

# 2025 Activity Report

RESEARCH CENTRE: Inria Centre at Université Grenoble Alpes  
IN PARTNERSHIP WITH: Université de Grenoble Alpes, CNRS

---

Project-Team

## POLARIS

Performance analysis and Optimization of LARge  
Infrastructures and Systems

---

*In collaboration with* Laboratoire d'Informatique de Grenoble (LIG)



## **Project-Team POLARIS**

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

Each year, Inria research teams publish an Activity Report presenting their work and results over the reporting period. These reports follow a common structure, with some optional sections depending on the specific team. They typically begin by outlining the overall objectives and research programme, including the main research themes, goals, and methodological approaches. They also describe the application domains targeted by the team, highlighting the scientific or societal contexts in which their work is situated. The reports then present the highlights of the year, covering major scientific achievements, software developments, or teaching contributions. When relevant, they include sections on software, platforms, and open data, detailing the tools developed and how they are shared. A substantial part is dedicated to new results, where scientific contributions are described in detail, often with subsections specifying participants and associated keywords. Finally, the Activity Report addresses funding, contracts, partnerships, and collaborations at various levels, from industrial agreements to international cooperations. It also covers dissemination and teaching activities, such as participation in scientific events, outreach, and supervision. The document concludes with a presentation of scientific production, including major publications and those produced during the year.

## 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. – HPC for machine learning
- 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 networks
  - B6.2.2. – wireless networks
- B6.4. – Internet of things
- B8.3. – Urbanism and urban planning
- B9.6.7. – Geography
- B9.7.2. – Open data
- B9.8. – Reproducibility

## Contents

<b>Project-Team POLARIS</b>	<b>1</b>
<b>1 Team members, visitors, external collaborators</b>	<b>5</b>
<b>2 Overall objectives</b>	<b>6</b>
2.1 Context	7
2.2 Objectives	7
2.3 Contribution to AI/Learning	8
<b>3 Research program</b>	<b>8</b>
3.1 Performance Evaluation	8
3.2 Asymptotic Methods	10
3.3 Distributed Online Optimization and Learning in Games	11
3.4 Responsible Computer Science	13
<b>4 Application domains</b>	<b>15</b>
4.1 Large Computing Infrastructures	15
4.2 Next-Generation Wireless Networks	16
4.3 Energy and Transportation	16
4.4 Social Computing Systems	16
<b>5 Social and environmental responsibility</b>	<b>16</b>
5.1 Footprint of research activities	16
5.2 Impact of research results	17
<b>6 Highlights of the year</b>	<b>17</b>
6.1 Awards	17
6.2 PhD defense	17
<b>7 Latest software developments, platforms, open data</b>	<b>17</b>
7.1 Latest software developments	17
7.1.1 SimGrid	17
7.1.2 PSI	18
7.1.3 marmoteCore	18
7.1.4 MarTO	19
7.1.5 GameSeer	19
7.1.6 rmf_tool	19
7.1.7 SCIP	20
<b>8 New results</b>	<b>21</b>
8.1 Scheduling in Data Centers	21
8.2 Performance evaluation of Large Systems	22
8.2.1 Experimental practices and Simulation	22
8.2.2 Mean Field	23
8.3 Reinforcement Learning and MDP	23
8.4 Optimization Techniques	24
8.4.1 Langevin Sampling and Large Deviation Techniques	24
8.4.2 Frank-Wolfe	24
8.4.3 Mixed Integer Optimization	25
8.4.4 Application to various domains	26
8.4.5 Optimization for Machine Learning	27
8.5 Learning in games	27
8.5.1 Learning in stochastic games	28
8.5.2 Efficient Learning in Complex Landscape Games	28

8.5.3	Uncertainty and Robustness	29
8.6	Random matrix analysis and Machine Learning	30
8.7	Fairness and equity in digital (recommendation, advertising, persistent storage) systems	31
<b>9</b>	<b>Bilateral contracts and grants with industry</b>	<b>32</b>
<b>10</b>	<b>Partnerships and cooperations</b>	<b>32</b>
10.1	International initiatives	32
10.1.1	Inria associate team not involved in an IIL or an international program	32
10.2	International research visitors	33
10.2.1	Visits of international scientists	33
10.2.2	Visits to international teams	33
10.3	European initiatives	33
10.3.1	Horizon Europe	33
10.4	National initiatives	33
10.5	Regional initiatives	34
<b>11</b>	<b>Dissemination</b>	<b>34</b>
11.1	Promoting scientific activities	34
11.1.1	Scientific events: selection	34
11.1.2	Journal	35
11.1.3	Invited talks	35
11.1.4	Leadership within the scientific community	35
11.1.5	Research administration	35
11.2	Teaching - Supervision - Juries - Educational and pedagogical outreach	36
11.2.1	Supervision	37
11.2.2	Juries	37
11.3	Popularization	37
11.3.1	Specific official responsibilities in science outreach structures	37
<b>12</b>	<b>Scientific production</b>	<b>37</b>
12.1	Publications of the year	37
12.2	Cited publications	41

# 1 Team members, visitors, external collaborators

## Research Scientists

- Arnaud Legrand [Team leader, CNRS, Senior Researcher, HDR]
- Jonatha Anselmi [INRIA, Researcher, until Jul 2025]
- Dorian Baudry [INRIA, Researcher, from Jul 2025 until Jul 2025]
- Mathieu Besancon [INRIA, Researcher, until Jul 2025]
- Nicolas Gast [INRIA, Researcher, until Jul 2025, HDR]
- Bruno Gaujal [INRIA, Senior Researcher, until Jul 2025, HDR]
- Panayotis Mertikopoulos [CNRS, Senior Researcher, until Jul 2025, HDR]
- Bary Pradelski [CNRS, Researcher, until Jul 2025]

## Faculty Members

- Romain Couillet [UGA, Professor, until Mar 2025]
- Vincent Danjean [UGA, Associate Professor, until Feb 2025]
- Kevin Marquet [UGA, Associate Professor]
- Florence Perronnin [UGA, Associate Professor]
- Jean-Marc Vincent [UGA, Associate Professor, until Jul 2025]
- Philippe Waille [UGA, Associate Professor, until Feb 2025]

## Post-Doctoral Fellows

- Mete Ahunbay [CNRS, Post-Doctoral Fellow, from Apr 2025]
- Victor Boone [CNRS, Post-Doctoral Fellow, until Jun 2025]
- Dheeraj Narasimha [INRIA, Post-Doctoral Fellow, until Jun 2025]

## PhD Students

- Rim Alhawal [CNRS, until Jul 2025]
- Helene Arvis [EDF, CIFRE]
- Achille Baucher [UGA, until Sep 2025]
- Pierre-Louis Cauvin [CNRS, until Jul 2025]
- Romain Cravic [INRIA, until Jul 2025]
- Valentin Girard [UGA]
- Galaad Langlois [ENS DE LYON, from Sep 2025]
- Hugo Lebeau [Univ. Grenoble Alpes, until Jan 2025]
- Davide Legacci [UGA, until Oct 2025]
- Hubert Villuendas [UGA, until Jul 2025]

## Interns and Apprentices

- Samuel Bounan [INRIA, Intern, from Mar 2025 until Jul 2025]
- Samuel Bounan [INRIA, Intern, until Feb 2025]
- Joel Charles-Rebuffe [ENS DE LYON, Intern, from Apr 2025 until Jul 2025]
- Leandre Cheruy [UGA, from Apr 2025 until Jul 2025]
- Hamidou Diallo [INRIA, Intern, from Jun 2025 until Jul 2025]
- Hamidou Diallo [INRIA, Intern, from Feb 2025 until May 2025]
- Karl Gottlieb [UGA, from Feb 2025 until Aug 2025]
- Djamel Rassem Lamouri [INRIA, Intern, from Feb 2025 until Jul 2025]
- Galaad Langlois [ENS DE LYON, Intern, until Mar 2025]
- Adrien Obrecht [ENS DE LYON, Intern, until Jul 2025]

## Administrative Assistants

- Luce Coelho [INRIA]
- Annie Simon [INRIA]

## 2 Overall objectives

**Note:** The POLARIS project started in January 2016 and has ended in December 2025. This 10 year project has been an amazing human and scientific journey. Over this period, this group has enriched with many new members, both full-time researchers (6), postdocs & PhD students (60), and interns. The scope of scientific interests in POLARIS has continued to broaden (from HPC to learning, including aspects of privacy, statistics, energy optimization, combinatorial optimization, environmental and sustainability concerns, etc.), leading to a restructuring into two new Inria teams: **GHOST** (which has been created on August 1st 2025) and **ADN** (which has been created on January 1st 2026). Some members of POLARIS have also joined other teams such as CORSE (compilers), DATAMOVE (HPC), and DRAKKAR (networking and security). For convenience, the activity the members of the GHOST project until the end of 2025 (August-December 2025) has been reattached to the 2025 POLARIS activity report.

The **GHOST** (Games, Mathematical Optimization, and Stochastic Systems) targets the study of random dynamical systems, which are ubiquitous in many fields of computer science and applied mathematics: In machine learning for example, they are used to analyze learning algorithms and provide concrete convergence and generalization guarantees; in queuing theory, they are used to model, characterize, and optimize the performance of distributed systems; in game theory, they model the behavior of autonomous agents that are competing or cooperating to improve their performance; etc. GHOST will work at the interface of stochastic modeling, dynamical systems, online learning, game theory and optimization, and their aim will be to (a) design mathematical and algorithmic methods for studying the dynamics of complex systems in the presence of randomness and uncertainty; and (b) to use these methods to optimize performance, design new learning algorithms and optimize decision-making in all its aspects. The applications of GHOST mostly revolve around complex resource allocation problems such as job allocation in distributed computing resources, improving the methods used to train complex machine learning models, and the areas of energy management and scheduling in electrical networks.

The research of the **ADN** (Anthropocene, Degrowth, and ICT) builds on the observation that the environmental impacts of human activities have increased so dramatically since the beginning of the Industrial Revolution that they now represent a major driver of the Earth system, prompting the use of the term Anthropocene to describe this new epoch. *Which role do Information and Communication Technologies (ICT) play in this, and how could they most significantly contribute to mitigation and adaptation strategies for*

*tackling these environmental impacts?* The ADN project team seeks to address this question by rethinking ICT through a strong sustainability via degrowth approach. For this, the members of this project will (1) study the place and contribution of digital technologies in prospective scenarios, (2) taking into account their political nature through a value in design approach, and (3) with a focus on key software technologies and infrastructures as commons and dedigitization.

## 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 [113], BSC [58], UIUC [121], Sandia Nat. Lab. [119], ORNL [59] or ETH Zürich [89]) 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* [66], 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 [92, 91] or the HPL benchmark [64, 40]. We have shown that the performance (both for time and energy consumption [88]) 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 [51]. 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 [95, 47].

**Trace Analysis and Visualization** Many monolithic visualization tools have been developed by renowned HPC groups since decades (e.g., BSC [106], Jülich and TU Dresden [103, 61], UIUC [87, 109, 90] and ANL [120]) but most of these tools build on the classical information visualization [110] 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 [73], 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 [73] and more recently to a sparse direct solver [44]. 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-IoTLab 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 [60, 52] (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 [41, 67] 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 [74], 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 [78], 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 [80, 75].

**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 [62, 82, 81], 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 [56, 57, 55], 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 [56, 55] 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 [57] 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 [69], 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 [68], mean field games can be used to study how much vaccination should be subsidized to encourage people to adapt a socially optimal behaviour [48].

### 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 [77][42], where we showed that Whittle indices (as well as other index policies) become optimal at a geometric rate under the same technical conditions used by Weber and Weiss to prove Whittle’s conjecture, plus a technical requirement on the non-singularity of the fixed point of the mean-field dynamics. We also propose the first sub-cubic algorithm to compute Whittle and Gittins indexes. As for reinforcement learning in Markovian bandits, we have shown that Bayesian and optimistic approaches do not use the structure of Markovian bandits similarly: While Bayesian learning has both a regret and a computational complexity that scales linearly with the number of arms, optimistic approaches all incur an exponential computation time, at least in their current versions [76].

