

RESEARCH CENTRE

Grenoble - Rhône-Alpes

IN PARTNERSHIP WITH:

CNRS, Université de Grenoble Alpes

2020

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 Sound and Reproducible Experimental Methodology	5
3.2 Multi-Scale Analysis and Visualization	5
3.3 Fast and Faithful Performance Prediction of Very Large Systems	6
3.4 Local Interactions and Transient Analysis in Adaptive Dynamic Systems	7
3.5 Distributed Learning in Games and Online Optimization	7
4 Application domains	8
4.1 Large Computing Infrastructures	8
4.2 Next-Generation Wireless Networks	8
4.3 Energy and Transportation	9
4.4 Social Computing Systems	9
5 Highlights of the year	9
5.1 Awards	9
6 New software and platforms	9
6.1 New software	9
6.1.1 Framesoc	9
6.1.2 Ocelotl	10
6.1.3 SimGrid	10
6.1.4 StarPU	11
6.1.5 PSI	12
6.1.6 marmoteCore	12
6.1.7 MarTO	13
6.1.8 GameSeer	13
6.1.9 taktuk	14
7 New results	14
7.1 System Analysis and Experiments	14
7.2 Performance Evaluation and Measurements of Distributed Systems and Networks	14
7.3 Mean Field Analysis and Mean Field Games	14
7.4 Game Theory	15
7.5 Privacy, Fairness and Transparency	15
7.6 Optimization Methods	15
7.7 Learning	15
7.8 Energy Optimization	15
7.9 Covid Deconfinement Policies and Testing	15
8 Bilateral contracts and grants with industry	16
9 Partnerships and cooperations	16
9.1 National initiatives	16
9.2 Inria associate team not involved in an ILL	16

10 Dissemination	17
10.1 Promoting scientific activities	17
10.1.1 Scientific events: organisation	17
10.1.2 Scientific events: selection	17
10.1.3 Journal	17
10.1.4 Invited talks	17
10.1.5 Research administration	18
10.2 Teaching - Supervision - Juries	18
10.2.1 Teaching	18
10.2.2 Supervision	18
10.2.3 Juries	19
10.3 Popularization	20
10.3.1 Internal or external Inria responsibilities	20
10.3.2 Articles and contents	20
10.3.3 Education	20
11 Scientific production	20
11.1 Publications of the year	20
11.2 Other	25
11.3 Cited publications	25

Project-Team POLARIS

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

Keywords

Computer sciences and digital sciences

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

Other research topics and application domains

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

1 Team members, visitors, external collaborators

Research Scientists

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

Faculty Members

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

Post-Doctoral Fellows

- Olivier Bilenne [CNRS, until Aug 2020]
- Mouhcine Mendil [INPG Entreprise SA, until Jan 2020]
- Dong Quan Vu [CNRS, from Oct 2020]

PhD Students

- Sebastian Allmeier [Inria, from Nov 2020]
- Kimon Antonakopoulos [Inria]
- Thomas Barzola [Univ Grenoble Alpes]
- Tom Cornebize [Univ Grenoble Alpes]
- Bruno De Moura Donassolo [Orange Labs, CIFRE, until Sep 2020]
- Vitalii Emelianov [Inria]
- Yu Guan Hsieh [Univ Grenoble Alpes, from Oct 2020]
- Alexis Janon [Univ Grenoble Alpes, until Nov 2020]
- Simon Philipp Jantschgi [Universität de Zurich, from Feb 2020]
- Baptiste Jonglez [Institut polytechnique de Grenoble, until Aug 2020]
- Kimang Khun [Inria]
- Till Kletti [Naver Labs, CIFRE, from Feb 2020]

- Dimitrios Moustakas [Inria]
- Louis Sebastien Rebuffi [Univ Grenoble Alpes, from Oct 2020]
- Pedro Rocha Bruel [Université de Sao Paulo - Brésil]
- Benjamin Roussillon [Univ Grenoble Alpes]
- Vera Sosnovik [Univ Grenoble Alpes]
- Chen Yan [Univ Grenoble Alpes]

Technical Staff

- Bruno De Moura Donassolo [Inria, Engineer, from Oct 2020]
- Eleni Gkiouzepe [CNRS, Engineer]

Interns and Apprentices

- Krishna Acharya [Univ Grenoble Alpes, from Feb 2020 until Jul 2020]
- Sebastian Allmeier [Inria, from Feb 2020 until Jul 2020]
- Tsothe Chakhvadze [Univ Grenoble Alpes, from Feb 2020 until Jun 2020]
- Matthias Lotta [Inria, from Jun 2020 until Jul 2020]
- Marius Monnier [Univ Grenoble Alpes, from Feb 2020 until Jun 2020]
- Vincent Ribot [Inria, from Feb 2020 until Aug 2020]
- Nicolas Rocher [Univ Grenoble Alpes, until Jan 2020]

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 Sound and Reproducible Experimental Methodology

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

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

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

3.2 Multi-Scale Analysis and Visualization

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

As explained in the previous section, the first difficulty encountered when modeling large scale computer systems is to observe these systems and extract information on the behavior of both the architecture, the middleware, the applications, and the users. The second difficulty is to *visualize and analyze* such *multi-level traces to understand how the performance of the application can be improved*. While a lot of efforts are put into visualizing scientific data, in comparison little effort have gone into to developing techniques specifically tailored for understanding the behavior of distributed systems. Many visualization tools have been developed by renowned HPC groups since decades (e.g., BSC [101], Jülich and TU Dresden [100, 71], UIUC [88, 104, 91] and ANL [117], Inria Bordeaux [76] and Grenoble [92], ...) but most of these tools build on the classical information visualization mantra [109] 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 [96]. Such traces are typically made of events that occur at very different time and space scales, which unfortunately hinders classical approaches. Such visualization tools have focused on easing interaction and navigation in the trace (through gantcharts, intuitive filters, pie charts and kiviats) but they are very

difficult to maintain and evolve and they require some significant experience to identify performance bottlenecks.

