

RESEARCH CENTRE

**Inria Centre
at Université Grenoble Alpes**

IN PARTNERSHIP WITH:

Université de Grenoble Alpes, CNRS

2023

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

Project-Team POLARIS	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	3
2.1 Context	3
2.2 Objectives	4
2.3 Contribution to AI/Learning	4
3 Research program	5
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
4 Application domains	12
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
5 Social and environmental responsibility	13
5.1 Footprint of research activities	13
5.2 Raising awareness on the climate crisis	13
5.3 Impact of research results	13
6 Highlights of the year	14
6.1 Awards	14
7 New software, platforms, open data	14
7.1 New software	14
7.1.1 SimGrid	14
7.1.2 PSI	15
7.1.3 marmoteCore	15
7.1.4 MarTO	15
7.1.5 GameSeer	16
7.1.6 rmf_tool	16
8 New results	16
8.1 Performance evaluation of Large Systems	16
8.2 Energy optimization	18
8.3 Restless multi-armed bandits	18
8.4 Reinforcement Learning and MDP	19
8.5 Learning in games	20
8.6 Quantum Games	21
8.7 Continuous optimization methods	22
8.8 Random matrix analysis and Machine Learning	23
8.9 Fairness and equity in digital (recommendation, advertising, persistent storage) systems	23
9 Bilateral contracts and grants with industry	25
10 Partnerships and cooperations	25
10.1 European initiatives	25
10.1.1 Other european programs/initiatives	25
10.2 National initiatives	25

11 Dissemination	26
11.1 Promoting scientific activities	26
11.1.1 Scientific events: organisation	26
11.1.2 Scientific events: selection	27
11.1.3 Journal	27
11.1.4 Invited talks	28
11.1.5 Leadership within the scientific community	29
11.1.6 Research administration	29
11.2 Teaching - Supervision - Juries	30
11.2.1 Teaching	30
11.2.2 Supervision	30
11.2.3 Juries	30
11.3 Popularization	31
11.3.1 Articles and contents	31
11.3.2 Education	32
12 Scientific production	32
12.1 Publications of the year	32
12.2 Cited publications	34

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

Post-Doctoral Fellow

- Dheeraj Narasimha [INRIA, Post-Doctoral Fellow, from Oct 2023]

PhD Students

- Sebastian Allmeier [INPG SA, from Nov 2023]
- Sebastian Allmeier [INRIA, until Oct 2023]
- Thomas Barzola [INPG SA, until Mar 2023]
- Achille Baucher [UGA]
- Victor Boone [UGA]
- Rémi Castera [UGA]
- Romain Cravic [INRIA]
- Yu-Guan Hsieh [UGA, until Sep 2023]
- Simon Jantschgi [CNRS, until May 2023]
- Kimang Khun [INRIA, until Mar 2023]
- Till Kletti [UGA, from Feb 2023 until Jun 2023]
- Till Kletti [NAVER LABS, until Feb 2023]
- Lucas Leandro Nesi [GOUV BRESIL]
- Hugo Lebeau [UGA]

- Davide Legacci [UGA]
- Victor Leger [UGA]
- Minh Toan Nguyen [INRIA, from Oct 2023]
- Minh-Toan Nguyen [INRIA]
- Louis-Sebastien Rebuffi [INRIA, from Oct 2023]
- Louis-Sebastien Rebuffi [UGA, until Sep 2023]
- Charles Sejourne [UGA]

Technical Staff

- Sacha Hodencq [FLORALIS, Engineer, until Mar 2023]

Interns and Apprentices

- Pierre-Louis Cauvin [INRIA, Intern, from Feb 2023 until Jul 2023]
- Mingming Dai [INRIA, Intern, from Feb 2023 until Aug 2023]
- Abel Douzal [INRIA, Intern, from Jun 2023 until Jul 2023]
- Apolline Rodary [INRIA, Intern, from Jun 2023 until Jul 2023]

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 capable 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 [105], BSC [51], UIUC [114], Sandia Nat. Lab. [112], ORNL [52] or ETH Zürich [81]) 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* [59], 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 [84, 83] or the HPL benchmark [57, 34]. We have shown that the performance (both for time and energy consumption [80]) 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 [44]. 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 [87, 40].

Trace Analysis and Visualization Many monolithic visualization tools have been developed by renowned HPC groups since decades (e.g., BSC [98], Jülich and TU Dresden [95, 54], UIUC [79, 101, 82] and ANL [113]) but most of these tools build on the classical information visualization [102] 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 [66], 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 [66] and more recently to a sparse direct solver [37]. 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 [53, 45] (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 [35, 60] 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 [67], 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 [70], 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 [72, 68].

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 [55, 74, 73], 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 [49, 50, 48], 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 [49, 48] 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 [50] 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 [62], 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 [61], mean field games can be used to study how much vaccination should be subsidized to encourage people to adapt a socially optimal behaviour [41].

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 [6][7], 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 [69].

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, [91, 89] showed that the (optimal) class of “follow-the-regularized-leader” (FTRL) learning algorithms leads to Poincaré recurrence even in simple, 2×2 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 [58, 92, 93] suggested that *strict* Nash equilibria play an important role in this question. This suspicion was recently confirmed in a series of papers [65, 42] 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 [92], 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 [43] 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 [104]. 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) [56]. 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 [86], 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 [64]. 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 [63], 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 PII¹-based targeting option, allowed attackers to discover some personal data of users [106]. 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) [47]. A follow-up work shed further light on the typical uses of the platform [46]. In another work, we proposed an innovative protocol based on randomized withdrawal to protect public posts deletion privacy [94]. Finally, in [71], 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, [90] 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 [100][36].

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 [39][97] 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 [99] and a vaccination policy [96]). 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 [38]. Overall this work has been characterized by close interactions with

¹Personally 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 [103]. In fact, this is the reason why modern processors are equipped with Dynamic Voltage and Frequency Scaling (DVFS) technology [111]. 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 [77, 75, 76]) as well as on the experimental side (showing the gains of optimal policies over classical solutions [78]).

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 [109, 108, 107], (2) how to co-simulate an electrical network and a communication network [85], and (3) what is the performance of the communication protocol (PLC G3) used by the Linky smart meters [88].

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

We try to keep the carbon footprint of the team as low as possible by a stricter laptop renewal policy and by reducing plane travels (e.g., using visioconference or sometimes by avoiding publishing our research in conferences that would take place on the other side of the planet).

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-INP, as well as to associations in Grenoble (FNE, AgirAlternatif). He is also co-responsible of the *Digital Transformation* DU. He has published several articles on the issue of "usability" of artificial intelligence. 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.

5.3 Impact of research results

Jean-Marc Vincent is heavily engaged since several years in the training of computer science teachers at the elementary/middle/high school levels [33, 2]. 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 11.2.1 for more details.

6 Highlights of the year

6.1 Awards

Victor Boone and Panayotis Mertikopoulos have received a Spotlight at the NIPS conference for their article on *The equivalence of dynamic and strategic stability under regularized learning in games* [13].

7 New software, platforms, open data

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 2 major releases in 2023. On modeling aspects, we released new plugins simulating chiller, photovoltaic and battery components of Fog/Edge infrastructures, as well as the disk arrays used in desegregated infrastructures. We improved the consistency of the simulation core: the new ActivitySet containers now make it easy to way for the completion of an heterogeneous set of activities (computation, communication, I/O, etc). The simulation of workflow and dataflow applications was also streamlined, with more examples, more documentation and less bugs. A new model of activities mixing disk I/O and network communication was introduced, to efficiently simulate accesses to remote disks. In addition, many efforts were put on the profiling of the software, leading to massive performance gains. We also pursued our efforts to improve the overall framework, through bug fixes, code refactoring and other software quality improvements. In particular, interfaces that were deprecated since almost a decade were removed to ease the maintenance burden on our community.

Many improvement occurred on the model-checker side too. We dropped the old experiments toward stateful verification of liveness properties to boost the development of stateless verification of safety properties. Our tool is simpler internally, and usable on all major operating systems. We modernized the reduction algorithms, implementing several recent algorithms of the literature and paving the way to the introduction of new ones. We also introduced a new module allowing to verify not only distributed applications, but also threaded applications.

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

Contact: Martin Quinson

Participants: Adrien Lebre, Anne-Cecile 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.

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.

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

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.