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, [99, 97] showed that the (optimal) class of “follow-the-regularized-leader” (FTRL) learning algorithms leads to Poincaré recurrence even in simple,  $2 \times 2$  min-max games, thus precluding convergence to Nash equilibrium in this context.

This negative result generated significant interest in the literature as it contributed in shifting the focus towards identifying *which* Nash equilibria may arise as stable limit points of FTRL algorithms and dynamics. Earlier work by POLARIS on the topic [65, 100, 101] suggested that *strict* Nash equilibria play an important role in this question. This suspicion was recently confirmed in a series of papers [72, 49] 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 [100], 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 [50] 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 Pradelski.

#### 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 [112]. 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) [63]. 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 [94], 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 [71]. 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 [70], we also study similar questions in the two-stage setting, and derive the optimal selector and the “price of local fairness” one pays in utility by imposing that the interim stage be fair.

**Privacy and transparency in social computing system** Online services in general, and social networks in particular, collect massive amounts of data about their users (both online and offline). It is critical that (i) the users’ data is protected so that it cannot leak and (ii) users can know what data the service has about them and understand how it is used—this is the transparency requirement. In this context, we did two kinds of work. First, we studied social networks through measurement, in particular using the use case of Facebook. We showed that their advertising platform, through the PII<sup>1</sup>-based targeting option, allowed attackers to discover some personal data of users [114]. 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) [54]. A follow-up work shed further light on the typical uses of the platform [53]. In another work, we proposed an innovative protocol based on randomized withdrawal to protect public posts deletion privacy [102]. Finally, in [79], 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, [98] considers

---

<sup>1</sup>Personally Identifiable Information

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 [108][43].

**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 [46][105] 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 [107] and a vaccination policy [104]). 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 [45]. Overall this work has been characterized by close interactions with 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 [111]. In fact, this is the reason why modern processors are equipped with Dynamic Voltage and Frequency Scaling (DVFS) technology [118]. 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 [85, 83, 84]) as well as on the experimental side (showing the gains of optimal policies over classical solutions [86]).

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 [117, 116, 115], (2) how to co-simulate an electrical network and a communication network [93], and (3) what is the performance of the communication protocol (PLC G3) used by the Linky smart meters [96].

## 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 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. Among one of his many activities, we can mention his involvement in the design of the *Numérique et Sciences Informatiques, NSI : les fondamentaux* MOOC.

# 6 Highlights of the year

## 6.1 Awards

The article *Multi-agent learning under uncertainty: Recurrence vs. concentration* [20] Panayotis Mertikopoulos and his co-authors was selected for a spotlight at NeurIPS 2025.

The article *Does Stochastic Gradient really succeed for Bandits?* [37] by Dorian Baudry and his coauthors has been admitted for oral presentation at the NeurIPS 2025 conference. There were 21575 valid paper submissions to the NeurIPS Main Track this year, of which the program committee accepted 5290 (24.52%) papers in total, with breakdown of 4525 as posters, 688 as spotlight and 77 as oral.

Mathieu Besançon has been a member of the team that won the *2025 Mixed-Integer Programming Workshop Computational Competition*.

Victor Boone received the *2025 UGA Academic Thesis Prize* for his research work among PhDs graduating in 2024.

## 6.2 PhD defense

Hugo Lebeau has defended his PhD thesis on March 2025 entitled *Random Matrix and Tensor Models for Large Data Processing* [29].

Romain Cravic has defended his PhD thesis on November 2025 entitled *Learning in stochastic games* [28].

# 7 Latest software developments, platforms, open data

## 7.1 Latest software developments

### 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, Fogs, 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. This tool can be used to assess safety 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 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.

**Release Contributions:** Breaking the seal: v4.0 was not the final release.

\* Allow one to unseal netzones to modify the platform even after the simulation start. \* The model-checker can now report memory race conditions (see tutorial). \* Pip builds should now work out of the box. \* (+ the usual bug fixes overall, and improvements to the Java/Python bindings).

**News of the Year:** There were 2 major releases in 2025. We released v4.0 in March, embodying 10 years of development. This turns SimGrid into a mature and stable research instrument. The users can easily extend this tool to adapt it to their specific research, while trusting its software implementation. Release v4.1 was published in November, showing that the development did not stall even if the framework is mostly in maintenance mode. The performance simulation mode was extended to allow modifications of the platform topology during the simulation.

Most of our work this year is related to the use of SimGrid for performance prediction of HPC applications and capacity planning of supercomputers.

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

**Publication:** [hal-04909441](#)

**Contact:** Martin Quinson

**Participants:** Mathieu Laurent, Anne-Cécile Orgerie, Arnaud Legrand, Augustin Degomme, Arnaud Giersch, Frédéric Suter, 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](https://github.com/ngast/rmf_tool)

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

**Contact:** Nicolas Gast

### 7.1.7 SCIP

**Name:** Solving Constraint Integer Programs

**Keywords:** Linear optimization, Mathematical Optimization, Mixed Integer Programming

**Functional Description:** SCIP is currently one of the fastest non-commercial solvers for mixed integer programming (MIP) and mixed integer nonlinear programming (MINLP). It is also a framework for constraint integer programming and branch-cut-and-price. It allows for total control of the solution process and the access of detailed information down to the guts of the solver.

**Release Contributions:** SCIP 10.0.0

Features and Performance Improvements

Exact Solving:

- added numerically exact solving mode for mixed-integer linear programs to the core framework including certification of branch-and-bound phase, - core extensions (new wrapper struct SCIP\_RATIONAL for rational arithmetic currently based on Boost, GMP, and MPFR, new data structure SCIP\_LPEXACT for handling rational LP relaxation and computing safe dual bounds, new interfaces to exact LP solvers SoPlex and QSopt\_ex, safe dualproof version of conflict analysis, new data structure SCIP\_CERTIFICATE for certificate printing/proof logging) - new plugins: new constraint handler "exactlinear" for handling linear constraints with rational data, new constraint handler "exactsol" to post-process and repair solutions from floating-point heuristics, - plugins revised for numerically exact solving mode: adjusted readers for MPS, LP, CIP, OPB/WBO, and ZIMPL files, extended presolver "milp" to perform rational presolving with PaPILO, adjusted constraint handler "integral" and default reliability pseudo-cost branching rule "relpscost", extended Gomory cut separator to separate and certify numerically safe MIR cuts, adjusted all primal heuristics (except for five dedicated MINLP heuristics), new interfaces to exact LP solvers SoPlex and QSopt\_ex

Symmetry Handling

- added more techniques to handle reflection symmetries, in particular, for orbitopes with column reflections and matrices whose rows and columns can be permuted by a symmetry - Dejavu can be used to compute symmetries, the source code is shipped with SCIP and incorporates sassy - implemented symmetry detection callbacks for disjunction and superindicator constraint handlers - detailed information about applied symmetry handling techniques can be printed to the terminal - improve memory usage by introducing different constraint handlers for full orbitopes and packing/partitioning orbitopes - symmetry detection no longer treats implicit integer variables separately, but computes symmetries based on the variable type inferred from variable bounds and implied integrality - extended the statistics to also include information about the number of variables (per type) affected by symmetry - implemented method to compute new permutations from a given list of symmetry group generators - cons\_orbisack, cons\_orbitope\_full, cons\_orbitope\_pp, and cons\_symresack now try to replace the stored aggregated variables by active ones at the end of presolving, this should reduce the size of copies of the resolved problem simplified symmetry detection graphs in case all edges have the same color

Presolve:

- distinguish implicit integrality of variables into strong and weak type, depending on whether integrality is implied for all feasible or only at least one optimal solution - added a new presolver "implint", which detects implied integral variables by detecting (transposed) network submatrices in the problem, for now, this plugin is disabled by default - added support for (transposed) network matrix detection allow multi-aggregation of unbounded slack variables, which may enable more bound tightening due to a reduction in the number of unbounded variables resolve all fixings in xor constraints also for an available integer variable

**URL:** <https://www.scipopt.org/>

**Contact:** Mathieu Besancon

**Partners:** TU Darmstadt, RWTH Aachen University, Friedrich-Alexander-Universität Erlangen-Nürnberg, Eindhoven University of Technology, University of Twente, University of Bayreuth, Forschungscampus Modal

## 8 New results

The new results produced by the team in 2025 can be grouped into the following categories.

### 8.1 Scheduling in Data Centers

**Participants:** Jonatha Anselmi, Bruno Gaujal, Nicolas Gast.

Queuing theory is a general modeling framework, originally developed by Erlang to model the system of calls at the Copenhagen Telephone Exchange Company, and which has later been extensively used to optimize telecommunications, traffic, the design of factories and shops, etc. It is particularly suited to the modeling and optimization of the operation of data-centers, clouds or HPC centers and has led to the development of a variety of effective and low-cost scheduling strategies. We contribute to this framework and extend it in relation with recent online learning techniques as well as with characteristics of modern workload.

**Non-Stationary Gradient Descent for Optimal Auto-Scaling in Serverless Platforms** To efficiently manage serverless computing platforms, a key aspect is the auto-scaling of services, i.e., the set of computational resources allocated to a service adapts over time as a function of the traffic demand. The objective is to find a compromise between user-perceived performance and energy consumption. In [3], we consider the scale-per-request auto-scaling pattern and investigate how many function instances (or servers) should be spawned each time an unfortunate job arrives, i.e., a job that finds all servers busy upon its arrival. We address this problem by following a stochastic optimization approach: we develop a stochastic gradient descent scheme of the Kiefer-Wolfowitz type that applies over a single run of the state evolution. At each iteration, the proposed scheme computes an estimate of the number of servers to spawn each time an unfortunate job arrives to minimize some cost function. Under natural assumptions, we show that the sequence of estimates produced by our scheme is asymptotically optimal almost surely. In addition, we prove that its convergence rate is  $O(n^{-2/3})$  where  $n$  is the number of iterations.

From a mathematical point of view, the stochastic optimization framework induced by auto-scaling exhibits non-standard aspects that we approach from a general point of view. We consider the setting where a controller can only get samples of the transient – rather than stationary – behavior of the underlying stochastic system. To handle this difficulty, we develop arguments that exploit properties of the mixing time of the underlying Markov chain. By means of numerical simulations, we validate the proposed approach and quantify its gain with respect to common existing scale-up rules.

**Autoscaling in Serverless Platforms via Online Learning with Convergence Guarantees** As the adoption of serverless computing platforms continue to grow, designing autoscaling policies that strike the right balance between energy efficiency and user-perceived performance has become a central challenge. In [35], we propose an online learning algorithm with theoretical convergence guarantees that dynamically tunes control parameters in a serverless autoscaling environment. The proposed algorithm, grounded in stochastic gradient descent, learns online-during the actual operation of the platform-the optimal values of three key control parameters: (i) the target stock size of prewarmed (idle) functions, (ii) the threshold triggering provisioning actions, and (iii) the expiration rate of idle resources. We prove that, under Markovian dynamics, the algorithm converges to the parameter set that minimizes a cost function capturing the tradeoff between energy consumption and response latency. In addition, we demonstrate that its structure naturally supports parallelization, significantly accelerating convergence.

Extensive numerical experiments show that our method outperforms existing baselines, including recent deep learning-based approaches, even under non-Markovian settings-highlighting both its robustness and practical viability for next-generation serverless infrastructures.

**Energy-Optimal Scheduling with Variable Processing Speed: The Role of Task Size Variability** In [2], we study the execution of a single task with an unknown size on a server with variable processing speed. Our goal is to analyze structural properties of the optimal energy consumption under the optimal speed profile that minimizes the expected energy consumption while meeting a hard deadline constraint. Specifically, we investigate how the task size probability distribution impacts the overall energy.

