

RESEARCH CENTRE

Grenoble - Rhône-Alpes

IN PARTNERSHIP WITH:

Université de Grenoble Alpes, CNRS

2021

ACTIVITY REPORT

Project-Team

POLARIS

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

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

DOMAIN

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

THEME

**Distributed and High Performance
Computing**

Contents

Project-Team POLARIS	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	3
2.1 Context	3
2.2 Objectives	4
2.3 Contribution to AI/Learning	4
3 Research program	5
3.1 Performance Evaluation	5
3.2 Asymptotic Methods	7
3.3 Distributed Online Optimization and Learning in Games	8
3.4 Responsible Computer Science	10
4 Application domains	12
4.1 Large Computing Infrastructures	12
4.2 Next-Generation Wireless Networks	13
4.3 Energy and Transportation	13
4.4 Social Computing Systems	13
5 Social and environmental responsibility	13
5.1 Footprint of research activities	13
5.2 Sens Workshop	14
5.3 Raising awareness on the climate crisis	14
5.4 Impact of research results	14
6 Highlights of the year	14
6.1 Awards	14
7 New software and platforms	15
7.1 New software	15
7.1.1 SimGrid	15
7.1.2 PSI	16
7.1.3 marmoteCore	16
7.1.4 MarTO	17
7.1.5 GameSeer	17
7.1.6 rmf_tool	17
8 New results	17
8.1 HPC application analysis and optimization	18
8.2 Large system analysis and optimization	19
8.3 Energy optimization	20
8.4 Learning in time varying systems	21
8.5 Advanced learning methods	22
8.6 Adaptive algorithms	23
8.7 Learning in Games	24
8.8 Bandits	25
8.9 Fairness and equity in digital (recommendation, advertising, persistent storage) systems	26
8.10 Policy responses to Covid-19	27
9 Bilateral contracts and grants with industry	28

10 Partnerships and cooperations	28
10.1 International initiatives	28
10.1.1 Associate Teams in the framework of an Inria International Lab or in the framework of an Inria International Program	28
10.2 National initiatives	29
11 Dissemination	30
11.1 Promoting scientific activities	30
11.1.1 Scientific events: organisation	30
11.1.2 Scientific events: selection	31
11.1.3 Journal	31
11.1.4 Invited talks	31
11.1.5 Research Administration	32
11.2 Teaching - Supervision - Juries	32
11.2.1 Teaching	33
11.3 Popularization	34
12 Scientific production	34
12.1 Publications of the year	34
12.2 Cited publications	37

Project-Team POLARIS

Creation of the Project-Team: 2018 January 01

Keywords

Computer sciences and digital sciences

- A1.2. – Networks
- A1.3.5. – Cloud
- A1.3.6. – Fog, Edge
- A1.6. – Green Computing
- A3.4. – Machine learning and statistics
- A3.5.2. – Recommendation systems
- A5.2. – Data visualization
- A6. – Modeling, simulation and control
- A6.2.3. – Probabilistic methods
- A6.2.4. – Statistical methods
- A6.2.6. – Optimization
- A6.2.7. – High performance computing
- A8.2. – Optimization
- A8.9. – Performance evaluation
- A8.11. – Game Theory
- A9.2. – Machine learning
- A9.9. – Distributed AI, Multi-agent

Other research topics and application domains

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

1 Team members, visitors, external collaborators

Research Scientists

- Arnaud Legrand [Team leader, CNRS, Senior Researcher, HDR]
- Jonatha Anselmi [Inria, Researcher]
- Nicolas Gast [Inria, Researcher, HDR]
- Bruno Gaujal [Inria, Senior Researcher, HDR]
- Patrick Loiseau [Inria, Researcher, HDR]
- Panayotis Mertikopoulos [CNRS, Researcher, HDR]
- Bary Pradelski [CNRS, Researcher]

Faculty Members

- Romain Couillet [Institut polytechnique de Grenoble, Professor, from Sep 2021]
- Vincent Danjean [Univ Grenoble Alpes, Associate Professor]
- Guillaume Huard [Univ Grenoble Alpes, Associate Professor]
- Florence Perronnin [Univ Grenoble Alpes, Associate Professor, HDR]
- Jean-Marc Vincent [Univ Grenoble Alpes, Associate Professor]
- Philippe Waille [Univ Grenoble Alpes, Associate Professor]

Post-Doctoral Fellows

- Henry Joseph Audeoud [INPG Entreprise SA]
- Dong Quan Vu [CNRS]

PhD Students

- Sebastian Allmeier [Inria]
- Kimon Antonakopoulos [CNRS]
- Thomas Barzola [Univ Grenoble Alpes]
- Victor Boone [École Normale Supérieure de Lyon, from Sep 2021]
- Remi Castera [Univ Grenoble Alpes, from Oct 2021]
- Tom Cornebize [Inria, until Mar 2021]
- Romain Cravic [Inria, from Oct 2021]
- Vitalii Emelianov [Inria]
- Yu Guan Hsieh [Univ Grenoble Alpes]
- Simon Philipp Jantschgi [Université de Zurich]
- Kimang Khun [Inria]
- Till Kletti [Naver Labs, CIFRE]

- Lucas Leandro Nesi [Federal University of the Rio Grande Do Sul (UFRGS), from Nov 2021]
- Hugo Lebeau [Univ Grenoble Alpes, from Oct 2021]
- Victor Leger [Univ Grenoble Alpes, from Oct 2021]
- Dimitrios Moustakas [Institut polytechnique de Grenoble]
- Louis Sebastien Rebuffi [Univ Grenoble Alpes]
- Pedro Rocha Bruel [Université de Sao Paulo - Brésil, until Oct 2021]
- Benjamin Roussillon [Univ Grenoble Alpes, until Sep 2021]
- Vera Sosnovik [Univ Grenoble Alpes]
- Chen Yan [Univ Grenoble Alpes]

Technical Staff

- Bruno De Moura Donassolo [Inria, Engineer]
- Eleni Gkiouzepe [Univ Grenoble Alpes, Engineer, until Nov 2021]

Interns and Apprentices

- Achille Baucher [Univ Grenoble Alpes, from Oct 2021]
- Victor Boone [École Normale Supérieure de Lyon, from Feb 2021 until Jul 2021]
- Remi Castera [Inria, from Apr 2021 until Sep 2021]
- Romain Cravic [Inria, from Feb 2021 until Jul 2021]
- Mael Delorme [Inria, from May 2021 until Jun 2021]
- Aurelien Gauffre [Univ Grenoble Alpes, from Feb 2021 until Jul 2021]
- Jeremy Guerin [Inria, from Apr 2021 until Sep 2021]
- Oumaima Hajji [Inria, from May 2021 until Jul 2021]
- Mathieu Molina [Inria, from May 2021 until Oct 2021]
- Julie Reynier [Inria, from May 2021 until Jul 2021]

Administrative Assistant

- Annie Simon [Inria]

2 Overall objectives

2.1 Context

Large distributed infrastructures are rampant in our society. Numerical simulations form the basis of computational sciences and high performance computing infrastructures have become scientific instruments with similar roles as those of test tubes or telescopes. Cloud infrastructures are used by companies in such an intense way that even the shortest outage quickly incurs the loss of several millions of dollars. But every citizen also relies on (and interacts with) such infrastructures via complex wireless mobile embedded devices whose nature is constantly evolving. In this way, the advent of digital

miniaturization and interconnection has enabled our homes, power stations, cars and bikes to evolve into smart grids and smart transportation systems that should be optimized to fulfill societal expectations.

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

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

2.2 Objectives

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

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

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

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

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

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

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 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 (see Figure). The resulting algorithms should not only learn an efficient (Nash) equilibrium but they should also be able of doing so quickly (low regret), even when facing the difficulties associated to a distributed context (lack of coordination, uncertain world, information delay, limited feedback, ...)

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

3 Research program

3.1 Performance Evaluation

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

Project-team positioning

Evaluating the scalability, robustness, energy consumption and performance of large infrastructures such as exascale platforms and clouds raises severe methodological challenges. The complexity of such platforms mandates empirical evaluation but direct experimentation via an application deployment on a real-world testbed is often limited by the few platforms available at hand and is even sometimes impossible (cost, access, early stages of the infrastructure design, etc.). Furthermore, such experiments are costly, difficult to control and therefore difficult to reproduce. Although many of these digital systems have been built by men, they have reached such a complexity level that we are no longer able to study them like artificial systems and have to deal with the same kind of experimental issues as natural sciences. The development of a sound experimental methodology for the evaluation of resource management solutions is among the most important ways to cope with the growing complexity of computing environments. Although computing environments come with their own specific challenges, we believe such general observation problems should be addressed by borrowing good practices and techniques developed in many other domains of science, 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[103], BSC[49], UIUC[111], Sandia Nat. Lab.[109], ORNL[50] or ETH Zürich[79]) but most of them builds 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*[57], 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 a legacy Geophysics application that suffers from spatial and temporal load balancing problem [82, 81] or the HPL benchmark [55][37]. We have shown that the performance (both for time and energy consumption [78]) 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 [29]. 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[85, 112].

Trace Analysis and Visualization Many monolithic visualization tools have been developed by renowned HPC groups since decades (e.g., BSC [96], Jülich and TU Dresden [93, 52], UIUC [77, 99, 80] and ANL [110]) but most of these tools build on the classical information visualization [100] 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/SMPD 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 [64], 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 [64] and more recently to a sparse direct solver [43]. 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 execution

Design of Experiments and Reproducibility Part of our work is devoted to the control of experiments on both classical (HPC) and novel (IoT/Fog in a smart home context) infrastructures. To this end, we heavily rely on experimental testbeds such as Grid5000 and FIT-IoT Lab that can be well-controlled but real experiments are nonetheless quite resource-consuming. *Design of experiments* has been successfully applied in many fields (e.g., agriculture, chemistry, industrial processes) where experiments are considered expensive. Building on concrete use cases, we explore how *Design of Experiments* and *Reproducible Research* techniques can be used to (1) design transparent auto-tuning strategies of scientific computation kernels [51][34] (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 degradations of more than 10%, and had yet been unnoticed by both Grid'5000 technical team and Grid'5000 users (3) design and evaluate the performance of service provisioning strategies [4][58] 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 form 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 lead 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 of the 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 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 [65], 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 [68], 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 [70, 66].

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 [53, 72, 71], 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 obtain 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 [47, 48, 46], 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 [47, 46] 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 [48] 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 [60], we develop the notion of discrete space mean field games, that is more amendable to study that the previously introduced notions of mean field games. This lead to two interesting contributions: mean field games are not always the limits of stochastic games as the number of players grow [59], mean field games can be used to study how much vaccination should be subsidized to encourage people to adapt an socially optimal behaviour [18].

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 [67][40], 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 [39].

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, [89, 87] showed that the (optimal) class of “follow-the-regularized-leader” (FTRL) learning algorithms leads to Poincaré recurrence even in simple, 2×2 min-max games, thus precluding convergence to Nash equilibrium in this context.