Therefore many groups have more recently proposed in combination to these tools some techniques to help identifying the structure of the application or regions (applicative, spatial or temporal) of interest. For example, researchers from the SDSC [99] propose some segment matching techniques based on clustering (Euclidean or Manhattan distance) of start and end dates of the segments that enables to reduce the amount of information to display. Researchers from the BSC use clustering, linear regression and Kriging techniques [108, 95, 87] to identify and characterize (in term of performance and resource usage) application phases and present aggregated representations of the trace [107]. Researchers from Jülich and TU Darmstadt have proposed techniques to identify specific communication patterns that incur wait states [114, 63]

3.3 Fast and Faithful Performance Prediction of Very Large Systems

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

Evaluating the scalability, robustness, energy consumption and performance of large infrastructures such as exascale platforms and clouds raises severe methodological challenges. The complexity of such platforms mandates empirical evaluation but direct experimentation via an application deployment on a real-world testbed is often limited by the few platforms available at hand and is even sometimes impossible (cost, access, early stages of the infrastructure design, ...). Unlike direct experimentation via an application deployment on a real-world testbed, simulation enables fully repeatable and configurable experiments that can often be conducted quickly for arbitrary hypothetical scenarios. In spite of these promises, current simulation practice is often not conducive to obtaining scientifically sound results. To date, most simulation results in the parallel and distributed computing literature are obtained with simulators that are ad hoc, unavailable, undocumented, and/or no longer maintained. For instance, Naicken et al. [62] point out that out of 125 recent papers they surveyed that study peer-to-peer systems, 52% use simulation and mention a simulator, but 72% of them use a custom simulator. As a result, most published simulation results build on throw-away (short-lived and non validated) simulators that are specifically designed for a particular study, which prevents other researchers from building upon it. There is thus a strong need for recognized simulation frameworks by which simulation results can be reproduced, further analyzed and improved.

The *SimGrid* simulation toolkit [74], whose development is partially supported by POLARIS, is specifically designed for studying large scale distributed computing systems. It has already been successfully used for simulation of grid, volunteer computing, HPC, cloud infrastructures and we have constantly invested on the software quality, the scalability [66] and the validity of the underlying network models [64, 112]. Many simulators of MPI applications have been developed by renowned HPC groups (e.g., at SDSC [110], BSC [60], UIUC [118], Sandia Nat. Lab. [113], ORNL [67] or ETH Zürich [89] for the most prominent ones). Yet, to scale most of them build on restrictive network and application modeling assumptions that make them difficult to extend to more complex architectures and to applications that do not solely build on the MPI API. Furthermore, simplistic modeling assumptions generally prevent to faithfully predict execution times, which limits the use of simulation to indication of gross trends at best. Our goal is to improve the quality of SimGrid to the point where it can be used effectively on a daily basis by practitioners to *reproduce the dynamic of real HPC systems*.

We also develop another simulation software, *PSI* (Perfect Simulator) [78, 72], dedicated to the simulation of very large systems that can be modeled as Markov chains. PSI provides a set of simulation kernels for Markov chains specified by events. It allows one to sample stationary distributions through the Perfect Sampling method (pioneered by Propp and Wilson [102]) or simply to generate trajectories with a forward Monte-Carlo simulation leveraging time parallel simulation (pioneered by Fujimoto [82], Lin and Lazowska [94]). One of the strength of the PSI framework is its expressiveness that allows us to easily study networks with finite and infinite capacity queues [73]. Although PSI already allows to simulate very large and complex systems, our main objective is to push its scalability even further and *improve its capabilities by one or several orders of magnitude*.

3.4 Local Interactions and Transient Analysis in Adaptive Dynamic Systems

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

Many systems can be effectively described by stochastic population models. These systems are composed of a set of n entities interacting together and the resulting stochastic process can be seen as a continuous-time Markov chain with a finite state space. Many numerical techniques exist to study the behavior of Markov chains, to solve stochastic optimal control problems [103] or to perform model-checking [61]. These techniques, however, are limited in their applicability, as they suffer from the *curse of dimensionality*: the state-space grows exponentially with n .

This results in the need for approximation techniques. Mean field analysis offers a viable, and often very accurate, solution for large n . The basic idea of the mean field approximation is to count the number of entities that are in a given state. Hence, the fluctuations due to stochasticity become negligible as the number of entities grows. For large n , the system becomes essentially deterministic. This approximation has been originally developed in statistical mechanics for very large systems composed of more than 10^{20} particles (called entities here). More recently, it has been claimed that, under some conditions, this approximation can be successfully used for stochastic systems composed of a few tens of entities. The claim is supported by various convergence results [83, 93, 116], and has been successfully applied in various domains: wireless networks [65], computer-based systems [86, 98, 111], epidemic or rumour propagation [75, 90] and bike-sharing systems [79]. It is also used to develop distributed control strategies [115, 97] or to construct approximate solutions of stochastic model checking problems [68, 70, 69].

Within the POLARIS project, we will continue developing both the theory behind these approximation techniques and their applications. Typically, these techniques require a homogeneous population of objects where the dynamics of the entities depend only on their state (the state space of each object must not scale with n the number of objects) but neither on their identity nor on their spatial location. Continuing our work in [83], we would like to be able to handle heterogeneous or uncertain dynamics. Typical applications are caching mechanisms [86] or bike-sharing systems [80]. A second point of interest is the use of mean field or large deviation asymptotics to compute the time between two regimes [105] or to reach an equilibrium state. Last, mean-field methods are mostly descriptive and are used to analyse the performance of a given system. We wish to extend their use to solve optimal control problems. In particular, we would like to implement numerical algorithms that use the framework that we developed in [84] to build distributed control algorithms [77] and optimal pricing mechanisms [85].

3.5 Distributed Learning in Games and Online Optimization

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

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

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

A key challenge here is that most of the literature on learning in games has focused on *static* games with a *finite number of actions* per player [81, 106]. While relatively tractable, such games are ill-suited to practical applications where players pick an action from a continuous space or when their payoff

functions evolve over time – this being typically the case in our target applications (e.g., routing in packet-switched networks or energy-efficient throughput maximization in wireless). On the other hand, the framework of online convex optimization typically provides worst-case performance bounds on the learner's *regret* that the agents can attain irrespectively of how their environment varies over time. However, if the agents' environment is determined chiefly by their interactions these bounds are fairly loose, so more sophisticated convergence criteria should be applied.