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

Publications: [hal-01622054](#), [tel-02509756](#), [hal-03485044](#)

Contact: Nicolas Gast

8 New results

The new results produced by the team in 2023 can be grouped into the following categories.

8.1 Performance evaluation of Large Systems

Participants: Sebastian Allmeier, Thomas Barzola, Vincent Danjean, Arnaud Legendre, Nicolas Gast, Guillaume Huard, Lucas Leandro Nesi, Jean-Marc Vincent.

Visualization strategies are a valuable tool in the performance evaluation of HPC applications. Although the traditional Gantt charts are a widespread and enlightening strategy, it presents scalability problems and may misguide the analysis by focusing on resource utilization alone. In [16], we propose an overview strategy to indicate nodes of interest for further investigation with classical visualizations like Gantt charts. For this, it uses a progression metric that captures work done per node inferred from the task-based structure, a time-step clustering of those metrics to decrease redundant information, and a more scalable visualization technique. We demonstrate with six scenarios and two applications that such a strategy can indicate problematic nodes more straightforwardly while using the same visualization

space. Also, we provide examples where it correctly captures application work progression, showing application problems earlier and as an easy way to compare nodes. At the same time that traditional methods are misleading.

This work completes our previous work on performance analysis of task-based applications on heterogeneous platforms and is part of the PhD thesis of Lucas Leandro Nesi [27]. It will be pursued in the WP5 (*Performance analysis and prediction*) of the ExaSoft pillar (*High Performance Computing software and tools*) of the PEPR NumPEX (Numérique Hautes Performances pour l'Exascale). The rest of the thesis of Lucas Leandro Nesi is more related to performance optimization (through algorithmic and reinforcement learning techniques) and evaluation (through predictive simulation and real experiments). A particular effort has been devoted to the reproducibility of the results through the opening of both the data, the code, and the underlying methodology.

Mean field approximation is a powerful technique which has been used in many settings to study large-scale stochastic systems. Some of our latest developments have been transferred in the open source project `rmf_tool` 7.1.6. In the case of two-timescale systems, the approximation is obtained by a combination of scaling arguments and the use of the averaging principle. In [1], we analyze the approximation error of this 'average' mean field model for a two-timescale model (\mathbf{X}, \mathbf{Y}) , where the slow component \mathbf{X} describes a population of interacting particles which is fully coupled with a rapidly changing environment \mathbf{X} . The model is parametrized by a scaling factor N , e.g. the population size, which as N gets large decreases the jump size of the slow component in contrast to the unchanged dynamics of the fast component. We show that under relatively mild conditions, the 'average' mean field approximation has a bias of order $O(1/N)$ compared to $\mathbb{E}[\mathbf{X}]$. This holds true under any continuous performance metric in the transient regime, as well as for the steady-state if the model is exponentially stable. To go one step further, we derive a bias correction term for the steady-state, from which we define a new approximation called the refined 'average' mean field approximation whose bias is of order $O(1/N^2)$. This refined 'average' mean field approximation allows computing an accurate approximation even for small scaling factors, i.e., $N \approx 10 - 50$. We illustrate the developed framework and accuracy results through an application to a random access CSMA model.

Finally, the PhD thesis of Thomas Barzola [22] presents a modular approach to compare optimization methods for bike sharing systems. Bike Sharing Systems (BSSs) are nowadays installed in many cities. In such a system, a user can take any available bike and return it to wherever there is an available parking spot. The Operations Research literature contains many papers that study optimization questions related to BSS, and in particular how to maximize the availability of bikes where and when the users need them. Yet, the optimization methods proposed by these papers are difficult to compare because most papers use their own problem instances and define their own metrics. This thesis aims to fill this gap by building a reproducible research methodology for BSSs. In this work, we divide this methodology in four modules: use of historical data, demand estimation, optimization methods, and performance evaluation. We study each module separately. In each case, we propose a prototype implementation and compare existing solutions when they are available. The first module handles the use of data from real systems. For many systems, two types of data are usually available: trips made by users, and records of the number of bikes available in each station. We observe that in general these data are inconsistent and we propose a method to correct this and detect relocation operations. The second module is the demand estimation one. In optimizing a BSS, it is essential to estimate the demand of the users for whom the system is designed. Most of the optimization works in the literature use historical demand to estimate the demand of the system. We experiment with the few existing methods of the literature along with a newly introduced method to detect censored demand. The third module is bike availability optimization. We implement a published optimization algorithm for this module as an example. We illustrate the challenges of reproducible research by trying to replicate the results. This chapter shows that, although the original authors made the data about their experiments available, we did not get the same quantitative results as the original publication. This difference highlights the need for better publication standards to produce more reproducible results. Finally, our fourth and last module is used to validate the optimization methods implemented in the 3rd module. We advocate that a simulator having all the requirements (user behavior models, demand scenarios, management strategies, etc.) can be a validation model. We use a third-party simulator to illustrate this module. We observe throughout this thesis that making research reproducible is not always handled with due diligence while being fundamental to produce valuable knowledge. In this work, we try our best efforts to specify and provide reproducible

tools to ensure that researchers could obtain the same results with the same data. We give links to the data, codes, environments and analyses needed to reproduce the experiments.

8.2 Energy optimization

Participants: Jonatha Anselmi, Bruno Gaujal.

In [4], we optimize the scheduling of Deep Learning training jobs from the perspective of a Cloud Service Provider running a data center, which efficiently selects resources for the execution of each job to minimize the average energy consumption while satisfying time constraints. To model the problem, we first develop a Mixed-Integer Non-Linear Programming formulation. Unfortunately, the computation of an optimal solution is prohibitively expensive, and to overcome this difficulty, we design a heuristic STochastic Scheduler (STS). Exploiting the probability distribution of early termination, STS determines how to adapt the resource assignment during the execution of the jobs to minimize the expected energy cost while meeting the job due dates. The results of an extensive experimental evaluation show that STS guarantees significantly better results than other methods in the literature, effectively avoiding due date violations and yielding a percentage total cost reduction between 32% and 80% on average. We also prove the applicability of our method in real-world scenarios, as obtaining optimal schedules for systems of up to 100 nodes and 400 concurrent jobs requires less than 5 seconds. Finally, we evaluated the effectiveness of GPU sharing, i.e., running multiple jobs in a single GPU. The obtained results demonstrate that depending on the workload and GPU memory, this further reduces the energy cost by 17-29% on average.

8.3 Restless multi-armed bandits

Participants: Nicolas Gast, Bruno Gaujal, Kimang Kuhn, Chen Yan.

Multi-Armed Bandits are a fundamental model for problems in which a decision maker has to iteratively select one of multiple fixed alternatives (i.e. arms or actions) when the reward of each choice is only partially known at the time of decision and is learned as the decision maker interacts with the bandits. The regret of a strategy is the expectation of the sum of the collected rewards minus the expectation of the optimal reward (the one corresponding to the arm with the larger reward). *Markov Decision Processes* (MDP) provide a framework for modeling situations where the state of a system (and its associated reward) evolves partly random and partly under the control of the decision maker. The reward depends on the current state of the machine, but good policies can be computed (e.g., using dynamic programming although it can be computationally unreasonable) when the system is fully known upfront. We have considered the intermediate situation of *restless and restless bandits* where each arm corresponds to an independent Markov chain but where neither the chain nor the associated reward is initially known. Each time a particular arm is played, the state of that chain advances to a new one, chosen according to the Markov state evolution probabilities. In the *restless bandits* problem, the states of non-played arms can also evolve over time.

Whittle index is a generalization of Gittins index that provides very efficient allocation rules for restless multi-armed bandits. In [5], 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 $O(2/3)n^3 + o(n^3)$ 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.

This work is part of the PhD thesis of Kimang Kuhn [25], where it was shown that no learning algorithms can perform uniformly well over the general class of restless bandits, and where several strategies for restful bandits have also been studied,

In [6], we evaluate the performance of Whittle index policy for restless Markovian bandit. It is shown in Weber and Weiss [110] that if the bandit is indexable and the associated deterministic system has a global attractor fixed point, then the Whittle index policy is asymptotically optimal in the regime where the arm population grows proportionally with the number of activation arms. In this paper, we show that, under the same conditions, this convergence rate is exponential in the arm population, unless the fixed point is singular, which almost never happens in practice. Our result holds for the continuous-time model of Weber and Weiss (1990) and for a discrete-time model in which all bandits make synchronous transitions. Our proof is based on the nature of the deterministic equation governing the stochastic system: We show that it is a piecewise affine continuous dynamical system inside the simplex of the empirical measure of the arms. Using simulations and numerical solvers, we also investigate the singular cases, as well as how the level of singularity influences the (exponential) convergence rate. We illustrate our theorem on a Markovian fading channel model.