This negative result generated significant interest in the literature as it contributed in shifting the focus towards identifying *which* Nash equilibria may arise as stable limit points of FTRL algorithms and dynamics. Earlier work by POLARIS on the topic [56, 90, 91] suggested that *strict* Nash equilibria play an important role in this question. This suspicion was recently confirmed in a series of papers [63][19] 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 [90], 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 [25] 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 [102]. 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) [54]. 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 [84], 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 [62]. 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 [61], we also study similar questions in the two-stage setting, and derive the optimal selector and the “price of local fairness” one pays in utility by imposing that the interim stage be fair.

Privacy and transparency in social computing system Online services in general, and social networks in particular, collect massive amounts of data about their users (both online and offline). It is critical that (i) the users’ data is protected so that it cannot leak and (ii) users can know what data the service has about them and understand how it is used—this is the transparency requirement. In this context, we did two kinds of work. First, we studied social networks through measurement, in particular using the use case of Facebook. We showed that their advertising platform, through the PII¹-based targeting option, allowed attackers to discover some personal data of users [104]. 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) [45]. A follow-up work shed further light on the typical uses of the platform [44]. In another work, we proposed an innovative protocol based on randomized withdrawal to protect public posts deletion privacy [92]. Finally, in [69], 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, [88] considers a matching market where buyers and sellers have heterogeneous preferences over each other. Traders arrive at random to the market and the market maker, having limited information, aims to optimize when to open the market for a clearing event to take place. There is a tradeoff between thickening the market (to achieve better matches) and matching quickly (to reduce waiting time of traders in the market). The tradeoff is made explicit for a wide range of underlying preferences. These works are adding to an ongoing effort to better understand and design markets [98][8].

¹Personally Identifiable Information

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 [10][95] 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 [97] and a vaccination policy [94]). 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 [9]. 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 [101]. In fact, this is the reason why modern processors are equipped with Dynamic Voltage and Frequency Scaling (DVFS) technology [108]. 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 [75, 73, 74]) as well as on the experimental side (showing the gains of optimal policies over classical solutions [76]).

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 [107, 106, 105], (2) how to co-simulate an electrical network and a communication network [83], and (3) what is the performance of the communication protocol (PLC G3) used by the Linky smart meters [86].

4 Application domains

4.1 Large Computing Infrastructures

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

4.2 Next-Generation Wireless Networks

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

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

4.3 Energy and Transportation

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

4.4 Social Computing Systems

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

5 Social and environmental responsibility

5.1 Footprint of research activities

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

Along this line, an interesting initiative lies in the PhD work of Tom Cornebize [33] who evaluated the carbon footprint of this thesis. The total amount of greenhouse gas emissions of this thesis due to airplane transportation is 7.6 t of CO₂eq and he estimates that the total amount of greenhouse gas emissions of this thesis due to computing is 10.6 t of CO₂eq (a total of 2,112,014 core hours on Grid5000). The initial goal

of the thesis was to evaluate (and possibly improve) the reliability of SimGrid to predict the performance of an MPI application like HPL. About 10% have been devoted to running SMPI simulations (to obtain predictions), 30% to run MPI applications (so as to compare the predictions with real executions), and 60% for a systematic performance measurement of a dozen of Grid5000 clusters over a long period of time (over two years). We are still unsure about whether this last part was worth the cost but this long and large-scale measurement was motivated by the need to characterize the variability of modern platforms, which is documented but largely underestimated by observational studies. Although we made our best to decrease the cost of this experiment from the start and proceeded gradually, the overall cost remains high. This study raises the question of the responsibility of the verification of the platform state (user vs. administrator) and we expect the lessons learned will be useful to the community.

5.2 Sens Workshop

On November 19th, a Sens Workshop was held, organized within the Inria DataMove and Polaris teams. We were ten participants. All participants were permanent members of one of the two teams. Participation in the workshop was on a voluntary basis. The day's proceedings followed four main axes: (1) Why do you do research? (2) Construction of a map of the expectations of everyone in the team. (3) Selection of two texts to be read and exchange around questions about the goal of research; (3) Prospective.

The first axe took place mainly in small groups. We were interested in why we work in the academic world and why this subject. It emerged that the motivations are very varied (ranging from intellectual curiosity to the desire to change the world), as well as the desire of several participants to change their object of study. The second session aimed to map the goals and constraints that bind us to our profession as researchers and to organize these different themes into a mural. Rigorous scientific production and education seem to be at the center of our priorities, while a question remains about the lack of group emulation and the too great part of individualism (and competition) in the current academic world. The last axes were the occasion to exchange around several texts on the impact of our profession of researcher in the current digital world, in particular linked to the fact that digital technology deeply modifies human activities and relationships, with a strong societal and environmental impact.

Without taking a concrete direction, this day was rich in learning. We think that it will be followed by other days of this type in the future.

5.3 Raising awareness on the climate crisis

Romain Couillet has organized several introductory seminars on the Anthropocene, which he has presented to students at the UGA and Grenoble-INP, as well as to associations in Grenoble (FNE, AgirAlter-natif). He is also co-responsible of the *Digital Transformation DU*. He has published three articles on the issue of "usability" of artificial intelligence, and is the organizer of a special session on "Signal processing and resilience" for the GRETSI 2022 conference. He is also co-creator of the *sustainable AI* transversal axis of the MIAI project in Grenoble. Finally, he is a trainer for the "Fresque du Climat" and a member of Adrastia and FNE Isère.

5.4 Impact of research results

The efforts of B. Pradelski on COVID policy has received lots of media coverage in Le Monde and other major newspapers. See Section 11.3 for more details.

6 Highlights of the year

6.1 Awards

- R. Couillet has received the IEEE/SEE Prix Glavieux 2021 for his work on *large dimension statistics for artificial intelligence*.

- Y.-P. Hsieh, P. Mertikopoulos, and V. Cevher have been accepted for a long talk to present their work on *The limits of min-max optimization algorithms: Convergence to spurious non-critical sets* at ICML 2021 [25].
- L. Nesi, A. Legrand, and L. Schnorr have received a best paper award from the ICPP'21 conference for their work on *Exploiting system level heterogeneity to improve the performance of a GeoStatistics multi-phase task-based application* [29].

7 New software and platforms

7.1 New software

7.1.1 SimGrid

Keywords: Large-scale Emulators, Grid Computing, Distributed Applications

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

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

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

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

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

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

News of the Year: There were 3 major releases in 2021. A new API was introduced to create the platform descriptions directly from the source code instead of XML, providing much more expressiveness to the experimenters. SMPI now reports memory leaks and correctly diagnoses API misuses, which makes it even more adapted to teaching settings. The documentation was thoroughly overhauled to ease the use of the framework. We also pursued our efforts to improve the overall framework, through bug fixes, code refactoring and other software quality improvement.

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

Contact: Martin Quinson

Participants: Adrien Lebre, Anne-Cécile Orgerie, Arnaud Legrand, Augustin Degomme, Emmanuelle Saillard, Frédéric Suter, Jean-Marc Vincent, Jonathan Pastor, Luka Stanisic, 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 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.

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

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

Contact: Jean-Marc Vincent

7.1.3 marmoteCore

Name: Markov Modeling Tools and Environments - the Core

Keywords: Modeling, Stochastic models, Markov model

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

This software was started within the ANR MARMOTE project: ANR-12-MONU-00019.

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

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

Publications: [hal-01651940](https://hal.archives-ouvertes.fr/hal-01651940), [hal-01276456](https://hal.archives-ouvertes.fr/hal-01276456)

Contact: Alain Jean-Marie

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

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

7.1.4 MarTO

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

Keywords: Perfect sampling, Markov model

Functional Description: MarTO is a simulation software of markovian models. It is able to simulate discrete and continuous time models to provide a perfect sampling of the stationary distribution or directly a sampling of functional of this distribution by using coupling from the past. The simulation kernel is based on the CFTP algorithm, and the internal simulation of transitions on the Aliasing method. This software is a rewrite, more efficient and flexible, of PSI

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

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

Contact: Vincent Danjean

7.1.5 GameSeer

Keyword: Game theory

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

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

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

Contact: Panayotis Mertikopoulos

7.1.6 rmf_tool

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

Keyword: Mean Field

Functional Description: The tool accepts three model types:

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

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

URL: https://github.com/ngast/rmf_tool

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

Contact: Nicolas Gast

8 New results

The new results produced by the team in 2020 can be grouped into the following categories.

8.1 HPC application analysis and optimization

Participants: Tom Cornebize, Vincent Danjean, Arnaud Legrand, Lucas Leandro Nesi, Jean-Marc Vincent.

Finely tuning applications and understanding the influence of key parameters (number of processes, granularity, collective operation algorithms, virtual topology, and process placement) is critical to obtain good performance on supercomputers. With the high consumption of running applications at scale, doing so solely to optimize their performance is particularly costly. We have shown in [37] that SimGrid and SMPI could be used to obtain inexpensive but faithful predictions of expected performance. The methodology we propose decouples the complexity of the platform, which is captured through statistical models of the performance of its main components (MPI communications, BLAS operations), from the complexity of adaptive applications by emulating the application and skipping regular non-MPI parts of the code. We demonstrate the capability of our method with High-Performance Linpack (HPL), the benchmark used to rank supercomputers in the TOP500, which requires careful tuning. This work presents an extensive (in)validation study that compares simulation with real experiments and demonstrates our ability to predict the performance of HPL within a few percent consistently. This study allows us to identify the main modeling pitfalls (e.g., spatial and temporal node variability or network heterogeneity and irregular behavior) that need to be considered. Our “surrogate” also allows studying several subtle HPL parameter optimization problems while accounting for uncertainty on the platform. This work is part of the PhD work of Tom Cornebize [33] and the spatial and temporal node variability has also been investigated and quantified through a systematically measurement of the performance of more than 450 nodes from a dozen of clusters of the Grid’5000 testbed for two years using a rigorous experimental discipline. Using a simple statistical test, we managed to detect many performance changes, from subtle ones of 1% to much more severe degradation of more than 10%, but which could significantly impact the outcome of experiments. The root cause behind detected performance changes ranges from BIOS and system upgrades to cooling issues, faulty memory, and power instability. These events went unnoticed by both Grid’5000 technical team and Grid’5000 users, yet they could greatly harm the reproducibility of experiments and lead to wrong scientific conclusions. All this work heavily builds on reproducible research methodology: the data and metadata collected for this work are permanently and publicly archived under an open license² and presented at through a collection of Jupyter notebooks at cornebize.net/g5k_test/.

Overall, over the last few years, the quality of SimGrid predictions for HPC applications has reach an unprecedented quality, which allows investigating and optimizing the performance of complex applications in a very controlled and reproducible yet realistic way. In [29], we study ExaGeoStat, a task-based machine learning framework specifically designed for geostatistics data. Every iteration of this application comprises several phases that do not scale in the same way, which makes the load particularly challenging to balance. In this work, we show how such applications with multiple phases with distinct resource necessities can take advantage of inter-node heterogeneity to improve performance and reduce resource idleness. We first show how to improve application phase overlap by optimizing runtime and scheduling decisions and then how to compute the optimal distribution for all the phases using a linear program leveraging node heterogeneity while limiting communication overhead. The performance gains of our phase overlap improvements are between 36% and 50% compared to the original base synchronous and homogeneous execution. We show that by adding some slow nodes to a homogeneous set of fast nodes, we can improve the performance by another 25% compared to a standard block-cyclic distribution, thereby harnessing any machine. Most of these algorithmic and scheduling improvements have been investigated in simulation with StarPU-SimGrid as it allows for controlled tracing and debugging on specific platform configurations before being confirmed through real experiments on real testbeds such as Grid 5000 and Santos Dumont.

