bound mechanisms for reducing the search space, and heterogeneous behavior modeling. We compare these methods with traditional exploration strategies in 16 different machines scenarios. In the end, the proposed strategies are able to gain up to  $\approx 51\%$  compared to the standard case of using all the nodes while having low overhead.

At the scale of the whole high-performance computing platform, job scheduling is also a hard problem that involves uncertainties on both the job arrival process and their execution times. Users typically provide only loose upper bounds for job execution times, which are not so useful for scheduling heuristics based on processing times. Previous studies focused on applying regression techniques to obtain better execution time estimates, which worked reasonably well and improved scheduling metrics. However, these approaches require a long period of training data. In [16], 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 a fast adaptation to changes in the workload characteristics. We evaluate our approach using four scheduling policies on seven HPC platform workload traces. We show that: first, incorporating such classification reduces the average bounded slowdown of jobs in all scenarios, second, in most considered scenarios, the improvements are comparable to the ideal hypothetical situation where the scheduler would know in advance the exact running time of jobs.

In [4], we evaluate the relevance of bandit-like strategies for the Fog computing context and explore the information-coordination trade-off. Fog computing emerges as a potential solution to handle the growth of traffic and processing demands, providing nearby resources to run IoT applications. In this paper, we consider the reconfiguration problem, i.e., how to dynamically adapt the placement of IoT applications running in the Fog, depending on application needs and evolution of resource usage. We propose and evaluate a series of reconfiguration algorithms, based on both online scheduling (dynamic packing) and online learning (bandit) approaches. Through an extensive set of experiments in a realistic testbed built on Grid5000 and FIT-IoT lab, we demonstrate that the performance strongly and mainly depends on the quality and availability of information from both Fog infrastructure and IoT applications. We show that a reactive and greedy strategy can overcome the performance of state-of-the-art online learning algorithms, as long as the strategy has access to a little extra information.

Finally, the high degree of variability present in current and emerging mobile wireless networks calls for mathematical tools and techniques that transcend classical (convex) optimization paradigms. In [20], we provide a gentle introduction to online learning and optimization algorithms that are able to provably cope with this variability and provide policies that are asymptotically optimal in hindsight-a property known as no regret. The focal point of this survey is to delineate the trade-off between the information available as feedback to the learner, and the achievable regret guarantees starting with the case of gradient-based (first-order) feedback, then moving on to value-based (zeroth-order) feedback, and, ultimately, pushing the envelope to the extreme case of a single bit of feedback. We illustrate our theoretical analysis with a series of practical wireless network examples that highlight the potential of this elegant toolbox.

## 8.4 Exploiting specific structures in reinforcement learning

**Participants:** Jonatha Anselmi, Yan Chen, Kimang Khun, Nicolas Gast, Bruno Gaujal, Louis-Sébastien Rebuffi, Panayotis Mertikopoulos.

The Multi-armed Stochastic Bandit framework is a classic reinforcement learning problem to study the exploration exploitation trade-off dilemma and for which several optimal algorithms like UCB<sup>2</sup> and Thompson sampling<sup>3</sup>, whose optimality has only recently been proved by Kaufmann et al.<sup>4</sup>, have been proposed. Although the first strategy is an optimistic strategy which systematically chooses the "most promising" arm, the second one build on a Bayesian perspective and samples the posterior to decide which arm to select. The Markovian Bandit allows to model situations where the reward distribution is modeled as a Markov chain and may thus exhibit temporal changes. A key challenge in this context is the curse of dimensionality, which basically says that the state size of the Markov process is exponential in the number of the system components so that the complexity of computing an optimal policy and its value

<sup>2</sup>P. Auer, N. Cesa-Bianchi, P. Fischer. *Finite-time analysis of the multiarmed bandit problem*. Machine Learning 47, 2002

<sup>3</sup>W. R. Thompson, *On the likelihood that one unknown probability exceeds another in view of the evidence of two samples*, Biometrika, 25:285–294, 1933

<sup>4</sup>E. Kaufmann, N. Korda, R. Munos. *Thompson Sampling: An Asymptotically Optimal Finite Time Analysis*. ALT 2012

are exponential. This is why specific algorithms should be designed to exploit the very specific structures that some state space exhibit. Based on earlier results, we have presented these lines of thought in two articles in special issue of Queueing Systems, *100 views on queues: Learning in Queues* [7], and *Why (and When) do Asymptotic Methods Work so well?* [6].

Restless bandits are a specific kind of bandits in which the state of each arm evolves according to a Markov process independently of the learner's actions. Most restless Markovian bandits problems in infinite horizon can be solved quasioptimally using the famous Whittle index, which is a generalization of the Gittins index. In [41], we develop an algorithm to test the indexability and compute the Whittle indices of any finite-state restless bandit arm. This algorithm works in the discounted and non-discounted cases, and can compute Gittins index. Our algorithm builds on three tools: (1) a careful characterization of Whittle index that allows one to compute recursively the  $k$ -th smallest index from the  $(k - 1)$ -th smallest, and to test indexability, (2) the use of the Sherman-Morrison formula to make this recursive computation efficient, and (3) a sporadic use of the fastest matrix inversion and multiplication methods to obtain a subcubic complexity. We show that an efficient use of the Sherman-Morrison formula leads to an algorithm that computes Whittle index in  $2/3n^3 + o(n^3)$  arithmetic operations, where  $n$  is the number of states of the arm. The careful use of fast matrix multiplication leads to the first subcubic algorithm to compute Whittle or Gittins index: By using the current fastest matrix multiplication, the theoretical complexity of our algorithm is  $O(n^{2.5286})$ . We also develop an efficient implementation of our algorithm that can compute indices of Markov chains with several thousands of states in less than a few seconds.

In [42], we provide a framework to analyse control policies for the restless Markovian bandit model, under both finite and infinite time horizon. We show that when the population of arms goes to infinity, the value of the optimal control policy converges to the solution of a linear program (LP). We provide necessary and sufficient conditions for a generic control policy to be: i) asymptotically optimal; ii) asymptotically optimal with square root convergence rate; iii) asymptotically optimal with exponential rate. We then construct the LP-index policy that is asymptotically optimal with square root convergence rate on all models, and with exponential rate if the model is non-degenerate in finite horizon, and satisfies a uniform global attractor property in infinite horizon. We next define the LP-update policy, which is essentially a repeated LP-index policy that solves a new linear program at each decision epoch. We provide numerical experiments to compare the efficiency of LP-based policies. We compare the performance of the LP-index policy and the LP-update policy with other heuristics. Our result demonstrates that the LP-update policy outperforms the LP-index policy in general, and can have a significant advantage when the transition matrices are wrongly estimated.

In [33], we revisit the regret of undiscounted reinforcement learning in MDPs with a birth and death structure, which are typical of queueing systems. Specifically, we consider a controlled queue with impatient jobs and the main objective is to optimize a trade-off between energy consumption and user-perceived performance. Within this setting, the diameter  $D$  of the MDP is  $\Omega(S^5)$ , where  $S$  is the number of states. Therefore, the existing lower and upper bounds on the regret at time  $T$ , of order  $O(\sqrt{DSAT})$  for MDPs with  $S$  states and  $A$  actions, may suggest that reinforcement learning is inefficient here. In our main result however, we exploit the structure of our MDPs to show that the regret of a slightly-tweaked version of the classical learning algorithm UCRL2 is in fact upper bounded by  $\tilde{O}(\sqrt{E_2AT})$  where  $E_2$  is related to the weighted second moment of the stationary measure of a reference policy. Importantly,  $E_2$  is bounded independently of  $S$ . Thus, our bound is asymptotically independent of the number of states and of the diameter. This result is based on a careful study of the number of visits performed by the learning algorithm to the states of the MDP, which is highly non-uniform.

Finally, in many online decision processes, the optimizing agent is called to choose between large numbers of alternatives with many inherent similarities; in turn, these similarities imply closely correlated losses that may confound standard discrete choice models and bandit algorithms. We study this question in the context of nested bandits, a class of adversarial multi-armed bandit problems where the learner seeks to minimize their regret in the presence of a large number of distinct alternatives with a hierarchy of embedded (non-combinatorial) similarities. In this setting, optimal algorithms based on the exponential weights blueprint (like Hedge, EXP3, and their variants) may incur significant regret because they tend to spend excessive amounts of time exploring irrelevant alternatives with similar, suboptimal costs. To account for this, we propose in [30] a nested exponential weights (NEW) algorithm that performs a layered exploration of the learner's set of alternatives based on a nested, step-by-step selection method. In so doing, we obtain a series of tight bounds for the learner's regret showing that online learning problems

with a high degree of similarity between alternatives can be resolved efficiently, without a red bus / blue bus paradox occurring.

## 8.5 Distributed learning and optimization

**Participants:** Yu-Guan Hsieh, Panayotis Mertikopoulos.

Many learning algorithms operate in centralized way, which raises many practical issues in terms of scalability, privacy, hence a high interest for designing efficient distributed and federated machine learning algorithms. In such context it is essential to design systems that can seamlessly adapt to the workload and to the evolving behaviour of its components (users, resources, network).

In [43] we consider decentralized optimization problems in which a number of agents collaborate to minimize the average of their local functions by exchanging over an underlying communication graph. Specifically, we place ourselves in an asynchronous model where only a random portion of nodes perform computation at each iteration, while the information exchange can be conducted between all the nodes and in an asymmetric fashion. For this setting, we propose an algorithm that combines gradient tracking and variance reduction over the entire network. This enables each node to track the average of the gradients of the objective functions. Our theoretical analysis shows that the algorithm converges linearly, when the local objective functions are strongly convex, under mild connectivity conditions on the expected mixing matrices. In particular, our result does not require the mixing matrices to be doubly stochastic. In the experiments, we investigate a broadcast mechanism that transmits information from computing nodes to their neighbors, and confirm the linear convergence of our method on both synthetic and real-world datasets.

One of the most widely used methods for solving large-scale stochastic optimization problems is distributed asynchronous stochastic gradient descent (DASGD), a family of algorithms that result from parallelizing stochastic gradient descent on distributed computing architectures (possibly) asynchronously. However, a key obstacle in the efficient implementation of DASGD is the issue of delays: when a computing node contributes a gradient update, the global model parameter may have already been updated by other nodes several times over, thereby rendering this gradient information stale. These delays can quickly add up if the computational throughput of a node is saturated, so the convergence of DASGD may be compromised in the presence of large delays. In [15], we show that, by carefully tuning the algorithm's step-size, convergence to the critical set is still achieved in mean square, even if the delays grow unbounded at a polynomial rate. We also establish finer results in a broad class of structured optimization problems (called variationally coherent), where we show that DASGD converges to a global optimum with probability 1 under the same delay assumptions. Together, these results contribute to the broad landscape of large-scale non-convex stochastic optimization by offering state-of-the-art theoretical guarantees and providing insights for algorithm design.

In [10], we provide a general framework for studying multi-agent online learning problems in the presence of delays and asynchronicities. Specifically, we propose and analyze a class of adaptive dual averaging schemes in which agents only need to accumulate gradient feedback received from the whole system, without requiring any between-agent coordination. In the single-agent case, the adaptivity of the proposed method allows us to extend a range of existing results to problems with potentially unbounded delays between playing an action and receiving the corresponding feedback. In the multi-agent case, the situation is significantly more complicated because agents may not have access to a global clock to use as a reference point; to overcome this, we focus on the information that is available for producing each prediction rather than the actual delay associated with each feedback. This allows us to derive adaptive learning strategies with optimal regret bounds, even in a fully decentralized, asynchronous environment. Finally, we also analyze an "optimistic" variant of the proposed algorithm which is capable of exploiting the predictability of problems with a slower variation and leads to improved regret bounds.