In [7], we also 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.

8.4 Reinforcement Learning and MDP

Participants: Jonatha Anselmi, Victor Boone, Romain Cravic, Nicolas Gast, Bruno Gaujal, Louis-Sébastien Rebuffi.

Although *regret* is a common objective in Reinforcement Learning, other criteria are relevant and allow to better understand or discriminate algorithms.

The first contribution of [12] is the introduction of a new performance measure of a RL algorithm that is more discriminating than the regret, that we call the regret of exploration that measures the asymptotic cost of exploration. The second contribution is a new performance test (PT) to end episodes in RL optimistic algorithms. This test is based on the performance of the current policy with respect to the best policy over the current confidence set. This is in contrast with all existing RL algorithms whose episode lengths are only based on the number of visits to the states. This modification does not harm the regret and brings an additional property. We show that while all current episodic RL algorithms have a linear regret of exploration, our method has a $O(\log T)$ regret of exploration for non-degenerate deterministic MDPs.

In [11], we investigate a new learning problem, the identification of Blackwell optimal policies on deterministic MDPs (DMDPs): A learner has to return a Blackwell optimal policy with fixed confidence using a minimal number of queries. First, we characterize the maximal set of DMDPs for which the identification is possible. Then, we focus on the analysis of algorithms based on product-form confidence regions. We minimize the number of queries by efficiently visiting the state-action pairs with respect to

the shape of confidence sets. Furthermore, these confidence sets are themselves optimized to achieve better performance. The performance of our method compares to the lower bound up to a factor n^2 in the worst case, where n is the number of states, and constant in certain classes of DMDPs.

In [14], we propose the first model-free algorithm that achieves low regret performance for decentralized learning in two-player zero-sum tabular stochastic games with infinite-horizon average-reward objective. In decentralized learning, the learning agent controls only one player and tries to achieve low regret performances against an arbitrary opponent. This contrasts with centralized learning where the agent tries to approximate the Nash equilibrium by controlling both players. In our infinite-horizon undiscounted setting, additional structure assumptions is needed to provide good behaviors of learning processes : here we assume for every strategy of the opponent, the agent has a way to go from any state to any other. This assumption is the analogous to the "communicating" assumption in the MDP setting. We show that our Decentralized Optimistic Nash Q-Learning (DONQ-learning) algorithm achieves both sublinear high probability regret of order $3/4$ and sublinear expected regret of order $2/3$. Moreover, our algorithm enjoys a low computational complexity and low memory space requirement compared to the previous works in the same setting.

Finally, in [30], we present an efficient reinforcement learning algorithm that learns the optimal admission control policy in a partially observable queueing network. Specifically, only the arrival and departure times from the network are observable, and optimality refers to the average holding/rejection cost in infinite horizon. While reinforcement learning in Partially Observable Markov Decision Processes (POMDP) is prohibitively expensive in general, we show that our algorithm has a regret that only depends *sub-linearly* on the maximal number of jobs in the network, S . In particular, in contrast with existing regret analyses, our regret bound does not depend on the *diameter* of the underlying Markov Decision Process (MDP), which in most queueing systems is at least exponential in S . The novelty of our approach is to leverage Norton's equivalent theorem for closed product-form queueing networks and an efficient reinforcement learning algorithm for MDPs with the structure of birth-and-death processes.

This work is part of the PhD thesis of Louis-Sebastien Rebuffi [29] and allows to propose reinforcement learning algorithms for controlled queueing systems that demonstrate a weak dependence on the state space compared to results obtained in the general case.

8.5 Learning in games

Participants: Victor Boone, Yu-Guan Hsieh, Panayotis Mertikopoulos.

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.

A wide array of modern machine learning applications – from adversarial models to multi-agent reinforcement learning – can be formulated as non-cooperative games whose Nash equilibria represent the system's desired operational states. Despite having a highly non-convex loss landscape, many cases of interest possess a latent convex structure that could potentially be leveraged to yield convergence to an equilibrium. Driven by this observation, we propose in [20] a flexible first-order method that successfully exploits such "hidden structures" and achieves convergence under minimal assumptions for the transformation connecting the players' control variables to the game's latent, convex-structured layer. The proposed method – which we call preconditioned hidden gradient descent (PHGD) – hinges on a judiciously chosen gradient preconditioning scheme related to natural gradient methods. Importantly, we make no separability assumptions for the game's hidden structure, and we provide explicit convergence rate guarantees for both deterministic and stochastic environments.

In [13], we show the equivalence of dynamic and strategic stability under regularized learning in games by examining the long-run behavior of regularized, no-regret learning in finite games. A well-known result in the field states that the empirical frequencies of no-regret play converge to the game's set of coarse correlated equilibria; however, our understanding of how the players' actual strategies evolve over time is much more limited - and, in many cases, non-existent. This issue is exacerbated

further by a series of recent results showing that only strict Nash equilibria are stable and attracting under regularized learning, thus making the relation between learning and pointwise solution concepts particularly elusive. In lieu of this, we take a more general approach and instead seek to characterize the *setwise* rationality properties of the players' day-to-day play. To that end, we focus on one of the most stringent criteria of setwise strategic stability, namely that any unilateral deviation from the set in question incurs a cost to the deviator - a property known as closedness under better replies (club). In so doing, we obtain a far-reaching equivalence between strategic and dynamic stability: a product of pure strategies is closed under better replies if and only if its span is stable and attracting under regularized learning. In addition, we estimate the rate of convergence to such sets, and we show that methods based on entropic regularization (like the exponential weights algorithm) converge at a geometric rate, while projection-based methods converge within a finite number of iterations, even with bandit, payoff-based feedback.

In [3], 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 [9], we develop a flexible stochastic approximation framework for analyzing the long-run behavior of learning in games (both continuous and finite). The proposed analysis template incorporates a wide array of popular learning algorithms, including gradient-based methods, the exponential / multiplicative weights algorithm for learning in finite games, optimistic and bandit variants of the above, etc. In addition to providing an integrated view of these algorithms, our framework further allows us to obtain several new convergence results, both asymptotic and in finite time, in both continuous and finite games. Specifically, we provide a range of criteria for identifying classes of Nash equilibria and sets of action profiles that are attracting with high probability, and we also introduce the notion of coherence, a game-theoretic property that includes strict and sharp equilibria, and which leads to convergence in finite time. Importantly, our analysis applies to both oracle-based and bandit, payoff-based methods – that is, when players only observe their realized payoffs.

This work is part of the PhD thesis of Yu Guan Hsieh [23] entitled *Decision-Making in multi-agent systems: delays, adaptivity, and learning in games*, and which has investigated separately two critical aspects of multi-agent systems: the impact of delays and the interactions among agents with non-aligned interests.

8.6 Quantum Games

Participants: Panayotis Mertikopoulos.

Although the games we generally consider for learning have nothing to do with the quantum world, they often involve probabilities (to account for the uncertainty of the agents or of the nature) and semi-definite programming (e.g., when dealing with the optimization of MIMO antennas). Quantum games have thus been a natural target for which we have proposed several contributions.

Recent developments in domains such as non-local games, quantum interactive proofs, and quantum generative adversarial networks have renewed interest in quantum game theory and, specifically, quantum zero-sum games. Central to classical game theory is the efficient algorithmic computation of Nash equilibria, which represent optimal strategies for both players. In 2008, Jain and Watrous proposed the first classical algorithm for computing equilibria in quantum zero-sum games using the Matrix Multiplicative Weight Updates (MMWU) method to achieve a convergence rate of $O(d/\epsilon^2)$ iterations to ϵ -Nash equilibria in the 4^d -dimensional spectraplex. In [21], we propose a hierarchy of quantum optimization algorithms that generalize MMWU via an extra-gradient mechanism. Notably, within this proposed hierarchy, we introduce the Optimistic Matrix Multiplicative Weights Update (OMMWU) algorithm and establish its average-iterate convergence complexity as $O(d/\epsilon)$ iterations to ϵ -Nash equilibria. This

quadratic speed-up relative to Jain and Watrous' original algorithm sets a new benchmark for computing ϵ -Nash equilibria in quantum zero-sum games.