Under mild assumptions, our main result shows that the expected energy consumption induced by the optimal speed profile preserves the convex increasing order with respect to the task size distribution. Then, we leverage this property to derive simple bounds and conduct a worst-case analysis. In particular, we derive a simple, general formula for the energy gap induced by the 'best' and 'worst' task size distributions, expressed in terms of the support and expectation of the task size.

**Time-Constrained Energy Minimization for Online Execution of a Stochastic DAG Task** In [30], We study the problem of energy-efficient online execution of a complex task on a server with variable processing speed. The task consists of a set of stochastic elementary jobs structured as a Directed Acyclic Graph (DAG), where each job's execution may reveal new information that influences future scheduling decisions. Our objective is to determine an online speed control policy that minimizes the expected energy consumption while ensuring that the task completes before a strict deadline. Leveraging tools from convex optimization, the optimality principle, and backward induction, we derive a structural characterization of the optimal policy. We find that this is linked to a set of second-order differential equations and exhibits a non-trivial form. Building on this result, we develop a discretization-based algorithm that efficiently approximates the optimal policy. The proposed algorithm is provably asymptotically exact in the discretization step and has computational complexity  $O(|A| + KN)$ , where  $|A|$  and  $K$  denote the number of edges and vertices (i.e., jobs) in the underlying DAG, and  $N$  is the discretization granularity. Our results offer a principled and computationally efficient solution framework for online execution of structured stochastic workloads under strict energy and timing constraints.

## 8.2 Performance evaluation of Large Systems

**Participants:** Nicolas Gast, Arnaud Legrand.

### 8.2.1 Experimental practices and Simulation

**Lowering entry barriers to developing custom simulators of distributed applications and platforms with SimGrid** Researchers in parallel and distributed computing (PDC) often resort to simulation because experiments conducted using a simulator can be for arbitrary experimental scenarios, are less resource-, labor-, and time-consuming than their real-world counterparts, and are perfectly repeatable and observable. Many frameworks have been developed to ease the development of PDC simulators, and these frameworks provide different levels of accuracy, scalability, versatility, extensibility, and usability. The SimGrid framework [39] has been used by many PDC researchers to produce a wide range of simulators for over two decades. Its popularity is due to a large emphasis placed on accuracy, scalability, and versatility, and is in spite of shortcomings in terms of extensibility and usability. Although SimGrid provides sensible simulation models for the common case, it was difficult for users to extend these models to meet domain-specific needs. Furthermore, SimGrid only provided relatively low-level simulation abstractions, making the implementation of a simulator of a complex system a labor-intensive undertaking. In [6] we describe developments in the last decade that have contributed to vastly improving extensibility and usability, thus lowering or removing entry barriers for users to develop custom SimGrid simulators.

**Journée thématique du GDR RSD : Experimental Practices in the Systems and Networks Community** [33] This working group met on September 14, 2024, in Paris to discuss experimental practices in the Systems and Networks communities. The purpose of this document is fourfold:

- To reflect on the practices for conducting experimental research in this community;

- To share the best practices we have identified with the community;
- List available resources (apart from platforms, which are the subject of a separate workshop);
- Propose recommendations for disseminating these best practices and creating an “experimental culture.”

### 8.2.2 Mean Field

As a system of stochastically interacting entities grows large, the size of its state-space quickly explodes (curse of dimensionality) and both its analysis and optimization become intractable. Fortunately, symmetries and regularities can be exploited, in particular through the so-called Mean Field approximation, which averages out the state over degrees of freedom. This technique from statistical physics is particularly effective when studying computer systems and over the years, we have refined such approximation, proposed extensions, and developed fine analysis methods.

**Accuracy of the Graphon Mean Field Approximation for Interacting Particle Systems** In [1] we consider a system of  $N$  particles whose interactions are characterized by a (weighted) graph  $G^N$ . Each particle is a node of the graph with an internal state. This state changes according to Markovian dynamics that depend on the states of neighboring particles. We study the limiting properties of the state dynamics, focusing on the dense graph regime, in which the average degree of a node grows linearly with  $N$ . We show that, when  $G^N$  converges to a piecewise Lipschitz graphon  $G$ , the behavior of the system converges to a deterministic limit, the graphon mean field approximation. We obtain convergence rates depending on the system size  $N$  and cut-norm distance between  $G^N$  and  $G$ . We apply these results for two subcases: when  $G^N$  is a discretization of the graph  $G$  with individually weighted edges; when  $G^N$  is a random graph obtained by sampling edges with probabilities obtained from  $G$ . In the case of weighted interactions, we obtain a bound of order  $O(1/N)$ . In the random graph case, the error is of order  $O(\sqrt{\log(N)/N})$  with high probability. We illustrate the applicability of our results and the numerical efficiency of the approximation through two examples: a graph-based load-balancing model and a heterogeneous bike-sharing system.

## 8.3 Reinforcement Learning and MDP

**Participants:** Victor Boone, Dorian Baudry, Nicolas Gast, Bruno Gaujal.

A limitation of the classical queuing theory framework is that policies may require knowledge of the system parameters whereas such parameters are rarely known in advance, hence our interest in reinforcement learning technique and Markov Decision Processes. Over the years, we have developed analysis and modeling techniques specifically suited to queuing systems whose state space quickly explodes (learning in queues), which makes classical RL approaches or results inappropriate.

**Logarithmic regret of exploration in average reward Markov decision processes** In average reward Markov decision processes, state-of-the-art algorithms for regret minimization follow a well-established framework: They are model-based, optimistic and episodic. First, they maintain a confidence region from which optimistic policies are computed using a well-known subroutine called Extended Value Iteration (EVI). Second, these policies are used over time windows called episodes, each ended by the Doubling Trick (DT) rule or a variant thereof. In [13], without modifying EVI, we show that there is a significant advantage in replacing (DT) by another simple rule, that we call the Vanishing Multiplicative (VM) rule. When managing episodes with (VM), the algorithm’s regret is, both in theory and in practice, as good if not better than with (DT), while the one-shot behavior is greatly improved. More specifically, the management of bad episodes (when sub-optimal policies are being used) is much better under (VM) than (DT) by making the regret of exploration logarithmic rather than linear. These results are made possible by a new in-depth understanding of the contrasting behaviors of confidence regions during good and bad episodes.

## 8.4 Optimization Techniques

**Participants:** Mathieu Besançon, Dorian Baudry, Panayotis Mertikopoulos.

Optimization arises in almost every domain of computer science and requires a variety of techniques, to which we contribute both from a theoretical and practical perspective.

### 8.4.1 Langevin Sampling and Large Deviation Techniques

**The global convergence of stochastic gradient descent in non-convex landscapes: Sharp estimates via large deviations** In [11], we examine the time it takes for stochastic gradient descent (SGD) to reach the global minimum of a general, non-convex loss function. We approach this question through the lens of randomly perturbed dynamical systems and large deviations theory, and we provide a tight characterization of the global convergence time of SGD via matching upper and lower bounds. These bounds are dominated by the most "costly" set of obstacles that the algorithm may need to overcome in order to reach a global minimizer from a given initialization, coupling in this way the global geometry of the underlying loss landscape with the statistics of the noise entering the process. Finally, motivated by applications to the training of deep neural networks, we also provide a series of refinements and extensions of our analysis for loss functions with shallow local minima.

**Tamed Langevin sampling under weaker conditions** Motivated by applications to deep learning which often fail standard Lipschitz smoothness requirements, we examine in [22] the problem of sampling from distributions that are not log-concave and are only weakly dissipative, with log-gradients allowed to grow superlinearly at infinity. In terms of structure, we only assume that the target distribution satisfies either a log-Sobolev or a Poincaré inequality and a local Lipschitz smoothness assumption with modulus growing possibly polynomially at infinity. This set of assumptions greatly exceeds the operational limits of the "vanilla" unadjusted Langevin algorithm (ULA), making sampling from such distributions a highly involved affair. To account for this, we introduce a taming scheme which is tailored to the growth and decay properties of the target distribution, and we provide explicit non-asymptotic guarantees for the proposed sampler in terms of the Kullback-Leibler (KL) divergence, total variation, and Wasserstein distance to the target distribution.

**Contractive kinetic Langevin samplers beyond global Lipschitz continuity** In [36], we examine the problem of sampling from log-concave distributions with (possibly) superlinear gradient growth under kinetic (underdamped) Langevin algorithms. Using a carefully tailored taming scheme, we propose two novel discretizations of the kinetic Langevin SDE, and we show that they are both contractive and satisfy a log-Sobolev inequality. Building on this, we establish a series of non-asymptotic bounds in 2-Wasserstein distance between the law reached by each algorithm and the underlying target measure.

### 8.4.2 Frank-Wolfe

**Frank-Wolfe Algorithms: Sparsity Guarantees and an Application to Robust Optimization** Frank-Wolfe (FW) methods are a class of nonlinear optimization algorithms over a compact constraint set leveraging first-order information of the objective and linear optimization on the constraints. They have risen in popularity in the last decade for their applications in operations research and learning. In [27], we first present a recently proposed enhancement of all Frank-Wolfe algorithms that ensure that the iterates remain sparse, in the sense that they are formed as a convex combination of a low number of vertices. We then view an application of FW to robust optimization. Sparsity Guarantees for Frank-Wolfe. In a first part, we present recent progress on the pivoting framework that modifies Frank-Wolfe variants, ensuring the sparsity of the iterate. We then introduce the so-called active set identification property of some FW variants and how they can be leveraged to ensure sparser iterates in a so-called pivoting framework. Frank-Wolfe for robust optimization under an oracle model. In the second part of the talk, we focus on recent applications of FW or FW-inspired methods to tackle robust optimization under an oracle setting. We highlight how different approaches from the literature are connected through the lens of these algorithms or are dual of each other, including a recent simplicial decomposition, a cutting plane, and a FW approach with Nesterov smoothing.

**Improved algorithms and novel applications of the FrankWolfe.jl library** Frank-Wolfe (FW) algorithms have emerged as an essential class of methods for constrained optimization, especially on large-scale problems. In [4], we summarize the algorithmic design choices and progress made in the last years of the development of FrankWolfe.jl, a Julia package gathering high-performance implementations of state-of-the-art FW variants. We review key use cases of the library in the recent literature, which match its original dual purpose: first, becoming the de-facto toolbox for practitioners applying FW methods to their problem, and second, offering a modular ecosystem to algorithm designers who experiment with their own variants and implementations of algorithmic blocks. Finally, we demonstrate the performance of several FW variants on important problem classes in several experiments, which we curated in a separate repository for continuous benchmarking.

**Efficient Quadratic Corrections for Frank-Wolfe Algorithms** In [16], we develop a Frank-Wolfe algorithm with corrective steps, generalizing previous algorithms including blended conditional gradients, blended pairwise conditional gradients, and fully-corrective Frank-Wolfe. For this, we prove tight convergence guarantees together with an optimal face identification property. Furthermore, we propose two highly efficient corrective steps for convex quadratic objectives based on linear optimization or linear system solving, akin to Wolfe’s minimum-norm point, and show that they converge in finite time under suitable conditions. Beyond optimization problems that are directly quadratic, we revisit two algorithms - split conditional gradient and second-order conditional gradient sliding - which can leverage quadratic corrections to accelerate their quadratic subproblems. We demonstrate improved convergence rates for the first and broader applicability for the second, which may be of independent interest. Finally, we show substantial computational speedups for Frank-Wolfe-based algorithms with quadratic corrections across the considered problem classes.