In decentralized optimization environments, each agent  $i$  in a network of  $n$  nodes has its own private function  $f_i$ , and nodes communicate with their neighbors to cooperatively minimize the aggregate objective  $\sum_{i=1}^n f_i$ . In this setting, synchronizing the nodes' updates incurs significant communication overhead and computational costs, so much of the recent literature has focused on the analysis and

design of asynchronous optimization algorithms, where agents activate and communicate at arbitrary times without needing a global synchronization enforcer. However, most works assume that when a node activates, it selects the neighbor to contact based on a fixed probability (e.g., uniformly at random), a choice that ignores the optimization landscape at the moment of activation. Instead, in [21] we introduce an optimization-aware selection rule that chooses the neighbor providing the highest dual cost improvement (a quantity related to a dualization of the problem based on consensus). This scheme is related to the coordinate descent (CD) method with the Gauss-Southwell (GS) rule for coordinate updates; in our setting however, only a subset of coordinates is accessible at each iteration (because each node can communicate only with its neighbors), so the existing literature on GS methods does not apply. To overcome this difficulty, we develop a new analytical framework for smooth and strongly convex  $f_i$  that covers the class of set-wise CD algorithms – a class that directly applies to decentralized scenarios, but is not limited to them – and we show that the proposed setwise GS rule achieves a speedup factor of up to the maximum degree in the network (which is in the order of  $\Theta(n)$  for highly connected graphs). The speedup predicted by our analysis is validated in numerical experiments with synthetic data.

## 8.6 Learning in games

**Participants:** Kimon Antonakopoulos, Yu Guan Hsieh, Panayotis Mertikopoulos, Bary Pradelski, Patrick Loiseau.

Learning in games naturally occurs in situations where the resources or the decision is distributed among several agents or even in situations where a centralised decision maker has to adapt to strategic users. Yet, it is considerably more difficult than in classical minimization games as the resulting equilibria may be attractive or not and the dynamic often exhibit cyclic behaviors.

In [24], we examine the problem of regret minimization when the learner is involved in a continuous game with other optimizing agents: in this case, if all players follow a no-regret algorithm, it is possible to achieve significantly lower regret relative to fully adversarial environments. We study this problem in the context of variationally stable games (a class of continuous games which includes all convex-concave and monotone games), and when the players only have access to noisy estimates of their individual payoff gradients. If the noise is additive, the game-theoretic and purely adversarial settings enjoy similar regret guarantees; however, if the noise is multiplicative, we show that the learners can, in fact, achieve constant regret. We achieve this faster rate via an optimistic gradient scheme with learning rate separation that is, the method's extrapolation and update steps are tuned to different schedules, depending on the noise profile. Subsequently, to eliminate the need for delicate hyperparameter tuning, we propose a fully adaptive method that smoothly interpolates between worst-and best-case regret guarantees.

In [9], we examine the long-run behavior of a wide range of dynamics for learning in nonatomic games, in both discrete and continuous time. The class of dynamics under consideration includes fictitious play and its regularized variants, the best reply dynamics (again, possibly regularized), as well as the dynamics of dual averaging / "follow the regularized leader" (which themselves include as special cases the replicator dynamics and Friedman's projection dynamics). Our analysis concerns both the actual trajectory of play and its time-average, and we cover potential and monotone games, as well as games with an evolutionarily stable state (global or otherwise). We focus exclusively on games with finite action spaces; nonatomic games with continuous action spaces are treated in detail in Part II of this work.

In [45], we develop a unified stochastic approximation framework for analyzing the long-run behavior of multi-agent online learning in games. Our framework is based on a "primal-dual", mirrored Robbins-Monro (MRM) template which encompasses a wide array of popular game-theoretic learning algorithms (gradient methods, their optimistic variants, the EXP3 algorithm for learning with payoff-based feedback in finite games, etc.). In addition to providing an integrated view of these algorithms, the proposed MRM blueprint allows us to obtain a broad range of new convergence results, both asymptotic and in finite time, in both continuous and finite games.

Learning in stochastic games is a notoriously difficult problem because, in addition to each other's strategic decisions, the players must also contend with the fact that the game itself evolves over time, possibly in a very complicated manner. Because of this, the convergence properties of popular learning algorithms - like policy gradient and its variants - are poorly understood, except in specific classes of

games (such as potential or two-player, zero-sum games). In view of this, we examine in [23] the long-run behavior of policy gradient methods with respect to Nash equilibrium policies that are second-order stationary (SOS) in a sense similar to the type of sufficiency conditions used in optimization. Our first result is that SOS policies are locally attracting with high probability, and we show that policy gradient trajectories with gradient estimates provided by the REINFORCE algorithm achieve an  $O(1/\sqrt{n})$  distance-squared convergence rate if the method's step-size is chosen appropriately. Subsequently, specializing to the class of deterministic Nash policies, we show that this rate can be improved dramatically and, in fact, policy gradient methods converge within a finite number of iterations in that case.

In [5], we examine the long-run behavior of multi-agent online learning in games that evolve over time. Specifically, we focus on a wide class of policies based on mirror descent, and we show that the induced sequence of play (a) converges to Nash equilibrium in time-varying games that stabilize in the long run to a strictly monotone limit; and (b) it stays asymptotically close to the evolving equilibrium of the sequence of stage games (assuming they are strongly monotone). Our results apply to both gradient-based and payoff-based feedback - i.e., when players only get to observe the payoffs of their chosen actions.

In [29], we also investigate the impact of feedback quantization on multi-agent learning. In particular, we analyze the equilibrium convergence properties of the well-known "follow the regularized leader" (FTRL) class of algorithms when players can only observe a quantized (and possibly noisy) version of their payoffs. In this information-constrained setting, we show that coarser quantization triggers a qualitative shift in the convergence behavior of FTRL schemes. Specifically, if the quantization error lies below a threshold value (which depends only on the underlying game and not on the level of uncertainty entering the process or the specific FTRL variant under study), then (i) FTRL is attracted to the game's strict Nash equilibria with arbitrarily high probability; and (ii) the algorithm's asymptotic rate of convergence remains the same as in the non-quantized case. Otherwise, for larger quantization levels, these convergence properties are lost altogether: players may fail to learn anything beyond their initial state, even with full information on their payoff vectors. This is in contrast to the impact of quantization in continuous optimization problems, where the quality of the obtained solution degrades smoothly with the quantization level.

Finally, the literature on evolutionary game theory suggests that pure strategies that are strictly dominated by other pure strategies always become extinct under imitative game dynamics, but they can survive under innovative dynamics. As we explain in [12], this is because innovative dynamics favour rare strategies while standard imitative dynamics do not. However, as we also show, there are reasonable imitation protocols that favour rare or frequent strategies, thus allowing strictly dominated strategies to survive in large classes of imitation dynamics. Dominated strategies can persist at nontrivial frequencies even when the level of domination is not small.

## 8.7 Advanced learning and optimization methods

**Participants:** Kimon Antonakopoulos, Yu Guan Hsieh, Panayotis Mertikopoulos.

Variational inequalities - and, in particular, stochastic variational inequalities - have recently attracted considerable attention in machine learning and learning theory as a flexible paradigm for "optimization beyond minimization", i.e., for problems where finding an optimal solution does not necessarily involve minimizing a loss function.

In [39], we examine the last-iterate convergence rate of Bregman proximal methods - from mirror descent to mirror-prox - in constrained variational inequalities. Our analysis shows that the convergence speed of a given method depends sharply on the Legendre exponent of the underlying Bregman regularizer (Euclidean, entropic, or other), a notion that measures the growth rate of said regularizer near a solution. In particular, we show that boundary solutions exhibit a clear separation of regimes between methods with a zero and non-zero Legendre exponent respectively, with linear convergence for the former versus sublinear for the latter. This dichotomy becomes even more pronounced in linearly constrained problems where, specifically, Euclidean methods converge along sharp directions in a finite number of steps, compared to a linear rate for entropic methods.

Universal methods for optimization are designed to achieve theoretically optimal convergence rates without any prior knowledge of the problem's regularity parameters or the accuracy of the gradient oracle employed by the optimizer. In this regard, existing state-of-the-art algorithms achieve an  $O(1/T^2)$  value convergence rate in Lipschitz smooth problems with a perfect gradient oracle, and an  $O(1/\sqrt{T})$  convergence rate when the underlying problem is non-smooth and/or the gradient oracle is stochastic. On the downside, these methods do not take into account the problem's dimensionality, and this can have a catastrophic impact on the achieved convergence rate, in both theory and practice. In [19] we aim to bridge this gap by providing a scalable universal gradient method - dubbed *UnderGrad* - whose oracle complexity is almost dimension-free in problems with a favorable geometry (like the simplex, linearly constrained semidefinite programs and combinatorial bandits), while retaining the order-optimal dependence on  $T$  described above. These "best-of-both-worlds" results are achieved via a primal-dual update scheme inspired by the dual exploration method for variational inequalities.

Adaptive first-order methods in optimization are prominent in machine learning and data science owing to their ability to automatically adapt to the landscape of the function being optimized. However, their convergence guarantees are typically stated in terms of vanishing gradient norms, which leaves open the issue of converging to undesirable saddle points (or even local maximizers). In [18], we focus on the *AdaGrad* family of algorithms-with scalar, diagonal or full-matrix preconditioning-and we examine the question of whether the method's trajectories avoid saddle points. A major challenge that arises here is that *AdaGrad*'s step-size (or, more accurately, the method's preconditioner) evolves over time in a filtration-dependent way, i.e., as a function of all gradients observed in earlier iterations; as a result, avoidance results for methods with a constant or vanishing step-size do not apply. We resolve this challenge by combining a series of step-size stabilization arguments with a recursive representation of the *AdaGrad* preconditioner that allows us to employ stable manifold techniques and ultimately show that the induced trajectories avoid saddle points from almost any initial condition.

Many important learning algorithms, such as stochastic gradient methods, are often deployed to solve nonlinear problems on Riemannian manifolds. Motivated by these applications, we propose in [25] a family of Riemannian algorithms generalizing and extending the seminal stochastic approximation framework of Robbins and Monro. Compared to their Euclidean counterparts, Riemannian iterative algorithms are much less understood due to the lack of a global linear structure on the manifold. We overcome this difficulty by introducing an extended Fermi coordinate frame which allows us to map the asymptotic behavior of the proposed Riemannian Robbins-Monro (RRM) class of algorithms to that of an associated deterministic dynamical system under very mild assumptions on the underlying manifold. In so doing, we provide a general template of almost sure convergence results that mirrors and extends the existing theory for Euclidean Robbins-Monro schemes, albeit with a significantly more involved analysis that requires a number of new geometric ingredients. We showcase the flexibility of the proposed RRM framework by using it to establish the convergence of a retraction-based analogue of the popular optimistic / extra-gradient methods for solving minimization problems and games, and we provide a unified treatment for their convergence.

Last, in [44], we propose and analyze exact and inexact regularized Newton-type methods for finding a global saddle point of a convex-concave unconstrained min-max optimization problem. Compared to their first-order counterparts, investigations of second-order methods for min-max optimization are relatively limited, as obtaining global rates of convergence with second-order information is much more involved. In this paper, we highlight how second-order information can be used to speed up the dynamics of dual extrapolation methods despite inexactness. Specifically, we show that the proposed algorithms generate iterates that remain within a bounded set and the averaged iterates converge to an  $\epsilon$ -saddle point within  $O(\epsilon^{-2/3})$  iterations in terms of a gap function. Our algorithms match the theoretically established lower bound in this context and our analysis provides a simple and intuitive convergence analysis for second-order methods without requiring any compactness assumptions. Finally, we present a series of numerical experiments on synthetic and real data that demonstrate the efficiency of the proposed algorithms.

## 8.8 Random matrix analysis and Machine Learning

**Participants:** Romain Couillet, Hugo Lebeau.

Random matrix theory has recently proven to be a very effective tool to understand Machine Learning challenges. In particular, concentration results can be used to derive more efficient and frugal algorithms.