From an algorithmic standpoint, a further challenge occurs when players can only observe their own payoffs (or a perturbed version thereof). In this bandit-like setting regret-matching or trial-and-error procedures guarantee convergence to an equilibrium in a weak sense in certain classes of games. However, these results apply exclusively to static, finite games: learning in games with continuous action spaces and/or nonlinear payoff functions cannot be studied within this framework. Furthermore, even in the case of finite games, the complexity of the algorithms described above is not known, so it is impossible to decide a priori which algorithmic scheme can be applied to which application.

4 Application domains

4.1 Large Computing Infrastructures

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

4.2 Next-Generation Wireless Networks

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

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

4.3 Energy and Transportation

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

4.4 Social Computing Systems

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

5 Highlights of the year

P. Mertikopoulos is a CNRS bronze medal finalist: <https://members.loria.fr/SPerdrix/files/cn6/>

5.1 Awards

- Spotlight award at NeurIPS 2020 for the paper "No-regret learning and mixed Nash equilibria: They do not mix" [22]
- Spotlight award at NeurIPS 2020 for the paper "Explore aggressively, update conservatively: Stochastic extragradient methods with variable stepsize scaling" [26]
- Spotlight award at ICLR 2020 for the paper "Online and stochastic optimization beyond Lipschitz continuity: A Riemannian approach" [20]

6 New software and platforms

6.1 New software

6.1.1 Framesoc

Keywords: HPC, Embedded systems

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

News of the Year: No new development. Maintenance is ensured by Damien Dosimont, now at Barcelona Supercomputer Center.

URL: <http://soctrace-inria.github.io/framesoc/>

Contacts: Jean-Marc Vincent, Guillaume Huard

Participants: Arnaud Legrand, Jean-Marc Vincent

6.1.2 Ocelotl

Name: Multidimensional Overviews for Huge Trace Analysis

Keywords: HPC, Embedded systems

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

News of the Year: No new development. Maintenance is ensured by Damien Dosimont, now at Barcelona Supercomputer Center.

URL: <http://soctrace-inria.github.io/ocelotl/>

Contacts: Jean-Marc Vincent, Arnaud Legrand

Participants: Arnaud Legrand, Jean-Marc Vincent

6.1.3 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 2 major releases in 2020. SMPI is now regularly tested on medium scale benchmarks of the exascale suite. The Wifi support was improved, through more example and documentation, and an energy model of wifi links was proposed. Many bugs were fixed in the bindings to the ns-3 packet-level network simulator, which now allows to simulate Wifi links using ns-3 too. We enriched the API expressiveness to allow the construction of activity tasks. We also pursued our efforts to improve the documentation of the software, simplified the web site, and made a lot of bug fixing and code refactoring.

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

Contacts: Arnaud Legrand, Martin Quinson, Frédéric Suter

Participants: Adrien Lèbre, Arnaud Legrand, Augustin Degomme, Frédéric Suter, Jean-Marc Vincent, Jonathan Pastor, Luka Stanisic, Martin Quinson, Samuel Thibault, Emmanuelle Saillard

Partners: CNRS, ENS Rennes

6.1.4 StarPU

Name: The StarPU Runtime System

Keywords: Multicore, GPU, Scheduling, HPC, Performance

Scientific Description: Traditional processors have reached architectural limits which heterogeneous multicore designs and hardware specialization (eg. coprocessors, accelerators, ...) intend to address. However, exploiting such machines introduces numerous challenging issues at all levels, ranging from programming models and compilers to the design of scalable hardware solutions. The design of efficient runtime systems for these architectures is a critical issue. StarPU typically makes it much easier for high performance libraries or compiler environments to exploit heterogeneous multicore machines possibly equipped with GPGPUs or Cell processors: rather than handling low-level issues, programmers may concentrate on algorithmic concerns. Portability is obtained by the means of a unified abstraction of the machine. StarPU offers a unified offloadable task abstraction named "codelet". Rather than rewriting the entire code, programmers can encapsulate existing functions within codelets. In case a codelet may run on heterogeneous architectures, it is possible to specify one function for each architectures (eg. one function for CUDA and one function for CPUs). StarPU takes care to schedule and execute those codelets as efficiently as possible over the entire machine. In order to relieve programmers from the burden of explicit data transfers, a high-level data management library enforces memory coherency over the machine: before a codelet starts (eg. on an accelerator), all its data are transparently made available on the compute resource. Given its expressive interface and portable scheduling policies, StarPU obtains portable performances by efficiently (and easily) using all computing resources at the same time. StarPU also takes advantage of the heterogeneous nature of a machine, for instance by using scheduling strategies based on auto-tuned performance models.

StarPU is a task programming library for hybrid architectures

The application provides algorithms and constraints: - CPU/GPU implementations of tasks - A graph of tasks, using either the StarPU's high level GCC plugin pragmas or StarPU's rich C API

StarPU handles run-time concerns - Task dependencies - Optimized heterogeneous scheduling - Optimized data transfers and replication between main memory and discrete memories - Optimized cluster communications

Rather than handling low-level scheduling and optimizing issues, programmers can concentrate on algorithmic concerns!

Functional Description: StarPU is a runtime system that offers support for heterogeneous multicore machines. While many efforts are devoted to design efficient computation kernels for those architectures (e.g. to implement BLAS kernels on GPUs), StarPU not only takes care of offloading such kernels (and implementing data coherency across the machine), but it also makes sure the kernels are executed as efficiently as possible.

URL: <https://starpu.gitlabpages.inria.fr/>

Publications: hal-02403109, hal-02421327, hal-02872765, hal-02914793, hal-02933803, hal-01473475, hal-01474556, tel-01538516, hal-01718280, hal-01618526, tel-01816341, hal-01410103, hal-01616632, hal-01353962, hal-01842038, hal-01181135, tel-01959127, hal-01355385, hal-01284004, hal-01502749, hal-01502749, hal-01332774, hal-01372022, tel-01483666, hal-01147997, hal-01182746, hal-01120507, hal-01101045, hal-01081974, hal-01101054, hal-01011633, hal-01005765, hal-01283949, hal-00987094, hal-00978364, hal-00978602, hal-00992208, hal-00966862, hal-00925017, hal-00920915, hal-00824514, hal-00926144, hal-00773610, hal-01284235, hal-00853423, hal-00807033, tel-00948309, hal-00772742, hal-00725477, hal-00773114, hal-00697020, hal-00776610, hal-01284136, inria-00550877, hal-00648480, hal-00661320, inria-00606200, hal-00654193, inria-00547614, hal-00643257, inria-00606195, hal-00803304, inria-00590670, tel-00777154, inria-00619654, inria-00523937, inria-00547616, inria-00467677, inria-00411581, inria-00421333, inria-00384363, inria-00378705, hal-01517153, tel-01162975, hal-01223573, hal-01361992, hal-01386174, hal-01409965, hal-02275363, hal-02296118