Finally, we have shown in [43] how the structure of complex applications such as a multifrontal sparse linear solvers could be exploited to detect and correct non-trivial performance problems. Efficiently exploiting computational resources in heterogeneous platforms is a real challenge which has motivated

²T. Cornebize and A. Legrand. “Grid’5000 performance data.” (Mar. 2021), [Online]. Available: [doi:10.5281/zenodo.5024655](https://doi.org/10.5281/zenodo.5024655).

the adoption of the task-based programming paradigm where resource usage is dynamic and adaptive. Unfortunately, classical performance visualization techniques used in routine performance analysis often fail to provide any insight in this new context, especially when the application structure is irregular. We propose and implement in StarVZ several performance visualization techniques tailored for the analysis of task-based multifrontal sparse linear solvers and show that by building on both a performance model of irregular tasks and on structure of the application (in particular the elimination tree), we can detect and highlight anomalies and understand resource utilization from the application point-of-view in a very insightful way. We validate these novel performance analysis techniques with the QR_mumps sparse parallel solver by describing a series of case studies where we identify and address non trivial performance issues thanks to our visualization methodology.

8.2 Large system analysis and optimization

Participants: Jonatha Anselmi, Nicolas Gast, Olivier Bilenne, Sebastian Allmeier, Jean-Marc Vincent.

Large systems can be (1) particularly difficult to analyze because of inherent state-space explosion and (2) require robust and scalable scheduling techniques. In this series of work, we contribute to a better understanding of both aspects.

In [6], we study the impact of communication latency on the classical Work Stealing load balancing algorithm by extending the reference model. By using a theoretical analysis and simulation, we study the overall impact of latency on the Makespan (maximum completion time) and we derive a new expression of the expected running time of a bag of independent tasks scheduled by Work Stealing. This expression enables us to predict under which conditions a given run will yield acceptable performance. For instance, we can easily calibrate the maximal number of processors to use for a given work/platform combination. All our results are validated through simulation on a wide range of parameters.

A complementary approach to work stealing in such system is replication but it is a double-edged weapon that must be handled with caution as the resource overhead may be detrimental when used too aggressively. In [1], we provide a queueing-theoretic framework for job replication schemes based on the principle "replicate a job as soon as the system detects it as a straggler". This is called job speculation. Recent works have analyzed replication on arrival, which we refer to as replication. Replication is motivated by its implementation in Google's BigTable. However, systems such as Apache Spark and Hadoop MapReduce implement speculative job execution. The performance and optimization of speculative job execution is not well understood. To this end, we propose a queueing network model for load balancing where each server can speculate on the execution time of a job. Specifically, each job is initially assigned to a single server by a frontend dispatcher. Then, when its execution begins, the server sets a timeout. If the job completes before the timeout, it leaves the network, otherwise the job is terminated and relaunched or resumed at another server where it will complete. We provide a necessary and sufficient condition for the stability of speculative queueing networks with heterogeneous servers, general job sizes and scheduling disciplines. We find that speculation can increase the stability region of the network when compared with standard load balancing models and replication schemes. We provide general conditions under which timeouts increase the size of the stability region and derive a formula for the optimal speculation time, i.e., the timeout that minimizes the load induced through speculation. We compare speculation with redundant-d and redundant-to-idle-queued rules under an S&X model. For light loaded systems, redundancy schemes provide better response times. However, for moderate to heavy loadings, redundancy schemes can lose capacity and have markedly worse response times when compared with the proposed speculative scheme.

A key challenge in such systems comes from the structure and from the high variability of the execution time distribution. We also study the dispatching to parallel servers problem in [2] where we seek to minimize the average cost experienced by the system over an infinite time horizon. A standard approach for solving this problem is through policy iteration techniques, which relies on the computation of value functions. In this context, we consider the continuous-space $M/G/1 - FCFS$ queue endowed with an arbitrarily-designed cost function for the waiting times of the incoming jobs. The associated relative value function is a solution of Poisson's equation for Markov chains, which in this work we solve in

the Laplace transform domain by considering an ancillary, underlying stochastic process extended to (imaginary) negative backlog states. This construction enables us to issue closed-form relative value functions for polynomial and exponential cost functions and for piecewise compositions of the latter, in turn permitting the derivation of interval bounds for the relative value function in the form of power series or trigonometric sums. We review various cost approximation schemes and assess the convergence of the interval bounds these induce on the relative value function. Namely: Taylor expansions (divergent, except for a narrow class of entire functions with low orders of growth), and uniform approximation schemes (polynomials, trigonometric), which achieve optimal convergence rates over finite intervals. This study addresses all the steps to implementing dispatching policies for systems of parallel servers, from the specification of general cost functions towards the computation of interval bounds for the relative value functions and the exact implementation of the first-policy improvement step.

Finally, when the number of entities is large, most computations are untractable but mean field approximation is a powerful technique to study the performance of very large stochastic systems represented as systems of interacting objects. Applications include load balancing models, epidemic spreading, cache replacement policies, or large-scale data centers, for which mean field approximation gives very accurate estimates of the transient or steady-state behaviors. In a series of recent papers, a new and more accurate approximation, called the refined mean field approximation has been presented. A key strength of this technique lies in its applicability to not-so-large systems. Yet, computing this new approximation can be cumbersome. In [13], we present a tool, called `rmf_tool` and available at github.com/ngast/rmf_tool, that takes the description of a mean field model, and can numerically compute its mean field approximations and refinement.

8.3 Energy optimization

Participants: Jonatha Anselmi, Bruno Gaujal, Panayotis Mertikopoulos, Stéphan Plassart, Louis-Sébastien Rebuffi.

Energy consumption is a major concern in modern architectures. In [7], we consider the classical problem of minimizing off-line the total energy consumption required to execute a set of n real-time jobs on a single processor with a finite number of available speeds. Each real-time job is defined by its release time, size, and deadline (all bounded integers). The goal is to find a processor speed schedule, such that no job misses its deadline and the energy consumption is minimal. We propose a pseudo-linear time algorithm that checks the schedulability of the given set of n jobs and computes an optimal speed schedule. The time complexity of our algorithm is in $O(n)$, to be compared with $O(n \log(n))$ for the best known solution. Besides the complexity gain, the main interest of our algorithm is that it is based on a completely different idea: instead of computing the critical intervals, it sweeps the set of jobs and uses a dynamic programming approach to compute an optimal speed schedule. Our linear time algorithm is still valid (with some changes) when arbitrary (non-convex) power functions and when switching costs are taken into account.

In [36], we consider this problem in a dynamic setting where a Dynamic Voltage and Frequency Scaling (DVFS) processor should execute jobs with obsolescence deadlines: A job becomes obsolete and is removed from the system if it is not completed before its deadline. The objective is to design a dynamic speed policy for the processor that minimizes its average energy consumption plus an obsolescence cost per deadline miss. Under Poisson arrivals and exponentially distributed deadlines and job sizes, we show that this problem can be modeled as a continuous time Markov decision process (MDP) with unbounded state space and unbounded rates. While this MDP admits a continuous time optimality equation for its average cost, the standard uniformization approach is not applicable. Inspired by the scaling method introduced by Blok and Spieksma, we first define a family of truncated MDPs and we then show that the optimal speed profiles are increasing in the number of jobs in the system and are uniformly bounded by a constant. Finally, we show that these properties are inherited from the original (infinite) system. The proposed upper bound on the optimal speed profile is tight and is used to develop an extremely simple policy that accurately approximates the optimal average cost in heavy traffic conditions.

Finally, we study power management in a distributed and online context in [12] through learning and game design. We consider the target-rate power management problem for wireless networks and we

propose two simple, distributed power management schemes that regulate power in a provably robust manner by efficiently leveraging past information. Both schemes are obtained via a combined approach of learning and “game design” where we (1) design a game with suitable payoff functions such that the optimal joint power profile in the original power management problem is the unique Nash equilibrium of the designed game; and (2) derive distributed power management algorithms by directing the networks’ users to employ a no-regret learning algorithm to maximize their individual utility over time. To establish convergence, we focus on the well-known online eager gradient descent learning algorithm in the class of weighted strongly monotone games. In this class of games, we show that when players only have access to imperfect stochastic feedback, multi-agent online eager gradient descent converges to the unique Nash equilibrium in mean square at an $O(1/T)$ rate. In the context of power management in static networks, we show that the designed games are weighted strongly monotone if the network is feasible (i.e. when all users can concurrently attain their target rates). This allows us to derive geometric convergence rate to the joint optimal transmission power. More importantly, in stochastic networks where channel quality fluctuates over time, the designed games are also weighted strongly monotone and the proposed algorithms converge in mean square to the joint optimal transmission power at a $O(1/T)$ rate, even when the network is only feasible on average (i.e. users may be unable to meet their requirements with positive probability). This comes in stark contrast to existing algorithms (like the seminal Foschini–Miljanic algorithm and its variants) that may fail to converge altogether.

8.4 Learning in time varying systems

Participants: Yu Guan Hsieh, Panayotis Mertikopoulos, Barry Pradeliski, Patrick Loiseau.

Some form of stationarity and time synchronism is generally required to guarantee the efficiency of distributed algorithms in large systems. In this series of work, we work on releasing some of these properties.

One of the most widely used methods for solving large-scale stochastic optimization problems is distributed asynchronous stochastic gradient descent (DASGD), a family of algorithms that result from parallelizing stochastic gradient descent on distributed computing architectures (possibly) asynchronously. However, a key obstacle in the efficient implementation of DASGD is the issue of delays: when a computing node contributes a gradient update, the global model parameter may have already been updated by other nodes several times over, thereby rendering this gradient information stale. These delays can quickly add up if the computational throughput of a node is saturated, so the convergence of DASGD may be compromised in the presence of large delays. In [11], we show that, by carefully tuning the algorithm’s step-size, convergence to the critical set is still achieved in mean square, even if the delays grow unbounded at a polynomial rate. We also establish finer results in a broad class of structured optimization problems (called variationally coherent), where we show that DASGD converges to a global optimum with probability 1 under the same delay assumptions. Together, these results contribute to the broad landscape of large-scale non-convex stochastic optimization by offering state-of-the-art theoretical guarantees and providing insights for algorithm design.

In [42], we provide a general framework for studying multi-agent online learning problems in the presence of delays and asynchronicities. Specifically, we propose and analyze a class of adaptive dual averaging schemes in which agents only need to accumulate gradient feedback received from the whole system, without requiring any between-agent coordination. In the single-agent case, the adaptivity of the proposed method allows us to extend a range of existing results to problems with potentially unbounded delays between playing an action and receiving the corresponding feedback. In the multi-agent case, the situation is significantly more complicated because agents may not have access to a global clock to use as a reference point; to overcome this, we focus on the information that is available for producing each prediction rather than the actual delay associated with each feedback. This allows us to derive adaptive learning strategies with optimal regret bounds, even in a fully decentralized, asynchronous environment. Finally, we also analyze an “optimistic” variant of the proposed algorithm which is capable of exploiting the predictability of problems with a slower variation and leads to improved regret bounds.