Several machine learning problems such as latent variable model learning and community detection can be addressed by estimating a low-rank signal from a noisy tensor. Despite recent substantial progress on the fundamental limits of the corresponding estimators in the large-dimensional setting, some of the most significant results are based on spin glass theory, which is not easily accessible to non-experts. In [8], we propose a sharply distinct and more elementary approach, relying on tools from random matrix theory. The key idea is to study random matrices arising from contractions of a random tensor, which give access to its spectral properties. In particular, for a symmetric  $d$ -th order rank-one model with Gaussian noise, our approach yields a novel characterization of maximum likelihood (ML) estimation performance in terms of a fixed-point equation valid in the regime where weak recovery is possible. For  $d = 3$ , the solution to this equation matches the existing results. We conjecture that the same holds for any order  $d$ , based on numerical evidence for  $d \in 4, 5$ . Moreover, our analysis illuminates certain properties of the large-dimensional ML landscape. Our approach can be extended to other models, including asymmetric and non-Gaussian ones.

In [11], we introduce a random matrix framework for the analysis of clustering on high-dimensional data streams, a particularly relevant setting for a more sober processing of large amounts of data with limited memory and energy resources. Assuming data  $x_1, x_2, \dots$  arrives as a continuous flow and a small number  $L$  of them can be kept in the learning pipeline, one has only access to the diagonal elements of the Gram kernel matrix:  $[K_L]_{i,j} = \frac{1}{p} x_i^T x_j \mathbb{1}_{|i-j| < L}$ . Under a large-dimensional data regime, we derive the limiting spectral distribution of the banded kernel matrix  $K_L$  and study its isolated eigenvalues and eigenvectors, which behave in an unfamiliar way. We detail how these results can be used to perform efficient online kernel spectral clustering and provide theoretical performance guarantees. Our findings are empirically confirmed on image clustering tasks. Leveraging on optimality results of spectral methods for clustering, this work offers insights on efficient online clustering techniques for high-dimensional data. This work has also been presented at the GRETSI [28].

## 8.9 Fairness and equity in digital (recommendation, advertising, persistent storage) systems

**Participants:** Nicolas Gast, Patrick Loiseau, Rémi Molina, Bary Pradelski, Benjamin Roussillon, Till Kletti.

The general deployment of machine-learning systems in many domains ranging from security to recommendation and advertising to guide strategic decisions leads to an interesting line of research from a game theory perspective. In this context, fairness, discrimination, and privacy are particularly important issues.

Discrimination in machine learning often arises along multiple dimensions (a.k.a. protected attributes); it is then desirable to ensure intersectional fairness-i.e., that no subgroup is discriminated against. It is known that ensuring marginal fairness for every dimension independently is not sufficient in general. Due to the exponential number of subgroups, however, directly measuring intersectional fairness from data is impossible. In [31], our primary goal is to understand in detail the relationship between marginal and intersectional fairness through statistical analysis. We first identify a set of sufficient conditions under which an exact relationship can be obtained. Then, we prove bounds (easily computable through marginal fairness and other meaningful statistical quantities) in high probability on intersectional fairness in the general case. Beyond their descriptive value, we show that these theoretical bounds can be leveraged to derive a heuristic improving the approximation and bounds of intersectional fairness by choosing, in a relevant manner, protected attributes for which we describe intersectional subgroups. Finally, we test the performance of our approximations and bounds on real and synthetic data-sets.

To better understand discriminations and the effect of affirmative actions in selection problems (e.g., college admission or hiring), a recent line of research proposed a model based on differential variance. This model assumes that the decision-maker has a noisy estimate of each candidate's quality and puts forward the difference in the noise variances between different demographic groups as a key factor to explain discrimination. The literature on differential variance, however, does not consider the strategic behavior of candidates who can react to the selection procedure to improve their outcome, which is well-known to happen in many domains. In [22], we study how the strategic aspect affects fairness in selection problems. We propose to model selection problems with strategic candidates as a contest game: A population of rational candidates compete by choosing an effort level to increase their quality. They incur a cost-of-effort but get a (random) quality whose expectation equals the chosen effort. A Bayesian decision-maker observes a noisy estimate of the quality of each candidate (with differential variance) and selects the fraction  $\alpha$  of best candidates based on their posterior expected quality; each selected candidate receives a reward  $S$ . We characterize the (unique) equilibrium of this game in the different parameters' regimes, both when the decision-maker is unconstrained and when they are constrained to respect the fairness notion of demographic parity. Our results reveal important impacts of the strategic behavior on the discrimination observed at equilibrium and allow us to understand the effect of imposing demographic parity in this context. In particular, we find that, in many cases, the results contrast with the non-strategic setting. We also find that, when the cost-of-effort depends on the demographic group (which is reasonable in many cases), then it entirely governs the observed discrimination (i.e., the noise becomes a second-order effect that does not have any impact on discrimination). Finally we find that imposing demographic parity can sometimes increase the quality of the selection at equilibrium; which surprisingly contrasts with the optimality of the Bayesian decision-maker in the non-strategic case. Our results give a new perspective on fairness in selection problems, relevant in many domains where strategic behavior is a reality.

In [34], we consider the problem of linear regression from strategic data sources with a public good component, i.e., when data is provided by strategic agents who seek to minimize an individual provision cost for increasing their data's precision while benefiting from the model's overall precision. In contrast to previous works, our model tackles the case where there is uncertainty on the attributes characterizing the agents' data – a critical aspect of the problem when the number of agents is large. We provide a characterization of the game's equilibrium, which reveals an interesting connection with optimal design. Subsequently, we focus on the asymptotic behavior of the covariance of the linear regression parameters estimated via generalized least squares as the number of data sources becomes large. We provide upper and lower bounds for this covariance matrix and we show that, when the agents' provision costs are super-linear, the model's covariance converges to zero but at a slower rate relative to virtually all learning problems with exogenous data. On the other hand, if the agents' provision costs are linear, this covariance fails to converge. This shows that even the basic property of consistency of generalized least squares estimators is compromised when the data sources are strategic.

In [26], we consider the problem of computing a sequence of rankings that maximizes consumer-side utility while minimizing producer-side individual unfairness of exposure. While prior work has addressed this problem using linear or quadratic programs on bistochastic matrices, such approaches, relying on Birkhoff-von Neumann (BvN) decompositions, are too slow to be implemented at large scale. In this paper we introduce a geometrical object, a polytope that we call *expohedron*, whose points represent all achievable exposures of items for a Position Based Model (PBM). We exhibit some of its properties and lay out a Carathéodory decomposition algorithm with complexity  $O(n^2 \log(n))$  able to express any point inside the expohedron as a convex sum of at most  $n$  vertices, where  $n$  is the number of items to rank. Such a decomposition makes it possible to express any feasible target exposure as a distribution over at most  $n$  rankings. Furthermore we show that we can use this polytope to recover the whole Pareto frontier of the multi-objective fairness-utility optimization problem, using a simple geometrical procedure with complexity  $O(n^2 \log(n))$ . Our approach compares favorably to linear or quadratic programming baselines in terms of algorithmic complexity and empirical runtime and is applicable to any merit that is a non-decreasing function of item relevance. Furthermore our solution can be expressed as a distribution over only  $n$  permutations, instead of the  $(n - 1)^2 + 1$  achieved with BvN decompositions. We perform experiments on synthetic and real-world datasets, confirming our theoretical results.

In recent years, it has become clear that rankings delivered in many areas need not only be useful to the users but also respect fairness of exposure for the item producers. We consider the problem of finding

ranking policies that achieve a Pareto-optimal tradeoff between these two aspects. Several methods were proposed to solve it; for instance a popular one is to use linear programming with a Birkhoffvon Neumann decomposition. These methods, however, are based on a classical Position Based exposure Model (PBM), which assumes independence between the items (hence the exposure only depends on the rank). In many applications, this assumption is unrealistic and the community increasingly moves towards considering other models that include dependencies, such as the Dynamic Bayesian Network (DBN) exposure model. For such models, computing (exact) optimal fair ranking policies remains an open question. In [27], we answer this question by leveraging a new geometrical method based on the so-called *expohedron* proposed recently for the PBM. We lay out the structure of a new geometrical object (the *DBN-expohedron*), and propose for it a Carathéodory decomposition algorithm of complexity  $O(n^3)$ , where  $n$  is the number of documents to rank. Such an algorithm enables expressing any feasible expected exposure vector as a distribution over at most  $n$  rankings; furthermore we show that we can compute the whole set of Pareto-optimal expected exposure vectors with the same complexity  $O(n^3)$ . Our work constitutes the first exact algorithm able to efficiently find a Pareto-optimal distribution of rankings. It is applicable to a broad range of fairness notions, including classical notions of meritocratic and demographic fairness. We empirically evaluate our method on the TREC2020 and MSLR datasets and compare it to several baselines in terms of Pareto-optimality and speed.

## 9 Bilateral contracts and grants with industry

**Participants:** Henry-Joseph Audéoud, Till Kletti, Nicolas Gast, Patrick Loiseau.

Patrick Loiseau has a Cifre contract with Naver labs (2020-2023) on "Fairness in multi-stakeholder recommendation platforms", which supports the PhD student Till Kletti.

Nicolas Gast obtained a grant from Enedis to evaluate the performance of the PLC-G3 protocol. This grant supported the post-doc of Henry-Joseph Audeoud.

## 10 Partnerships and cooperations

### 10.1 International initiatives

#### 10.1.1 Inria associate team not involved in an IIL or an international program

**Participants:** Vincent Danjean, Guillaume Huard, Arnaud Legrand, Lucas Leandro Nesi, Jean-Marc Vincent.

#### ReDaS

**Title:** Reproducible Data Science

**Duration:** 2019 - 2022

**Coordinator:** Lucas Mello Schnorr (schnorr@inf.ufrgs.br)

#### Partners:

- Universidade Federal do Rio Grande do Sul Porto Alegre (Brésil)

**Inria contact:** Guillaume Huard

**Summary:** Data science builds on a variety of techniques 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.

## 10.2 International research visitors

### 10.2.1 Visits of international scientists

#### **Anshul Gandhi**

**Status** Full Professor

**Institution of origin:** Stony Brook University

**Country:** USA

**Dates:** 20 - 29 Sep.

**Context of the visit:** Collaboration with Nicolas Gast

**Mobility program/type of mobility:** Research stay

#### **Nicolas Maillard**

**Status** Full Professor

**Institution of origin:** Univ. Federale do Rio Grande do Sul

**Country:** Brazil

**Dates:** 16 - 26 Oct.

**Context of the visit:** ReDAS associated team

**Mobility program/type of mobility:** Research stay

#### **Bryce Ferguson**

**Status** PhD Student

**Institution of origin:** Univ. of California, Santa Barbara

**Country:** USA

**Dates:** 9 Oct- 23 Nov.

**Context of the visit:** Collaboration with Panayotis Mertikopoulos

**Mobility program/type of mobility:** Research stay

#### **Swann Perarnau**

**Status** Research scientist

**Institution of origin:** Argonne National Laboratory

**Country:** USA

**Dates:** 1 May - 10 July

**Context of the visit:** Joint Laboratory on Extreme Scale Computing

**Mobility program/type of mobility:** Research stay

### 10.2.2 Visits to international teams

**Research stays abroad**

### Panayotis Mertikopoulos

**Visited institution:** Simons Institute for the Theory of Computing, Berkeley, CA

**Country:** USA

**Dates:** March-April 2022

**Context of the visit:** Visiting scientist

**Mobility program/type of mobility:** Research stay

## 10.3 European initiatives

### 10.3.1 Other european programs/initiatives

**Participants:** Arnaud Legrand.

**Unite!** Arnaud Legrand is involved in the WP6 (Open Science) of the Unite! (University Network for Innovation, Technology and Engineering) project, which aims to create a large European campus from Finland to Portugal. Unite! brings together 7 partners, recognized for the quality of their education and research: Technische Universität Darmstadt (Germany), Aalto University (Finland), Kunglia Tekniska Hoegskolan (Sweden), Politecnico di Torino (Italy), Universitat Politecnica de Catalunya (Spain), Universidade de Lisboa (Portugal) and Grenoble INP, Graduate Schools of Engineering and Management, Université Grenoble Alpes.