**Secant Line Search for Frank-Wolfe Algorithms** In [17], we present a new step-size strategy based on the secant method for Frank-Wolfe algorithms. This strategy, which requires mild assumptions about the function under consideration, can be applied to any Frank-Wolfe algorithm. It is as effective as full line search and, in particular, allows for adapting to the local smoothness of the function but comes with a significantly reduced computational cost, leading to higher effective rates of convergence. We provide theoretical guarantees and demonstrate the effectiveness of the strategy through numerical experiments.

**The Pivoting Framework: Frank-Wolfe Algorithms with Active Set Size Control** In [25], we propose the pivoting meta algorithm (PM) to enhance optimization algorithms that generate iterates as convex combinations of vertices of a feasible region  $C \subseteq \mathbb{R}^n$ , including Frank-Wolfe (FW) variants. PM guarantees that the active set (the set of vertices in the convex combination) of the modified algorithm remains as small as  $\dim(C) + 1$  as stipulated by Carathéodory’s theorem. PM achieves this by reformulating the active set expansion task into an equivalent linear program, which can be efficiently solved using a single pivot step akin to the primal simplex algorithm; the convergence rate of the original algorithms are maintained. Furthermore, we establish the connection between PM and active set identification, in particular showing under mild assumptions that PM applied to the away-step Frank-Wolfe algorithm (AFW) or the blended pairwise Frank-Wolfe algorithm (BPFW) bounds the active set size by the dimension of the optimal face plus 1. We provide numerical experiments to illustrate practicality and efficacy on active set size reduction.

### 8.4.3 Mixed Integer Optimization

**Convex mixed-integer optimization with Frank-Wolfe methods** Mixed-integer nonlinear optimization encompasses a broad class of problems that present both theoretical and computational challenges. We propose in [8] a new type of method to solve these problems based on a branch-and-bound algorithm with convex node relaxations. These relaxations are solved with a Frank-Wolfe algorithm over the convex hull of mixed-integer feasible points instead of the continuous relaxation via calls to a mixed-integer linear solver as the linear minimization oracle. The proposed method computes feasible solutions while working on a single representation of the polyhedral constraints, leveraging the full extent of mixed-integer linear solvers without an outer approximation scheme and can exploit inexact solutions of node subproblems.

**The SCIP Optimization Suite 10.0** The SCIP Optimization Suite provides a collection of software packages for mathematical optimization, centered around the constraint integer programming (CIP) framework SCIP. [34] discusses the enhancements and extensions included in SCIP Optimization Suite 10.0. The updates in SCIP 10.0 include a new solving mode for exactly solving rational mixed-integer linear programs, a new presolver for detecting implied integral variables, a novel cut-based conflict analysis and separator for flower inequalities, two new heuristics, a novel tool for explaining infeasibility, a new interface for nonlinear solvers as well as improvements in symmetry handling, branching strategies, and SCIP's Benders' decomposition framework. SCIP Optimization Suite 10.0 also includes new and improved features in the the presolving library PaPILO, the parallel framework UG, and the decomposition framework GCG. Moreover, the SCIP Optimization Suite 10.0 contains MIP-DD, the first open-source delta debugger for mixed-integer programming solvers. These additions and enhancements have resulted in an overall performance improvement of SCIP in terms of solving time, number of nodes in the branch-and-bound tree, as well as the reliability of the solver.

#### 8.4.4 Application to various domains

**Integrating Aggregated Electric Vehicle Flexibilities in Unit Commitment Models using Submodular Optimization** The Unit Commitment (UC) problem consists in controlling a large fleet of heterogeneous electricity production units in order to minimize the total production cost while satisfying consumer demand. Electric Vehicles (EVs) are used as a source of flexibility and are often aggregated for problem tractability. In [32], we develop a new approach to integrate EV flexibilities in the UC problem and exploit the generalized polymatroid structure of aggregated flexibilities of a large population of users to develop an exact optimization algorithm, combining a cutting-plane approach and submodular optimization. We show in particular that the UC can be solved exactly in a time which scales linearly, up to a logarithmic factor, in the number of EV users when each production unit is subject to convex constraints. We illustrate our approach by solving a real instance of a long-term UC problem, combining open-source data of the European grid (European Resource Adequacy Assessment project) and data originating from a survey of user behavior of the French EV fleet.

**Efficient Sparse Flow Decomposition Methods for RNA Multi-Assembly** Decomposing a flow on a Directed Acyclic Graph (DAG) into a weighted sum of a small number of paths is an essential task in operations research and bioinformatics. This problem, referred to as Sparse Flow Decomposition (SFD), has gained significant interest, in particular for its application in RNA transcript multi-assembly, the identification of the multiple transcripts corresponding to a given gene and their relative abundance. Several recent approaches cast SFD variants as integer optimization problems, motivated by the NP-hardness of the formulations they consider. In [12], we propose an alternative formulation of SFD as a data fitting problem on the conic hull of the flow polytope. By reformulating the problem on the flow polytope for compactness and solving it using specific variants of the Frank-Wolfe algorithm, we obtain a method converging rapidly to the minimizer of the chosen loss function while producing a parsimonious decomposition. Our approach subsumes previous formulations of SFD with exact and inexact flows and can model different priors on the error distributions. Computational experiments show that our method outperforms recent integer optimization approaches in runtime, but is also highly competitive in terms of reconstruction of the underlying transcripts, despite not explicitly minimizing the solution cardinality.

**Mixed-Integer Optimization for Loopless Flux Distributions in Metabolic Networks** Constraint-based metabolic models can be used to investigate the intracellular physiology of microorganisms. These models couple genes to reactions, and typically seek to predict metabolite fluxes that optimize some biologically important metric. Classical techniques, like Flux Balance Analysis (FBA), formulate the metabolism of a microbe as an optimization problem where growth rate is maximized. While FBA has found widespread use, it often leads to thermodynamically infeasible solutions that contain internal cycles (loops). To address this shortcoming, Loopless-Flux Balance Analysis (ll-FBA) seeks to predict flux distributions that do not contain these loops. ll-FBA is a disjunctive program, usually reformulated as a mixed-integer program, and is challenging to solve for biological models that often contain thousands of reactions and metabolites. In [24], we compare various reformulations of ll-FBA and different solution approaches. Overall, the combinatorial Benders' decomposition is the most promising of the tested approaches with which we could solve most

instances. However, the model size and numerical instability pose a challenge to the combinatorial Benders' method.

#### 8.4.5 Optimization for Machine Learning

Many learning algorithms operate in centralized way, which raises many practical issues in terms of scalability, privacy, hence a high interest for designing efficient distributed and federated machine learning algorithms. Furthermore generally, the optimization space is also quite particular, which calls for specific regularization techniques and optimization algorithms.

**Non-stationary Bandit Convex Optimization: A Comprehensive Study** Bandit Convex Optimization is a fundamental class of sequential decision-making problems, where the learner selects actions from a continuous domain and observes a loss (but not its gradient) at only one point per round. In [38], we study this problem in non-stationary environments, and aim to minimize the regret under three standard measures of non-stationarity: the number of switches  $S$  in the comparator sequence, the total variation  $\Delta$  of the loss functions, and the path-length  $P$  of the comparator sequence. We propose a polynomial-time algorithm, Tilted Exponentially Weighted Average with Sleeping Experts (TEWA-SE), which adapts the sleeping experts framework from online convex optimization to the bandit setting. For strongly convex losses, we prove that TEWA-SE is minimax-optimal with respect to known  $S$  and  $\Delta$  by establishing matching upper and lower bounds. By equipping TEWA-SE with the Bandit-over-Bandit framework, we extend our analysis to environments with unknown non-stationarity measures. For general convex losses, we introduce a second algorithm, clipped Exploration by Optimization (cExO), based on exponential weights over a discretized action space. While not polynomial-time computable, this method achieves minimax-optimal regret with respect to known  $S$  and  $\Delta$ , and improves on the best existing bounds with respect to  $P$ .

**Model Predictive Control is Almost Optimal for Restless Bandit** In [15], we consider the discrete time infinite horizon average reward restless markovian bandit (RMAB) problem. We propose a *model predictive control* based non-stationary policy with a rolling computational horizon  $\tau$ . At each time-slot, this policy solves a  $\tau$  horizon linear program whose first control value is kept as a control for the RMAB. Our solution requires minimal assumptions and quantifies the loss in optimality in terms of  $\tau$  and the number of arms,  $N$ . We show that its sub-optimality gap is  $O(1/\sqrt{N})$  in general, and  $\exp(-\Omega(N))$  under a local-stability condition. Our proof is based on a framework from dynamic control known as *dissipativity*. Our solution easy to implement and performs very well in practice when compared to the state of the art. Further, both our solution and our proof methodology can easily be generalized to more general constrained MDP settings and should thus, be of great interest to the burgeoning RMAB community.

**Does Stochastic Gradient really succeed for Bandits?** Recent works of Mei et al. have deepened the theoretical understanding of the Stochastic Gradient Bandit (SGB) policy, showing that using a constant learning rate guarantees asymptotic convergence to the optimal policy, and that sufficiently small learning rates can yield logarithmic regret. However, whether logarithmic regret holds beyond small learning rates remains unclear. In [37], we take a step towards characterizing the regret regimes of SGB as a function of its learning rate. For two-armed bandits, we identify a sharp threshold, scaling with the suboptimality gap  $\Delta$ , below which SGB achieves logarithmic regret on all instances, and above which it can incur polynomial regret on some instances. This result highlights the necessity of knowing (or estimating)  $\Delta$  to ensure logarithmic regret with a constant learning rate. For general  $K$ -armed bandits, we further show the learning rate must additionally scale inversely with  $K$  to avoid polynomial regret. We introduce novel techniques to derive regret upper bounds for SGB, laying the groundwork for future advances in the theory of gradient-based bandit algorithms.

## 8.5 Learning in games

**Participants:** Davide Legacci, Panayotis Mertikopoulos, Bary Pradelski.

Learning in games naturally occurs in situations where the resources or the decision is distributed among several agents or even in situations where a centralized 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. Understanding and characterizing the geometry of such spaces is thus the key to propose efficient algorithms. This line of work has led to the defense of one PhD thesis in 2025.

### 8.5.1 Learning in stochastic games

**Learning in stochastic games** The thesis of Romain Cravic [28] addresses two learning problems in two-player zero-sum stochastic games. The first problem concerns a learner who embodies one of the two players and learns to play efficiently against an arbitrary opponent via a simulator of the game, whose internal parameters are unknown to the learner at the beginning of the learning process. The second problem concerns learning the Nash equilibrium of the game through self-play, assuming perfect knowledge of the considered model. For the first problem, we consider games with perfect information, meaning that players fully observe the current state of the game when making decisions. In contrast, we do not make this assumption for the second problem and instead consider games with imperfect information, where players must rely on the history of their past observations to infer the current state of the game and the information obtained by their opponent. For each of these two problems, we propose an algorithm that handles infinite-horizon games by linking them to the finite-horizon case through the introduction of a discount factor on the rewards produced by the game over time. More precisely, we propose the DONQ-learning algorithm for the first problem, which we refer to as black-box learning in games with perfect information, and the DOS-CFR algorithm for the second problem, which we refer to here as the solving of games with imperfect information. In both cases, the theoretical guarantees established in our analysis of these algorithms translate into probabilistic, sublinear regret bounds in the time horizon  $T$  as it tends to infinity. Finally, the thesis includes an experimental component consisting of developing an artificial intelligence capable of playing efficiently in a memory-bluff game: Robin Wood.

**A Quadratic Speedup in Finding Nash Equilibria of Quantum Zero-Sum Games** 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  $\epsilon$ -Nash equilibria in the  $4d$ -dimensional spectraplex. In [10], 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  $\epsilon$ -Nash equilibria. This quadratic speed-up relative to Jain and Watrous' original algorithm sets a new benchmark for computing  $\epsilon$ -Nash equilibria in quantum zero-sum games.