We also use the dual averaging technique in [24], where we address an open network (agents can join and leave the network at any time) context. In networks of autonomous agents (e.g., fleets of vehicles, scattered sensors), the problem of minimizing the sum of the agents' local functions has received a lot of interest. Leveraging recent online optimization techniques, we propose and analyze the convergence of a decentralized asynchronous optimization method for open networks.

Finally, we examine in [38] and [27] the long-run behavior of multi-agent online learning in games that evolve over time. Specifically, we examine the equilibrium tracking and convergence properties of no-regret learning algorithms in continuous games that evolve over time. We focus on learning via "mirror descent", a widely used class of noregret learning schemes where players take small steps along their individual payoff gradients and then "mirror" the output back to their action sets, and we show that the induced sequence of play (a) converges to Nash equilibrium in time-varying games that stabilize in the long run to a strictly monotone limit; and (b) it stays asymptotically close to the evolving equilibrium of the sequence of stage games (assuming they are strongly monotone). Our results apply to both gradient-based and payoff-based feedback, i.e., the "bandit" case where players only observe the payoffs of their chosen actions.

8.5 Advanced learning methods

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

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

In [17], we analyze the convergence rate of optimistic mirror descent methods in stochastic variational inequalities. Our analysis reveals an intricate relation between the algorithm's rate of convergence and the local geometry induced by the method's underlying Bregman function. We quantify this relation by means of the Legendre exponent, a notion that we introduce to measure the growth rate of the Bregman divergence relative to the ambient norm near a solution. We show that this exponent determines both the optimal step-size policy of the algorithm and the optimal rates attained, explaining in this way the differences observed for some popular Bregman functions (Euclidean projection, negative entropy, fractional power, etc.).

[3], we develop a new stochastic algorithm for solving pseudo-monotone stochastic variational inequalities. Our method builds on Tseng's forward-backward-forward (FBF) algorithm, which is known in the deterministic literature to be a valuable alternative to Korpelevich's extragradient method when solving variational inequalities over a convex and closed set governed by pseudo-monotone, Lipschitz continuous operators. The main computational advantage of Tseng's algorithm is that it relies only on a single projection step and two independent queries of a stochastic oracle. Our algorithm incorporates a mini-batch sampling mechanism and leads to almost sure (a.s.) convergence to an optimal solution. To the best of our knowledge, this is the first stochastic look-ahead algorithm achieving this by using only a single projection at each iteration.

In [16], we examine a flexible algorithmic framework for solving monotone variational inequalities in the presence of randomness and uncertainty. The proposed template encompasses a wide range of popular first-order methods, including dual averaging, dual extrapolation and optimistic gradient algorithms – both adaptive and non-adaptive. Our first result is that the algorithm achieves the optimal rates of convergence for cocoercive problems when the profile of the randomness is known to the optimizer: $O(1/\sqrt{T})$ for absolute noise profiles, and $O(1/T)$ for relative ones. Subsequently, we drop all prior knowledge requirements (the absolute/relative variance of the randomness affecting the problem, the operator's cocoercivity constant, etc.), and we analyze an adaptive instance of the method that gracefully interpolates between the above rates, i.e., it achieves $O(1/\sqrt{T})$ and $O(1/T)$ in the absolute and relative cases, respectively. To our knowledge, this is the first universality result of its kind in the literature and, somewhat surprisingly, it shows that an extra-gradient proxy step is not required to achieve optimal rates.

An other challenging and promising problem motivated by applications in machine learning and operations research is stochastic regret minimization in non-convex problems:

In [21], we study regret minimization with stochastic first-order oracle feedback in online constrained, and possibly non-smooth, non-convex problems. In this setting, the minimization of external regret is beyond reach, so we focus on a local regret measure defined via a proximal-gradient mapping. To achieve no (local) regret in this setting, we develop a prox-grad method based on stochastic first-order feedback, and a simpler method for when access to a perfect first-order oracle is possible. Both methods are min-max order-optimal, and we also establish a bound on the number of prox-grad queries these methods require. As an important application of our results, we also obtain a link between online and offline non-convex stochastic optimization manifested as a new prox-grad scheme with complexity guarantees matching those obtained via variance reduction techniques.

In [22], we propose a hierarchical version of dual averaging for zeroth-order online non-convex optimization-i.e., learning processes where, at each stage, the optimizer is facing an unknown non-convex loss function and only receives the incurred loss as feedback. The proposed class of policies relies on the construction of an online model that aggregates loss information as it arrives, and it consists of two principal components: (a) a regularizer adapted to the Fisher information metric (as opposed to the metric norm of the ambient space); and (b) a principled exploration of the problem's state space based on an adapted hierarchical schedule. This construction enables sharper control of the model's bias and variance, and allows us to derive tight bounds for both the learner's static and dynamic regret-i.e., the regret incurred against the best dynamic policy in hindsight over the horizon of play.

8.6 Adaptive algorithms

Participants: Kimon Antonakopoulos, Yu Guan Hsieh, Panayotis Mertikopoulos, Dong Quan Vu.

Designing algorithms that perform well in a variety of regime is particularly challenging. In a series of work, we study how to get the best of both worlds in a variety of contexts:

In [30], we examine an adaptive learning framework for nonatomic congestion games where the players' cost functions may be subject to exogenous fluctuations (e.g., due to disturbances in the network, variations in the traffic going through a link). In this setting, the popular multiplicative/ exponential weights algorithm enjoys an $O(1/\sqrt{T})$ equilibrium convergence rate. However, this rate is suboptimal in static environments, i.e., when the network is not subject to randomness. In this static regime, accelerated algorithms achieve an $O(1/T^2)$ convergence speed, but they fail to converge altogether in stochastic problems. To fill this gap, we propose a novel, adaptive exponential weights method that seamlessly interpolates between the $O(1/T^2)$ and $O(1/\sqrt{T})$ in the static and stochastic regimes respectively. Importantly, this "best-of-both-worlds" guarantee does not require any prior knowledge of the problem's parameters or tuning by the optimizer. In addition, the method's convergence speed depends subquadratically on the size of the network (number of vertices and edges), so it scales gracefully to large, real-life urban networks.

In [14], we present a new family of min-max optimization algorithms that automatically exploit the geometry of the gradient data observed at earlier iterations to perform more informative extra-gradient steps in later ones. Thanks to this adaptation mechanism, the proposed method automatically detects whether the problem is smooth or not, without requiring any prior tuning by the optimizer. As a result, the algorithm simultaneously achieves order-optimal convergence rates, i.e., it converges to an ε -optimal solution within $O(1/\varepsilon)$ iterations in smooth problems, and within $O(1/\varepsilon^2)$ iterations in non-smooth ones. Importantly, these guarantees do not require any of the standard boundedness or Lipschitz continuity conditions that are typically assumed in the literature; in particular, they apply even to problems with singularities (such as resource allocation problems and the like). This adaptation is achieved through the use of a geometric apparatus based on Finsler metrics and a suitably chosen mirror-prox template that allows us to derive sharp convergence rates for the methods at hand.

In [15], we propose a new family of adaptive first-order methods for a class of convex minimization problems that may fail to be Lipschitz continuous or smooth in the standard sense. Specifically, motivated by a recent flurry of activity on non-Lipschitz (NoLips) optimization, we consider problems that are

continuous or smooth relative to a reference Bregman function—as opposed to a global, ambient norm (Euclidean or otherwise). These conditions encompass a wide range of problems with singular objective, such as Fisher markets, Poisson tomography, D-design, and the like. In this setting, the application of existing order-optimal adaptive methods—like UnixGrad or AcceleGrad—is not possible, especially in the presence of randomness and uncertainty. The proposed method, adaptive mirror descent (AdaMir), aims to close this gap by concurrently achieving min-max optimal rates in problems that are relatively continuous or smooth, including stochastic ones.

Finally, we study how such no-regret strategies fare in a multi-agent context. In game-theoretic learning, several agents are simultaneously following their individual interests, so the environment is non-stationary from each player’s perspective. In this context, the performance of a learning algorithm is often measured by its regret. However, no-regret algorithms are not created equal in terms of game-theoretic guarantees: depending on how they are tuned, some of them may drive the system to an equilibrium, while others could produce cyclic, chaotic, or otherwise divergent trajectories. To account for this, we propose in [23] a range of no-regret policies based on optimistic mirror descent, with the following desirable properties: i) they do not require any prior tuning or knowledge of the game; ii) they all achieve $O(\sqrt{T})$ regret against arbitrary, adversarial opponents; and iii) they converge to the best response against convergent opponents. Also, if employed by all players, then iv) they guarantee $O(1)$ social regret; while v) the induced sequence of play converges to Nash equilibrium with $O(1)$ individual regret in all variationally stable games (a class of games that includes all monotone and convex-concave zero-sum games).

Exploiting past information is thus essential. In [8], we study the dynamics of price discovery in decentralized two-sided markets. We show that there exist memoryless dynamics that converge to the core of the underlying assignment game in which agents’ actions depend only on their current payoff. However, we show that for any such dynamic the convergence time can grow exponentially in relation to the population size. We present a natural dynamic in which a player’s reservation value provides a summary of his past information and show that this dynamic converges to the core in polynomial time in homogeneous markets.

8.7 Learning in Games

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

Learning in Games is considerably more difficult than in classical minimization games as the resulting equilibria may be attractive or not and the dynamic often exhibit cyclic behaviors.

In this context, we examine in [19] the Nash equilibrium convergence properties of no-regret learning in general N -player games. For concreteness, we focus on the archetypal "follow the regularized leader" (FTRL) family of algorithms, and we consider the full spectrum of uncertainty that the players may encounter—from noisy, oracle-based feedback, to bandit, payoff-based information. In this general context, we establish a comprehensive equivalence between the stability of a Nash equilibrium and its support: a Nash equilibrium is stable and attracting with arbitrarily high probability if and only if it is strict (i.e., each equilibrium strategy has a unique best response). This equivalence extends existing continuous-time versions of the "folk theorem" of evolutionary game theory to a *bona fide* algorithmic learning setting, and it provides a clear refinement criterion for the prediction of the day-today behavior of no-regret learning in games.

We study convergence rate of a wide range of regularized methods for learning in games in [20]. To that end, we propose a unified algorithmic template that we call "follow the generalized leader" (FTGL), and which includes as special cases the canonical "follow the regularized leader" algorithm, its optimistic variants, extra-gradient schemes, and many others. The proposed framework is also sufficiently flexible to account for several different feedback models – from full information to bandit feedback. In this general setting, we show that FTGL algorithms converge locally to strict Nash equilibria at a rate which does not depend on the level of uncertainty faced by the players, but only on the geometry of the regularizer near the equilibrium. In particular, we show that algorithms based on entropic regularization—like the exponential weights algorithm—enjoy a linear convergence rate, while Euclidean projection methods converge to equilibrium in a finite number of iterations, even with bandit feedback.

In [41], we examine the long-run behavior of a wide range of dynamics for learning in nonatomic games with finite action spaces and population games, in both discrete and continuous time. The class of dynamics under consideration includes fictitious play and its regularized variants, the best reply dynamics (again, possibly regularized), as well as the dynamics of dual averaging / "follow the regularized leader" (which themselves include as special cases the replicator dynamics and Friedman's projection dynamics). Our analysis concerns both the actual trajectory of play and its time-average, and we cover potential and monotone games, as well as games with an evolutionarily stable state (global or otherwise). Nonatomic games with continuous action spaces will be treated in detail in a second part.