## 10.4 National initiatives

Projects indicated with a ★ are projects coordinated by members of the POLARIS team.

### ANR

**ANR ALIAS (PRCI 2020-2023)★** *Adaptive Learning for Interactive Agents and Systems* [284K€]

**Partners:** Singapore University of Technology and Design (SUTD).

ALIAS is a bilateral PRCI (collaboration internationale) project joint with Singapore University of Technology and Design (SUTD), coordinated by Bary Pradelski (PI) and involving P. Mertikopoulos and P. Loiseau. The Singapore team consists of G. Piliouras and G. Panageas. The goal of the project is to provide a unified answer to the question of stability in multi-agent systems: for systems that can be controlled (such as programmable machine learning models), prescriptive learning algorithms can steer the system towards an optimum configuration; for systems that cannot (e.g., online assignment markets), a predictive learning analysis can determine whether stability can arise in the long run. We aim at identifying the fundamental limits of learning in multi-agent systems and design novel, robust algorithms that achieve convergence in cases where conventional online learning methods fail.

**ANR REFINO (JCJC 2020-2024)★** *Refined Mean Field Optimization* [250K€]

REFINO is an ANR starting grant (JCJC) coordinated by Nicolas Gast. The main objective on this project is to provide an innovative framework for optimal control of stochastic distributed agents. Restless bandit allocation is one particular example where the control that can be sent to each arm is restricted to an on/off signal. The originality of this framework is the use of refined mean field approximation to develop control heuristics that are asymptotically optimal as the number of arms goes to infinity and that also have a better performance than existing heuristics for a moderate number of arms. As an example, we will use this framework in the context of smart grids, to develop control policies for distributed electric appliances.

**ANR FAIRPLAY (JCJC 2021-2025)★** *Fair algorithms via game theory and sequential learning* [245K€]  
 FAIRPLAY is an ANR starting grant (JCJC) coordinated by Patrick Loiseau. Machine learning algorithms are increasingly used to optimize decision making in various areas, but this can result in unacceptable discrimination. The main objective of this project is to propose an innovative framework for the development of learning algorithms that respect fairness constraints. While the literature mostly focuses on idealized settings, the originality of this framework and central focus of this project is the use of game theory and sequential learning methods to account for constraints that appear in practical applications: strategic and decentralized aspects of the decisions and the data provided and absence of knowledge of certain parameters key to the fairness definition.

## IRS/UGA

**UGA MIAI Chaire (2019-2023)★** [365K€] Patrick Loiseau is in charge of the *Explainable and Responsible AI* chaire of the MIAI institute. To build more trustworthy AI systems, we investigate how to produce explanations for results returned by AI systems and how to build AI algorithms with guarantees of fairness and privacy, in the setting of varied tasks such as classification, recommendation, resource allocation or matching.

**IRS DISCMAN (IRS 2020-2022)★** (Distributed Control for Multi-Agent systems and Networks) is a joint IRS project funded by IDEX Université Grenoble-Alpes. Its main objectives is to develop distributed equilibrium convergence algorithms for large-scale control and optimization problems, both offline and online. It is being coordinated by Panayotis Mertikopoulos (POLARIS), and it involves a joint team of researchers from the LIG and LJK laboratories in Grenoble.

# 11 Dissemination

## 11.1 Promoting scientific activities

### 11.1.1 Scientific events: organisation

- Bruno Gaujal, Nicolas Gast, and Annie Simon have organized the **12th Atelier d'Évaluation de Performance** in Grenoble, France, July 2022, which has attracted about 70 researchers of the domain.
- Arnaud Legrand has organized the **3rd Workshop of the LIG SRCPR Axis**, Grenoble, June 2022, which is an internal event to the LIG laboratory.
- Panayotis Mertikopoulos has been co-organizer of the 7th workshop on "Stochastic Methods in Game Theory" in Erice, Italy, May 2022.

**General chair, scientific chair** Panayotis Mertikopoulos has been area Chair for ICLR 2022 and NeurIPS 2022.

### Member of the conference program committees

- Jonatha Anselmi has been a member of the IEEE MASCOTS 2022 Program Committee.
- Nicolas Gast has been a member of the Sigmetrics 2022 and ICLR 2023 Program Committees
- Bruno Gaujal has been a member of the NeurIPS 2022 Program Committee.
- Arnaud Legrand has been a member of the Europar 2022 Program Committee.

### Member of the editorial boards

- Nicolas Gast is member of the editorial boards of the journals "Performance Evaluation" and "Stochastic Models".
- Panayotis Mertikopoulos has been Guest editor of the special issue on "Optimization Challenges in Data Science" for the EURO Journal on Computational Optimization (EJCO).
- Panayotis Mertikopoulos has been guest editor of special issue on "Population Games and Evolutionary Dynamics in Memory of William H. Sandholm" for the Journal of Dynamics and Games (JDG).

### 11.1.2 Invited talks

- Arnaud Legrand has been invited to present his latest results at the "15th Scheduling for Large Scale Systems" workshop in Fréjus, June 2022 and at the Scheduling workshop in Aussois, June 2022.
- Romain Couillet has been invited to present "Pourquoi et comment démanteler l'IA et le numérique?" at many seminars (with up to 120 participants) and round tables.
- Nicolas Gast has been invited to present his latest results to the GDR COSMOS, at Séminaire Inria Paris, at the DDQC workshop, and at the Stochastic Networks conference.
- Nicolas Gast has been invited to present his latest results at the Workshop scheduling, Fréjus, June 2022.
- Nicolas Gast has been invited to present his latest results at the Gipsa Lab seminar, Grenoble, March 2022.
- Nicolas Gast has been invited to present his latest results at the MASCOTS conference, Nice, October 2022.
- Panayotis Mertikopoulos has given "A crash course in optimization for machine learning" at the Archimedes research center for Artificial Intelligence and Data Science, Athens, June-July 2022.
- Panayotis Mertikopoulos has been invited to present his latest results "On the limits – and limitations – of learning in games" at the Eccellenza workshop on algorithmic game theory, mechanism design, and learning, Turin, IT, Nov. 2022
- Panayotis Mertikopoulos has been invited to present his latest results "Equilibrium and optimality under uncertainty" at the Amazon Science Summit, Barcelona, ES, September 2022
- Panayotis Mertikopoulos has been invited to present his latest results "The dynamics of artificial intelligence" at the Second Congress of Greek Mathematicians, Athens, GR, July 2022
- Panayotis Mertikopoulos has been invited to present his latest results "The limits of game-theoretic learning" at the Controversies in Game Theory Symposium, ETH Zürich, CH, June 2022
- Panayotis Mertikopoulos has been invited to present his latest results "Regularized learning in games" at the Learning with Strategic Agents Keynote at AAMAS 2022, Auckland, NZ, June 2022
- Panayotis Mertikopoulos has been invited to present his latest results "The limits of regularized learning", at the Learning in the Presence of Strategic Behavior workshop, Berkeley, CA, April 2022
- Panayotis Mertikopoulos has been invited to present his latest results "Min-max optimization from a dynamical systems viewpoint" at the Adversarial Approaches in Machine Learning workshop, Berkeley, CA, March 2022
- Panayotis Mertikopoulos has been invited to present his latest results "The long-run limit of online learning in games" at the Purdue University, Lafayette, IN, September 2022

- Panayotis Mertikopoulos has been invited to present his latest results "The evolution of learning in games" at the University of Athens, Athens, GR, March 2022
- Bary Pradelski has been invited to present his experience on COVID policy at Bruegel, at the OECD Economics Brown Bag Seminar, at Terra Nova, at the Institute for Interdisciplinary Innovation in healthcare (Université Libre de Bruxelles) and at the Institut Pasteur.
- Bary Pradelski has been invited to present his latest research results at the Multidisciplinary Institute in Artificial Intelligence (Univ. Grenoble Alpes) and at the Séminaire parisien de Théorie des Jeux.

### 11.1.3 Scientific expertise

- Barry Pradelski was appointed to share his expertise on COVID policy to the *Centre européen de prévention et de contrôle des maladies* (ECDC).
- Jean-Marc Vincent is vice-head of the SIF, adjunct on teaching.
  - He participated to the coordination of the *trophées NSI* and has been a member of the national committee.
  - He also co-organised the *teaching days* of the SIF
  - He participated to the mediation school organized by the SIF and the Blaise Pascal foundation.
- Jean-Marc Vincent has been a member of the NSI (high school computer science teachers) CAPES committee.

### 11.1.4 Research administration

- Arnaud Legrand is a member of the Section 6 of the CoNRS.
- Arnaud Legrand is head of the SRCPR axis of the LIG and a member of LIG bureau.
- Arnaud Legrand is a member of *Comité Scientifique* of the Inria Grenoble.
- Nicolas Gast has been a member of the hiring committee of a *Maître de Conférences* at Centrale/Supélec.
- Arnaud Legrand has been a member of the hiring committee of a *Professeur* at University of Bordeaux.

## 11.2 Teaching - Supervision - Juries

### 11.2.1 Teaching

- Jonatha Anselmi teaches the *Probability and Simulation* and the *Performance Evaluation* lectures in M1, Polytech Grenoble.
- Arnaud Legrand and Jean-Marc Vincent teach the transversal *Scientific Methodology and Empirical Evaluation* lecture (36h) at the M2 MOSIG.
- Bruno Gaujal teaches the M2 course on *Optimization under Uncertainties* in M2 ORCO Grenoble
- Bruno Gaujal taught the *Reinforcement Learning* course at the special week of the computer science department a ENS Lyon
- Nicolas Gast is responsible of the course *Reinforcement Learning* in Master MOSIG/MSIAM (Grenoble) and of the course « Introduction to machine learning » (License 3, Grenoble).
- Guillaume Huard is responsible of the courses *UNIX & C programming* in the L1 and L3 INFO, of *Object Oriented and Event-Driven Programming* in the L3 INFO, and of the *Objet Oriented Design* in M1 INFO.

- Vincent Danjean teaches the *Operating Systems, Programming Languages, Algorithms, Computer Science and Mediation* lectures in L3, M1 and Polytech Grenoble.
- Romain Couillet is the initiator of a new lecture on the *Introduction to Artificial Intelligence* in the L1 INFO.
- Romain Couillet has created an *Écosophie* workshop with Yoan Svejcar (écopsychologue), which has been taught 4 times, in particular in the context of the *Penser la crise écologique* transversal lecture (Aurélien Barrau, Éléonore Cartelier) and in the PISTE program of Grenoble-INP.
- Romain Couillet has created and managed the continuing education program *Transformation numérique* in collaboration with Denis Trystram (5 students this year, including 2 academics)
- Romain Couillet has created the *Incarner le changement: le numérique à l'ère des lowtechs* lecture in M1 MoSIG (practical sessions world dynamics, "fresque du numérique", ecosophy workshop, surveying Illich's "conviviality", etc.)
- Panayotis Mertikopoulos teaches the *Online Optimization and Learning in Games* M2 lecture at Univ. Limoges
- Panayotis Mertikopoulos teaches the *Reinforcement and Online Learning* in M2 MOSIG/MSIAM (Grenoble)
- The 3rd edition of the MOOC of Arnaud Legrand, K. Hinsén and C. Pouzat on *Reproducible Research: Methodological Principles for a Transparent Science* is still running. Over the 3 editions (Oct.-Dec. 2018, Apr.-June 2019, March 2020 - end of 2023), about 18,800 persons have followed this MOOC and 1900 certificates of achievement have been delivered. 54% of participants are PhD students and 12% are undergraduates.
- Jean-Marc Vincent teaches *Algorithms and Probabilities* at the L3, UGA.
- Jean-Marc Vincent participates to the *Histoire de l'informatique* lecture at the ENSIMAG.