Authors: Simon Archipoff, Cédric Augonnet, Olivier Aumage, Guillaume Beauchamp, William Braik, Bérenger Bramas, Alfredo Buttari, Adrien Cassagne, Arthur Chevalier, Jérôme Clet-Ortega, Terry Cojean, Nicolas Collin, Ludovic Courtès, Yann Courtois, Jean-Marie Couteyen, Vincent Danjean, Alexandre Denis, Lionel Eyraud-Dubois, Nathalie Furmento, Brice Goglin, David Antonio Gomez Jauregui, Sylvain Henry, Andra Hugo, Mehdi Juhour, Thibaud Lambert, Erwan Leria, Xavier Lacoste, Mathieu Lirzin, Benoît Lize, Benjamin Lorendeau, Antoine Lucas, Brice Mortier, Stojce Nakov, Raymond Namyst, Lucas Leandro Nesi, Joris Pablo, Damien Pasqualinotto, Samuel Pitoiset, Nguyen Quôc-Dinh, Cyril Roelandt, Anthony Roy, Chiheb Sakka, Corentin Salingue, Lucas Schnorr, Marc Sergent, Anthony Simonet, Luka Stanisic, Ludovic Stordeur, Guillaume Sylvand, Francois Tessier, Samuel Thibault, Leo Villeveygoux, Pierre-André Wacrenier

Contacts: Samuel Thibault, Nathalie Furmento, Olivier Aumage

Participants: Corentin Salingue, Andra Hugo, Benoît Lize, Cédric Augonnet, Cyril Roelandt, Francois Tessier, Jérôme Clet-Ortega, Ludovic Courtès, Ludovic Stordeur, Marc Sergent, Mehdi Juhour, Nathalie Furmento, Nicolas Collin, Olivier Aumage, Pierre-André Wacrenier, Raymond Namyst, Samuel Thibault, Simon Archipoff, Xavier Lacoste, Terry Cojean, Yanis Khorsi, Philippe Virouleau, Loïc Jouans, Leo Villeveygoux

6.1.5 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: <http://psi.gforge.inria.fr/>

Contacts: Jean-Marc Vincent, Florence Perronnin, Vincent Danjean

6.1.6 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. The next generations of PSI and marmoteCore lie in the MARTO project.

URL: <http://marmotecore.gforge.inria.fr/>

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

Contacts: Alain Jean-Marie, Jean-Marc Vincent

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

6.1.7 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>

Contacts: Vincent Danjean, Jean-Marc Vincent

6.1.8 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://mescal.imag.fr/membres/panayotis.mertikopoulos/publications.html>

Contact: Panayotis Mertikopoulos

6.1.9 taktuk

Name: TakTuk: Adaptive large scale remote executions deployment

Keyword: Deployment

Functional Description: TakTuk is a tool for deploying parallel remote executions of commands to a potentially large set of remote nodes. It spreads itself using an adaptive algorithm and sets up an interconnection network to transport commands and perform I/Os multiplexing/demultiplexing. The TakTuk engine dynamically adapts to environment (machine performance and current load, network contention) by using a reactive work-stealing algorithm that mixes local parallelization and work distribution.

News of the Year: No new development but the software is maintained to follow architecture and software upgrades.

URL: <http://taktuk.gforge.inria.fr/>

Contacts: Pierre Neyron, Guillaume Huard

Participants: Benoît Claudel, Guillaume Huard, Johann Bourcier, Olivier Richard, Pierre Neyron, Thierry Gautier

Partner: LIG

7 New results

The new results produced by the team in 2020 can be grouped into the following categories; for each new result, see the corresponding reference for further details.

7.1 System Analysis and Experiments

- Measuring phenology uncertainty with large scale image processing [1]
- F 3 ARIoT: A Framework for Autonomic Resilience of IoT Applications in the Fog [14]
- Verification of a Failure Management Protocol for Stateful IoT Applications [32]

7.2 Performance Evaluation and Measurements of Distributed Systems and Networks

- Faithful Performance Prediction of a Dynamic Task-based Runtime System, an Opportunity for Task Graph Scheduling [38]
- Communication-Aware Load Balancing of the LU Factorization over Heterogeneous Clusters [31]
- Fast Optimization with Zeroth-Order Feedback in Distributed, Multi-User MIMO Systems [3]
- SRPT-ECF: challenging Round-Robin for stream-aware multipath scheduling [27]
- Derivative-Free Optimization over Multi-User MIMO Networks [34]

7.3 Mean Field Analysis and Mean Field Games

- Power-of-d-Choices with Memory: Fluid Limit and Optimality [2]
- Vaccination in a Large Population: Mean Field Equilibrium versus Social Optimum [23]
- Performance analysis methods for list-based caches with non-uniform access [5]
- A Mean Field Game Analysis of SIR Dynamics with Vaccination [7]

7.4 Game Theory

- When is selfish routing bad? The price of anarchy in light and heavy traffic [6]
- No-regret learning and mixed Nash equilibria: They do not mix [22]
- Market sentiments and convergence dynamics in decentralized assignment economies [15]
- Quick or cheap? Breaking points in dynamic markets [30]
- The importance of memory for price discovery in decentralized markets [13]

7.5 Privacy, Fairness and Transparency

- Linear Regression from Strategic Data Sources [8]
- On Fair Selection in the Presence of Implicit Variance [21]

7.6 Optimization Methods

- On the convergence of mirror descent beyond stochastic convex programming [18]
- Mini-batch forward-backward-forward methods for solving stochastic variational inequalities
Explore Aggressively, Update Conservatively: Stochastic Extragradient Methods with Variable Step-size Scaling [4]
- Online and Stochastic Optimization beyond Lipschitz Continuity: A Riemannian Approach [20]
- On the almost sure convergence of stochastic gradient descent in non-convex problems [29]
- Online non-convex optimization with imperfect feedback [24]