### 8.5.2 Efficient Learning in Complex Landscape Games

**Characterizing the Convergence of Game Dynamics via Potentialness** Understanding the convergence landscape of multi-agent learning is a fundamental problem of great practical relevance in many applications of artificial intelligence and machine learning. While it is known that learning dynamics converge to Nash equilibrium in potential games, the behavior of dynamics in many important classes of games that do not admit a potential is poorly understood. To measure how “close” a game is to being potential, we consider in [5] a distance function, that we call “potentialness”, and which relies on a strategic decomposition of games introduced by Candogan et al. (2011). We introduce a numerical framework enabling the computation of this metric, which we use to calculate the degree of “potentialness” in generic matrix games, as well as (non-generic) games that are important in economic applications, namely auctions and contests. Understanding learning in the latter games has become increasingly important due to the wide-spread automation of bidding and pricing with no-regret learning algorithms. We empirically show that potentialness decreases and concentrates with an increasing number of agents or actions; in addition, potentialness turns out to be a good

predictor for the existence of pure Nash equilibria and the convergence of no-regret learning algorithms in matrix games. In particular, we observe that potentialness is very low for complete-information models of the all-pay auction where no pure Nash equilibrium exists, and much higher for Tullock contests, first- and second-price auctions, explaining the success of learning in the latter. In the incomplete-information version of the all-pay auction, a pure Bayes-Nash equilibrium exists and it can be learned with gradient-based algorithms. Potentialness nicely characterizes these differences to the complete-information version.

**Efficient kernelized learning in polyhedral games beyond full-information: From Colonel Blotto to congestion games** In [18], we examine the problem of efficiently learning coarse correlated equilibria (CCE) in polyhedral games, that is, normal-form games with an exponentially large number of actions per player and an underlying combinatorial structure. Prominent examples of such games are the classical Colonel Blotto and congestion games. To achieve computational efficiency, the learning algorithms must exhibit regret and per-iteration complexity that scale polylogarithmically in the size of the players' action sets. This challenge has recently been addressed in the full-information setting, primarily through the use of kernelization. However, in the case of the realistic, but mathematically challenging, partial-information setting, existing approaches result in suboptimal and impractical runtime complexity to learn CCE. We tackle this limitation by building a framework based on the kernelization paradigm. We apply this framework to prominent examples of polyhedral games—namely the Colonel Blotto, graphic matroid and network congestion games—and provide computationally efficient payoff-based learning algorithms, which significantly improve upon prior works in terms of the runtime for learning CCE in these settings.

**Invariance and concentration properties of gradient-based learning in games** In [19], we study the long-run behavior of learning in strongly monotone games with stochastic, gradient-based feedback. For concreteness, we focus on the stochastic projected gradient (SPG) algorithm, and we examine the asymptotic distribution of its iterates when the method is run with constant step-size updates (the de facto choice for practical deployments of the algorithm). In contrast to variants of the method with a vanishing step-size case, SPG with a constant step-size does not converge: instead, it reaches a neighborhood of the game's Nash equilibrium at an exponential rate, and then, due to persistent noise, it fluctuates in its vicinity without converging (occasionally moving away on rare occasions). We provide a theoretical quantification of this behavior by analyzing the Markovian structure of the process. Namely, we show that, regardless of the algorithm's initialization, the distribution of its iterates converges at a geometric rate to a unique invariant measure which is concentrated in a neighborhood of the game's Nash equilibrium. More explicitly, we quantify the degree of this concentration and the rate of convergence of the algorithm's empirical frequency of play to the invariant measure of the process in Wasserstein distance, and we provide explicit bounds in terms of the method's step-size, the variance of the noise entering the process, and the geometric features of the game's payoff landscape.

### 8.5.3 Uncertainty and Robustness

**The impact of uncertainty on regularized learning in games** In [14], we investigate how randomness and uncertainty influence learning in games. Specifically, we examine a perturbed variant of the dynamics of "follow the regularized leader" (FTRL), where the players' payoff observations and strategy updates are continually impacted by random shocks. Our findings reveal that, in a fairly precise sense, "uncertainty favors extremes": in any game, regardless of the noise level, every player's trajectory of play reaches an arbitrarily small neighborhood of a pure strategy in finite time (which we estimate). Moreover, even if the player does not ultimately settle at this strategy, they return arbitrarily close to some (possibly different) pure strategy infinitely often. This prompts the question of which sets of pure strategies emerge as robust predictions of learning under uncertainty. We show that (a) the only possible limits of the FTRL dynamics under uncertainty are pure Nash equilibria; and (b) a span of pure strategies is stable and attracting if and only if it is closed under better replies. Finally, we turn to games where the deterministic dynamics are recurrent—such as zero-sum games with interior equilibria—and show that randomness disrupts this behavior, causing the stochastic dynamics to drift toward the boundary on average.

**Multi-agent learning under uncertainty: Recurrence vs. concentration** In [20], we examine the convergence landscape of multi-agent learning under uncertainty. Specifically, we analyze two stochastic

models of regularized learning in continuous games—one in continuous and one in discrete time—with the aim of characterizing the long-run behavior of the induced sequence of play. In stark contrast to deterministic, full-information models of learning (or models with a vanishing learning rate), we show that the resulting dynamics do not converge in general. In lieu of this, we ask instead which actions are played more often in the long run, and by how much. We show that, in strongly monotone games, the dynamics of regularized learning may wander away from equilibrium infinitely often, but they always return to its vicinity in finite time (which we estimate), and their long-run distribution is sharply concentrated around a neighborhood thereof. We quantify the degree of this concentration, and we show that these favorable properties may all break down if the underlying game is not strongly monotone—underscoring in this way the limits of regularized learning in the presence of persistent randomness and uncertainty.

**Robust equilibria in continuous games: From strategic to dynamic robustness** In [21], we examine the robustness of Nash equilibria in continuous games, under both strategic and dynamic uncertainty. Starting with the former, we introduce the notion of a robust equilibrium as those equilibria that remain invariant to small-but otherwise arbitrary-perturbations to the game’s payoff structure, and we provide a crisp geometric characterization thereof. Subsequently, we turn to the question of dynamic robustness, and we examine which equilibria may arise as stable limit points of the dynamics of "follow the regularized leader" (FTRL) in the presence of randomness and uncertainty. Despite their very distinct origins, we establish a structural correspondence between these two notions of robustness: strategic robustness implies dynamic robustness, and, conversely, the requirement of strategic robustness cannot be relaxed if dynamic robustness is to be maintained. Finally, we examine the rate of convergence to robust equilibria as a function of the underlying regularizer, and we show that entropically regularized learning converges at a geometric rate in games with affinely constrained action spaces.

**On the discrete-time origins of the replicator dynamics: From convergence to instability and chaos** In [7], we consider three distinct discrete-time models of learning and evolution in games: a biological model based on intra-species selective pressure, the dynamics induced by pairwise proportional imitation, and the exponential / multiplicative weights (EW) algorithm for online learning. Even though these models share the same continuous-time limit – the replicator dynamics – we show that second-order effects play a crucial role and may lead to drastically different behaviors in each model, even in very simple, symmetric  $2 \times 2$  games. Specifically, we study the resulting discrete-time dynamics in a class of parametrized congestion games, and we show that (i) in the biological model of intra-species competition, the dynamics remain convergent for any parameter value; (ii) the dynamics of pairwise proportional imitation exhibit an entire range of behaviors for larger time steps and different equilibrium configurations (stability, instability, and even Li-Yorke chaos); while (iii) in the EW algorithm, increasing the time step (almost) inevitably leads to chaos (again, in the formal, Li-Yorke sense). This divergence of behaviors comes in stark contrast to the globally convergent behavior of the replicator dynamics, and serves to delineate the extent to which the replicator dynamics provide a useful predictor for the long-run behavior of their discrete-time origins.

## 8.6 Random matrix analysis and Machine Learning

**Participants:** Romain Couillet, Hugo Lebeau, Victor Leger, Charles SejourneLeger.

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. This line of work has led to the defense of one PhD thesis in 2025.

**Random Matrix and Tensor Models for Large Data Processing** The exponential growth in computing power has enabled the widespread deployment of machine learning, which has in turn given rise to new challenges in data processing. The sheer volume of data now being generated means that the standard statistical assumption of a number of samples far greater than their dimension is no longer tenable. In the paradigm of the Big Data era, datasets are typically of very large dimension and may also comprise several

modes, indicating a variety of sources, modalities, domains, and so on. Furthermore, the advancement of technologies required to develop models capable of processing vast quantities of data results in significant environmental and human costs. In light of these concerns, it is imperative to promote a more clever and prudent use of our resources. Random matrix theory provides powerful tools to precisely study the statistical and computational limitations associated with the processing of large and multidimensional data. Through this lens, the thesis of Hugo Lebeau [29] examines several learning approaches to identify the relevant parameters influencing the success of a task and thereby facilitate an informed use.

Firstly, we examine an extension of spectral clustering to data streams. This approach enables the clustering of a potentially very large dataset with a controlled and limited memory usage. Our findings demonstrate that, with an astute management of the available memory, it is possible to achieve performance levels comparable to those obtained without resource constraints. We then turn our attention to the computational limits to tensor estimation, with a particular focus on low-rank approximation. The study the reconstruction performance of the truncated MLSVD (which generalizes the concept of truncated SVD to tensors) as well as the HOOI algorithm precisely describes the conditions required for the reconstruction of a noisy signal in multidimensional data. Additionally, we utilize a similar approach to investigate the multi-view clustering problem from the perspective of a rank-one tensor approximation. Our findings shed light on and precisely quantifies the pivotal role of view informativeness in the quality of the estimation. Lastly, we investigate the statistical limits to tensor estimation by introducing a matrix model associated to the maximum likelihood problem. This approach allows us to characterize the reconstruction performance of the corresponding estimator through a spectral analysis that can be performed with the standard tools of random matrix theory. To illustrate this method, we examine the best rank-one approximation of a tensor, where a given proportion of entries is randomly removed to reduce its memory cost. This allows us to quantify the impact of such a procedure on the quality of the resulting estimate. Finally, we propose to extend the presented method to a more general tensor estimation framework, which reveals attractive new challenges for the study of large random tensors.

**A Random Matrix Approach to Low-Multilinear-Rank Tensor Approximation** [9] presents a comprehensive understanding of the estimation of a planted low-rank signal from a general spiked tensor model near the computational threshold. Relying on standard tools from the theory of large random matrices, we characterize the large-dimensional spectral behavior of the unfoldings of the data tensor and exhibit relevant signal-to-noise ratios governing the detectability of the principal directions of the signal. These results allow to accurately predict the reconstruction performance of truncated multilinear SVD (MLSVD) in the non-trivial regime. This is particularly important since it serves as an initialization of the higher-order orthogonal iteration (HOOI) scheme, whose convergence to the best low-multilinear-rank approximation depends entirely on its initialization. We give a sufficient condition for the convergence of HOOI and show that the number of iterations before convergence tends to 1 in the large-dimensional limit.

**Performance of Rank-One Tensor Approximation on Incomplete Data** In [26], we are interested in the estimation of a rank-one tensor signal when only a portion  $\epsilon$  of its noisy observation is available. We show that the study of this problem can be reduced to that of a random matrix model whose spectral analysis gives access to the reconstruction performance. These results shed light on and specify the loss of performance induced by an artificial reduction of the memory cost of a tensor via the deletion of a random part of its entries.