In [18], we study the problem of learning in quantum games and other classes of semidefinite games with scalar, payoff-based feedback. For concreteness, we focus on the widely used matrix multiplicative weights (MMW) algorithm and, instead of requiring players to have full knowledge of the game (and/or each other's chosen states), we introduce a suite of minimal-information matrix multiplicative weights (3MW) methods tailored to different information frameworks. The main difficulty to attaining convergence in this setting is that, in contrast to classical finite games, quantum games have an infinite continuum of pure states (the quantum equivalent of pure strategies), so standard importance-weighting techniques for estimating payoff vectors cannot be employed. Instead, we borrow ideas from bandit convex optimization and we design a zeroth-order gradient sampler adapted to the semidefinite geometry of the problem at hand. As a first result, we show that the 3MW method with deterministic payoff feedback retains the $O(1/\sqrt{T})$ convergence rate of the vanilla, full information MMW algorithm in quantum min-max games, even though the players only observe a single scalar. Subsequently, we relax the algorithm's information requirements even further and we provide a 3MW method that only requires players to observe a random realization of their payoff observable, and converges to equilibrium at an $O(T^{-1/4})$ rate. Finally, going beyond zero-sum games, we show that a regularized variant of the proposed 3MW method guarantees local convergence with high probability to all equilibria that satisfy a certain first-order stability condition.

Finally, in [17], we study the equilibrium convergence and stability properties of the widely used matrix multiplicative weights (MMW) dynamics for learning in general quantum games. A key difficulty in this endeavor is that the induced quantum state dynamics decompose naturally into (i) a classical, commutative component which governs the dynamics of the system's eigenvalues in a way analogous to the evolution of mixed strategies under the classical replicator dynamics; and (ii) a non-commutative component for the system's eigenvectors. This non-commutative component has no classical counterpart and, as a result, requires the introduction of novel notions of (asymptotic) stability to account for the nonlinear geometry of the game's quantum space. In this general context, we show that (i) only pure quantum equilibria can be stable and attracting under the MMW dynamics; and (ii) as a partial converse, pure quantum states that satisfy a certain "variational stability" condition are always attracting. This allows us to fully characterize the structure of quantum Nash equilibria that are stable and attracting under the MMW dynamics, a fact with significant implications for predicting the outcome of a multi-agent quantum learning process.

8.7 Continuous optimization methods

Participants: 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.

Many modern machine learning applications – from online principal component analysis to covariance matrix identification and dictionary learning – can be formulated as minimization problems on Riemannian manifolds, and are typically solved with a Riemannian stochastic gradient method (or some variant thereof). However, in many cases of interest, the resulting minimization problem is not geodesically convex, so the convergence of the chosen solver to a desirable solution – i.e., a local minimizer – is by no means guaranteed. In [15], we study precisely this question, that is, whether stochastic Riemannian optimization algorithms are guaranteed to avoid saddle points with probability 1. For generality, we study a family of retraction-based methods which, in addition to having a potentially much lower per-iteration cost relative to Riemannian gradient descent, include other widely used algorithms, such as natural policy gradient methods and mirror descent in ordinary convex spaces. In this general setting, we show that, under mild assumptions for the ambient manifold and the oracle providing gradient information, the policies under study avoid strict saddle points / submanifolds with probability 1, from any initial

condition. This result provides an important sanity check for the use of gradient methods on manifolds as it shows that, almost always, the limit state of a stochastic Riemannian algorithm can only be a local minimizer.

In [31], we examine the last-iterate convergence rate of Bregman proximal methods - from mirror descent to mirror-prox and its optimistic variants - as a function of the local geometry induced by the prox-mapping defining the method. For generality, we focus on local solutions of constrained, non-monotone variational inequalities, and we show that the convergence rate of a given method depends sharply on its associated Legendre exponent, a notion that measures the growth rate of the underlying Bregman function (Euclidean, entropic, or other) near a solution. In particular, we show that boundary solutions exhibit a stark separation of regimes between methods with a zero and non-zero Legendre exponent: the former converge at a linear rate, while the latter converge, in general, sublinearly. This dichotomy becomes even more pronounced in linearly constrained problems where methods with entropic regularization achieve a linear convergence rate along sharp directions, compared to convergence in a finite number of steps under Euclidean regularization.

8.8 Random matrix analysis and Machine Learning

Participants: Romain Couillet, Minh-Toan Nguyen.

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.

The PhD thesis of Minh-Toan Nguyen [28] has provided a nice overview with a deep perspective on replica method and asymptotic equivalence. Replica method is a favorite tool of physicists for studying large disordered systems. Although the method is highly non-rigorous, it can solve difficult problems across various domains: random matrix theory, convex optimization, combinatorial optimization, Bayesian inference, etc. The method has been successfully used to analyze theoretical models in wireless communication and machine learning. The rigorous alternatives for the replica method include the method of deterministic equivalents in random matrix theory, the objective method in combinatorial optimization, and the CGMT (convex Gaussian min-max theorem) in random convex optimization. Although these methods work in different domains, they offer one common insight: the asymptotic equivalence, which tells us that the large system under study is equivalent to a simpler system. As a result, many difficult computations on the original system can be done more easily on the equivalent system. In contrast, with the replica method, the insights come after the calculations. We start by writing down what we want to compute and then proceed to get the answer at the end. After calculating various quantities related to the system, with some observations and a good intuition, we may uncover the equivalent system. In this thesis, we show that the asymptotic equivalent of a disordered system can be obtained directly through the replica formalism by paying attention to the large-deviation computations lurking behind the replica computations. In other words, we develop a version of the replica method that can directly compute the asymptotic equivalent of a disordered system. This version of the replica method, which fits into the same framework of deterministic equivalence as the rigorous methods above, can compute the deterministic equivalents of random matrices, formally derive the CGMT, and solve problems in high-dimensional Bayesian statistics. Moreover, it can derive results on the Sherrington-Kirkpatrick model in a clear and simple manner. In this version of the replica method, each disordered system is associated with an object called “the replica density”. By De Finetti’s theorem, a disordered system can be recovered from its replica density. To compute the asymptotic equivalent of a disordered system, we compute the equivalent of its replica density, using a result that we derive from the fundamental Gibbs principle. We thus obtain another replica density, which corresponds to another disordered system. This system is asymptotically equivalent to the original disordered system.

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

Participants: Rémi Castera, Nicolas Gast, Till Kletti, Simon Jantschgi, Mathieu Molina, Bary Pradelski.

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.

In [32], we study statistical discrimination in matching, where multiple decision-makers are simultaneously facing selection problems from the same pool of candidates. We propose a model where decision-makers observe different, but correlated estimates of each candidate's quality. The candidate population consists of several groups that represent gender, ethnicity, or other attributes. The correlation differs across groups and may, for example, result from noisy estimates of candidates' latent qualities, a weighting of common and decision-maker specific evaluations, or different admission criteria of each decision maker. We show that lower correlation (e.g., resulting from higher estimation noise) for one of the groups worsens the outcome for all groups, thus leading to efficiency loss. Further, the probability that a candidate is assigned to their first choice is independent of their group. In contrast, the probability that a candidate is assigned at all depends on their group, and — against common intuition — the group that is subjected to lower correlation is better off. The resulting inequality reveals a novel source of statistical discrimination.

In [8], we conducted a large number of controlled continuous double auction experiments to reproduce and stress-test the phenomenon of convergence to competitive equilibrium under private information with decentralized trading feedback. Our main finding is that across a total of 104 markets (involving over 1,700 subjects), convergence occurs after a handful of trading periods. Initially, however, there is an inherent asymmetry that favors buyers, typically resulting in prices below equilibrium levels. Analysis of over 80,000 observations of individual bids and asks helps identify empirical ingredients contributing to the observed phenomena including higher levels of aggressiveness initially among buyers than sellers.

This work is part of the PhD thesis of Simon Jantschgi [24] on market design for double auctions.

Individual behavior such as choice of fashion, adoption of new products, and selection of means of transport is influenced by taking account of others' actions. In [10], we study social influence in a heterogeneous population and analyze the behavior of the dynamic processes. We distinguish between two information regimes: (i) agents are influenced by the adoption ratio, (ii) agents are influenced by the usage history. We identify the stable equilibria and long-run frequencies of the dynamics. We then show that the two processes generate qualitatively different dynamics, leaving characteristic 'footprints'. In particular, (ii) favors more extreme outcomes than (i).