7.7 Learning

- Towards Designing Cost-Optimal Policies to Utilize IaaS Clouds with Online Learning [17]
- Path Planning Problems with Side Observations—When Colonels Play Hide-and-Seek [33]
- Gradient-free Online Learning in Games with Delayed Rewards [25]
- A new regret analysis for Adam-type algorithms [19]
- Finite-time last-iterate convergence for multi-agent learning in games [28]

7.8 Energy Optimization

- Dynamic Speed Scaling Minimizing Expected Energy Consumption for Real-Time Tasks [10]
- Feasibility of on-line speed policies in real-time systems [11]
- A Pseudo-Linear Time Algorithm for the Optimal Discrete Speed Minimizing Energy Consumption [9]
- Discrete and Continuous Optimal Control for Energy Minimization in Real-Time Systems [35]

7.9 Covid Deconfinement Policies and Testing

- Focus mass testing: How to overcome low test accuracy [58]
- Green bridges: Reconnecting Europe to avoid economic disaster [40]
- A vaccination policy by zones [56]

8 Bilateral contracts and grants with industry

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.

9 Partnerships and cooperations

9.1 National initiatives

ANR Bary Pradelski (PI), P. Mertikopoulos and P. Loiseau obtained funding from the ANR for the project ALIAS (Adaptive Learning for Interactive Agents and Systems). This is a bilateral PRCI (collaboration internationale) project joint with Singapore University of Technology and Design (SUTD). The Singapore team consists of G. Piliouras and G. Panageas.

ORACLESS (2016–2021) 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.

Nicolas Gast obtained a funding from the ANR for the JCJC project REFINO (Refined Mean Field Optimization). The main objective of this project is to leverage our expertise on mean field and refined mean field approximation to solve distributed optimization problems.

Patrick Loiseau obtained a funding from the ANR for FairPlay, a starting grant (JCJC) obtained in September 2020 (covering the period 2021-2025). The goal of the project is to develop fair algorithms via game theory and sequential learning techniques, in particular for problems of auctions and of matching.

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.

9.2 Inria associate team not involved in an IIL

Title: ReDaS

Coordinator: Guillaume Huard

Partners:

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

Inria contact: Guillaume Huard

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

10 Dissemination

10.1 Promoting scientific activities

10.1.1 Scientific events: organisation

General chair, scientific chair A. Legrand was scientific chair of the "Performance and Power Modeling, Prediction and Evaluation" track for the EuroPar 2020 conference.

10.1.2 Scientific events: selection

Chair of conference program committees P. Mertikopoulos: Area chair at NeurIPS 2020; Area chair at ICLR 2021 (paper selection in 2020, conference taking place in 2021)

Member of the conference program committees J. Anselmi: IFIP Performance

N. Gast: SIGMETRICS, ICML, ICLR

B. Gaujal: SIGMETRICS

A. Legrand: EuroPar, PRECS

P. Loiseau: NeurIPS, ICML, AAAI, IJCAI, PETS, NetEcon

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

10.1.3 Journal

Member of the editorial boards P. Mertikopoulos is associate editor for JDG (Journal of Dynamics and Games).

P. Mertikopoulos is associate editor for MCAP (Methodology and Computing in Applied Probability).

P. Mertikopoulos is associate editor for RAIRO Operations Research

N. Gast is associate editor for PEVA (Performance Evaluation) and for Stochastic Models.

P. Loiseau is an associate editor at ACM Transactions on Internet Technology (TOIT)

P. is an associate editor at IEEE Transactions on Big Data (TBD)

Reviewer - reviewing activities All members of the team are active reviewers for several international journals.

10.1.4 Invited talks

A. Legrand: Two invited talks at the JDEV (<http://devlog.cnrs.fr/jdev2020>): "Reproducibility Crisis, Open Science, and Software Factories" and "Software Factories for Reproducible Research in Big Data/DL/AI [Online; invited talk]

P. Mertikopoulos:

- National Technical University of Athens "Games, Dynamics, and Spurious Attractors" [Online; invited talk]
- French Days on Optimization and Decision Science "Algorithmic game theory: from multi-agent optimization to online learning" [Online; invited course]
- One World Optimization Seminar / One World Game Theory Seminar "Games, Dynamics, and Optimization" [Online; invited talk]
- GDO 2020 "Learning in time-varying games" [Rome, Feb. 2020; invited talk]

10.1.5 Research administration

N. Gast is co-responsible of the Doctoral-School "MSTII" (maths and computer science)

B. Gaujal is a member of the scientific committee of GDR-IM and a member of the council of 'pole MSTIC' Grenoble

A. Legrand is responsible for the SRCPR ("Systèmes Répartis, Calcul Parallèle et Réseaux") research axis of the LIG.

A. Legrand is leading the HAC-SPECIS ("High-performance Application and Computers, Studying Performance and Correctness In Simulation") Inria Project Laboratory.

P. Mertikopoulos is responsible for the "Noeud Est" of the GDR Jeux (RT 2932)

P. Mertikopoulos is the working group coordinator, core group member and management committee (MC) representative for France in the European Network for Game Theory (GAMENET).

P. Loiseau is chair of the steering committee of NetEcon (since 2013)

P. Loiseau is the co-holder (with Marie-Christine Rousset from LIG) of a chair of the 3IA institute MIAI at Grenoble Alpes on "Explainable and Responsible AI".