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

**Participants:** Rémi Castera, Nicolas Gast, 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.

**Prophet Inequalities: Competing with the Top  $\ell$  Items is Easy** In [23], we explore a prophet inequality problem, where the values of a sequence of items are drawn i.i.d. from some distribution, and an online decision maker must select one item irrevocably. We establish that  $CR_\ell$  the worst-case competitive ratio between the expected optimal performance of an online decision maker compared to that of a prophet who uses the average of the top  $\ell$  items is exactly the solution to an integral equation. This quantity  $CR_\ell$  is larger than  $1 - e^{-\ell}$ . This implies that the bound converges exponentially fast to 1 as  $\ell$  grows. In particular for  $\ell = 2$ ,  $CR_2 \approx 0.966$ , which is much closer to 1 than the classical bound of 0.745 for  $\ell = 1$ . Additionally, we prove asymptotic lower bounds for the competitive ratio of a more general scenario, where the decision maker is permitted to select  $k$  items. This subsumes the  $k$  multi-unit i.i.d. prophet problem and provides the current best asymptotic guarantees, as well as enables broader understanding in the more general framework. Finally, we prove a tight asymptotic competitive ratio when only static threshold policies are allowed.

## 9 Bilateral contracts and grants with industry

**Participants:** Nicolas Gast.

Nicolas Gast participates to the "Defi EDF" and is currently supervising a PhD student (Hélène Arvis) via a "CIFRE" contract.

## 10 Partnerships and cooperations

### 10.1 International initiatives

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

##### AIRBA

**Participants:** Nicolas Gast.

**Title:** AI for restless bandits and its application

**Duration:** 2023 – 2025

**Coordinator:** Gupta Manu Kumar

**Partners:**

- Indian Institute of Technology Roorkee Roorkee (Inde)

**Inria contact:** Nicolas Gast

**Summary:** Multi-armed restless bandit problems (MARBPs) are Markov decision process models for optimal dynamic priority allocation to a collection of stochastic binary-action (active/passive) projects evolving over time. Typical applications include maintenance problems, in which a collection of agents must be send to various objects subject to failures, or stochastic scheduling problems. If MARBPs are in general intractable, there exists efficient relaxation when the problems parameters are known. The goal of this project is to build on recent progress on Reinforcement Learning to create new tools to solve this problem when the parameters are unknown. Problems of interests are: how to define online indices and how to learn them; What is the performance of such online policies; Application of these to real-life examples such as machine-repairman problem of dynamic asset allocations.

## 10.2 International research visitors

### 10.2.1 Visits of international scientists

**Other international visits to the team** We have hosted Deborah Hendrych (Zuse Institute Berlin) from July 1st to August 1st 2025, Mohamed Ghannan (Zuse Institute Berlin) in July, and Abednego Kambale (Politecnico di Milano) from April to June.

### 10.2.2 Visits to international teams

#### Research stays abroad

- Panayotis Mertikopoulos taught a two-week invited course at the recently inaugurated Moroccan Center for Game Theory (MCGT) at UM6P, Rabat. As a result of this course, R. Laraki (CNRS/UM6P) and Panayotis Mertikopoulos will jointly co-supervise the PhD of Omar Abbadi (co-tutelle between UM6P and UGA)
- Panayotis Mertikopoulos spent three weeks as a research visitor at the Archimedes AI research center in Athens, Greece

## 10.3 European initiatives

### 10.3.1 Horizon Europe

Panayotis Mertikopoulos participated in the submission of the QU-METRIC proposal to the call HORIZON-EIC-2025-PATHFINDEROPEN. The proposal was unsuccessful; unclear if there will be a follow-up with the same consortium.

## 10.4 National initiatives

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

### ANR

**ANR REFINO (JCJC 2020-2025)★** *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.

## 10.5 Regional initiatives

**Participants:** Bruno Gaujal.

**MIAI Cluster chair:** Fundamentals of Reinforcement Learning

- PI: Pierre Gaillard and Bruno Gaujal
- Members: Nicolas Gast and Jean-Philippe GAYON (UCA)
- The goal of a reinforcement learning algorithm is to gather information about the unknown system being explored by the learner to better understand its dynamical properties and exploit them to optimize its behavior. Whenever the learner has an a priori offline information about the system, it can leverage this knowledge to be more efficient in learning its optimal behavior. This approach is coined by the global concept of structured learning.

This leads us to the research question that we want to tackle with the FunRL project: *How to design algorithms with optimal theoretical guarantees that exploit a (known or unknown) structure of the problem to solve?* This question will be developed in three directions. First, we will tackle the online control of queueing networks, which raises the important issue of stability and rarely visited states. The as Markov decision processes (MDPs), which are stochastic dynamical systems that can be controlled. The main originality of this axe with respect to the others is that these dynamical systems are constrained by the structure of the problem, the challenge being to efficiently use our knowledge of such a structure. Third we will study parametric learning, where a learner adapts its policy to a problem with a known structure but whose parameters are unknown. This has applications to auto-scaling problems in cloud computing, resource allocation, and sequential decisions.

## 11 Dissemination

### 11.1 Promoting scientific activities

#### 11.1.1 Scientific events: selection

Panayotis Mertikopoulos has been senior area chair at NeurIPS 2025 and area chair at ICML 2025.

#### Chair of conference program committees

- Nicolas Gast is TPC co-chair (ACM SIGMETRICS 2026)

#### Member of the conference program committees

- Jonatha Anselmi has been a PC member of [MASCOTS 2025](#).
- Mathieu Besancon has been a PC member of [JuMP-dev 2025](#), [JuliaCon Local Paris 2025](#), [European Mixed Integer Programming Workshop 2025](#), and [AAAI 2026](#)
- Bruno Gaujal has been a PC member of Sigmetrics, NeurIPS

#### Reviewer

- Mathieu Besancon has been a reviewer for Mathematical Programming and Journal of Global Optimization.
- Jonatha Anselmi has been a reviewer for IEEE transactions on parallel and distributed systems, IEEE transactions on cloud computing, Journal of Applied probability, IEEE transactions on networking.

### 11.1.2 Journal

#### Member of the editorial boards

- Nicolas Gast is associate editor of the journals *Performance Evaluation* and *Stochastic Models*.
- Panayotis Mertikopoulos is managing co-editor of the *Open Journal of Mathematical Optimization* (OJMO) and associate editor at *Mathematics of Operations Research* (MOR), the *International Journal of Game Theory* (IJGT), the *EURO Journal on Computational Optimization* (EJCO), and *Operations Research Letters* (ORL).

### 11.1.3 Invited talks

- Arnaud Legrand was invited to give lectures and keynotes on Reproducible Research and Open Science on the following occasions:
  - Congrès de la ROADEF, Champ-sur-Marne (Feb. 2025)
  - École thématique *Science Ouverte pour les SHS*, Oléron (June 2025)
  - M1 students in computer science at UGA (Dec. 2025): *Reproducible Research and Computer Science*
  - 1st year Inria PhD students (Nov. 2025)
- Bruno Gaujal and Nicolas Gast were invited to give a tutorial on *Stochastic bandits* at Performance 2025 (Amsterdam).
- Bruno Gaujal was invited to a roundtable on *the future of performance evaluation* at Atelier Evaluation des Performances (Toulouse).
- Panayotis Mertikopoulos was invited to present his work at the following events:
  - Invited conference talk: 75 years of Nash equilibrium Oxford, UK
  - Invited workshop talk: 3rd Paris Workshop on Game Theory and Language
  - Tutorial talk at the 7th Annual Conference on Learning for Dynamics & Control
  - Invited seminar talks at: LSE, MCGT

### 11.1.4 Leadership within the scientific community

- Mathieu Besancon is a member of the MOBIDEC PEPR in the **ACME** project.
- 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).
- Jean-Marc Vincent is a member of the scientific committee of the CIST.
- Jean-Marc Vincent is vice-head of the SIF, adjunct on teaching. In this context he is in charge of the organization of the annual meeting on education in computer science.

### 11.1.5 Research administration

- Mathieu Besancon has been a member of the hiring committee for an Assistant Professor at Univ. Clermont Auvergne.
- Nicolas Gast is vice-head of the Labex *EnergyAlps* that federates the community working on electrical energy in Grenoble.
- Nicolas Gast is vice-head of the *école doctorale MSTII* (the doctoral school managing PhD students in computer science and mathematics at Univ. Grenoble Alpes).

- Bruno Gaujal was president of the hiring community for and Assistant Professor at Grenoble INP-GI.
- 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. He has been particularly involved in the HCERES evaluation process of the LIG.
- Arnaud Legrand has been a member of the HCERES evaluation committee of the IRIT.
- Arnaud Legrand is a member of *Comité Scientifique* of the Inria Grenoble.
- Florence Perronin is a member of the QVT team of the LIG.
- Jean-Marc Vincent was a member of the jury for Mcf hiring in Univ. Nancy.

## 11.2 Teaching - Supervision - Juries - Educational and pedagogical outreach

- Jonatha Anselmi teaches Probabilités et simulation (32h) and Évaluation de performances (32h) at PolyTech Grenoble.
- Mathieu Besançon teaches half of the Advanced Models and Methods of Operations Research at the M2 ORCO (UGA) and in the Master de Mathématiques et Applications (UGA).
- Nicolas Gast teaches the *Reinforcement learning* part of the M2 course *Mathematical foundations of machine learning* at the M2 MOSIG (Grenoble).
- Bruno Gaujal teaches *Optimisation under uncertainties* (18h) at the M2 ORCO (UGA).
- Bruno Gaujal and Nicolas Gast teach *Markov Decision Process and Reinforcement Learning* (32h in total) at the M2 Info (ENS Lyon).
- Nicolas Gast is responsible of the course "Introduction au Machine Learning" that is an optional course of the third year or the "licence informatique" at Univ. Grenoble Alpes.
- Nicolas Gast is co-responsible of the courses "MDP and reinforcement learning" (M2, Ens Lyon), "Mathematical Foundations of machine learning"(M2, UGA), and "Optimization under uncertainties" (M2, UGA).
- Arnaud Legrand and Jean-Marc Vincent teach the transversal *Scientific Methodology and Empirical Evaluation lecture* (18h) at the M2 MOSIG (UGA).
- The 3rd edition of the MOOC of Arnaud Legrand, K. Hinsin 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 2026), more than 25,600 persons have followed this MOOC and about 3200 certificates of achievement have been delivered (about 14% for the last session, which is very high for a MOOC). More than half of participants are PhD students and about 10% are undergraduates.
- The 2nd edition of the MOOC of Arnaud Legrand, K. Hinsin and C. Pouzat on *Reproducible Research II: Practices and tools for managing computations and data* has been launched from May 2025 to October 2025. It has attracted 3,000 persons.
- Florence Perronin teaches *Programming Languages* in L1.
- Florence Perronin is a member of the *conseil de perfectionnement* of the Mathematics license.
- Jean-Marc Vincent is in charge of the coordination of the training of high school teachers in computer science (NSI) in Grenoble.
- Jean-Marc Vincent teaches *Algorithms and Probabilities* at the L3, UGA.
- Jean-Marc Vincent teaches *Statistical Models and Literate Programming* at the L3 MIAGE, UGA.

- Jean-Marc Vincent teaches *Mathematics for Computer Science* at the M1 MOSIG, UGA.
- Jean-Marc Vincent participates to the *Histoire de l'informatique* lecture at the ENSIMAG.
- Panayotis Mertikopoulos organized and taught a PhD seminar course on Game Theory at MCGT (Moroccan Center for Game Theory (MCGT) at UM6P, Rabat).

### 11.2.1 Supervision

Bruno Gaujal is member of the CSI of Vittorio Puricelli (Toulouse) and Alessia Rigonat (Paris).