In [19], we consider the problem of online allocation subject to a long-term fairness penalty. Contrary to existing works, however, we do not assume that the decision-maker observes the protected attributes — which is often unrealistic in practice. Instead they can purchase data that help estimate them from sources of different quality; and hence reduce the fairness penalty at some cost. We model this problem as a multi-armed bandit problem where each arm corresponds to the choice of a data source, coupled with the online allocation problem. We propose an algorithm that jointly solves both problems and show that it has a regret bounded by $O(\sqrt{T})$. A key difficulty is that the rewards received by selecting a source are correlated by the fairness penalty, which leads to a need for randomization (despite a stochastic setting). Our algorithm takes into account contextual information available before the source selection, and can adapt to many different fairness notions. We also show that in some instances, the estimates used can be learned on the fly.

Finally, fairness has also been studied in the PhD thesis of Till Kletti [26], in the context of multistakeholder recommendation platforms. The object of study of this thesis is the ranking of potentially relevant objects in response to an information request, for example when using a search engine or in the case of online content recommendation. Such a ranking brings together two groups: users searching for relevant information, and content producers, whose goal is to make the produced information visible. For example, when searching for restaurants, the user is interested in seeing good restaurants, while the interest of the restaurant owners is to be seen by many people, in order to attract customers. The objects to be ranked are thus competing with each other and it is in the interest of the platform generating

the rankings to ensure that the exposure allocated to the objects is fairly distributed. Obviously there are many possibilities of defining what fair means and none of them will be unanimously agreed upon. Therefore in this thesis the definition of fairness is taken as a parameter represented by a vector of merit, which determines the proportion with which visibility should be distributed amongst the items. This will make our method applicable to a wide range of possible definitions. Two things then become apparent. First, there does not in general exist ranking that is fair in the sense of proportionality of exposure to merit. It is therefore necessary to produce several rankings that compensate each other in order to give, on average, fair exposures to the items. Secondly, these rankings do not generally give maximum utility to the user. Indeed, to guarantee fairness, less relevant objects could potentially be shown to him. These two objectives, fairness and utility, are thus not simultaneously optimizable. The contribution of this thesis is to develop methods to determine Pareto optimal ranking sequences, i.e. such that it is not possible to improve one of the two objectives without deteriorating the other. The idea is that this would make it possible for a qualified decision maker to make an informed choice about an adequate trade-off between user utility and fairness amongst items. The determination of these optimal sequences is accomplished via the introduction of a geometric object, a polytope named expohedron. This polytope expresses the set of average exposures attainable with ranking sequences and is therefore a good decision space for both fairness and utility. The expohedron makes it possible to compute these optimal ranking sequences using only mathematically exact geometric constructions inside it, and this in a significantly faster way than previous methods based on linear programming. Moreover, the proposed method is applicable to two large classes of exposure models including Position Based Model (PBM) and Dynamic Bayesian Network (DBN) models to which linear programming is not applicable.

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 European initiatives

10.1.1 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), Kungliga Tekniska Högskolan (Sweden), Politecnico di Torino (Italy), Universitat Politècnica de Catalunya (Spain), Universidade de Lisboa (Portugal) and Grenoble INP, Graduate Schools of Engineering and Management, Université Grenoble Alpes.

10.2 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.

11 Dissemination

11.1 Promoting scientific activities

11.1.1 Scientific events: organisation

1. Member of the organizing committees

- Nicolas Gast, Aditya Mahajan, and Annie Simon have organized the **Workshop on restless bandits, index policies and applications in reinforcement learning** in Grenoble, France, November 2023, which has attracted about 40 researchers of the domain from all over Europe. This event was co-organized with the support of the GDR COSMOS.

- Arnaud Legrand co-organized with five other colleagues the first **Meeting days of the French network for reproducible research** in Paris, France, March 2023, which has attracted about a hundred of researchers. The aim of these days was to provide an overview of the state of reproducibility in scientific research in France, taking into account the diversity of disciplines and practices.
- Jean-Marc Vincent has co-organized the *Colloque de la SIF: le système dans tous ses états* in Grenoble, April 2023.

2. General chair, scientific chair

- Panayotis Mertikopoulos co-organized the **2023 thematic program on Games, Learning, and Networks** in Singapore, from April 3 to April 21, 2023. This program brought together around 80 researchers in total, with the aim of studying emerging questions on game theory and its applications.
- Panayotis Mertikopoulos, Bruno Gaujal, and Annie Simon co-organized the **2023 Alpine Game Theory Symposium** in Grenoble, from June 26 to June 30, 2023. This symposium brought together about 100 researchers working on all aspects of game theory (mathematical, economic, algorithmic, etc.), with the aim of fostering collaborations and interactions between the various communities. It was primarily sponsored by the French Game Theory Society, Inria, MIAI, and the ANR project ALIAS.

11.1.2 Scientific events: selection

1. Member of the conference program committees

- Jonatha Anselmi was a PC member of the Algorithms Track for ICPP 2023.
- Bruno Gaujal was member of the following TPC: ICML, Mascots, and NeurIPS.
- Nicolas Gast was member of the Technical Program Committee of Sigmetrics 2024 and ICLR 2023.
- Panayotis Mertikopoulos served as an area chair for ICLR 2023, ICML 2023, and NeurIPS 2023.

2. Reviewer

- Panayotis Mertikopoulos served as a reviewer for COLT 2023.

11.1.3 Journal

1. Member of the editorial boards

- Nicolas Gast is associate editor of Performance Evaluation, Stochastic Models and TMLR.
- Panayotis Mertikopoulos served as an associate editor for Operations Research Letters, EURO Journal on Computational Optimization, Methodology and Computing in Applied Probability, and the Journal on Dynamics and Games.
- Since December 2023, Panayotis Mertikopoulos is an associate editor for Mathematics of Operations Research, and the managing co-editor for the Open Journal of Mathematical Optimization.

2. Reviewer - reviewing activities

- Jonatha Anselmi was a reviewer for Mathematics of Operation Research, IEEE Transactions on Parallel and Distributed Systems, Queueing Systems, Performance Evaluation.
- Bruno Gaujal was a reviewer for MOR.
- Panayotis Mertikopoulos reviewed around 10 papers for various journals in mathematical optimization and game theory.

11.1.4 Invited talks

- Bruno Gaujal was invited to present his work at the following events:
 - ROADEF as a plenary speaker, at Rennes (Feb. 2023): *Reinforcement Learning and Markovian Bandits*
 - Séminaire Parisien d'Optimisation (April 2023): *Higher Order Reinforcement Learning*
 - OR Seminar, at Univ. of Porto (June 2023): *Reinforcement Learning in Markov Decision Processes*
 - Informs Applied Probability Society, at Nancy (June 2023): *Reinforcement Learning*
 - EDF R&D Seminar, (Oct. 2023): *Learning Indexes in Rested Bandits*
 - LAAS Seminar, Toulouse, (Nov. 2023): *Learning Optimal Admission Control in Partially Observable Queueing Networks*
- Nicolas was invited to present his work at the following events:
 - EDF R&D: *Mean-Field Control for Restless Bandits and Weakly Coupled MDPs*
 - Inria Paris, workshop "8 days on Network Mathematics": *Approximations for dynamics on graphs*
 - Cornell OREI: *Mean-Field Control for Restless Bandits and Weakly Coupled MDPs*
 - Online CNI seminar series: *Mean-Field Control for Restless Bandits and Weakly Coupled MDPs*
 - Université Laval (Québec): *The bias of mean field approximation*
 - McGill university (Montreal): *How to Use Mean-Field Control for Restless Bandits and Weakly Coupled MDPs*
- Arnaud Legrand was invited to give lectures and keynotes on Reproducible Research and Open Science on the following occasions:
 - M1 students in computer science at UGA (Dec. 2023): *Reproducible Research and Computer Science*
 - Neurocampus Open Science Workshop at Bordeaux (Oct. 2023): *Good practices in the lab: Research documentation and electronic notebooks*
 - New Inria PhD students at Grenoble (Oct. 2023): *Doing a PhD, good practice and pitfalls to avoid*
 - Inria foresight seminar of the *Optimization, Machine Learning and Statistical Methods* theme at Rungis (Oct. 2023): *Reproducible Research and Benchmarking*
 - Colloque de la MITI *Répliquabilité/reproductibilité de la recherche : enjeux et propositions* at Paris (Sep. 2023): *Répliquabilité et reproductibilité de la recherche: Impact sur les pratiques des chercheuses et chercheurs*
 - Journée GitLab, GT "Données" de la MITI du CNRS at Paris (June 2023): *Scientific Data Management with Git and Git-Annex*
 - MaiMosine, GRICAD, SARI network, remote conference (June 2023): *Laboratory notebook, computational document, reproducible article. Emacs/Org-mode: One ring to rule them all?*
 - Journées scientifiques de l'équipe AVALON at Le Sapey (June 2023): *Reproducible Research and Computer Science*
 - 20th Anniversary of Grid'5000 at Lyon (May 2023): *Reproducible Research and Computer Science*
 - Réseau des référents science ouverte à la CPU, remote conference (March 2023): *Formation à grande échelle à la Recherche Reproductible*
 - Master2 Mathématiques, Vision, Apprentissage at Saclay (March 2023): *Reproducibility Crisis, Open Science, ... and Computer Science*