Finally, we study the limits of min-max optimization algorithms in [25]. Compared to minimization problems, the min-max landscape in machine learning applications is considerably more convoluted because of the existence of cycles and similar phenomena. Such oscillatory behaviors are well-understood in the convex-concave regime, and many algorithms are known to overcome them. In this paper, we go beyond the convex-concave setting and we characterize the convergence properties of a wide class of zeroth-, first-, and (scalable) second-order methods in non-convex/nonconcave problems. In particular, we show 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. 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.

8.8 Bandits

Participants: Yan Chen, Bruno Donassolo, Kimang Khun, Nicolas Gast, Bruno Gauthier, Arnaud Legrand, Panayotis Mertikopoulos.

The Multi-armed Stochastic Bandit framework is a classic reinforcement learning problem to study the exploration exploitation trade-off dilemma and for which several optimal algorithms like UCB³ and Thompson sampling⁴, whose optimality has only recently been proved by Kaufmann et al.⁵, have been proposed. Although the first strategy is an optimistic strategy which systematically chooses the "most promising" arm, the second ones build on a Bayesian perspective and samples the posterior to decide which arm to select. The Markovian Bandit allows to model situations where the reward distribution is modeled as a Markov chain and may thus exhibit temporal changes. A key challenge in this context is the curse of dimensionality, which basically says that the state size of the Markov process is exponential in the number of the system components so that the complexity of computing an optimal policy and its value are exponential.

In [39], we study learning algorithms for the classical restless Markovian bandit (in which the state of each arm evolves only when it is chosen) problem with discount and compare posterior sampling strategies with optimistic strategies in terms of scalability. We explain how to adapt PSRL and UCRL2 to exploit the problem structure (MB-PSRL and MB-UCRL2). While the regret bound and runtime of vanilla implementations of PSRL and UCRL2 are exponential in the number of bandits, we show that the episodic regret of MB-PSRL and MB-UCRL2 is $O(S\sqrt{nK})$, where K is the number of episodes, n is the number of bandits, and S is the number of states of each bandit. Up to a factor \sqrt{S} , this matches the lower bound of $\Omega(\sqrt{SnK})$ that we also derive in the paper. MB-PSRL is also computationally efficient: its runtime is linear in the number of bandits. We further show that this linear runtime cannot be achieved by adapting classical non-Bayesian algorithms such as UCRL2 or UCBVI to Markovian bandit problems. Finally, we perform numerical experiments that confirm that MB-PSRL outperforms other existing algorithms in practice, both in terms of regret and of computation time.

In [40], we further develop complexity results for finite horizon restless bandits (in which the state of each arm evolves according to a Markov process independently of the learner's actions). Most restless Markovian bandits problems in infinite horizon can be solved quasioptimally: the famous Whittle index policy is known to become asymptotically optimal exponentially fast as the number of arms grows, at

³P. Auer, N. Cesa-Bianchi, P. Fischer. *Finite-time analysis of the multiarmed bandit problem*. Machine Learning 47, 2002

⁴W. R. Thompson, *On the likelihood that one unknown probability exceeds another in view of the evidence of two samples*, Biometrika, 25:285–294, 1933

⁵E. Kaufmann, N. Korda, R. Munos. *Thompson Sampling: An Asymptotically Optimal Finite Time Analysis*. ALT 2012

least under certain conditions (including having a so-called indexable problem). For restless Markovian bandits problems in finite horizons no such optimal policy is known. In this paper, we define a new policy, based on a linear program relaxation of the finite horizon problem (called the LP-filling policy), that is asymptotically optimal under no condition. Furthermore we show that for regular problems (defined in the paper) the LP-filling policy becomes an index policy (called the LP-regular policy) and becomes optimal exponentially fast in the number of arms. We also introduce the LP-update policy that significantly improves the performance compared to the LP-filling policy for large time horizons. We provide numerical studies that show the prevalence of our LP-policies over previous solutions.

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

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

Participants: Dong Quan Vu, Vitalii Emelianov, Nicolas Gast, Patrick Loiseau, Benjamin Roussillon.

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.

A first line of work in this context is related to fairness and adversarial classification. Discrimination in selection problems such as hiring or college admission is often explained by implicit bias from the decision maker against disadvantaged demographic groups. In [5], we consider a model where the decision maker receives a noisy estimate of each candidate's quality, whose variance depends on the candidate's group—we argue that such differential variance is a key feature of many selection problems. We analyze two notable settings: in the first, the noise variances are unknown to the decision maker who simply picks the candidates with the highest estimated quality independently of their group; in the second, the variances are known and the decision maker picks candidates having the highest expected quality given the noisy estimate. We show that both baseline decision makers yield discrimination, although in opposite directions: the first leads to under-representation of the low-variance group while the second leads to under-representation of the high-variance group. We study the effect on the selection utility of imposing a fairness mechanism that we term the γ -rule (it is an extension of the classical four-fifths rule and it also includes demographic parity). In the first setting (with unknown variances), we prove that under mild conditions, imposing the γ -rule increases the selection utility—here there is no trade-off between fairness and utility. In the second setting (with known variances), imposing the γ -rule decreases the utility but we prove a bound on the utility loss due to the fairness mechanism.

We also consider classification in adversarial context. In [31], we consider the problem of finding optimal classifiers in an adversarial setting where the class-1 data is generated by an attacker whose objective is not known to the defender – an aspect that is key to realistic applications but has so far been overlooked in the literature. To model this situation, we propose a Bayesian game framework where the defender chooses a classifier with no a priori restriction on the set of possible classifiers. The key difficulty in the proposed framework is that the set of possible classifiers is exponential in the set of possible data, which is itself exponential in the number of features used for classification. To counter this, we first show that Bayesian Nash equilibria can be characterized completely via functional threshold classifiers

with a small number of parameters. We then show that this low-dimensional characterization enables us to develop a training method to compute provably approximately optimal classifiers in a scalable manner; and to develop a learning algorithm for the online setting with low regret (both independent of the dimension of the set of possible data). We illustrate our results through simulations.

The Colonel Blotto game is a well-known resource allocation games introduced by Borel (1921) that finds application in many domains like politics (where political parties distribute their budgets to compete over voters), cybersecurity (where effort is distributed to attack/defend targets), online advertising (where marketing campaigns allocate the time to broadcast ads to attract web users), or telecommunication (where network service providers distribute and lease their spectrum to the users). In [32], we introduce the Colonel Blotto game with favoritism, an extension where the winner-determination rule is generalized to include pre-allocations and asymmetry of the players' resources effectiveness on each battlefield. Such favoritism is found in many classical applications of the Colonel Blotto game. We focus on the Nash equilibrium. First, we consider the closely related model of all-pay auctions with favoritism and completely characterize its equilibrium. Based on this result, we prove the existence of a set of optimal univariate distributions-which serve as candidate marginals for an equilibrium-of the Colonel Blotto game with favoritism and show an explicit construction thereof. In several particular cases, this directly leads to an equilibrium of the Colonel Blotto game with favoritism. In other cases, we use these optimal univariate distributions to derive an approximate equilibrium with well-controlled approximation error. Finally, we propose an algorithm-based on the notion of winding number in parametric curves-to efficiently compute an approximation of the proposed optimal univariate distributions with arbitrarily small error.

Finally, we propose in [28], a game-theoretic analysis of the transaction ordering protocol in the Bitcoin blockchain. Most public blockchain protocols, including the popular Bitcoin and Ethereum blockchains, do not formally specify the order in which miners should select transactions from the pool of pending (or uncommitted) transactions for inclusion in the blockchain. Over the years, informal conventions or "norms" for transaction ordering have, however, emerged via the use of shared software by miners, e.g., the GetBlockTemplate (GBT) mining protocol in Bitcoin Core. Today, a widely held view is that Bitcoin miners prioritize transactions based on their offered "transaction fee-per-byte." Bitcoin users are, consequently, encouraged to increase the fees to accelerate the commitment of their transactions, particularly during periods of congestion. In this paper, we audit the Bitcoin blockchain and present statistically significant evidence of mining pools deviating from the norms to accelerate the commitment of transactions for which they have (i) a selfish or vested interest, or (ii) received dark-fee payments via opaque (non-public) side-channels. As blockchains are increasingly being used as a record-keeping substrate for a variety of decentralized (financial technology) systems, our findings call for an urgent discussion on defining neutrality norms that miners must adhere to when ordering transactions in the chains.

8.10 Policy responses to Covid-19

Participants: Nicolas Gast, Bruno Gaujal, Bary Pradelksi.

The Covid-19 pandemic has deeply impacted our lives and caused more than 5.49 million deaths worldwide, making it one of the deadliest pandemic in history. Several policies have been proposed to respond this illness and mitigate both its spread and its impact on both populations' health and economy. We have been studying and supporting some of policies through our expertise in large system analysis (game theory, mean field, etc.)

Bary Pradelksi has been among the first researchers to actively promote Green zoning which has emerged as a widely used policy response to tackle the Covid-19 pandemic [10]. '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. Subsequently, more and more countries followed suit and the European Commission advocated for the implementation of a

European green zoning strategy, which has been supported by the EU member states. While there remain coordination problems, green zoning has proven to be an effective strategy for containing the spread of the virus and limiting its negative economic and social impact. This strategy should provide important lessons and prove useful in future outbreaks. Research in epidemiology indicates that thoroughly implemented and operationalised green zoning can prevent the spread of a transmittable disease that is poorly understood, highly virulent, and potentially highly lethal. Finally, there is strong evidence that green zoning can reduce economic and societal damage as it avoids worst-in-class measures.

Unfortunately, locking down entire regions is not satisfactory and has dramatic consequences on health, the economy, and civil liberties, and several countries have responded in very different ways. Some countries have consistently aimed for elimination – i.e., maximum action to control SARS-CoV-2 and stop community transmission as quickly as possible – while others have opted for mitigation – i.e., action increased in a step-wise, targeted way to reduce cases so as not to overwhelm health-care systems. In [9], we show that the former ones have generally fared much better than the latter ones by comparing deaths, gross domestic product (GDP) growth, and strictness of lock-down measures during the first 12 months of the pandemic. Furthermore, the mitigation has favored the proliferation of new SARS-CoV-2 variants and countries that opt to live with the virus will likely pose a threat to other countries, notably those that have less access to COVID-19 vaccines. Although many scientists have called for a coordinated international strategy to eliminate SARS-CoV-2, it has unfortunately not been heard yet.

Finally, we analyze in [18] a virus propagation dynamics in a large population of agents (or nodes) with three possible states (Susceptible, Infected, Recovered) where agents may choose to vaccinate. We show that this system admits a unique symmetric equilibrium when the number of agents goes to infinity. We also show that the vaccination strategy that minimizes the social cost has the same threshold structure as the mean field equilibrium, but with a shorter threshold. This implies that, to encourage optimal vaccination behavior, vaccination should always be subsidized.

9 Bilateral contracts and grants with industry

Participants: Till Kletti, Patrick Loiseau.

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

10 Partnerships and cooperations

Participants: Nicolas Gast, Guillaume Huard, Patrick Loiseau, Panayotis Mertikopoulos, Bary Pradelski, Benjamin Roussillon.