### 11.2.2 Juries

- Nicolas Gast was president of the PhD jury of Bianca Marin Moreno (Univ. Grenoble Alpes) entitled *Apprentissage par renforcement convexe en ligne et applications aux problèmes de gestion de l'énergie*.
- Nicolas Gast was reviewer for the PhD of Thomas Le Corre (ENS PSL) entitled *Distributed control of flexible loads in power grids*.
- Bruno Gaujal was a member of the jury of the PhD defense of Lucas Weber (Paris) entitled *Exploiting Partial System Knowledge in Reinforcement Learning for Admission Control and Electricity Storage Optimization*, president of the jury of the PhD of Andrea Fox (Avignon) entitled *Reinforcement Learning for Resource Allocation in Edge/Fog Systems*, and reviewer for the PhD of Maria Cherifa (Saclay) entitled *Dynamics and learning in some online allocation problems* and of the PhD of Chiara Mignacco (Saclay) entitled *A mathematical study of policy orchestration for reinforcement learning*.
- Bruno Gaujal was president of the HDR jury of Pierre Gaillard (Grenoble)
- Arnaud Legrand was a reviewer for the PhD of Léo Cosseron (ÉNS Rennes): *Time-accurate network simulation interconnecting virtual machines with hardware virtualization towards stealth analysis*.
- Arnaud Legrand was president of the PhD jury of Eduardo Tomasi Ribeiro (Univ. Grenoble Alpes): *Single address space for 128-bit massively parallel computers*.

## 11.3 Popularization

### 11.3.1 Specific official responsibilities in science outreach structures

- Jean-Marc Vincent is the head of the evaluation committee for *education in AI* of the *MIAI* chaire.

## 12 Scientific production

### 12.1 Publications of the year

#### International journals

- [1] S. Allmeier and N. Gast. 'Accuracy of the Graphon Mean Field Approximation for Interacting Particle Systems'. In: *Stochastic Systems* 15.4 (Dec. 2025), pp. 273–290. DOI: [10.1287/stsy.2024.0070](https://doi.org/10.1287/stsy.2024.0070). URL: <https://inria.hal.science/hal-05457094> (cit. on p. 23).
- [2] J. Anselmi and B. Gaujal. 'Energy-Optimal Scheduling with Variable Processing Speed: The Role of Task Size Variability'. In: *Operations Research Letters* 64 (2nd Dec. 2025), p. 107379. DOI: [10.1016/j.orl.2025.107379](https://doi.org/10.1016/j.orl.2025.107379). URL: <https://hal.science/hal-05136334> (cit. on p. 22).
- [3] J. Anselmi, B. Gaujal and L.-S. Rebuffi. 'Non-Stationary Gradient Descent for Optimal Auto-Scaling in Serverless Platforms'. In: *IEEE/ACM Transactions on Networking* (2025), pp. 1–14. DOI: [10.1109/TON.2025.3538982](https://doi.org/10.1109/TON.2025.3538982). URL: <https://hal.science/hal-04942350> (cit. on p. 21).

- [4] M. Besançon, S. Designolle, J. Halbey, D. Hendrych, D. Kuzinowicz, S. Pokutta, H. Troppens, D. Viladrich Herrmannsdoerfer and E. Wirth. ‘Improved algorithms and novel applications of the FrankWolfe.jl library’. In: *ACM Transactions on Mathematical Software* (3rd Sept. 2025), pp. 1–31. DOI: [10.1145/3765626](https://doi.org/10.1145/3765626). URL: <https://inria.hal.science/hal-05250102> (cit. on p. 25).
- [5] M. Bichler, D. Legacci, M. Oberlechner, P. Mertikopoulos and B. Pradelski. ‘Characterizing the Convergence of Game Dynamics via Potentialness’. In: *Transactions on Machine Learning Research Journal* (Mar. 2025), pp. 1–22. URL: <https://hal.science/hal-05133709> (cit. on p. 28).
- [6] H. Casanova, A. Giersch, A. Legrand, M. Quinson and F. Suter. ‘Lowering entry barriers to developing custom simulators of distributed applications and platforms with SimGrid’. In: *Parallel Computing* 123 (Mar. 2025), p. 103125. DOI: [10.1016/j.parco.2025.103125](https://doi.org/10.1016/j.parco.2025.103125). URL: <https://hal.science/hal-04909441> (cit. on p. 22).
- [7] F. Falniowski and P. Mertikopoulos. ‘On the discrete-time origins of the replicator dynamics: From convergence to instability and chaos’. In: *International Journal of Game Theory* 54.7 (27th Feb. 2025). DOI: [10.1007/s00182-025-00929-3](https://doi.org/10.1007/s00182-025-00929-3). URL: <https://hal.science/hal-04629311> (cit. on p. 30).
- [8] D. Hendrych, H. Troppens, M. Besançon and S. Pokutta. ‘Convex mixed-integer optimization with Frank–Wolfe methods’. In: *Mathematical Programming Computation* (28th June 2025), pp. 1–49. DOI: [10.1007/s12532-025-00288-w](https://doi.org/10.1007/s12532-025-00288-w). URL: <https://hal.science/hal-05136465> (cit. on p. 25).
- [9] H. Lebeau, F. Chatelain and R. Couillet. ‘A Random Matrix Approach to Low-Multilinear-Rank Tensor Approximation’. In: *Journal of Machine Learning Research* 26.7 (2025), pp. 1–64. URL: <https://hal.science/hal-04673321> (cit. on p. 31).
- [10] 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: *Quantum* 9 (6th May 2025), p. 1737. DOI: [10.22331/q-2025-05-06-1737](https://doi.org/10.22331/q-2025-05-06-1737). URL: <https://inria.hal.science/hal-05136745> (cit. on p. 28).

#### International peer-reviewed conferences

- [11] W. Azizian, F. Iutzeler, J. Malick and P. Mertikopoulos. ‘The global convergence of stochastic gradient descent in non-convex landscapes: Sharp estimates via large deviations’. In: *Proceedings of the 42nd International Conference on Machine Learning*. ICML 2025 - 42nd International Conference on Machine Learning. Vancouver, Canada, 2025, pp. 1–67. URL: <https://inria.hal.science/hal-05136755> (cit. on p. 24).
- [12] M. Besançon. ‘Efficient Sparse Flow Decomposition Methods for RNA Multi-Assembly’. In: *ECAI 2025 - 28th European Conference on Artificial Intelligence*. ECAI 2025 - 28th European Conference on Artificial Intelligence. Bologna, Italy, 2025, pp. 1–8. URL: <https://hal.science/hal-05173160> (cit. on p. 26).
- [13] V. Boone and B. Gaujal. ‘Logarithmic regret of exploration in average reward Markov decision processes’. In: *Proceedings of Thirty Eighth Conference on Learning Theory*. Proceedings of Thirty Eighth Conference on Learning Theory. Vol. 291. 291. Lyon, France, 2025, pp. 454–533. URL: <https://hal.science/hal-05352833> (cit. on p. 23).
- [14] P.-L. Cauvin, D. Legacci and P. Mertikopoulos. ‘The impact of uncertainty on regularized learning in games’. In: *Proceedings of the 42nd International Conference on Machine Learning*. ICML 2025 - 42nd International Conference on Machine Learning. Vancouver, Canada, 2025, pp. 1–50. URL: <https://inria.hal.science/hal-05136772> (cit. on p. 29).
- [15] N. Gast and D. Narasimha. ‘Model Predictive Control is Almost Optimal for Restless Bandit’. In: *COLT 2025 - 38th Annual Conference on Learning Theory*. Lyon, France, 2025, pp. 1–35. DOI: [10.48550/arXiv.2410.06307](https://doi.org/10.48550/arXiv.2410.06307). URL: <https://inria.hal.science/hal-05194338> (cit. on p. 27).

- [16] J. Halbey, S. Rakotomandimby, M. Besançon, S. Designolle and S. Pokutta. ‘Efficient Quadratic Corrections for Frank-Wolfe Algorithms’. In: NeurIPS. San Diego, United States, 2nd Dec. 2025. URL: <https://hal.science/hal-05426705> (cit. on p. 25).
- [17] D. Hendrych, S. Pokutta, M. Besançon and D. Martínez-Rubio. ‘Secant Line Search for Frank-Wolfe Algorithms’. In: *Proceedings of Machine Learning Research*. ICML 2025 - Forty-Second International Conference on Machine Learning. Vancouver (BC), Canada, 13th July 2025, pp. 1–22. URL: <https://hal.science/hal-05086780> (cit. on p. 25).
- [18] A. Kontogiannis, V. Pollatos, G. Farina, P. Mertikopoulos and I. Panageas. ‘Efficient kernelized learning in polyhedral games beyond full-information: From Colonel Blotto to congestion games’. In: NeurIPS 2025 - Proceedings of the 39th International Conference on Neural Information Processing Systems, 2025. San Diego, United States, Sept. 2025. URL: <https://hal.science/hal-05410395> (cit. on p. 29).
- [19] K. Lotidis, P. Mertikopoulos, N. Bambos and J. Blanchet. ‘Invariance and concentration properties of gradient-based learning in games’. In: CDC 2025 - 64th IEEE Annual Conference on Decision and Control. Rio de Janeiro, Brazil: IEEE, 2025, pp. 1–8. URL: <https://hal.science/hal-05410392> (cit. on p. 29).
- [20] K. Lotidis, P. Mertikopoulos, N. Bambos and J. Blanchet. ‘Multi-agent learning under uncertainty: Recurrence vs. concentration’. In: NeurIPS 2025 - Proceedings of the 39th International Conference on Neural Information Processing Systems, 2025. San Diego, United States, 9th Dec. 2025. URL: <https://hal.science/hal-05410394> (cit. on pp. 17, 29).
- [21] K. Lotidis, P. Mertikopoulos, N. Bambos and J. Blanchet. ‘Robust equilibria in continuous games: From strategic to dynamic robustness’. In: NeurIPS 2025 - Proceedings of the 39th International Conference on Neural Information Processing Systems, 2025. San Diego, United States, 2025, pp. 1–33. URL: <https://hal.science/hal-05410393> (cit. on p. 30).
- [22] I. Lytras and P. Mertikopoulos. ‘Tamed Langevin sampling under weaker conditions’. In: AISTATS 2025 - 28th International Conference on Artificial Intelligence and Statistics. Mai Khao, Thailand, 2025, pp. 1–32. URL: <https://hal.science/hal-04629313> (cit. on p. 24).
- [23] M. Molina, N. Gast, P. Loiseau and V. Perchet. ‘Prophet Inequalities: Competing with the Top  $\ell$  Items is Easy’. In: SODA 2025 - ACM-SIAM Symposium on Discrete Algorithms. New Orleans, United States: ACM, 2025, pp. 1–45. URL: <https://hal.science/hal-04912640> (cit. on p. 32).
- [24] H. Troppens, M. Besançon, S. E. Wilken and S. Pokutta. ‘Mixed-Integer Optimization for Loopless Flux Distributions in Metabolic Networks’. In: *23rd International Symposium on Experimental Algorithms (SEA 2025)*. SEA 2025 - 23rd International Symposium on Experimental Algorithms. Vol. 338. Venice, Italy: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. DOI: [10.4230/LIPIcs.SEA.2025.26](https://doi.org/10.4230/LIPIcs.SEA.2025.26). URL: <https://inria.hal.science/hal-05432362> (cit. on p. 26).
- [25] E. Wirth, M. Besançon and S. Pokutta. ‘The Pivoting Framework: Frank-Wolfe Algorithms with Active Set Size Control’. In: *Proceedings of Machine Learning Research*. AISTATS 2025 - 28th International Conference on Artificial Intelligence and Statistics. Mai Khao, Thailand, 3rd May 2025, pp. 1–24. URL: <https://hal.science/hal-04913437> (cit. on p. 25).

#### National peer-reviewed Conferences

- [26] H. Lebeau. ‘Performance of Rank-One Tensor Approximation on Incomplete Data’. In: GRETSI 2025 – XXXème Colloque Francophone de Traitement du Signal et des Images. Strasbourg, France, 2025, pp. 1–4. URL: <https://hal.science/hal-05247960> (cit. on p. 31).

#### Conferences without proceedings

- [27] M. Besançon. ‘Frank-Wolfe Algorithms: Sparsity Guarantees and an Application to Robust Optimization’. In: ROADEF 2025 - 26ème Congrès annuel de la Société française de recherche opérationnelle et d’aide à la décision. Champs-Sur-Marne, France, 25th Feb. 2025, pp. 1–1. URL: <https://hal.science/hal-04979715> (cit. on p. 24).