- 1ères Journées du Réseau National Recherche reproductible at Paris (March 2023): *Formation à grande échelle à la Recherche Reproductible*
- Unite! Dialogue, remotely (March 2023): *CNRS Open Science policy*
- Seminar at the AG of the LISTIC, at Annecy (Feb. 2023): *Reproducibility Crisis, Open Science, ... and Computer Science*
- Panayotis Mertikopoulos was invited to give a tutorial at the following events:
 - NTUA Summer School on the Mathematics of Machine Learning, Athens, GR: *Optimization algorithms for machine learning*
- Panayotis Mertikopoulos was invited to give a talk in the following conferences:
 - Conference in honor of Roberto Cominetti's 60th birthday, Viña del Mar, CL: *Adaptive routing in large-scale networks*
 - Workshop in honor of Elias Koutsoupias' 60th birthday, Athens, GR: *The attractors of regularized learning in games*
 - 2023 Athens Colloquium on Algorithms and Complexity, Athens, GR: *Strategic stability under regularized learning*
- Panayotis Mertikopoulos was invited to give a talk in the following universities and research institutes:
 - University of Athens, Athens, GR: *Adaptive routing under uncertainty*
 - DeepMind, London, UK: *Accelerated and optimistic methods for learning*
 - University of Athens, Athens, GR: *From Robbins–Monro to artificial intelligence: 70 years of stochastic approximation*
 - TII, online: *Training models with a min-max landscape*
 - Mannheim University, Mannheim, DE: *A stochastic approximation framework for multi-agent learning*

11.1.5 Leadership within the scientific community

- Vincent Danjean, Guillaume Huard, Arnaud Legrand, and Jean-Marc Vincent, are involved in the WP5 (*Performance analysis and prediction*) of the ExaSoft pillar (*High Performance Computing software and tools*) of the PEPR NumPEX (Numérique Hautes Performances pour l'Exascale).
- Panayotis Mertikopoulos is the scientific coordinator of the PEPR IA projet ciblé (acceleration grant) FOUNDRY: *Foundations of robustness and reliability in artificial intelligence*. The project has a total budget of 5M€ and involves five research teams across France (POLARIS in Grenoble, the Inria teams SCOOOL and FAIRPLAY, Dauphine and IMT in Paris, and ENS Lyon).

11.1.6 Research administration

- Vincent Danjean is a member of the *Conseil d'Administration* of Grenoble-INP.
- Nicolas Gast is vice-director of the *école doctorale MSTII*.
- 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.
- Florence Perronin is a member of the QVT team of the LIG.

11.2 Teaching - Supervision - Juries

11.2.1 Teaching

- Jonatha Anselmi teaches *Probabilités et simulation* (32h) and *Évaluation de performances* (32h) at PolyTech Grenoble.
- Vincent Danjean teaches the *Operating Systems, Programming Languages, Algorithms, Computer Science and Mediation* lectures in L3, M1 and Polytech Grenoble.
- Vincent Danjean organized with J.M. Vincent a complementary training for high school professors to teach computer science.
- Nicolas Gast teaches the *Reinforcement learning* part of the M2 course *Mathematical foundations of machine learning* at the M2 MOSIG (Grenoble).
- Bruno Gaujal and Nicolas Gast teach *Markov Decision Process and Reinforcement Learning* (32h in total) at the M2 Info (ENS Lyon).
- Bruno Gaujal teaches *Optimisation under uncertainties* (18h) at the M2 ORCO (UGA).
- Guillaume Huard is responsible of L3 Info and of the Licence Info.
- 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.
- Arnaud Legrand and Jean-Marc Vincent teach the transversal *Scientific Methodology and Empirical Evaluation lecture* (36h) at the M2 MOSIG (UGA).
- The 3rd edition of the MOOC of Arnaud Legrand, K. Hinsen 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 2024), more than 20,800 persons have followed this MOOC and about 2100 certificates of achievement have been delivered. More than half of participants are PhD students and about 10% are undergraduates.
- Florence Perronin teaches *Programming Languages* in L1.
- Florence Perronin is a member of the *conseil de perfectionnement* of the Mathematics license.
- Jean-Marc Vincent contributed to the MOOC NSI: Introduction à la préparation au CAPES NSI, pour les futurs enseignants de lycée en informatique.
- Panayotis Mertikopoulos teaches a graduate course in game theory in the University of Athens.

11.2.2 Supervision

- Arnaud Legrand was a member of the *Comité de Suivi Individuel* of Adeyemi Adetula (UGA)
- Panayotis Mertikopoulos is co-supervising two PhD students in addition of the ones in POLARIS (Victor Boone and Davide Legacci): : Waïss Azizian and Pierre-Louis Cauvin, both with Jérôme Malick (LJK).

11.2.3 Juries

- Nicolas Gast was a reviewer of the PhD committee of Aymen Al Mariani (Ens de Lyon): *Adaptive Pure Exploration in Markov Decision Processes and Bandits*.
- Nicolas Gast was a member of the PhD jury of Michel Davydov (Ens Paris): *Point-process-based Markovian dynamics and their applications*.

- Bruno Gaujal was a member of the PhD thesis committee of Aymen Al Marjani (ENS Lyon): *Adaptive Pure Exploration in Markov Decision Processes and Bandits*.
- Bruno Gaujal was a reviewer for the HDR of Balakrishna Prabhu (Univ. Toulouse): *Some applications of asymptotic analysis in communication networks*.
- Bruno Gaujal was a member of the PhD thesis committee of Fabien Pesquerel (Univ. Lille): *Information par unité d'interaction dans la prise de décisions séquentielles en environnement stochastique*.
- Arnaud Legrand was a reviewer for the HDR of Guillaume Pallez (Univ. Bordeaux): *Model Design and Accuracy for Resource Management in HPC*.
- Arnaud Legrand was a member of the PhD thesis committee of Julien Emmanuel (Univ. Lyon): *Un simulateur pour le calcul haute performance : modélisation multi-niveau de l'interconnect BXI pour prédire les performances d'applications MPI*.
- Panayotis Mertikopoulos was a reviewer / external examiner for the PhD theses of:
 - Le Cong Dinh (U. Southampton): *Online Learning in the Presence of Strategic Adversary*
 - Simon Jantschgi (ETH Zürich / UGA): *Market Design for Double Auctions*
 - Lucas Baudin (Paris-Dauphine): *Contributions to Online Learning in Stochastic Games*
 - Étienne de Montbrun (Toulouse 1 Capitole): *Game Theory and Online Learning: Gradient Descent Ascent, Optimization, Approximation and Certification*
 - Maurizio d'Andrea (Toulouse 1 Capitole): *Learning Algorithms in Game-Theoretic Contexts*
- Panayotis Mertikopoulos was an examiner for the PhD theses of:
 - Rémi Leluc (IP Paris): *Monte Carlo Methods and Stochastic Approximation: Theory and Applications to Machine Learning*
 - Vicky Kouni (U. Athens): *Model-based and data-driven approaches meet redundancy in signal processing*

11.3 Popularization

11.3.1 Articles and contents

We have contributed to a French-speaking teacher training center dedicated to computer science education in high school.