10.1 International initiatives

10.1.1 Associate Teams in the framework of an Inria International Lab or in the framework of an Inria International Program

ReDaS

Title: Reproducible Data Science

Duration: 2019 – 2021

Coordinator: Guillaume Huard and Lucas Mello Schnorr

Partners: Universidade Federal do Rio Grande do Sul

Inria contact: Guillaume Huard

Summary: The main scientific context of this project is to develop novel analysis techniques and workflow methodologies to support reproducible data science. We focus our efforts along three axes:

1. **Analysis Techniques:** large volumes of data are hard to summarize using simple statistics that hides important behavior in the data. Therefore, raw information visualization plays a key role to explore such data, in particular when curating data and trying to develop intuition about the mathematical models underlying data. Yet, such visualizations require data aggregation, which may lead to significant information loss. It is thus essential to investigate adaptive data aggregation schemes that enable the reduction of the data while controlling the information loss.
2. **Workflow Methodologies:** the analysis process often involves a mix of tools to produce the end result. The data has to be filtered before it can be passed to some standard statistical tool to, eventually, produce some projection of the transformed data that can be visualized and studied by the analyst. Furthermore, the process is interactive: when the analyst is unsatisfied with the end result, a part of the analysis has to be changed to produce a new visualization. These adaptations of the whole analysis typically start from intermediate data and only a part of the analysis has to be rerun. The issue comes with the increasing size of these analysis, the disparity of the analysis tools and the large space of analysis parameters.
3. **Evaluation:** In the previous work packages we will propose both a theoretical and practical methodology whose relevance should be evaluated with real case studies. We will build our evaluation on well identified and quite different datasets originating from the following three areas, on which we already have some past experience:

Performance analysis of HPC applications These applications and their underlying run-times tend to be increasingly complex and dynamic. As a consequence, their execution traces become too large and impossible to analyze with classical tools.

Long-term phenology behavior analysis and correlation with climate change The phenology is the study of plant grow through the use of digital cameras attached to towers installed in the middle of the natural environments. These cameras take photos every a certain number of minutes and enable the researcher to verify how certain species grow, including their relation with the climate.

General public datasets from government transparency reports All public Brazilian institutions are obliged by law to provide datasets about any publicly-financed data measurements. The city of Porto Alegre has long-term weather datasets that contain temperature, pressure and other indicators from different parts of the city. The goal in this case study is very exploratory, for example to envision a way to represent such data in a geographical manner to verify if certain parts of the city may suffer from flash flood more than others.

10.2 National initiatives

ANR

ANR ORACLESS (JCJC 2016-2021) ★ *Online Resource Allocation, unprediCtable wireLESS Systems* [207K€]
ORACLESS is an ANR starting grant (JCJC) coordinated by Panayotis Mertikopoulos. The goal of the project is to develop highly adaptive resource allocation methods for wireless communication networks that are provably capable of adapting to unpredictable changes in the network. In particular, the project will focus on the application of online optimization and online learning methodologies to multi-antenna systems and cognitive radio networks.

ANR ALIAS (PRCI 2020-2021) ★ *Adaptive Learning for Interactive Agents and Systems* [284K€]

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

ALIAS is a bilateral PRCI (collaboration internationale) project joint with Singapore University of Technology and Design (SUTD), coordinated by Bary Pradelski (PI) and involving P. Mertikopoulos and P. Loiseau. The Singapore team consists of G. Piliouras and G. Panageas. The goal of the

project is to provide a unified answer to the question of stability in multi-agent systems: for systems that can be controlled (such as programmable machine learning models), prescriptive learning algorithms can steer the system towards an optimum configuration; for systems that cannot (e.g., online assignment markets), a predictive learning analysis can determine whether stability can arise in the long run. We aim at identifying the fundamental limits of learning in multi-agent systems and design novel, robust algorithms that achieve convergence in cases where conventional online learning methods fail.

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

REFINO is an ANR starting grant (JCJC) coordinated by Nicolas Gast. The main objective on this project is to provide an innovative framework for optimal control of stochastic distributed agents. Restless bandit allocation is one particular example where the control that can be sent to each arm is restricted to an on/off signal. The originality of this framework is the use 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.

DGA Grants Patrick Loiseau and Panayotis Mertikopoulos have a grant from DGA (2018-2021) that complements the funding of PhD student (Benjamin Roussillon) to work on game theoretic models for adversarial classification.

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

11 Dissemination

Participants: Jonatha Anselmi, Romain Couillet, Vincent Danjean, Nicolas Gast, Bruno Gaujal, Guillaume Huard, Arnaud Legrand, Patrick Loiseau, Panayotis Mertikopoulos, Bary Pradelski.

11.1 Promoting scientific activities

11.1.1 Scientific events: organisation

General chair, scientific chair

- P. Loiseau has been chair of the steering committee of NetEcon;

11.1.2 Scientific events: selection

Chair of conference program committees

- P. Mertikopoulos has been a TPC co-chair for NetGCoOp 2020 (online, postponed from 2020), area chair for NeurIPS 2021, and area chair for ICLR 2021

Member of the conference program committees

- P. Loiseau has been a PC member of AAAI 2021, ECML-PKDD 2021 (Area Chair), FAccT 2021 (Area Chair), and NetEcon;
- J. Anselmi has been a PC member of IFIP Performance 2021 and IEEE MASCOTS 2021
- N. Gast has been a PC member of Sigmetrics 2021 and NeurIPS 2021.
- B. Pradelski has been a PC member of FAT* 2021.

Reviewer

- B. Gaujal has been a reviewer for ACM Sigmetrics

11.1.3 Journal

Member of the editorial boards

- N. Gast is associate editor for *Performance Evaluation* and *Stochastic Models*.
- R. Couillet is associate editor for *Random Matrix Theory and Applications*.
- P. Mertikopoulos is an associate editor for *Operations Research Letters*, *RAIRO Operations Research*, *EURO Journal on Computational Optimization*, *Journal on Dynamics and Games*, *Methodology and Computing in Applied Probability*, *Reviewer for*., *Games and Economic Behavior*, *Journal of Economic Theory*, *Journal of Optimization Theory and Applications*, *Mathematics of Operations Research*, *Mathematical Programming*, *Operations Research*, *SIAM Journal on Control and Optimization*, and *SIAM Journal on Optimization*

Reviewer - reviewing activities

- P. Loiseau has been a reviewer for *Games and Economics Behavior*
- B. Gaujal has been a reviewer for *CHAOS* and *Games and Economic Behavior*.
- B. Pradelski has been a reviewer for *Games and Economic Behavior*, *Operations Research*, and *International Journal of Game Theory*

11.1.4 Invited talks

- B. Gaujal has been invited to give a talk on *Discrete Mean Field Games* to the monthly seminar of the Gipsa Lab, Grenoble, in October 2021;
- J. Anselmi has been invited to give a talk on *Stability and Optimization of Speculative Queueing Networks* at the 2021 INFORMS Annual Meeting;
- R. Couillet has been invited several time to expose his work on sustainable AI and the anthropocene (ENS Paris, Ecole Polytechnique, TU Berlin, etc.).
- B. Pradelsky has been invited to present his work on *The effect of COVID certificates on vaccine uptake, health outcomes, and the economy* to the Brussels-based economic think tank Bruegel.
- P. Mertikopoulos has been invited to present his work on:

- *Adaptive Routing in Large-scale Networks: Optimality Under Uncertainty* at NYU Operations Management Series
 - *Online learning in games* at NTUA (Computation and Reasoning Laboratory)
 - *Generalized Robbins-Monro algorithms for min-min and min-max optimization* at RWTH Aachen (Mathematics and Information Processing Seminar)
 - *Online optimization: A unified view through the lens of stochastic approximation* at Télécom ParisTech (Signal, Statistics & Learning Seminar)
 - *Dynamics, (min-max) optimization, and games* at TSE (MAD-Stat Seminar)
 - *Spurious attractors in min-max optimization* at Montréal Machine Learning and Optimization Seminar
 - *Optimization, games, and dynamical systems* at TOUTELIA 2021, Toulouse
- A. Legrand has been invited to give lectures on Open Science and Reproducible Research at:
 - *Laboratory Notebooks and Reproducible Research* at the DKM department of the IRISA laboratory, Feb 2021.
 - *Reproducibility Crisis and Open Science* at the UGA Reproducible Research seminar, Apr 2021.
 - *Obtaining Faithful/Reproducible Measurements on Modern CPUs* at the Journée SIF on research reproducibility, May 2021.
 - *Generating a Controlled Software Environment with Debian Snapshot Archive* at the GUIX workshop on software environment reproducibility, May 2021.
 - *Reproducibility Crisis and Open Science* at the reproducible research workshop of the Montpellier Bio-Stats network
 - *Reproducibility Crisis and Open Science* at the *Sciences de l'information géographique reproductibles* thematic school of the CIST-CNRS.

11.1.5 Research Administration

- P. Loiseau is an expert for the FRS-FNRS
- A. Legrand has participated to the *Electronic Laboratory Notebook* working group of the French CoSO (Comité pour la Science Ouverte), whose **report** has recently been produced.
- A. Legrand is a member of the Inria Grenoble COS (conseil scientifique);
- A. Legrand is a member of the CoSO UGA board (commission consultative science ouverte);
- A. Legrand is in charge of the *Distributed Systems, Parallel Computing, and Networks* research area theme of the LIG;
- N. Gast is co-responsible of the PhD program in Computer Science at the University Grenoble-Alpes (vice director of ED MSTII).

11.2 Teaching - Supervision - Juries

- B. Gaujal has been a member of the jury for CR2 selection in Inria Grenoble
- A. Legrand is a member of the *section 6* of the CoNRS (Comité National de la Recherche Scientifique)
- N. Gast served in the « Comité de spécialistes » (COS) for a MCF position in the G-SCOP laboratory (Grenoble)
- P. Mertikopoulos served in the « Comité de spécialistes » (COS) for a MCF position in the Université Gustave Eiffel.
- A. Legrand served in the « Comité de spécialistes » (COS) for a Professor position in the ENS Rennes.