### Doctoral dissertations and habilitation theses

- [28] R. Cravic. ‘Learning in stochastic games’. UGA - Université Grenoble Alpes, 3rd Nov. 2025. URL: <https://theses.hal.science/tel-05489070> (cit. on pp. 17, 28).
- [29] H. Lebeau. ‘Random Matrix and Tensor Models for Large Data Processing’. Université Grenoble Alpes [2020-....], 11th Mar. 2025. URL: <https://theses.hal.science/tel-05222408> (cit. on pp. 17, 31).

### Reports & preprints

- [30] J. ANSEMI, B. Gaujal and K. Gottlieb. *Time-Constrained Energy Minimization for Online Execution of a Stochastic DAG Task*. 2025. URL: <https://cnrs.hal.science/hal-05245021> (cit. on p. 22).
- [31] N. Appel, J. Bourdon, N. Bousquet, J. Cohen, A. Genitrini, P. Georgeon, Y. Grandvalet, K. Jaffrès-Runser, A. Legrand, D. Markham, A. Muscholl, A. Paparrizou, L. Paulevé, M. Poss, M. Gradinariu Potop-Butucaru, J.-F. Raymond, R. Rouvoy, Y. Sallent, P. Senellart, T. Seiller, Y.-Q. Song, A. Tchana and H. Waeselync. *Section 06 Sciences de l’information : fondements de l’informatique, calculs, algorithmes, représentations, exploitations: Rapport de conjoncture 2024*. CNRS, 2025, pp. 1–20. URL: <https://inria.hal.science/hal-05238890>.
- [32] H. Arvis, O. Beaude, N. Gast, S. Gaubert and B. Gaujal. *Integrating Aggregated Electric Vehicle Flexibilities in Unit Commitment Models using Submodular Optimization*. 14th Nov. 2025. URL: <https://inria.hal.science/hal-05399441> (cit. on p. 26).
- [33] M. Bacou, D. Beserra, E. Dedu, L. Desgeorges, D. Donsez, A. Guitton, B. Jonglez, A. Legrand, G. Papadopoulos, O. Richard, S. Si-Mohammed, N. Tamdrari and F. Theoleyre. *Journée thématique du GDR RSD : pratiques expérimentales de la communauté systèmes et réseaux*. 31st Jan. 2025. URL: <https://hal.science/hal-04924273> (cit. on p. 22).
- [34] C. Hojny, M. Besançon, K. Bestuzheva, S. Borst, A. Chmiela, J. Dionísio, J. Ehls, L. Eifler, M. Ghannam, A. Gleixner, A. Göß, A. Hoen, J. von Holly-Ponientzietz, R. van der Hulst, D. Kamp, T. Koch, K. Kofler, J. Lentz, M. Lübbecke, S. J. Maher, P. M. Meinhold, G. Mexi, T. Mohr, E. Mühmer, K. K. Patel, M. E. Pfetsch, S. Pokutta, C. R. Groba, F. Serrano, Y. Shinano, M. Turner, S. Vigerske, M. Walter, D. Weninger and L. Xu. *The SCIP Optimization Suite 10.0*. Zuse Institute Berlin, 23rd Nov. 2025. URL: <https://inria.hal.science/hal-05397675> (cit. on p. 26).
- [35] A. W. Kambale, J. ANSEMI, D. Ardagna and B. Gaujal. *Autoscaling in Serverless Platforms via Online Learning with Convergence Guarantees*. 2025. URL: <https://cnrs.hal.science/hal-05245047> (cit. on p. 21).
- [36] I. Lytras and P. Mertikopoulos. *Contractive kinetic Langevin samplers beyond global Lipschitz continuity*. Sept. 2025. URL: <https://hal.science/hal-05410397> (cit. on p. 24).

### Other scientific publications

- [37] D. Baudry, E. Johnson, S. Vary, C. Pike-Burke and P. Rebeschini. ‘Does Stochastic Gradient really succeed for Bandits?’ In: *NeurIPS 2025 - 39th Conference on Neural Information Processing Systems*. San Diego (Californie - EU), United States, 2nd Dec. 2025. URL: <https://hal.science/hal-05446148> (cit. on pp. 17, 27).
- [38] X. Liu, D. Baudry, J. Zimmert, P. Rebeschini and A. Akhavan. ‘Non-stationary Bandit Convex Optimization: A Comprehensive Study’. In: *NeurIPS 2025 - 39th Conference on Neural Information Processing Systems*. San Diego (Californie - EU), United States, 2nd Dec. 2025. URL: <https://hal.science/hal-05446174> (cit. on p. 27).

## Software

- [39] [SW] H. Casanova, A. Degomme, A. Giersch, A. Legrand, M. Quinson and F. Suter, *SimGrid* version 4.0, 18th Apr. 2025. LIC: GNU Lesser General Public License v2.1 or later. HAL: ([hal-05040017](https://hal.inria.fr/hal-05040017)), URL: <https://inria.hal.science/hal-05040017>, vcs: <https://framagit.org/simgrid/simgrid/>, SWHID: ([swh:1:rev:3714440fcff4a244f3d51dc0495bab153a5c298c](https://sw.hicet.fr/rev/3714440fcff4a244f3d51dc0495bab153a5c298c);origin=<https://framagit.org/simgrid/simgrid.git>;visit=swh:1:snp:f2414fd9d467641d7a695ed64e3d2fd779087095) (cit. on p. 22).

## 12.2 Cited publications

- [40] 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> (cit. on p. 9).
- [41] 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> (cit. on p. 10).
- [42] 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> (cit. on p. 12).
- [43] 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> (cit. on p. 15).
- [44] 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> (cit. on p. 10).
- [45] M. Oliu-Barton, B. Pradelski, P. Aghion, P. Artus, I. Kickbusch, J. Lazarus, D. Sridhar and S. Vanderslott. ‘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> (cit. on p. 15).
- [46] 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> (cit. on p. 15).
- [47] 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> (cit. on p. 9).
- [48] 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> (cit. on p. 11).
- [49] 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> (cit. on p. 13).
- [50] 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> (cit. on p. 13).

- [51] 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> (cit. on p. 9).
- [52] 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> (cit. on p. 10).
- [53] 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> (cit. on p. 14).
- [54] 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> (cit. on p. 14).
- [55] 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> (cit. on p. 11).
- [56] 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> (cit. on p. 11).
- [57] 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> (cit. on p. 11).
- [58] 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 (cit. on p. 9).
- [59] S. Böhm and C. Engelmann. ‘xSim: The Extreme-Scale Simulator’. In: *HPCS*. Istanbul, Turkey, 2011 (cit. on p. 9).
- [60] 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.2019.000026](https://doi.org/10.1109/CCGRID.2019.000026). URL: <https://hal.inria.fr/hal-02110868> (cit. on p. 10).
- [61] 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](http://dx.doi.org/10.1007/978-3-642-11261-4_2) (cit. on p. 9).
- [62] 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> (cit. on p. 11).
- [63] 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> (cit. on p. 14).

- [64] 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> (cit. on p. 9).
- [65] 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> (cit. on p. 13).
- [66] 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> (cit. on p. 9).
- [67] 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> (cit. on p. 10).
- [68] 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> (cit. on p. 11).
- [69] 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> (cit. on p. 11).
- [70] 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> (cit. on p. 14).
- [71] 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> (cit. on p. 14).
- [72] L. Flokas, E. V. Vlatakis-Gkaragkounis, T. Lianas, 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> (cit. on p. 13).
- [73] 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> (cit. on p. 10).
- [74] 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> (cit. on p. 11).
- [75] 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> (cit. on p. 11).
- [76] 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> (cit. on p. 12).

- [77] N. Gast, B. Gaujal and C. Yan. ‘Exponential Convergence Rate for the Asymptotic Optimality of Whittle Index Policy’. Dec. 2020. URL: <https://hal.inria.fr/hal-03041176> (cit. on p. 12).
- [78] 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> (cit. on p. 11).
- [79] 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> (cit. on p. 14).
- [80] 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> (cit. on p. 11).
- [81] 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> (cit. on p. 11).
- [82] 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> (cit. on p. 11).
- [83] 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> (cit. on p. 15).
- [84] B. Gaujal, A. Girault and S. Plassart. ‘Discrete and Continuous Optimal Control for Energy Minimization in Real-Time Systems’. In: *EBCSP 2020 - 6th International Conference on Event-Based Control, Communication, and Signal Processing*. Krakow, Poland: IEEE, Sept. 2020, pp. 1–8. DOI: [10.1109/EBCSP51266.2020.9291353](https://doi.org/10.1109/EBCSP51266.2020.9291353). URL: <https://hal.archives-ouvertes.fr/hal-03020330> (cit. on p. 15).
- [85] 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> (cit. on p. 15).
- [86] 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> (cit. on p. 15).
- [87] M. Heath and J. Etheridge. ‘Visualizing the performance of parallel programs’. In: *IEEE software* 8.5 (1991) (cit. on p. 9).
- [88] 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> (cit. on p. 9).
- [89] 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 (cit. on p. 9).
- [90] 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) (cit. on p. 9).
- [91] 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> (cit. on p. 9).

- [92] 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> (cit. on p. 9).
- [93] 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> (cit. on p. 15).
- [94] 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 (cit. on p. 14).
- [95] 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> (cit. on p. 9).
- [96] 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> (cit. on p. 15).
- [97] 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> (cit. on p. 13).
- [98] 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> (cit. on p. 14).
- [99] 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> (cit. on p. 13).
- [100] 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> (cit. on p. 13).
- [101] 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> (cit. on p. 13).
- [102] 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> (cit. on p. 14).
- [103] W. Nagel, A. Arnold, M. Weber, H. Hoppe and K. Solchenbach. ‘VAMPIR: Visualization and Analysis of MPI Resources’. In: *Supercomputer* 12.1 (1996) (cit. on p. 9).
- [104] 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> (cit. on p. 15).
- [105] 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> (cit. on p. 15).
- [106] 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 (cit. on p. 9).

- [107] 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> (cit. on p. 15).
- [108] 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> (cit. on p. 15).
- [109] 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 (cit. on p. 9).
- [110] B. Shneiderman. ‘The eyes have it: A task by data type taxonomy for information visualizations’. In: *IEEE Symposium on Visual Languages*. IEEE. 1996 (cit. on p. 9).
- [111] 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 (cit. on p. 15).
- [112] 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> (cit. on p. 14).
- [113] 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 (cit. on p. 9).
- [114] 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> (cit. on p. 14).
- [115] B. Vinot, F. Cadoux and R. Héliot. ‘Decentralized Optimization of Energy Exchanges in an Electricity Microgrid’. In: *ACM e-Energy 2016 - 7th ACM International Conference on Future Energy Systems*. Waterloo, Canada, June 2016. DOI: [10.1145/2939912.2939915](https://doi.org/10.1145/2939912.2939915). URL: <https://hal.archives-ouvertes.fr/hal-01340394> (cit. on p. 15).
- [116] B. Vinot, F. Cadoux and R. Héliot. ‘Decentralized optimization of energy exchanges in an electricity microgrid’. In: *2016 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe)*. Ljubljana, Slovenia: IEEE, Oct. 2016, pp. 1–6. DOI: [10.1109/ISGTEurope.2016.7856319](https://doi.org/10.1109/ISGTEurope.2016.7856319). URL: <https://hal.inria.fr/hal-01784148> (cit. on p. 15).
- [117] 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> (cit. on p. 15).
- [118] 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 (cit. on p. 15).
- [119] 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 (cit. on p. 9).
- [120] 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> (cit. on p. 9).
- [121] G. Zheng, G. Kakulapati and L. Kalé. ‘BigSim: A Parallel Simulator for Performance Prediction of Extremely Large Parallel Machines’. In: *IPDPS*. 2004 (cit. on p. 9).