### 11.2.2 Supervision

- Arnaud Legrand has been a member of the *Comité de Suivi Individuel* of Amaury Maillé (ENS Lyon)
- Arnaud Legrand has been a member of the *Comité de Suivi Individuel* of Yishu Du (ENS Lyon)
- Arnaud Legrand has been a member of the *Comité de Suivi Individuel* of Adeyemi Adetula (Univ. Grenoble Alpes)
- Arnaud Legrand has been a member of the *Comité de Suivi Individuel* of Julien Emmanuel (ENS Lyon)
- Panayotis Mertikopoulos is co-supervising the PhD thesis of Waïss Azizian with Jérôme Malick.

### 11.2.3 Juries

- Nicolas Gast has been reviewer for the PhD thesis of Thomas Tournaire (Telecom Sud Paris): *Model-based reinforcement learning for dynamic resource allocation in cloud environments*
- Bruno Gaujal has been reviewer for the PhD thesis of Leonardo Cianfanelli (Polytechnico di Torino): *Learning dynamics in congestion games, intervention in traffic networks, and epidemics control*
- Bruno Gaujal has been reviewer for the PhD thesis of Leonardo Massai (Univ. degli studi di Torino): *Bankruptcy cascades in interbank market*
- Bruno Gaujal has been reviewer for the PhD thesis of Marin Boyet (école polytechnique): *Systèmes dynamiques affines par morceaux appliqués à l'évaluation de performance de centres d'appels d'urgence*

- Bruno Gaujal has been reviewer for the PhD thesis of Guilherme Espindola-Winck (Univ. Angers): *On the stochastic filtering of max-plus linear systems*
- Arnaud Legrand has been reviewer for the PhD thesis of Philippe Swartvagher (Univ. Bordeaux): *On the Interactions between HPC Task-based Runtime Systems and Communication Libraries*
- Arnaud Legrand has been member of the PhD thesis committee of Mathieu Vérité (Univ. Bordeaux): *Algorithmes d'allocation statique pour la planification d'applications haute performance*
- Arnaud Legrand has been member of the PhD thesis committee of Alexis Colin (Télécom Sud Paris): *De la collecte de trace à la prédiction du comportement d'applications parallèles*
- Nicolas Gast has been a member of the committee of the *prix de thèse Paul Casseau*
- Panayotis Mertikopoulos has been reviewer for the PhD thesis of Geovani Rizk (Univ. Paris-Dauphine): *Stochastic graphical bilinear bandits*
- Panayotis Mertikopoulos has been reviewer for the PhD thesis of Laurent Meunier (Univ. Paris-Dauphine): *Adversarial attacks: A theoretical journey*
- Panayotis Mertikopoulos has been reviewer for the PhD thesis of Juliette Achddou (École Normale Supérieure): *Zeroth-order optimization for real-time bidding: A mathematical perspective*

## 11.3 Popularization

### 11.3.1 Articles and contents

- Romain Couillet and his PhD student Achille Baucher have designed an introduction to “World dynamics and limits to the growth” as a Python practical session which has been taught in several occasions (in the PISTE program of Grenoble INP and in several awareness-raising groups).

### 11.3.2 Education

- Jean-Marc Vincent is particularly involved in the teaching of computer science at the national level.
  - He co-animates the national group *infosansordi* (unplugged computer science).
  - He participates to the training of high school computer science teachers (NSI) and he is preparing a new training on programming methods.
  - He also co-organizes series of seminars for high school computer science teachers.
  - He is a member of the national *inter-IREM* group on computer science and writes leaflets for high school computer science teachers.
  - He contributes to the MOOC « Enseigner l'Informatique au Lycée » and is responsible of the algorithms part.
- Jean-Marc Vincent has participated to the *Module d'Initiative Nationale: École inclusive et outils numériques au service des apprentissages*, in November 2022.

### 11.3.3 Interventions

- Arnaud Legrand has presented Open Science and Reproducible Research Challenges at the Université Inter Ages du Dauphiné, May 2022, at the séminaire des élèves de l'ENS Lyon, October 2022, and at the M1 MOSIG students, Univ. Grenoble Alpes, December 2022.
- Jean-Marc Vincent has organized 4 conferences at the *Kateb Yacine* library for general audience: *Le numérique en questions*, with Inria, UGA and the Grenoble city.
- Romain Couillet has organized a "Ingénieur lowtech" workshop at the *forum des métiers* of the *Lycée du Grésivaudan à Meylan*.

- Romain Couillet connects his professional activity with public action: Lowtechlab de Grenoble, Université Autogérée, Arche des Innovateurs, etc.

## 12 Scientific production

### 12.1 Publications of the year

#### International journals

- [1] S. Allmeier and N. Gast. ‘Mean Field and Refined Mean Field Approximations for Heterogeneous Systems: It Works!’ In: *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 6.1 (24th Feb. 2022), pp. 1–43. DOI: [10.1145/3508033](https://doi.org/10.1145/3508033). URL: <https://hal.inria.fr/hal-03600672>.
- [2] J. Anselmi. ‘Replication vs speculation for load balancing’. In: *Queueing Systems* 100.3 (Apr. 2022), pp. 389–391. DOI: [10.1007/s11134-022-09809-z](https://doi.org/10.1007/s11134-022-09809-z). URL: <https://hal.archives-ouvertes.fr/hal-03859228>.
- [3] T. Cornebize and A. Legrand. ‘Simulation-based Optimization and Sensibility Analysis of MPI Applications: Variability Matters’. In: *Journal of Parallel and Distributed Computing* (20th Apr. 2022). DOI: [10.1016/j.jpdc.2022.04.002](https://doi.org/10.1016/j.jpdc.2022.04.002). URL: <https://hal.inria.fr/hal-03141988>.
- [4] B. Donassolo, A. Legrand, P. Mertikopoulos and I. Fajjari. ‘Online Reconfiguration of IoT Applications in the Fog: The Information-Coordination Trade-off’. In: *IEEE Transactions on Parallel and Distributed Systems* 33.5 (2022), pp. 1156–1172. DOI: [10.1109/TPDS.2021.3097281](https://doi.org/10.1109/TPDS.2021.3097281). URL: <https://hal.inria.fr/hal-02636987>.
- [5] B. Duvocelle, P. Mertikopoulos, M. Staudigl and D. Vermeulen. ‘Multi-agent online learning in time-varying games’. In: *Mathematics of Operations Research* (2022). DOI: [10.1287/moor.2022.1283](https://doi.org/10.1287/moor.2022.1283). URL: <https://hal.inria.fr/hal-01891545>.
- [6] N. Gast. ‘Why (and When) do Asymptotic Methods Work so well?’ In: *Queueing Systems* (May 2022), pp. 1–3. DOI: [10.1007/s11134-022-09834-y](https://doi.org/10.1007/s11134-022-09834-y). URL: <https://hal.inria.fr/hal-03638310>.
- [7] B. Gaujal. ‘Learning in Queues’. In: *Queueing Systems* 100 (1st Apr. 2022), pp. 521–523. DOI: [10.1007/s11134-022-09806-2](https://doi.org/10.1007/s11134-022-09806-2). URL: <https://hal.inria.fr/hal-03850698>.
- [8] J. H. d. M. Goulart, R. Couillet and P. Comon. ‘A Random Matrix Perspective on Random Tensors’. In: *Journal of Machine Learning Research* 23.264 (Sept. 2022), pp. 1–36. URL: <https://hal.science/hal-03793940>.
- [9] S. Hadikhanloo, R. Laraki, P. Mertikopoulos and S. Sorin. ‘Learning in nonatomic games, Part I: Finite action spaces and population games’. In: *Journal of Dynamics and Games* 9.4 (Oct. 2022), pp. 433–460. DOI: [10.3934/jdg.2022018](https://doi.org/10.3934/jdg.2022018). URL: <https://hal.inria.fr/hal-03342992>.
- [10] Y.-G. Hsieh, F. Iutzeler, J. Malick and P. Mertikopoulos. ‘Multi-agent online optimization with delays: Asynchronicity, adaptivity, and optimism’. In: *Journal of Machine Learning Research* 23.78 (2022), pp. 1–49. URL: <https://hal.inria.fr/hal-03410422>.
- [11] H. Lebeau, R. Couillet and F. Chatelain. ‘A Random Matrix Analysis of Data Stream Clustering: Coping With Limited Memory Resources’. In: *Proceedings of Machine Learning Research* (June 2022), pp. 1–29. URL: <https://hal.archives-ouvertes.fr/hal-03755939>.
- [12] P. Mertikopoulos and Y. Viossat. ‘Survival of dominated strategies under imitation dynamics’. In: *Journal of Dynamics and Games* 9.4 (Oct. 2022), pp. 499–528. DOI: [10.3934/jdg.2022021](https://doi.org/10.3934/jdg.2022021). URL: <https://hal.archives-ouvertes.fr/hal-03873966>.
- [13] M. C. Miletto, L. L. Nesi, L. Mello Schnorr and A. Legrand. ‘Performance Analysis of Irregular Task-Based Applications on Hybrid Platforms: Structure Matters’. In: *Future Generation Computer Systems* 135 (1st Oct. 2022). URL: <https://hal.inria.fr/hal-03298021>.

- [14] M. Oliu-Barton, B. Pradelski, Y. Algan, M. Baker, A. Binagwaho, G. Dore, A. El-Mohandes, A. Fontanet, A. Peichl, V. Priesemann, G. Wolff, G. Yamey and J. Lazarus. ‘Elimination versus mitigation of SARS-CoV-2 in the presence of effective vaccines’. In: *The Lancet global health* 10.1 (Jan. 2022), e142–e147. DOI: [10.1016/S2214-109X\(21\)00494-0](https://doi.org/10.1016/S2214-109X(21)00494-0). URL: <https://hal.science/hal-03934247>.
- [15] Z. Zhou, P. Mertikopoulos, N. Bambos, P. W. Glynn and Y. Ye. ‘Distributed stochastic optimization with large delays’. In: *Mathematics of Operations Research* 47.3 (Aug. 2022), pp. 2082–2111. DOI: [10.1287/moor.2021.1200](https://doi.org/10.1287/moor.2021.1200). URL: <https://hal.inria.fr/hal-03342384>.
- [16] S. Zrigui, R. y. de Camargo, A. Legrand and D. Trystram. ‘Improving the Performance of Batch Schedulers Using Online Job Runtime Classification’. In: *Journal of Parallel and Distributed Computing* 164 (2nd Feb. 2022), pp. 83–95. DOI: [10.1016/j.jpdc.2022.01.003](https://doi.org/10.1016/j.jpdc.2022.01.003). URL: <https://hal.archives-ouvertes.fr/hal-03023222>.