- A. Legrand was a reviewer for the PhD thesis of Mrs Yiqin Gao, from ENS Lyon: *Scheduling independent tasks under budget and time constraints*
- P. Mertikopoulos was a member of the PhD thesis committee of Mr. Yassine Laguel from UGA: *Optimisation convexe pour l'apprentissage robuste au risque*
- A. Legrand is a member of the mid-term PhD committee of Adeyemi Adetula (UGA)
- P. Loiseau is a member of the mid-term PhD committee of O. Boufous (U. Avignon) and C. Pinzon (Ecole Polytechnique)
- B. Gaujal is a member of the mid-term PhD committee of Michel Davydov (ENS Paris)

11.2.1 Teaching

- J. Anselmi teaches the *Probability and Simulation* and the *Performance Evaluation* lectures in M1, Polytech Grenoble.
- P. Loiseau teaches the M1 course at Ecole Polytechnique on Advanced Machine Learning and Autonomous Agents
- A. Legrand and J.-M. Vincent teach the transversal *Scientific Methodology and Empirical Evaluation* lecture (36h) at the M2 MOSIG.
- B. Gaujal teaches the M2 course on *Optimization under uncertainties* in M2 ORCO Grenoble
- N. Gast is responsible of the course *Reinforcement Learning* in Master MOSIG/MSIAM (Grenoble) and of the course « Introduction to machine learning » (License 3, Grenoble).
- G. Huard is responsible of the courses *UNIX & C programming* in the L1 and L3 INFO, of *Object oriented and event-driven programming* in the L3 INFO, and of the *Objet oriented design* in M1 INFO.
- During the COVID crisis, teaching as been particularly difficult and has required major involvements and adaptations of teachers. G. Huard has implemented the following developments related to teaching activities:
 - Moodle VPL backend & engine for local execution
 - VSCode extensions to use VPLs and help with C programming ([here](#) and [here](#))
 - Several VPL activities for the course in L1 INFO on Moodle
- V. Danjean teaches the *Operating Systems, Programming Languages, Algorithms, Computer Science and Mediation* lectures in L3, M1 and Polytech Grenoble.
- V. Danjean organized with J.M. Vincent the end of the DIU-EIL that have been setup for training high school professors to teach computer science.
- R. Couillet is the initiator of a new lecture on the *introduction to artificial intelligence* in the L1 INFO.
- The 3rd edition of the MOOC of A. Legrand, K. Hinsens 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 2022), about 17,800 persons have followed this MOOC and 1620 certificates of achievement have been delivered. 54% of participants are PhD students and 12% are undergraduates.
- A. Legrand has organized with F. Theoleyre, B. Dnnassolo, M. Simonin, and G. Schreiner the 16th edition of the RSD GDR research school on *Reproducibility and experimental research in networks and systems*. This school has attrated 40 PhD students and has offered them to experiment with both the Grid5000 and FIT-IoTLab platforms throughout the whole week using reproducible research technologies (Jupyter notebooks, docker containers, git, the EnOSLib experiment engine, R for data analysis and experiment design, etc.). See [here](#) and [here](#) for more details.

11.3 Popularization

- V. Danjean organized a Linux Install Party for all the students of INFO and IESE3 department of Polytech'Grenoble
- R. Couillet has been interviewed by France Bleu Isère in the *Retour sur Terre* show (Dec. 9th 2021)

The efforts of B. Pradelski on COVID policy has received lots of media coverage in Le Monde and other major newspapers. See for example:

- Le Monde: [Couverture vaccinale: nombre de décès et PIB, le très fort impact du passe sanitaire](#)
- [Les Echos: CoViD: le passe sanitaire a sauvé 4000 vies en France entre mi-juillet et fin 2021](#)
- [Financial Times: Covid passes boosted economies and vaccine uptake, study shows](#)
- The Lancet article *SARS-CoV-2 elimination, not mitigation, creates best outcomes for health, the economy, and civil liberties* [9] has been covered in the *New York Times*, *Financial Times*, *Daily Mail*, *Süddeutsche Zeitung*, *Le Monde*, *Libération*, *Le Temps*, *Euronews*, *Le Devoir*, *The Hill*, *20 minutos*, *Die Welt*, *The Guardian* (five times), *New Scientist*, *The Atlantic*, *Exame*, *L'Express*, *Die Zeit*, *Science*, *Politico*
- Several opinion articles have also been published in media:
 - ‘Covid-19 : « Le passe sanitaire doit servir à accélérer les doses de rappel »’ with Miquel Oliu-Barton, *Le Monde* (24 November 2021).
 - ‘Covid-19 : le faux dilemme entre santé, économie et libertés’ with Philippe Aghion, Patrick Artus, and Miquel Oliu-Barton, *Commentaire* (June 2021).
 - ‘Le recours temporaire à un passe sanitaire est nécessaire si nous voulons une sortie de crise durable’ with Philippe Aghion, Philippe Martin, and Miquel Oliu-Barton, *Le Monde* (2 June 2021).
 - ‘Covid-19 : « Nous avons besoin d’un “passe sanitaire” fiable, temporaire et accessible pour tout le monde »’ with Anne Bucher and Miquel Oliu-Barton, *Le Monde* (22 March 2021).
 - ‘Aiming for zero Covid-19: Europe needs to take action’ with a group of international scientists, published by Bruegel, *de Volkskrant*, *El Pais*, *la Repubblica*, *Le Monde*, *Rzeczpospolita*, *Süddeutsche Zeitung* (15 February 2021).

12 Scientific production

12.1 Publications of the year

International journals

- [1] J. Anselmi and N. Walton. ‘Stability and Optimization of Speculative Queueing Networks’. In: *IEEE/ACM Transactions on Networking* (2021), pp. 1–12. DOI: [10.1109/TNET.2021.3128778](#). URL: <https://hal.archives-ouvertes.fr/hal-03349775>.
- [2] O. Bilenne. ‘Dispatching to parallel servers: Solutions of Poisson’s equation for first-policy improvement’. In: *Queueing Systems*. *Queueing Systems* 99.3 (11th Oct. 2021), pp. 199–230. DOI: [10.1007/s11134-021-09713-y](#). URL: <https://hal.archives-ouvertes.fr/hal-02925284>.
- [3] R. I. Bot, P. Mertikopoulos, M. Staudigl and P. T. Vuong. ‘Minibatch forward-backward-forward methods for solving stochastic variational inequalities’. In: *Stochastic Systems* 11.2 (June 2021), pp. 112–139. DOI: [10.1287/stsy.2019.0064](#). URL: <https://hal.inria.fr/hal-02405776>.
- [4] B. Donassolo, A. Legrand, P. Mertikopoulos and I. Fajjari. ‘Online Reconfiguration of IoT Applications in the Fog: The Information-Coordination Trade-off’. In: *IEEE Transactions on Parallel and Distributed Systems* 33.5 (2022), pp. 1156–1172. DOI: [10.1109/TPDS.2021.3097281](#). URL: <https://hal.inria.fr/hal-02636987>.

- [5] V. Emelianov, N. Gast, K. P. Gummadi and P. Loiseau. ‘On fair selection in the presence of implicit and differential variance’. In: *Artificial Intelligence* 302 (Oct. 2021), pp. 1–20. DOI: [10.1016/j.artint.2021.103609](https://doi.org/10.1016/j.artint.2021.103609). URL: <https://hal.inria.fr/hal-03398739>.
- [6] N. Gast, M. Khatiri, D. Trystram and F. Wagner. ‘Analysis of Work Stealing with latency’. In: *Journal of Parallel and Distributed Computing* 153 (July 2021), pp. 119–129. DOI: [10.1016/j.jpdc.2021.03.010](https://doi.org/10.1016/j.jpdc.2021.03.010). URL: <https://hal.inria.fr/hal-03356234>.
- [7] B. Gaujal, A. Girault and S. Plassart. ‘A Pseudo-Linear Time Algorithm for the Optimal Discrete Speed Minimizing Energy Consumption’. In: *Discrete Event Dynamic Systems* (2021). DOI: [10.1007/s10626-020-00327-9](https://doi.org/10.1007/s10626-020-00327-9). URL: <https://hal.archives-ouvertes.fr/hal-03030416>.
- [8] 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>.
- [9] M. Oliu-Barton, B. Pradelski, P. Aghion, P. Artus, I. Kickbusch, J. Lazarus, D. Sridhar and S. Vander-slott. ‘SARS-CoV-2 elimination, not mitigation, creates best outcomes for health, the economy, and civil liberties’. In: *The Lancet* 397.10291 (June 2021), pp. 2234–2236. DOI: [10.1016/S0140-6736\(21\)00978-8](https://doi.org/10.1016/S0140-6736(21)00978-8). URL: <https://hal.archives-ouvertes.fr/hal-03355143>.
- [10] 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>.
- [11] Z. Zhou, P. Mertikopoulos, N. Bambos, P. W. Glynn and Y. Ye. ‘Distributed stochastic optimization with large delays’. In: *Mathematics of Operations Research* (2021), pp. 1–41. DOI: [10.1287/moor.2021.1200](https://doi.org/10.1287/moor.2021.1200). URL: <https://hal.inria.fr/hal-03342384>.
- [12] Z. Zhou, P. Mertikopoulos, A. L. Moustakas, N. Bambos and P. W. Glynn. ‘Robust power management via learning and game design’. In: *Operations Research* 69.1 (Jan. 2021), pp. 331–345. DOI: [10.1287/opre.2020.1996](https://doi.org/10.1287/opre.2020.1996). URL: <https://hal.archives-ouvertes.fr/hal-01643340>.

International peer-reviewed conferences

- [13] S. Allmeier and N. Gast. ‘rmftool - A library to Compute (Refined) Mean Field Approximation(s)’. In: TOSME 2021. Online conference, France, 12th Nov. 2021. URL: <https://hal.inria.fr/hal-03485044>.
- [14] K. Antonakopoulos, V. E. Belmega and P. Mertikopoulos. ‘Adaptive extra-gradient methods for min-max optimization and games’. In: ICLR 2021 - 9th International Conference on Learning Representations. Virtual, Unknown Region, 3rd May 2021, pp. 1–28. URL: <https://hal.inria.fr/hal-03342601>.
- [15] K. Antonakopoulos and P. Mertikopoulos. ‘Adaptive first-order methods revisited: Convex optimization without Lipschitz requirements’. In: NeurIPS 2021 - 35th International Conference on Neural Information Processing Systems. Virtual, Unknown Region, 2021. URL: <https://hal.inria.fr/hal-03342998>.
- [16] K. Antonakopoulos, T. Pethick, A. Kavis, P. Mertikopoulos and V. Cevher. ‘Sifting through the noise: Universal first-order methods for stochastic variational inequalities’. In: NeurIPS 2021 - 35th International Conference on Neural Information Processing Systems. Virtual, Unknown Region, 6th Dec. 2021, pp. 1–39. URL: <https://hal.inria.fr/hal-03357714>.
- [17] W. Azizian, F. Iutzeler, J. Malick and P. Mertikopoulos. ‘The last-iterate convergence rate of optimistic mirror descent in stochastic variational inequalities’. In: COLT 2021 - 34th Annual Conference on Learning Theory. Boulder, United States, 15th Aug. 2021, pp. 1–32. URL: <https://hal.inria.fr/hal-03342583>.
- [18] 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>.