10.2 Teaching - Supervision - Juries

10.2.1 Teaching

- B. Gaujal was involved in M1 exercise sessions(Ensimag) in applied probability
- V. Danjean was involved in INFO3 and INFO4 at Polytech Grenoble (System Architecture, Internship supervising, ...) and in M1 Info (Operationg systems and Parallel Programming course, Operating System project)
- A. Legrand was involved in Scientific Methodology and Performance Evaluation (M2 MOSIG, UGA),Parallel Systems (M2 MOSIG, UGA), Probability and Simulation (M1, Polytech/UGA),Performance Evaluation (M1, Polytech/UGA), Reproducible Research (Doctoral School MSTII, UGA)
- J. Anselmi: Probability and Simulation (M1, Polytech Grenoble), Performance Evaluation (M1, Polytech Grenoble).
- N. Gast is responsible of the master course Optimization under Uncertainties (Master ORCO [Operations Research in Grenoble]), L3 course Introduction to Machine Learning.
- P. Mertikopoulos gave an invited course for PhD and MSc students as part of the SMAI-MODE programme in September 2020
- J.-M. Vincent teaches Probability for Informatics and Performance Evaluation at Ensimag, and Mathematics for Computer Science (1st year) and Scientific Methodology and Performance Evaluation (2nd year) at the Master of Computer Science.
- G. Huard taught the course Object Oriented Design class for the M1 INFO, UGA.
- P. Loiseau: Introduction to Data Analysis (M1 MOSIG, UGA), INF421: Conception et analyse d'algorithmes (Ecole Polytechnique, 2A), and INF581: Advanced Topics in Artificial Intelligence (Ecole Polytechnique, 3A/M1)
- B. Pradelski: Introduction to Game Theory – ETH Zurich (Spring 2020)

10.2.2 Supervision

Supervision of PhD students and postdocs:

- B. Jonglez (Bruno Gaujal and Martin Heusse)[43]
- S. Plassart (Bruno Gaujal and Alain Girault)[44]
- B. Donassolo (P. Mertikopoulos and A. Legrand)[41]

- Dong Quan Vu (P. Loiseau)[45]
 - P. Rocha Bruel (A. Legrand and Alfredo Goldman)
 - T. Cornebize (A. Legrand)
 - S. Zrigui (A. Legrand and D. Trystram)
 - K. Khun (Bruno Gaujal and Nicolas Gast)
 - C. Yan (Bruno Gaujal and Nicolas Gast)
 - Y. G. Hsieh (P. Mertikopoulos, F. Iutzeler and J. Malick, LJK)
 - K. Antonakopoulos (P. Mertikopoulos and E. V. Belmega, ETIS/ENSEA)
 - B. Roussillon (P. Mertikopoulos and P. Loiseau)
 - A. Janon (G. Huard and A. Legrand)
 - V. Emelianov (N. Gast and P. Loiseau)
 - T. Barzolla (N. Gast with Vincent Jost and Van-Dat Cung from G-SCOP laboratory)
 - Lucas Leandro Nesi (A. Legrand and Lucas Mello Schnorr)
 - Till Kletti (Patrick Loiseau and Sihem Amer-Yahia from CNRS/LIG, Cifre with Jean-Michel Renders from Naver Labs)
 - Sebastian Allmeier (Nicolas Gast)
 - Vera Sosnovik (O. Goga and P. Loiseau)
 - Eleni Gkiouzepi (P. Loiseau)
 - Dimitrios Moustakas (B. Pradelski and P. Loiseau, with H. Nax from UZH)
 - Louis-Sebastien Rebuffi (J. Anselmi and B. Gaujal)
 - Simon Jantscheg (B. Pradelski and P. Loiseau, with H. Nax from UZH)
- Supervision of M2 Students:
- Victor Boone (B. Gaujal)
 - W. Azizian (P. Mertikopoulos, F. Iutzeler and J. Malick, LJK)
 - A. Giannou (P. Mertikopoulos, D. Fotakis, National Technical University of Athens)
 - Krishna Virendra Acharya (Patrick Loiseau and Nicolas Gast)

10.2.3 Juries

- N. Gast was member of the PhD Jury of Santi Duran.
- A. Legrand was member of the PhD Jury of Arthur Chevalier (ENS Lyon)
- P. Mertikopoulos was a reviewer for the PhD thesis of R. Pinot (U. Dauphine)
- P. Mertikopoulos was a reviewer for the PhD thesis of X. Fontaine (U. Paris-Saclay)
- P. Mertikopoulos was a reviewer for the PhD thesis of Y. P. Hsieh (EPFL; recipient of the EPFL EDEE thesis award)

10.3 Popularization

10.3.1 Internal or external Inria responsibilities

J-M. Vincenti is coordinating all the "Mediation Scientifique" activities for Inria Grenoble Rhône-Alpes.

10.3.2 Articles and contents

Bary Pradelski has been particularly active during the COVID-19 pandemic by promoting the Green zoning strategy to exit lockdown.

"Aiming for zero Covid-19: Europe needs to take action" with collective of ca. 30 academics, published in deVolkskrant, El Pais, la Repubblica, Le Monde, Rzeczpospolita, Sueddeutsche Zeitung.

"Vacunación: igualdad, fraternidad... y eficacia" with Miquel Oliu-Barton, El Mundo (15 December 2020).

"Covid-19 : « Qui vacciner en priorité ? Selon quels critères ? Comment hiérarchiser tout cela ? »" with Miquel Oliu-Barton, Le Monde (21 November 2020).

"Covid-19 : sanctuarisons les « zones vertes » !" with Miquel Oliu-Barton, Les Echos (14 October 2020).

"Más allá de las fronteras nacionales" with Miquel Oliu-Barton, El Pais (17 September 2020).

"Coronavirus : il faut « un plan de reconfinements ciblés réaliste, intelligible et commun »" with Miquel Oliu-Barton, Le Monde (26 August 2020).

"Sauver la saison touristique européenne" with Miquel Oliu-Barton, Le Monde (9 May 2020).

"Conectando las 'zonas verdes' de Europa: una propuesta para salvar el turismo" with Miquel Oliu-Barton, El Mundo (6 May 2020).

"Il faut une méthode de déconfinement efficace et sécurisée" with Miquel Oliu-Barton and Luc Attia, Le Monde (27 April 2020).

"Green zones: a mathematical proposal for how to exit from the COVID-19 lockdown" with Miquel Oliu-Barton, The Conversation (17 April 2020).

10.3.3 Education

- J-M. Vincent is elected member of the executive board of "Société Informatique de France"
- J-M. Vincent is member of the national organization committee of the Diplôme Inter-Universitaire "Enseigner l'Informatique au Lycée", and head of the cursus DIU EIL in UGA
- J-M Vincent is President of the committee of Baccalauréat subject committee in Academy of Grenoble
- V. Danjean is the head of the DU ISN formation (Diplôme Universitaire Informatique et Sciences du Numérique)
- A. Legrand has participated to the design of the 3rd edition of the MOOC on "Reproducible research: Methodological principles for a transparent science" <https://learninglab.inria.fr/mooc-recherche-reproductible-principes-methodologiques-pour-une-science-transparente/>. This 3rd edition is opened in self pace for a year and has attracted more than 7,100 persons.