### International peer-reviewed conferences

- [17] J. Anselmi and B. Gaujal. ‘Energy Optimal Activation of Processors for the Execution of a Single Task with Unknown Size’. In: 30th International Symposium on the Modeling, Analysis, and Simulation of Computer and Telecommunication Systems. Nice, France, 18th Oct. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03682485>.
- [18] K. Antonakopoulos, P. Mertikopoulos, G. Piliouras and X. Wang. ‘AdaGrad avoids saddle points’. In: ICML 2022 - 39th International Conference on Machine Learning. Baltimore, United States, 2022, pp. 1–41. URL: <https://hal.archives-ouvertes.fr/hal-03874036>.
- [19] K. Antonakopoulos, D. Q. Vu, V. Cevher, K. Y. Levy and P. Mertikopoulos. ‘UnderGrad: A universal black-box optimization method with almost dimension-free convergence rate guarantees’. In: ICML 2022 - 39th International Conference on Machine Learning. Baltimore, United States, 2022, pp. 1–31. URL: <https://hal.archives-ouvertes.fr/hal-03874041>.
- [20] E. V. Belmega, P. Mertikopoulos and R. Negrel. ‘Online convex optimization in wireless networks and beyond: The feedback -performance trade-off’. In: RAWNET 2022 - International Workshop on Resource Allocation and Cooperation in Wireless Networks. Turin, Italy, 2022, pp. 1–8. URL: <https://hal.archives-ouvertes.fr/hal-03737125>.
- [21] M. Costantini, N. Liakopoulos, P. Mertikopoulos and T. Spyropoulos. ‘Pick your neighbor: Local Gauss-Southwell rule for fast asynchronous decentralized optimization’. In: CDC 2022 - 61st IEEE Annual Conference on Decision and Control. Cancun, Mexico, 2022. URL: <https://hal.archives-ouvertes.fr/hal-03790278>.
- [22] V. Emelianov, N. Gast and P. Loiseau. ‘Fairness in Selection Problems with Strategic Candidates’. In: EC 2022 - ACM Conference on Economics and Computation. Boulder, Colorado, United States: ACM, 11th July 2022, pp. 1–29. DOI: [10.1145/3490486.3538287](https://doi.org/10.1145/3490486.3538287). URL: <https://hal.inria.fr/hal-03677966>.
- [23] A. Giannou, K. Lotidis, P. Mertikopoulos and E.-V. Vlatakis-Gkaragkounis. ‘On the convergence of policy gradient methods to Nash equilibria in general stochastic games’. In: NeurIPS 2022 - 36th International Conference on Neural Information Processing Systems. New Orleans, United States, 2022, pp. 1–43. URL: <https://hal.archives-ouvertes.fr/hal-03874018>.
- [24] Y.-G. Hsieh, K. Antonakopoulos, V. Cevher and P. Mertikopoulos. ‘No-Regret Learning in Games with Noisy Feedback: Faster Rates and Adaptivity via Learning Rate Separation’. In: NeurIPS 2022 - 36th International Conference on Neural Information Processing Systems. New Orleans, United States, 2022. URL: <https://hal.science/hal-03694134>.
- [25] M. R. Karimi, Y.-P. Hsieh, P. Mertikopoulos and A. Krause. ‘The dynamics of Riemannian Robbins-Monro algorithms’. In: COLT 2022 - 35th Annual Conference on Learning Theory. London, United Kingdom, 2022, pp. 1–31. URL: <https://hal.archives-ouvertes.fr/hal-03874052>.

- [26] T. Kletti, J.-M. Renders and P. Loiseau. ‘Introducing the Expohedron for Efficient Pareto-optimal Fairness-Utility Amortizations in Repeated Rankings’. In: WSDM 2022 - 15th ACM International Conference on Web Search and Data Mining. Phoenix (virtual), United States: ACM, Feb. 2022, pp. 1–10. DOI: [10.1145/3488560.3498490](https://doi.org/10.1145/3488560.3498490). URL: <https://hal.inria.fr/hal-03551061>.
- [27] T. Kletti, J.-M. Renders and P. Loiseau. ‘Pareto-Optimal Fairness-Utility Amortizations in Rankings with a DBN Exposure Model’. In: SIGIR 2022 - 45th International ACM SIGIR Conference on Research and Development in Information Retrieval. Madrid, Spain: ACM, 11th July 2022, pp. 1–12. DOI: [10.1145/3477495.3532036](https://doi.org/10.1145/3477495.3532036). URL: <https://hal.inria.fr/hal-03691743>.
- [28] H. Lebeau, R. Couillet and F. Chatelain. ‘Une analyse par matrices aléatoires du clustering en ligne : comprendre l’impact des limitations en mémoire’. In: GRETSI 2022 - XXVIIIème Colloque Francophone de Traitement du Signal et des Images. Nancy, France, 6th Sept. 2022, pp. 1–4. URL: <https://hal.archives-ouvertes.fr/hal-03756033>.
- [29] K. Lotidis, P. Mertikopoulos and N. Bambos. ‘Learning in games with quantized payoff observations’. In: CDC 2022 - 61st IEEE Annual Conference on Decision and Control. Cancun, Mexico: IEEE, 2022, pp. 1–8. URL: <https://hal.archives-ouvertes.fr/hal-03874022>.
- [30] M. Martin, P. Mertikopoulos, T. Rahier and H. Zenati. ‘Nested bandits’. In: ICML 2022 - 39th International Conference on Machine Learning. Baltimore, United States, 17th July 2022. URL: <https://hal.archives-ouvertes.fr/hal-03874048>.
- [31] M. Molina and P. Loiseau. ‘Bounding and Approximating Intersectional Fairness through Marginal Fairness’. In: NeurIPS 2022 - 36th Conference on Neural Information Processing Systems. New Orleans, United States, 28th Nov. 2022, pp. 1–32. URL: <https://hal.inria.fr/hal-03827777>.
- [32] L. Nesi, L. Mello Schnorr and A. Legrand. ‘Multi-Phase Task-Based HPC Applications: Quickly Learning how to Run Fast’. In: IPDPS 2022 - 36th IEEE International Parallel & Distributed Processing Symposium. Lyon, France: IEEE, 30th May 2022, pp. 1–11. URL: <https://hal.inria.fr/hal-03608579>.

#### Conferences without proceedings

- [33] J. Anselmi, B. Gaujal and L.-S. Rebuffi. ‘Reinforcement Learning in a Birth and Death Process: Breaking the Dependence on the State Space’. In: NeurIPS 2022 - 36th Conference on Neural Information Processing Systems. La Nouvelle Orléans, United States, 28th Nov. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03799394>.
- [34] B. Roussillon, N. Gast, P. Loiseau and P. Mertikopoulos. ‘Asymptotic Degradation of Linear Regression Estimates with Strategic Data Sources’. In: ALT 2022 - 33rd International Conference on Algorithmic Learning Theory. Paris, France, 29th Mar. 2022, pp. 1–31. URL: <https://hal.univ-grenoble-alpes.fr/hal-03593516>.

#### Scientific books

- [35] R. Couillet and Z. Liao. *Random Matrix Methods for Machine Learning*. 1. Cambridge University Press, 30th June 2022. DOI: [10.1017/9781009128490](https://doi.org/10.1017/9781009128490). URL: <https://hal.inria.fr/hal-03946535>.

#### Doctoral dissertations and habilitation theses

- [36] K. Antonakopoulos. ‘Adaptive Algorithms for Optimization Beyond Lipschitz Requirements’. Université Grenoble Alpes [2020-....], 12th Jan. 2022. URL: <https://theses.hal.science/tel-03716178>.
- [37] V. Emelianov. ‘Fairness in selection problems’. Université Grenoble Alpes [2020-....], 13th June 2022. URL: <https://theses.hal.science/tel-03783665>.
- [38] C. Yan. ‘Close-to-optimal policies for Markovian bandits’. Université Grenoble Alpes (UGA), 15th Dec. 2022. URL: <https://hal.inria.fr/tel-03931714>.

## Reports & preprints

- [39] W. Azizian, F. Iutzeler, J. Malick and P. Mertikopoulos. *On the rate of convergence of Bregman proximal methods in constrained variational inequalities*. Nov. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03873992>.
- [40] N. Gast, B. Gaujal and K. Khun. *Reinforcement Learning for Markovian Bandits: Is Posterior Sampling more Scalable than Optimism?* 2022. URL: <https://hal.inria.fr/hal-03262006>.
- [41] N. Gast, B. Gaujal and K. Khun. *Testing Indexability and Computing Whittle and Gittins Index in Subcubic Time*. 8th Dec. 2022. URL: <https://hal.inria.fr/hal-03602458>.
- [42] N. Gast, B. Gaujal and C. Yan. *LP-based policies for restless bandits: necessary and sufficient conditions for (exponentially fast) asymptotic optimality*. 13th Sept. 2022. URL: <https://hal.inria.fr/hal-03262307>.
- [43] Y.-G. Hsieh, Y. Laguel, F. Iutzeler and J. Malick. *Push-Pull with Device Sampling*. 8th June 2022. URL: <https://hal.archives-ouvertes.fr/hal-03691122>.
- [44] T. Lin, P. Mertikopoulos and M. I. Jordan. *Explicit second-order min-max optimization methods with optimal convergence guarantees*. Oct. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03873989>.
- [45] P. Mertikopoulos, Y.-P. Hsieh and V. Cevher. *Learning in games from a stochastic approximation viewpoint*. June 2022. URL: <https://hal.archives-ouvertes.fr/hal-03874012>.

## Other scientific publications

- [46] J. Lazarus, D. Romero, C. Kopka, S. A. Karim, L. Abu-Raddad, G. Almeida, R. Baptista-Leite, J. Barocas, M. Barreto, Y. Bar-Yam et al. ‘A multinational Delphi consensus to end the COVID-19 public health threat’. In: *Nature* 611.7935 (10th Nov. 2022), pp. 332–345. DOI: [10.1038/s41586-022-05398-2](https://doi.org/10.1038/s41586-022-05398-2). URL: <https://hal.science/hal-03934072>.

## 12.2 Other

### Scientific popularization

- [47] J.-M. Vincent and T. Viéville. *Des professionnel-le-s de l’informatique deviennent enseignant-e-s*. 7th Oct. 2022. URL: <https://hal.inria.fr/hal-03844358>.

## 12.3 Cited publications

- [48] J. D. Leshno and B. S. R. Pradelski. ‘The importance of memory for price discovery in decentralized markets’. In: *Games and Economic Behavior* 125 (Jan. 2021), pp. 62–78. DOI: [10.1016/j.geb.2020.11.002](https://doi.org/10.1016/j.geb.2020.11.002). URL: <https://hal.archives-ouvertes.fr/hal-03100097>.
- [49] M. Oliu-Barton, B. Pradelski, P. Aghion, P. Artus, I. Kickbusch, J. Lazarus, D. Sridhar and S. Vander-slott. ‘SARS-CoV-2 elimination, not mitigation, creates best outcomes for health, the economy, and civil liberties’. In: *The Lancet* 397.10291 (June 2021), pp. 2234–2236. DOI: [10.1016/S0140-6736\(21\)00978-8](https://doi.org/10.1016/S0140-6736(21)00978-8). URL: <https://hal.archives-ouvertes.fr/hal-03355143>.
- [50] B. Pradelski and M. Oliu-Barton. ‘Green zoning: An effective policy tool to tackle the Covid-19 pandemic’. In: *Health Policy* 125.8 (Aug. 2021), pp. 981–986. DOI: [10.1016/j.healthpol.2021.06.001](https://doi.org/10.1016/j.healthpol.2021.06.001). URL: <https://hal.archives-ouvertes.fr/hal-03355146>.
- [51] Z. Zhou, P. Mertikopoulos, A. L. Moustakas, N. Bambos and P. W. Glynn. ‘Robust power management via learning and game design’. In: *Operations Research* 69.1 (Jan. 2021), pp. 331–345. DOI: [10.1287/opre.2020.1996](https://doi.org/10.1287/opre.2020.1996). URL: <https://hal.archives-ouvertes.fr/hal-01643340>.
- [52] B. Gaujal, J. Doncel and N. Gast. ‘Vaccination in a Large Population: Mean Field Equilibrium versus Social Optimum’. In: *NETGCOOP 2020 - 10th International Conference on NETwork Games, CONTROL and OPTimization*. Cargèse, France, 22nd Sept. 2021, pp. 1–9. URL: <https://hal.inria.fr/hal-02938850>.