- [19] 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>.
- [20] A. Giannou, E. V. Vlatakis-Gkaragkounis and P. Mertikopoulos. ‘The convergence rate of regularized learning in games: From bandits and uncertainty to optimism and beyond’. In: NeurIPS 2021 - 35th International Conference on Neural Information Processing Systems. Virtual, Unknown Region, 6th Dec. 2021, pp. 1–28. URL: <https://hal.inria.fr/hal-03357715>.
- [21] N. Hallak, P. Mertikopoulos and V. Cevher. ‘Regret minimization in stochastic non-convex learning via a proximal-gradient approach’. In: ICML 2021 - 38th International Conference on Machine Learning. Vienna, Austria, July 2021. URL: <https://hal.archives-ouvertes.fr/hal-03043872>.
- [22] A. Héliou, M. Martin, T. Rahier and P. Mertikopoulos. ‘Zeroth-order non-convex learning via hierarchical dual averaging’. In: ICML 2021 - 38th International Conference on Machine Learning. Vienna, Austria, 18th July 2021, pp. 1–34. URL: <https://hal.inria.fr/hal-03341912>.
- [23] Y.-G. Hsieh, K. Antonakopoulos and P. Mertikopoulos. ‘Adaptive learning in continuous games: Optimal regret bounds and convergence to Nash equilibrium’. In: COLT 2021 - the 34th Annual Conference on Learning Theory. Boulder, United States, 15th Aug. 2021, pp. 1–34. URL: <https://hal.inria.fr/hal-03342410>.
- [24] Y.-G. Hsieh, F. Iutzeler, J. Malick and P. Mertikopoulos. ‘Optimization in open networks via dual averaging’. In: CDC 2021 - 60th IEEE Annual Conference on Decision and Control. Austin, United States: IEEE, 13th Dec. 2021, pp. 1–7. URL: <https://hal.inria.fr/hal-03342395>.
- [25] 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>.
- [26] T. Kletti, J.-M. Renders and P. Loiseau. ‘Introducing the Expohedron for Efficient Pareto-optimal Fairness-Utility Amortizations in Repeated Rankings’. In: WSDM 2022 - 15th ACM International Conference on Web Search and Data Mining. Phoenix (virtual), United States: ACM, Feb. 2022, pp. 1–10. DOI: [10.1145/3488560.3498490](https://doi.org/10.1145/3488560.3498490). URL: <https://hal.inria.fr/hal-03551061>.
- [27] P. Mertikopoulos and M. Staudigl. ‘Equilibrium tracking and convergence in dynamic games’. In: CDC 2021 - 60th IEEE Annual Conference on Decision and Control. Austin, United States, 13th Dec. 2021, pp. 1–8. URL: <https://hal.inria.fr/hal-03342397>.
- [28] J. Messias, M. Alzayat, B. Chandrasekaran, K. P. Gummadi, P. Loiseau and A. Mislove. ‘Selfish & Opaque Transaction Ordering in the Bitcoin Blockchain: The Case for Chain Neutrality’. In: IMC 2021 - ACM Internet Measurement Conference. Virtual Event, France, 2nd Nov. 2021, pp. 1–16. DOI: [10.1145/3487552.3487823](https://doi.org/10.1145/3487552.3487823). URL: <https://hal.inria.fr/hal-03361860>.
- [29] 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>.
- [30] D. Quan Vu, K. Antonakopoulos and P. Mertikopoulos. ‘Fast routing under uncertainty: Adaptive learning in congestion games with exponential weights’. In: NeurIPS 2021 - 35th International Conference on Neural Information Processing Systems. Virtual, Unknown Region, 6th Dec. 2021, pp. 1–36. URL: <https://hal.inria.fr/hal-03357716>.
- [31] B. Roussillon and P. Loiseau. ‘Scalable Optimal Classifiers for Adversarial Settings under Uncertainty’. In: GameSec 2021 - 12th Conference on Decision and Game Theory for Security. Prague, Czech Republic, 2021, pp. 1–20. URL: <https://hal.inria.fr/hal-03360526>.

- [32] D. Q. Vu and P. Loiseau. ‘Colonel Blotto Games with Favoritism: Competitions with Pre-allocations and Asymmetric Effectiveness’. In: *Proceedings of the 22nd ACM Conference on Economics and Computation (ACM EC '21)*. Budapest, Hungary, 18th July 2021, pp. 862–863. URL: <https://hal.archives-ouvertes.fr/hal-03359800>.

Doctoral dissertations and habilitation theses

- [33] T. Cornebize. ‘High Performance Computing : towards better Performance Predictions and Experiments’. Université Grenoble Alpes [2020-....], 2nd June 2021. URL: <https://tel.archives-ouvertes.fr/tel-03328956>.
- [34] 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>.
- [35] B. Roussillon. ‘Learning in the Presence of Strategic Data Sources: Models and Solutions’. Université Grenoble Alpes, 15th Sept. 2021. URL: <https://tel.archives-ouvertes.fr/tel-03528550>.

Reports & preprints

- [36] J. Anselmi, B. Gaujal and L. Sébastien Rebuffi. *Optimal Speed Profile of a DVFS Processor under Soft Deadlines*. 5th Oct. 2021. URL: <https://hal.archives-ouvertes.fr/hal-03364880>.
- [37] T. Cornebize and A. Legrand. *Simulation-based Optimization and Sensibility Analysis of MPI Applications: Variability Matters*. 6th Jan. 2022. URL: <https://hal.inria.fr/hal-03141988>.
- [38] B. Duvocelle, P. Mertikopoulos, M. Staudigl and D. Vermeulen. *Multi-agent online learning in time-varying games*. 2021. URL: <https://hal.inria.fr/hal-01891545>.
- [39] N. Gast, B. Gaujal and K. Khun. *Reinforcement Learning for Markovian Bandits: Is Posterior Sampling more Scalable than Optimism?* 16th June 2021. URL: <https://hal.inria.fr/hal-03262006>.
- [40] N. Gast, B. Gaujal and C. Yan. *(Close to) Optimal Policies for Finite Horizon Restless Bandits*. 18th June 2021. URL: <https://hal.inria.fr/hal-03262307>.
- [41] S. Hadikhannloo, R. Laraki, P. Mertikopoulos and S. Sorin. *Learning in nonatomic games, Part I: Finite action spaces and population games*. 13th Sept. 2021. URL: <https://hal.inria.fr/hal-03342992>.
- [42] Y.-G. Hsieh, F. Iutzeler, J. Malick and P. Mertikopoulos. *Multi-agent online optimization with delays: Asynchronicity, adaptivity, and optimism*. 2021. URL: <https://hal.inria.fr/hal-03410422>.
- [43] M. C. Miletto, L. L. Nesi, L. Mello Schnorr and A. Legrand. *Performance Analysis of Irregular Task-Based Applications on Hybrid Platforms: Structure Matters*. 23rd July 2021. URL: <https://hal.inria.fr/hal-03298021>.

12.2 Cited publications

- [44] 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>.
- [45] 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>.

- [46] 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>.
- [47] 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>.
- [48] 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>.
- [49] 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.
- [50] S. Böhm and C. Engelmann. ‘xSim: The Extreme-Scale Simulator’. In: *HPCS*. Istanbul, Turkey, 2011.
- [51] P. Bruel, S. Quinito Masnada, B. Videau, A. Legrand, J.-M. Vincent and A. Goldman. ‘Autotuning under Tight Budget Constraints: A Transparent Design of Experiments Approach’. In: *CCGrid 2019 - International Symposium in Cluster, Cloud, and Grid Computing*. Larcana, Cyprus: IEEE, May 2019, pp. 1–10. DOI: [10.1109/CCGRID.2019.00026](https://doi.org/10.1109/CCGRID.2019.00026). URL: <https://hal.inria.fr/hal-02110868>.
- [52] H. Brunst, D. Hackenberg, G. Juckeland and H. Rohling. ‘Comprehensive Performance Tracking with VAMPIR 7’. In: *Tools for High Performance Computing 2009*. Springer Berlin Heidelberg, 2010. DOI: [http://dx.doi.org/10.1007/978-3-642-11261-4_2](https://doi.org/10.1007/978-3-642-11261-4_2).
- [53] 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>.
- [54] 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>.
- [55] 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>.
- [56] 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>.
- [57] 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>.
- [58] 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>.
- [59] 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>.
- [60] 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>.

- [61] 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>.
- [62] 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>.
- [63] L. Flokas, E. V. Vlatakis-Gkaragkounis, T. Lianes, 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>.
- [64] 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>.
- [65] 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>.
- [66] 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>.
- [67] 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>.
- [68] 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>.
- [69] 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>.
- [70] 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>.
- [71] 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>.
- [72] 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>.
- [73] 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>.
- [74] B. Gaujal, A. Girault and S. Plassart. ‘Discrete and Continuous Optimal Control for Energy Minimization in Real-Time Systems’. In: *EBCCSP 2020 - 6th International Conference on Event-Based Control, Communication, and Signal Processing*. Krakow, Poland: IEEE, Sept. 2020, pp. 1–8. DOI: [10.1109/EBCCSP51266.2020.9291353](https://doi.org/10.1109/EBCCSP51266.2020.9291353). URL: <https://hal.archives-ouvertes.fr/hal-03020330>.

- [75] 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>.
- [76] 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>.
- [77] M. Heath and J. Etheridge. ‘Visualizing the performance of parallel programs’. In: *IEEE software* 8.5 (1991).
- [78] 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>.
- [79] T. Hoefler, T. Schneider and A. Lumsdaine. ‘LogGOPSim - Simulating Large-Scale Applications in the LogGOPSim Model’. In: *ACM Workshop on Large-Scale System and Application Performance*. 2010.
- [80] 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).
- [81] 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>.
- [82] 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>.
- [83] 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, Macao SAR China: IEEE, Dec. 2019, pp. 1–5. URL: <https://hal.archives-ouvertes.fr/hal-02384051>.
- [84] 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.
- [85] 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>.
- [86] 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>.
- [87] 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>.
- [88] 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>.

- [89] 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>.
- [90] 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>.
- [91] 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>.
- [92] 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>.
- [93] W. Nagel, A. Arnold, M. Weber, H. Hoppe and K. Solchenbach. ‘VAMPIR: Visualization and Analysis of MPI Resources’. In: *Supercomputer* 12.1 (1996).
- [94] 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>.
- [95] 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>.
- [96] 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.
- [97] 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>.
- [98] 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>.
- [99] 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.
- [100] B. Shneiderman. ‘The eyes have it: A task by data type taxonomy for information visualizations’. In: *IEEE Symposium on Visual Languages*. IEEE. 1996.
- [101] 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.
- [102] 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>.
- [103] M. Tikir, M. Laurenzano, L. Carrington and A. Snively. ‘PSINS: An Open Source Event Tracer and Execution Simulator for MPI Applications’. In: *Euro-Par*. 2009.
- [104] 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>.

- [105] B. Vinot, F. Cadoux and R. Hélot. ‘Decentralized Optimization of Energy Exchanges in an Electricity Microgrid’. In: *ACM e-Energy 2016 - 7th ACM International Conference on Future Energy Systems*. Waterloo, Canada, June 2016. DOI: [10.1145/2939912.2939915](https://doi.org/10.1145/2939912.2939915). URL: <https://hal.archives-ouvertes.fr/hal-01340394>.
- [106] B. Vinot, F. Cadoux and R. Hélot. ‘Decentralized optimization of energy exchanges in an electricity microgrid’. In: *2016 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe)*. Ljubljana, Slovenia: IEEE, Oct. 2016, pp. 1–6. DOI: [10.1109/ISGTEurope.2016.7856319](https://doi.org/10.1109/ISGTEurope.2016.7856319). URL: <https://hal.inria.fr/hal-01784148>.
- [107] 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>.
- [108] 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.
- [109] 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.
- [110] 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>.
- [111] G. Zheng, G. Kakulapati and L. Kalé. ‘BigSim: A Parallel Simulator for Performance Prediction of Extremely Large Parallel Machines’. In: *IPDPS*. 2004.
- [112] S. Zrigui, R. Y. De Camargo, A. Legrand and D. Trystram. ‘Improving the Performance of Batch Schedulers Using Online Job Runtime Classification’. Under review in JPDC. Oct. 2020. URL: <https://hal.archives-ouvertes.fr/hal-03023222>.