- The introduction of computer science education in high school will allow the next generations to master and participate in the development of digital technology. The main issue is therefore the training of teachers. We are helping to meet this challenge by forming a community of learning and practice with the welcome and support of hundreds of colleagues in activity or in training and by offering two online training courses, one regarding the fundamentals of computer science, with resources for initiation and improvement the other to learn to teach by doing, by co-preparing the educational activities of the courses to come, by sharing didactic practices and by taking a pedagogical step back, including from the point of view of the pedagogy of equality, supplemented by hybrid initiatives. We share in [33] the approach and the analysis from the point of view of educational sciences of the first results obtained. In terms of research, what we are presenting here falls within the framework of what is called "research within action".
- In many European countries, we're starting to teach our children not just how to use digital technology, but how to understand its fundamentals so they can master it. At the same time, we need to train part of our future population in the discipline of computing, which is both a science and a technique, and which can be found as a fundamental skill in all fields that have become digital. The key issue here is teacher training. We help to meet this challenge with a dual online training program, one for the fundamentals of computing, with introductory and advanced resources, and

the other for learning to teach by doing, by co-preparing the pedagogical activities of future courses, and sharing teaching practices. This includes pedagogical hindsight, including from the point of view of equality pedagogy. This resource enables a community of learning and practice to be formed, with hundreds of colleagues in training or already in work welcoming and helping each other. This is complemented by hybrid initiatives across the region. In [2], we share our approach and analyze the initial results from an educational science perspective, showing just how valuable the links between teaching and research can be at this level.

11.3.2 Education

- Florence Perronin has set up a program to promote mental health for L3 Maths students.
- Florence Perronin has been a member of the committee of Journées de l'innovation en promotion de la santé mentale (sep. 2023).
- Guillaume Huard is a member of the development team of [Caseine](#).

12 Scientific production

12.1 Publications of the year

International journals

- [1] S. Allmeier and N. Gast. 'Bias and Refinement of Multiscale Mean Field Models'. In: *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 7.1 (2023), pp. 1–29. DOI: [10.1145/3579336](https://doi.org/10.1145/3579336). URL: <https://inria.hal.science/hal-03981002>.
- [2] M.-H. Comte, S. Djeballah, M. Fourny, S. Hoarau, A. Juton, M. Khaneboubi, A. Lagarrigue, T. Massart, C. Poulmaire, V. Prince, S. Renouf, T. Viéville and J.-M. Vincent. 'Un espace de formation francophone des enseignants, dédié à l'apprentissage de l'informatique, dans le secondaire'. In: *Adjectif: analyses et recherches sur les TICE* (Nov. 2023). URL: <https://hal.science/hal-04350362>.
- [3] B. Duvocelle, P. Mertikopoulos, M. Staudigl and D. Vermeulen. 'Multi-agent online learning in time-varying games'. In: *Mathematics of Operations Research* 48.2 (2023), pp. 914–941. DOI: [10.1287/moor.2022.1283](https://doi.org/10.1287/moor.2022.1283). URL: <https://inria.hal.science/hal-01891545>.
- [4] F. Filippini, J. Anselmi, D. Ardagna and B. Gaujal. 'A Stochastic Approach for Scheduling AI Training Jobs in GPU-based Systems'. In: *IEEE Transactions on Cloud Computing* (2023), pp. 1–17. DOI: [10.1109/TCC.2023.3336540](https://doi.org/10.1109/TCC.2023.3336540). URL: <https://hal.science/hal-04337856>.
- [5] N. Gast, B. Gaujal and K. Khun. 'Testing Indexability and Computing Whittle and Gittins Index in Subcubic Time'. In: *Mathematical Methods of Operations Research* (13th June 2023). DOI: [10.1007/s00186-023-00821-4](https://doi.org/10.1007/s00186-023-00821-4). URL: <https://inria.hal.science/hal-03602458>.
- [6] N. Gast, B. Gaujal and C. Yan. 'Exponential Asymptotic Optimality of Whittle Index Policy.' In: *Queueing Systems* 104 (June 2023), pp. 1–44. DOI: [10.1007/s11134-023-09875-x](https://doi.org/10.1007/s11134-023-09875-x). URL: <https://inria.hal.science/hal-03041176>.
- [7] N. Gast, B. Gaujal and C. Yan. 'LP-based policies for restless bandits: necessary and sufficient conditions for (exponentially fast) asymptotic optimality'. In: *Mathematics of Operations Research* (1st Dec. 2023). DOI: [10.1287/moor.2022.0101](https://doi.org/10.1287/moor.2022.0101). URL: <https://inria.hal.science/hal-03262307>.
- [8] B. Ikica, S. Jantschgi, H. H. Nax, D. G. N. Duran and B. S. R. Pradelski. 'COMPETITIVE MARKET BEHAVIOR: CONVERGENCE AND ASYMMETRY IN THE EXPERIMENTAL DOUBLE AUCTION'. In: *International Economic Review* 64.3 (2023), pp. 1087–1126. DOI: [10.1111/iere.12630](https://doi.org/10.1111/iere.12630). URL: <https://hal.science/hal-04367512>.
- [9] P. Mertikopoulos, Y.-P. Hsieh and V. Cevher. 'A unified stochastic approximation framework for learning in games'. In: *Mathematical Programming* (4th Aug. 2023), pp. 1–40. DOI: [10.1007/s10107-023-02001-y](https://doi.org/10.1007/s10107-023-02001-y). URL: <https://hal.science/hal-03874012>.

- [10] B. S. R. Pradelski. ‘Social influence: The Usage History heuristic’. In: *Mathematical Social Sciences* (May 2023). URL: <https://hal.science/hal-04367513>.

International peer-reviewed conferences

- [11] V. Boone and B. Gaujal. ‘Identification of Blackwell Optimal Policies for Deterministic MDPs’. In: *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics, PMLR. AISTATS 2023 - 26th International Conference on Artificial Intelligence and Statistics*. Vol. 206. Valencia, Spain, 26th Apr. 2023, p. 32. URL: <https://inria.hal.science/hal-04082136>.
- [12] V. Boone and B. Gaujal. ‘The Regret of Exploration and the Control of Bad Episodes in Reinforcement Learning’. In: *PMLR. ICML 2023 - 40th International Conference on Machine Learning*. Vol. 202. Hawaii-Honolulu, United States, 10th July 2023, pp. 2824–2856. URL: <https://inria.hal.science/hal-04161584>.
- [13] V. Boone and P. Mertikopoulos. ‘The equivalence of dynamic and strategic stability under regularized learning in games’. In: *Proceedings of the 37th International Conference on Neural Information Processing Systems. NeurIPS 2023 - 37th Conference on Neural Information Processing Systems*. New Orleans (LA), United States, 4th Nov. 2023, pp. 1–31. URL: <https://inria.hal.science/hal-04307883>.
- [14] R. Cravic, N. Gast and B. Gaujal. ‘Decentralized model-free reinforcement learning in stochastic games with average-reward objective’. In: *AAMAS 2023 - International Conference on Autonomous Agents and Multiagent Systems*. London (U.K.), United Kingdom, 13th Jan. 2023, pp. 1–13. URL: <https://inria.hal.science/hal-04161628>.
- [15] Y.-P. Hsieh, M. R. Karimi, A. Krause and P. Mertikopoulos. ‘Riemannian stochastic optimization methods avoid strict saddle points’. In: *Riemannian stochastic optimization methods avoid strict saddle points*. *NeurIPS 2023 - 37th Conference on Neural Information Processing Systems*. New Orleans (LA), United States, 4th Nov. 2023, pp. 1–27. URL: <https://inria.hal.science/hal-04307876>.
- [16] L. Leandro Nesi, V. Garcia Pinto, L. M. Schnorr and A. Legrand. ‘Summarizing task-based applications behavior over many nodes through progression clustering’. In: *PDP 2023 - 31st Euromicro International Conference on Parallel, Distributed, and Network-Based Processing*. Naples, Italy, 1st Mar. 2023, pp. 1–8. URL: <https://hal.science/hal-04005071>.
- [17] K. Lotidis, P. Mertikopoulos and N. Bambos. ‘The stability of matrix multiplicative weights dynamics in quantum games’. In: *Proceedings of the 62nd IEEE Annual Conference on Decision and Control*. CDC 2023 - 62nd IEEE Conference on Decision and Control. Singapore, Singapore: IEEE, 2023, pp. 1–8. URL: <https://inria.hal.science/hal-04307803>.
- [18] K. Lotidis, P. Mertikopoulos, N. Bambos and J. Blanchet. ‘Payoff-based learning with matrix multiplicative weights in quantum games’. In: *Proceedings of the 37th International Conference on Neural Information Processing Systems. NeurIPS 2023 - 37th Conference on Neural Information Processing Systems*. New Orleans (LA), United States, Dec. 2023, pp. 1–39. URL: <https://inria.hal.science/hal-04307824>.
- [19] M. Molina, N. Gast, P. Loiseau and V. Perchet. ‘Trading-off price for data quality to achieve fair online allocation’. In: *NeurIPS 2023 - 37th Conference on Neural Information Processing Systems*. New Orleans, USA, United States, 10th Dec. 2023. URL: <https://inria.hal.science/hal-04360656>.
- [20] I. Sakos, E.-V. Vlatakis-Gkaragkounis, P. Mertikopoulos and G. Piliouras. ‘Exploiting hidden structures in non-convex games for convergence to Nash equilibrium’. In: *Proceedings of the 37th International Conference on Neural Information Processing Systems. NeurIPS 2023 - 37th Conference on Neural Information Processing Systems*. New Orleans (LA), United States, 2023, pp. 1–32. URL: <https://inria.hal.science/hal-04312981>.