11 Scientific production

11.1 Publications of the year

International journals

- [1] G. R. Alles, J. L. D. Comba, J.-M. Vincent, S. Nagai and L. M. Schnorr. 'Measuring phenology uncertainty with large scale image processing'. In: *Ecological Informatics* 59 (Sept. 2020), p. 101109. DOI: [10.1016/j.ecoinf.2020.101109](https://doi.org/10.1016/j.ecoinf.2020.101109). URL: <https://hal.inria.fr/hal-03130476>.
- [2] 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. URL: <https://hal.archives-ouvertes.fr/hal-02394147>.

- [3] O. Bilenne, P. Mertikopoulos and E. V. Belmega. ‘Fast Optimization with Zeroth-Order Feedback in Distributed, Multi-User MIMO Systems’. In: *IEEE Transactions on Signal Processing* 68 (2020), pp. 6085–6100. DOI: [10.1109/TSP.2020.3029983](https://doi.org/10.1109/TSP.2020.3029983). URL: <https://hal.archives-ouvertes.fr/hal-02861460>.
- [4] R. I. Bot, P. Mertikopoulos, M. Staudigl and P. T. Vuong. ‘Mini-batch forward-backward-forward methods for solving stochastic variational inequalities’. In: *Stochastic Systems* (2020). URL: <https://hal.inria.fr/hal-02405776>.
- [5] G. Casale and N. Gast. ‘Performance analysis methods for list-based caches with non-uniform access’. In: *IEEE/ACM Transactions on Networking* (18th Dec. 2020), p. 18. DOI: [10.1109/TNET.2020.3042869](https://doi.org/10.1109/TNET.2020.3042869). URL: <https://hal.inria.fr/hal-03102188>.
- [6] R. Colini-Baldeschi, R. Cominetti, P. Mertikopoulos and M. Scarsini. ‘When is selfish routing bad? The price of anarchy in light and heavy traffic’. In: *Operations Research* 68.2 (2020), pp. 411–434. URL: <https://hal.inria.fr/hal-01891558>.
- [7] J. Doncel, N. Gast and B. Gaujal. ‘A Mean Field Game Analysis of SIR Dynamics with Vaccination’. In: *Probability in the Engineering and Informational Sciences* (11th Dec. 2020), pp. 1–18. DOI: [10.1017/S0269964820000522](https://doi.org/10.1017/S0269964820000522). URL: <https://hal.inria.fr/hal-01496885>.
- [8] N. Gast, S. Ioannidis, P. Loiseau and B. Roussillon. ‘Linear Regression from Strategic Data Sources’. In: *ACM Transactions on Economics and Computation* 8.2 (14th May 2020), pp. 1–24. DOI: [10.1145/3391436](https://doi.org/10.1145/3391436). URL: <https://hal.inria.fr/hal-02982682>.
- [9] 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* (2020). URL: <https://hal.archives-ouvertes.fr/hal-03030416>.
- [10] B. Gaujal, A. Girault and S. Plassart. ‘Dynamic Speed Scaling Minimizing Expected Energy Consumption for Real-Time Tasks’. In: *Journal of Scheduling* (4th 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>.
- [11] B. Gaujal, A. Girault and S. Plassart. ‘Feasibility of on-line speed policies in real-time systems’. In: *Real-Time Systems* (27th Apr. 2020). DOI: [10.1007/s11241-020-09347-y](https://doi.org/10.1007/s11241-020-09347-y). URL: <https://hal.inria.fr/hal-02557148>.
- [12] R. Kusters, D. Misevic, H. Berry, A. Cully, Y. Le Cunff, L. Dandoy, N. Díaz-Rodríguez, M. Ficher, J. Grizou, A. Othmani, T. Palpanas, M. Komorowski, P. Loiseau, C. Moulin-Frier, S. Nanini, D. Quercia, M. Sebag, F. Soulié Fogelman, S. Taleb, L. Tupikina, V. Sahu, J.-J. Vie and F. Wehbi. ‘Interdisciplinary Research in Artificial Intelligence: Challenges and Opportunities’. In: *Frontiers in Big Data* 3 (23rd Nov. 2020). DOI: [10.3389/fdata.2020.577974](https://doi.org/10.3389/fdata.2020.577974). URL: <https://hal.inria.fr/hal-0311148>.
- [13] 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>.
- [14] U. Ozeer, L. Letondeur, G. Salaün, F.-G. Ottogalli and J.-M. Vincent. ‘F 3 ARIoT: A Framework for Autonomic Resilience of IoT Applications in the Fog’. In: *Internet of Things* (1st Dec. 2020), pp. 1–54. DOI: [10.1016/j.iot.2020.100275](https://doi.org/10.1016/j.iot.2020.100275). URL: <https://hal.inria.fr/hal-02933365>.
- [15] 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>.

- [16] K. Uguen, C. Jubin, Y. Duffourd, C. Bardel, V. Malan, J.-M. Dupont, L. El Khattabi, N. Chatron, A. Vitobello, P.-A. Rollat-Farnier, C. E. Baulard, M. Lelorch, A. A. Leduc, F. Tran Mau-Them, V. Danjean, M. Délépine, M. Till, V. Meyer, S. Lyonnet, A.-I. Mosca-Boidron, J. Thevenon, L. Faivre, C. Thauvin-Robinet, C. Schluth-Bolard, A. Boland, R. E. Olaso, P. Callier, S. Romana, J.-F. Deleuze and D. Sanlaville. ‘Genome sequencing in cytogenetics: Comparison of short-read and linked-read approaches for germline structural variant detection and characterization’. In: *Molecular Genetics & Genomic Medicine* (27th Jan. 2020), pp. 1–10. DOI: [10.1002/mgg3.1114](https://doi.org/10.1002/mgg3.1114). URL: <https://hal.archives-ouvertes.fr/hal-02467795>.
- [17] X. Wu, P. Loiseau and E. Hyytiä. ‘Towards Designing Cost-Optimal Policies to Utilize IaaS Clouds with Online Learning’. In: *IEEE Transactions on Parallel and Distributed Systems* 31.3 (Mar. 2020), pp. 501–514. DOI: [10.1109/TPDS.2019.2935199](https://doi.org/10.1109/TPDS.2019.2935199). URL: <https://hal.inria.fr/hal-02303480>.
- [18] Z. Zhou, P. Mertikopoulos, N. Bambos, S. Boyd and P. W. Glynn. ‘On the convergence of mirror descent beyond stochastic convex programming’. In: *SIAM Journal on Optimization* 30.1 (2020), pp. 687–716. DOI: [10.1137/17M1134925](https://doi.org/10.1137/17M1134925). URL: <https://hal.archives-ouvertes.fr/hal-01643346>.