- [53] A. Giannou, E. V. Vlatakis-Gkaragkounis and P. Mertikopoulos. ‘Survival of the strictest: Stable and unstable equilibria under regularized learning with partial information’. In: COLT 2021 - 34th Annual Conference on Learning Theory. Boulder, United States, 15th Aug. 2021, pp. 1–30. URL: <https://hal.inria.fr/hal-03342404>.
- [54] Y.-P. Hsieh, P. Mertikopoulos and V. Cevher. ‘The limits of min-max optimization algorithms: Convergence to spurious non-critical sets’. In: ICML 2021 - 38th International Conference on Machine Learning. Vienna, Austria, July 2021. URL: <https://hal.archives-ouvertes.fr/hal-03043862>.
- [55] L. L. Nesi, A. Legrand and L. Mello Schnorr. ‘Exploiting system level heterogeneity to improve the performance of a GeoStatistics multi-phase task-based application’. In: ICPP 2021 - 50th International Conference on Parallel Processing. Lemont, United States, 9th Aug. 2021, pp. 1–10. DOI: [10.1145/3472456.3472516](https://doi.org/10.1145/3472456.3472516). URL: <https://hal.inria.fr/hal-03280459>.
- [56] P. H. Rocha Bruel. ‘Toward transparent and parsimonious methods for automatic performance tuning’. UGA (Université Grenoble Alpes); USP (Universidade de São Paulo), 9th July 2021. URL: <https://tel.archives-ouvertes.fr/tel-03526483>.
- [57] A. Andreou, M. Silva, F. Benevenuto, O. Goga, P. Loiseau and A. Mislove. ‘Measuring the Facebook Advertising Ecosystem’. In: NDSS 2019 - Proceedings of the Network and Distributed System Security Symposium. San Diego, United States, Feb. 2019, pp. 1–15. DOI: [10.14722/ndss.2019.23280](https://doi.org/10.14722/ndss.2019.23280). URL: <https://hal.archives-ouvertes.fr/hal-01959145>.
- [58] A. Andreou, G. Venkatadri, O. Goga, K. P. Gummadi, P. Loiseau and A. Mislove. ‘Investigating Ad Transparency Mechanisms in Social Media: A Case Study of Facebook’s Explanations’. In: NDSS 2018 - Network and Distributed System Security Symposium. San Diego, United States, Feb. 2018, pp. 1–15. DOI: [10.14722/ndss.2018.23204](https://doi.org/10.14722/ndss.2018.23204). URL: <https://hal.archives-ouvertes.fr/hal-01955309>.
- [59] 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. DOI: [10.1109/TPDS.2019.2950621](https://doi.org/10.1109/TPDS.2019.2950621). URL: <https://hal.archives-ouvertes.fr/hal-02276789>.
- [60] J. Anselmi and J. Doncel. ‘Asymptotically Optimal Size-Interval Task Assignments’. In: IEEE Transactions on Parallel and Distributed Systems 30.11 (2019), pp. 2422–2433. DOI: [10.1109/TPDS.2019.2920121](https://doi.org/10.1109/TPDS.2019.2920121). URL: <https://hal.archives-ouvertes.fr/hal-02318576>.
- [61] J. Anselmi and F. Dufour. ‘Power-of-d-Choices with Memory: Fluid Limit and Optimality’. In: Mathematics of Operations Research 45.3 (2020), pp. 862–888. DOI: [10.1287/moor.2019.1014](https://doi.org/10.1287/moor.2019.1014). URL: <https://hal.archives-ouvertes.fr/hal-02394147>.
- [62] R. M. Badia, J. Labarta, J. Giménez and F. Escalé. ‘Dimemas: Predicting MPI Applications Behaviour in Grid Environments’. In: Proc. of the Workshop on Grid Applications and Programming Tools. 2003.
- [63] S. Böhm and C. Engelmann. ‘xSim: The Extreme-Scale Simulator’. In: HPCS. Istanbul, Turkey, 2011.
- [64] P. Bruel, S. Quinito Masnada, B. Videau, A. Legrand, J.-M. Vincent and A. Goldman. ‘Autotuning under Tight Budget Constraints: A Transparent Design of Experiments Approach’. In: CCGrid 2019 - International Symposium in Cluster, Cloud, and Grid Computing. Larcana, Cyprus: IEEE, May 2019, pp. 1–10. DOI: [10.1109/CCGRID.2019.00026](https://doi.org/10.1109/CCGRID.2019.00026). URL: <https://hal.inria.fr/hal-02110868>.
- [65] H. Brunst, D. Hackenberg, G. Juckeland and H. Rohling. ‘Comprehensive Performance Tracking with VAMPIR 7’. In: Tools for High Performance Computing 2009. Springer Berlin Heidelberg, 2010. DOI: [http://dx.doi.org/10.1007/978-3-642-11261-4\\_2](https://dx.doi.org/10.1007/978-3-642-11261-4_2).
- [66] G. Casale and N. Gast. ‘Performance analysis methods for list-based caches with non-uniform access’. In: IEEE/ACM Transactions on Networking (Dec. 2020), pp. 1–18. DOI: [10.1109/TNET.2020.3042869](https://doi.org/10.1109/TNET.2020.3042869). URL: <https://hal.inria.fr/hal-03102188>.

- [67] A. Chakraborty, G. K. Patro, N. Ganguly, K. P. Gummadi and P. Loiseau. ‘Equality of Voice: Towards Fair Representation in Crowdsourced Top-K Recommendations’. In: *FAT\* 2019 - ACM Conference on Fairness, Accountability, and Transparency*. Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAT\*). Atlanta, United States: ACM, Jan. 2019, pp. 129–138. DOI: [10.1145/3287560.3287570](https://doi.org/10.1145/3287560.3287570). URL: <https://hal.archives-ouvertes.fr/hal-01959135>.
- [68] T. Cornebize, A. Legrand and F. C. Heinrich. ‘Fast and Faithful Performance Prediction of MPI Applications: the HPL Case Study’. In: *2019 IEEE International Conference on Cluster Computing (CLUSTER)*. Albuquerque, United States, Sept. 2019. DOI: [10.1109/CLUSTER.2019.8891011](https://doi.org/10.1109/CLUSTER.2019.8891011). URL: <https://hal.inria.fr/hal-02096571>.
- [69] P. COUCHENEY, B. Gaujal and P. Mertikopoulos. ‘Penalty-Regulated Dynamics and Robust Learning Procedures in Games’. In: *Mathematics of Operations Research* 40.3 (2015), pp. 611–633. DOI: [10.1287/moor.2014.0687](https://doi.org/10.1287/moor.2014.0687). URL: <https://hal.inria.fr/hal-01235243>.
- [70] A. Degomme, A. Legrand, G. Markomanolis, M. Quinson, M. L. Stillwell and F. Suter. ‘Simulating MPI applications: the SMPI approach’. In: *IEEE Transactions on Parallel and Distributed Systems* 28.8 (Aug. 2017), p. 14. DOI: [10.1109/TPDS.2017.2669305](https://doi.org/10.1109/TPDS.2017.2669305). URL: <https://hal.inria.fr/hal-01415484>.
- [71] B. Donassolo, I. Fajjari, A. Legrand and P. Mertikopoulos. ‘Load Aware Provisioning of IoT Services on Fog Computing Platform’. In: *IEEE International Conference on Communications (ICC)*. Shanghai, China: IEEE, May 2019. DOI: [10.1109/ICC.2019.8762052](https://doi.org/10.1109/ICC.2019.8762052). URL: <https://hal.inria.fr/hal-02056743>.
- [72] J. Doncel, N. Gast and B. Gaujal. ‘Are mean-field games the limits of finite stochastic games?’ In: *The 18th Workshop on Mathematical performance Modeling and Analysis*. Nice, France, June 2016. URL: <https://hal.inria.fr/hal-01321020>.
- [73] J. Doncel, N. Gast and B. Gaujal. ‘Discrete Mean Field Games: Existence of Equilibria and Convergence’. In: *Journal of Dynamics and Games* 6.3 (2019), pp. 1–19. DOI: [10.3934/jdg.2019016](https://doi.org/10.3934/jdg.2019016). URL: <https://hal.inria.fr/hal-01277098>.
- [74] V. Emelianov, G. Arvanitakis, N. Gast, K. P. Gummadi and 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, Aug. 2019, pp. 5836–5842. DOI: [10.24963/ijcai.2019/809](https://doi.org/10.24963/ijcai.2019/809). URL: <https://hal.inria.fr/hal-02145071>.
- [75] V. Emelianov, N. Gast, K. P. Gummadi and P. Loiseau. ‘On Fair Selection in the Presence of Implicit Variance’. In: *EC 2020 - Twenty-First ACM Conference on Economics and Computation*. Budapest, Hungary: ACM, July 2020, pp. 649–675. DOI: [10.1145/3391403.3399482](https://doi.org/10.1145/3391403.3399482). URL: <https://hal.inria.fr/hal-02880578>.
- [76] L. Flokas, E. V. Vlatakis-Gkaragkounis, T. Lianeas, P. Mertikopoulos and G. Piliouras. ‘No-regret learning and mixed Nash equilibria: They do not mix’. In: *NeurIPS 2020 - 34th International Conference on Neural Information Processing Systems*. Vancouver, Canada, Dec. 2020, pp. 1–24. URL: <https://hal.inria.fr/hal-03043763>.
- [77] V. Garcia Pinto, L. M. Schnorr, L. Stanisic, A. Legrand, S. Thibault and V. Danjean. ‘A Visual Performance Analysis Framework for Task-based Parallel Applications running on Hybrid Clusters’. In: *Concurrency and Computation: Practice and Experience* 30.18 (Apr. 2018), pp. 1–31. DOI: [10.1002/cpe.4472](https://doi.org/10.1002/cpe.4472). URL: <https://hal.inria.fr/hal-01616632>.
- [78] N. Gast. ‘Expected Values Estimated via Mean-Field Approximation are  $1/N$ -Accurate’. In: *ACM SIGMETRICS / International Conference on Measurement and Modeling of Computer Systems, SIGMETRICS ’17*. Vol. 1. ACM SIGMETRICS / International Conference on Measurement and Modeling of Computer Systems, SIGMETRICS ’17. Urbana-Champaign, United States, June 2017, p. 26. DOI: [10.1145/2896377.2901463](https://doi.org/10.1145/2896377.2901463). URL: <https://hal.inria.fr/hal-01553133>.
- [79] N. Gast, L. Bortolussi and M. Tribastone. ‘Size Expansions of Mean Field Approximation: Transient and Steady-State Analysis’. In: *Performance Evaluation* (2018), pp. 1–15. URL: <https://hal.inria.fr/hal-01891632>.