- [21] F. Vasconcelos, E.-V. Vlatakis-Gkaragkounis, P. Mertikopoulos, G. Piliouras and M. I. Jordan. ‘A quadratic speedup in finding Nash equilibria of quantum zero-sum games’. In: *Proceedings of the 2023 Conference on Quantum Techniques in Machine Learning*. QTML 2023 - Annual international conference on Quantum Techniques in Machine Learning. CERN, Switzerland, 2023, pp. 1–54. URL: <https://inria.hal.science/hal-04312987>.

Doctoral dissertations and habilitation theses

- [22] T. Barzola-Poma Hild. ‘A Modular Approach to Compare Optimization Methods for Bike Sharing Systems’. Université Grenoble Alpes [2020-....], 20th Mar. 2023. URL: <https://theses.hal.science/tel-04122218>.
- [23] Y.-G. Hsieh. ‘Decision-Making in multi-agent systems: delays, adaptivity, and learning in games’. PhD thesis. Université Grenoble Alpes, 7th Nov. 2023. URL: <http://www.theses.fr/s245435>.
- [24] S. P. Jantschi. ‘Market design for double auctions’. PhD thesis. Université Grenoble Alpes, 31st May 2023. URL: <http://www.theses.fr/2023GRALM020>.
- [25] K. Khun. ‘Algorithms for Markovian bandits : Indexability and Learning’. Université Grenoble Alpes [2020-....], 30th Mar. 2023. URL: <https://theses.hal.science/tel-04190346>.
- [26] T. Kletti. ‘Fairness in multistakeholder recommendation platforms’. PhD thesis. Université Grenoble Alpes, 29th June 2023. URL: <http://www.theses.fr/2023GRALM031>.
- [27] L. Leandro Nesi. ‘Strategies for Distributing Task-Based Applications on Heterogeneous Platforms’. PhD thesis. Université Grenoble Alpes, 14th Sept. 2023. URL: <http://www.theses.fr/s300803>.
- [28] M.-T. Nguyen. ‘Replica method and asymptotic equivalence’. PhD thesis. Université Grenoble Alpes, 12th May 2023. URL: <http://www.theses.fr/s252318>.
- [29] L.-S. Rebuffi. ‘Reinforcement Learning Algorithms for Controlled Queueing Systems’. PhD thesis. Université Grenoble Alpes, 11th Dec. 2023. URL: <http://www.theses.fr/s276680>.

Reports & preprints

- [30] J. Anselmi, B. Gaujal and L.-S. Rebuffi. *Learning Optimal Admission Control in Partially Observable Queueing Networks*. 2023. URL: <https://hal.science/hal-04170992>.
- [31] W. Azizian, F. Iutzeler, J. Malick and P. Mertikopoulos. *The rate of convergence of Bregman proximal methods: Local geometry vs. regularity vs. sharpness*. 26th Nov. 2023. URL: <https://inria.hal.science/hal-04307796>.
- [32] R. Castera, P. Loiseau and B. Pradelski. *Statistical Discrimination in Stable Matching*. Apr. 2023. URL: <https://hal.science/hal-03672270>.
- [33] M.-H. Comte, S. Djeballah, M. Fourny, S. Hoarau, A. Juton, M. Khaneboubi, A. Lagarrigue, T. Massart, C. Poulmaire, V. Prince, S. Renouf, T. Viéville and J.-M. Vincent. *A French-speaking training platform for teachers, dedicated to learning computer science, in secondary school*. RR-9514. Inria, 26th June 2023, p. 15. URL: <https://inria.hal.science/hal-04140155>.

12.2 Cited publications

- [34] 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>.
- [35] 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>.
- [36] 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>.

- [37] 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>.
- [38] 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>.
- [39] 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>.
- [40] 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>.
- [41] 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>.
- [42] 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>.
- [43] 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>.
- [44] 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>.
- [45] 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>.
- [46] 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>.
- [47] 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>.
- [48] 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>.
- [49] 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>.
- [50] 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>.

- [51] 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.
- [52] S. Böhm and C. Engelmann. ‘xSim: The Extreme-Scale Simulator’. In: *HPCS*. Istanbul, Turkey, 2011.
- [53] 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>.
- [54] 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://doi.org/10.1007/978-3-642-11261-4_2).
- [55] 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>.
- [56] 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>.
- [57] 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>.
- [58] 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>.
- [59] 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>.
- [60] 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>.
- [61] 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>.
- [62] 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>.
- [63] 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>.
- [64] 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>.

- [65] 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>.
- [66] 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>.
- [67] 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>.
- [68] 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>.
- [69] N. Gast, B. Gaujal and K. Khun. ‘Learning algorithms for Markovian Bandits: Is Posterior Sampling more Scalable than Optimism?’ In: *Transactions on Machine Learning Research Journal* (Nov. 2022). URL: <https://inria.hal.science/hal-03262006>.
- [70] 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>.
- [71] 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>.
- [72] 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>.
- [73] 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>.
- [74] 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>.
- [75] 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>.
- [76] 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>.
- [77] 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>.
- [78] 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>.
- [79] M. Heath and J. Etheridge. ‘Visualizing the performance of parallel programs’. In: *IEEE software* 8.5 (1991).

- [80] 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>.
- [81] 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.
- [82] 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).
- [83] 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>.
- [84] 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 Over-decomposition 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>.
- [85] 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>.
- [86] 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.
- [87] 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>.
- [88] 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>.
- [89] 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>.
- [90] 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>.
- [91] 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>.
- [92] 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>.
- [93] 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>.

- [94] 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>.
- [95] W. Nagel, A. Arnold, M. Weber, H. Hoppe and K. Solchenbach. 'VAMPIR: Visualization and Analysis of MPI Resources'. In: *Supercomputer 12.1* (1996).
- [96] 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>.
- [97] 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>.
- [98] 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.
- [99] 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>.
- [100] 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>.
- [101] 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.
- [102] B. Shneiderman. 'The eyes have it: A task by data type taxonomy for information visualizations'. In: *IEEE Symposium on Visual Languages*. IEEE. 1996.
- [103] 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.
- [104] 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>.
- [105] M. Tikir, M. Laurenzano, L. Carrington and A. Snavely. 'PSINS: An Open Source Event Tracer and Execution Simulator for MPI Applications'. In: *Euro-Par*. 2009.
- [106] 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>.
- [107] B. Vinot, F. Cadoux and R. Hélot. '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>.
- [108] B. Vinot, F. Cadoux and R. Hélot. '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>.
- [109] 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>.

- [110] R. R. Weber and G. Weiss. 'On an Index Policy for Restless Bandits'. In: *Journal of Applied Probability* 27.3 (1990), pp. 637–648. URL: <http://www.jstor.org/stable/3214547>.
- [111] 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.
- [112] 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.
- [113] 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>.
- [114] G. Zheng, G. Kakulapati and L. Kalé. 'BigSim: A Parallel Simulator for Performance Prediction of Extremely Large Parallel Machines'. In: *IPDPS*. 2004.