- [80] N. Gast, B. Gaujal and C. Yan. ‘Exponential Convergence Rate for the Asymptotic Optimality of Whittle Index Policy’. Dec. 2020. URL: <https://hal.inria.fr/hal-03041176>.
- [81] N. Gast and B. V. Houdt. ‘A Refined Mean Field Approximation’. In: *ACM SIGMETRICS 2018*. Irvine, United States, June 2018, p. 1. URL: <https://hal.inria.fr/hal-01891642>.
- [82] N. Gast, S. Ioannidis, P. Loiseau and B. Roussillon. ‘Linear Regression from Strategic Data Sources’. In: *ACM Transactions on Economics and Computation* 8.2 (May 2020), pp. 1–24. DOI: [10.1145/3391436](https://doi.org/10.1145/3391436). URL: <https://hal.inria.fr/hal-02982682>.
- [83] N. Gast, D. Latella and M. Massink. ‘A Refined Mean Field Approximation for Synchronous Population Processes’. In: *MAMA 2018 Workshop on MAThematical performance Modeling and Analysis*. Irvine, United States, June 2018, pp. 1–3. URL: <https://hal.inria.fr/hal-01891629>.
- [84] N. Gast and B. Van Houdt. ‘Asymptotically Exact TTL-Approximations of the Cache Replacement Algorithms LRU(m) and h-LRU’. In: *28th International Teletraffic Congress (ITC 28)*. Würzburg, Germany, Sept. 2016. URL: <https://hal.inria.fr/hal-01292269>.
- [85] N. Gast and B. Van Houdt. ‘TTL Approximations of the Cache Replacement Algorithms LRU(m) and h-LRU’. In: *Performance Evaluation* (Sept. 2017). DOI: [10.1016/j.peva.2017.09.002](https://doi.org/10.1016/j.peva.2017.09.002). URL: <https://hal.inria.fr/hal-01622059>.
- [86] B. Gaujal, A. Girault and 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, Nov. 2019, pp. 1–14. URL: <https://hal.archives-ouvertes.fr/hal-02372136>.
- [87] B. Gaujal, A. Girault and S. Plassart. ‘Discrete and Continuous Optimal Control for Energy Minimization in Real-Time Systems’. In: *EBCCSP 2020 - 6th International Conference on Event-Based Control, Communication, and Signal Processing*. Krakow, Poland: IEEE, Sept. 2020, pp. 1–8. DOI: [10.1109/EBCCSP51266.2020.9291353](https://doi.org/10.1109/EBCCSP51266.2020.9291353). URL: <https://hal.archives-ouvertes.fr/hal-03020330>.
- [88] B. Gaujal, A. Girault and S. Plassart. ‘Dynamic Speed Scaling Minimizing Expected Energy Consumption for Real-Time Tasks’. In: *Journal of Scheduling* (July 2020), pp. 1–25. DOI: [10.1007/s10951-020-00660-9](https://doi.org/10.1007/s10951-020-00660-9). URL: <https://hal.inria.fr/hal-02888573>.
- [89] B. Gaujal, A. Girault and S. Plassart. *Exploiting Job Variability to Minimize Energy Consumption under Real-Time Constraints*. Research Report RR-9300. Inria Grenoble Rhône-Alpes, Université de Grenoble ; Université Grenoble - Alpes, Nov. 2019, p. 23. URL: <https://hal.inria.fr/hal-02371742>.
- [90] M. Heath and J. Etheridge. ‘Visualizing the performance of parallel programs’. In: *IEEE software* 8.5 (1991).
- [91] F. C. Heinrich, T. Cornebize, A. Degomme, A. Legrand, A. Carpen-Amarie, S. Hunold, A.-C. Orgerie and M. Quinson. ‘Predicting the Energy Consumption of MPI Applications at Scale Using a Single Node’. In: *Cluster 2017*. IEEE. Hawaii, United States, Sept. 2017. URL: <https://hal.inria.fr/hal-01523608>.
- [92] T. Hoefler, T. Schneider and A. Lumsdaine. ‘LogGOPSIm - Simulating Large-Scale Applications in the LogGOPS Model’. In: *ACM Workshop on Large-Scale System and Application Performance*. 2010.
- [93] L. V. Kalé, G. Zheng, C. W. Lee and S. Kumar. ‘Scaling applications to massively parallel machines using Projections performance analysis tool’. In: *Future Generation Comp. Syst.* 22.3 (2006).
- [94] R. Keller Tesser, L. Mello Schnorr, A. Legrand, F. Dupros and P. O. A. Navaux. ‘Using Simulation to Evaluate and Tune the Performance of Dynamic Load Balancing of an Over-decomposed Geophysics Application’. In: *Euro-Par 2017: 23rd International European Conference on Parallel and Distributed Computing*. Santiago de Compostela, Spain, Aug. 2017, p. 15. URL: <https://hal.inria.fr/hal-01567792>.

- [95] R. Keller Tesser, L. Mello Schnorr, A. Legrand, C. Heinrich, F. Dupros and P. O. Alexandre Navaux. ‘Performance Modeling of a Geophysics Application to Accelerate the Tuning of Overdecomposition Parameters through Simulation’. In: *Concurrency and Computation: Practice and Experience* (2018), pp. 1–21. DOI: [10.1002/cpe.5012](https://doi.org/10.1002/cpe.5012). URL: <https://hal.inria.fr/hal-01891416>.
- [96] T.-E. Kennouche, F. Cadoux, N. Gast and 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, Dec. 2019, pp. 1–5. URL: <https://hal.archives-ouvertes.fr/hal-02384051>.
- [97] J. M. Kleinberg and M. Raghavan. ‘Selection Problems in the Presence of Implicit Bias’. In: *Proceedings of the 9th Innovations in Theoretical Computer Science Conference (ITCS)*. 2018, 33:1–33:17.
- [98] A. Legrand, D. Trystram and 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://doi.org/10.1109/IPDPS.2019.00077). URL: <https://hal.archives-ouvertes.fr/hal-02044903>.
- [99] M. Mendil, N. Gast and H.-J. Audéoud. ‘Collisions groupées lors du mécanisme d’évitement de collisions de CPL-G3’. In: *CoRes 2020 - Rencontres Francophones sur la Conception de Protocoles, l’Évaluation de Performance et l’Expérimentation des Réseaux de Communication*. Lyon, France, Sept. 2020, pp. 1–4. URL: <https://hal.archives-ouvertes.fr/hal-02879294>.
- [100] P. Mertikopoulos, B. Lecouat, H. Zenati, C.-S. Foo, V. Chandrasekhar and 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. URL: <https://hal.inria.fr/hal-02111937>.
- [101] P. Mertikopoulos, H. Nax and B. Pradelski. ‘Quick or cheap? Breaking points in dynamic markets’. In: *EC 2020 - 21st ACM Conference on Economics and Computation*. Budapest, Hungary, July 2020, pp. 1–32. URL: <https://hal.inria.fr/hal-03043688>.
- [102] P. Mertikopoulos, C. H. Papadimitriou and G. Piliouras. ‘Cycles in adversarial regularized learning’. In: *SODA ’18 - Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*. New Orleans, United States, Jan. 2018, pp. 2703–2717. URL: <https://hal.archives-ouvertes.fr/hal-01643338>.
- [103] P. Mertikopoulos and W. H. Sandholm. ‘Learning in games via reinforcement learning and regularization’. In: *Mathematics of Operations Research* 41.4 (Nov. 2016), pp. 1297–1324. DOI: [10.1287/moor.2016.0778](https://doi.org/10.1287/moor.2016.0778). URL: <https://hal.inria.fr/hal-01073491>.
- [104] P. Mertikopoulos and W. H. Sandholm. ‘Riemannian game dynamics’. In: *Journal of Economic Theory* 177 (Sept. 2018), pp. 315–364. DOI: [10.1016/j.jet.2018.06.002](https://doi.org/10.1016/j.jet.2018.06.002). URL: <https://hal.archives-ouvertes.fr/hal-01382281>.
- [105] M. Minaei, M. Mondal, P. Loiseau, K. P. Gummadi and A. Kate. ‘Forgetting the Forgotten with Lethe: Conceal Content Deletion from Persistent Observers’. In: *PETS 2019 - 19th Privacy Enhancing Technologies Symposium*. Stockholm, Sweden, July 2019, pp. 1–21. URL: <https://hal.archives-ouvertes.fr/hal-01959119>.
- [106] W. Nagel, A. Arnold, M. Weber, H. Hoppe and K. Solchenbach. ‘VAMPIR: Visualization and Analysis of MPI Resources’. In: *Supercomputer* 12.1 (1996).
- [107] M. Oliu-Barton and B. Pradelski. *A vaccination policy by zones*. Other. Think tank Terra Nova, Oct. 2020. URL: <https://hal.archives-ouvertes.fr/hal-03100256>.
- [108] M. Oliu-Barton and B. Pradelski. ‘Green bridges: Reconnecting Europe to avoid economic disaster’. In: *Europe in the Time of Covid-19*. 2020. URL: <https://hal.archives-ouvertes.fr/hal-03100282>.
- [109] V. Pillet, J. Labarta, T. Cortes and S. Girona. ‘PARAVER: A tool to visualise and analyze parallel code’. In: *Proceedings of Transputer and Occam Developments, WOTUG-18*. Vol. 44. 1995.

- [110] B. Pradelski and M. Oliu-Barton. *Focus mass testing: How to overcome low test accuracy*. Other. Esade Centre for Economic Policy, Dec. 2020. URL: <https://hal.archives-ouvertes.fr/hal-03100250>.
- [111] B. S. R. Pradelski and H. H. Nax. ‘Market sentiments and convergence dynamics in decentralized assignment economies’. In: *International Journal of Game Theory* 49.1 (Mar. 2020), pp. 275–298. DOI: [10.1007/s00182-019-00694-0](https://doi.org/10.1007/s00182-019-00694-0). URL: <https://hal.archives-ouvertes.fr/hal-03100116>.
- [112] D. Reed, P. Roth, R. Aydt, K. Shields, L. Tavera, R. Noe and B. Schwartz. ‘Scalable performance analysis: the Pablo performance analysis environment’. In: *Scalable Parallel Libraries Conference*. 1993.
- [113] B. Shneiderman. ‘The eyes have it: A task by data type taxonomy for information visualizations’. In: *IEEE Symposium on Visual Languages*. IEEE. 1996.
- [114] D. C. Snowdon, S. Ruocco and G. Heiser. ‘Power Management and Dynamic Voltage Scaling: Myths and Facts’. In: *Proceedings of the 2005 Workshop on Power Aware Real-time Computing*. New Jersey, USA, Sept. 2005.
- [115] T. Speicher, M. Ali, G. Venkatadri, F. Ribeiro, G. Arvanitakis, F. Benevenuto, K. P. Gummadi, P. Loiseau and A. Mislove. ‘Potential for Discrimination in Online Targeted Advertising’. In: *FAT 2018 - Conference on Fairness, Accountability, and Transparency*. Vol. 81. New-York, United States, Feb. 2018, pp. 1–15. URL: <https://hal.archives-ouvertes.fr/hal-01955343>.
- [116] M. Tikir, M. Laurenzano, L. Carrington and A. Snavey. ‘PSINS: An Open Source Event Tracer and Execution Simulator for MPI Applications’. In: *Euro-Par*. 2009.
- [117] G. Venkatadri, A. Andreou, Y. Liu, A. Mislove, K. P. Gummadi, P. Loiseau and O. Goga. ‘Privacy Risks with Facebook’s PII-based Targeting: Auditing a Data Broker’s Advertising Interface’. In: *39th IEEE Symposium on Security and Privacy (S&P)*. Proceedings of the 39th IEEE Symposium on Security and Privacy (S&P). San Francisco, United States, 2018. URL: <https://hal.archives-ouvertes.fr/hal-01955327>.
- [118] B. Vinot, F. Cadoux and R. Héliot. ‘Decentralized Optimization of Energy Exchanges in an Electricity Microgrid’. In: *ACM e-Energy 2016 - 7th ACM International Conference on Future Energy Systems*. Waterloo, Canada, June 2016. DOI: [10.1145/2939912.2939915](https://doi.org/10.1145/2939912.2939915). URL: <https://hal.archives-ouvertes.fr/hal-01340394>.
- [119] B. Vinot, F. Cadoux and R. Héliot. ‘Decentralized optimization of energy exchanges in an electricity microgrid’. In: *2016 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe)*. Ljubljana, Slovenia: IEEE, Oct. 2016, pp. 1–6. DOI: [10.1109/ISGTEurope.2016.7856319](https://doi.org/10.1109/ISGTEurope.2016.7856319). URL: <https://hal.inria.fr/hal-01784148>.
- [120] B. Vinot, F. Cadoux and N. Gast. ‘Congestion Avoidance in Low-Voltage Networks by using the Advanced Metering Infrastructure’. In: *ePerf 2018 - IFIP WG PERFORMANCE - 36th International Symposium on Computer Performance, Modeling, Measurements and Evaluation*. Toulouse, France, Dec. 2018, pp. 1–3. URL: <https://hal.inria.fr/hal-01953386>.
- [121] M. Weiser, B. Welch, A. Demers and S. Shenker. ‘Scheduling for Reduced CPU Energy’. In: *Proceedings of the 1st USENIX Conference on Operating Systems Design and Implementation*. OSDI ’94. Monterey, California: USENIX Association, 1994, 2–es.
- [122] J. Wilke, K. Sargsyan, J. Kenny, B. Debusschere, H. Najm and G. Hendry. ‘Validation and Uncertainty Assessment of Extreme-Scale HPC Simulation through Bayesian Inference’. In: *Euro-Par*. 2013.
- [123] O. Zaki, E. Lusk, W. Gropp and D. Swider. ‘Toward Scalable Performance Visualization with Jumpshot’. In: *International Journal of High Performance Computing Applications* 13.3 (1999). DOI: <http://dx.doi.org/10.1177/109434209901300310>.
- [124] G. Zheng, G. Kakulapati and L. Kalé. ‘BigSim: A Parallel Simulator for Performance Prediction of Extremely Large Parallel Machines’. In: *IPDPS*. 2004.