International peer-reviewed conferences

- [19] A. Alacaoglu, Y. Malitsky, P. Mertikopoulos and V. Cevher. ‘A new regret analysis for Adam-type algorithms’. In: ICML ’20: The 37th International Conference on Machine Learning. Vienna, Austria, 2020. URL: <https://hal.inria.fr/hal-03043697>.
- [20] K. Antonakopoulos, E. V. Belmega and P. Mertikopoulos. ‘Online and Stochastic Optimization beyond Lipschitz Continuity: A Riemannian Approach’. In: ICLR 2020: The 2020 International Conference on Learning Representations. Addis Ababa, Ethiopia, 2020. URL: <https://hal.inria.fr/hal-03043671>.
- [21] V. Emelianov, N. Gast, K. P. Gummadi and P. Loiseau. ‘On Fair Selection in the Presence of Implicit Variance’. In: The Twenty-First ACM Conference on Economics and Computation (EC’20). Budapest, Hungary, 13th July 2020. DOI: [10.1145/3391403.3399482](https://doi.org/10.1145/3391403.3399482). URL: <https://hal.inria.fr/hal-02880578>.
- [22] L. Flokas, E. V. Vlatakis-Gkaragkounis, T. Lianas, P. Mertikopoulos and G. Piliouras. ‘No-regret learning and mixed Nash equilibria: They do not mix’. In: NeurIPS ’20 - 34th International Conference on Neural Information Processing Systems. Vancouver, Canada, 2020. URL: <https://hal.inria.fr/hal-03043763>.
- [23] B. Gaujal, J. Doncel and N. Gast. ‘Vaccination in a Large Population: Mean Field Equilibrium versus Social Optimum’. In: netgcoop’20. Cargèse, France, 22nd Sept. 2021. URL: <https://hal.inria.fr/hal-02938850>.
- [24] A. Héliou, M. Martin, P. Mertikopoulos and T. Rahier. ‘Online non-convex optimization with imperfect feedback’. In: NeurIPS ’20 - 34th International Conference on Neural Information Processing Systems. Vancouver, Canada, 2020. URL: <https://hal.inria.fr/hal-03043746>.
- [25] A. Héliou, P. Mertikopoulos and Z. Zhou. ‘Gradient-free Online Learning in Games with Delayed Rewards’. In: ICML ’20: The 37th International Conference on Machine Learning. Vienna, Austria, 2020. URL: <https://hal.inria.fr/hal-03043703>.
- [26] Y.-G. Hsieh, F. Iutzeler, J. Malick and P. Mertikopoulos. ‘Explore Aggressively, Update Conservatively: Stochastic Extragradient Methods with Variable Stepsize Scaling’. In: NeurIPS ’20 - 34th International Conference on Neural Information Processing Systems. Vancouver / Virtual, Canada, 6th Dec. 2020, pp. 1–29. URL: <https://hal.inria.fr/hal-03002844>.
- [27] B. Jonglez, M. Heusse and B. Gaujal. ‘SRPT-ECF: challenging Round-Robin for stream-aware multipath scheduling’. In: FIT 2020 - Second Workshop on the Future of Internet Transport. Paris, France: <https://networking.ifip.org/2020/workshops/fit-2020>, 22nd June 2020, pp. 1–7. URL: <https://hal.inria.fr/hal-02570686>.

- [28] T. Lin, Z. Zhou, P. Mertikopoulos and M. I. Jordan. ‘Finite-time last-iterate convergence for multi-agent learning in games’. In: ICML ’20: The 37th International Conference on Machine Learning. Vienna, Austria, 2020. URL: <https://hal.inria.fr/hal-03043711>.
- [29] P. Mertikopoulos, N. Hallak, A. Kavis and V. Cevher. ‘On the almost sure convergence of stochastic gradient descent in non-convex problems’. In: NeurIPS ’20 - 34th International Conference on Neural Information Processing Systems. Vancouver, Canada, 2020. URL: <https://hal.inria.fr/hal-03043771>.
- [30] P. Mertikopoulos, H. Nax and B. Pradelski. ‘Quick or cheap? Breaking points in dynamic markets’. In: EC ’20: The 21st ACM Conference on Economics and Computation. Budapest, Hungary, 2020. URL: <https://hal.inria.fr/hal-03043688>.
- [31] L. L. Nesi, L. Mello Schnorr and A. Legrand. ‘Communication-Aware Load Balancing of the LU Factorization over Heterogeneous Clusters’. In: IEEE International Conference on Parallel and Distributed Systems (ICPADS). Hong Kong, France: <https://icpads2020.comp.polyu.edu.hk/>, 2nd Dec. 2020. URL: <https://hal.inria.fr/hal-02633985>.
- [32] U. Ozeer, G. Salaün, L. Letondeur, F.-G. Ottogalli and J.-M. Vincent. ‘Verification of a Failure Management Protocol for Stateful IoT Applications’. In: Proc. of FMICS’20. Vienne, Austria, 2nd Sept. 2020. DOI: [10.1007/978-3-030-58298-2_12](https://doi.org/10.1007/978-3-030-58298-2_12). URL: <https://hal.inria.fr/hal-02930872>.
- [33] D. Quan Vu, P. Loiseau, A. Silva and L. Tran-Thanh. ‘Path Planning Problems with Side Observations—When Colonels Play Hide-and-Seek’. In: AAAI 2020 - Thirty-Fourth AAAI Conference on Artificial Intelligence. New-York, United States, 7th Feb. 2020, pp. 1–15. URL: <https://hal.inria.fr/hal-02375789>.

Conferences without proceedings

- [34] O. Bilenne, P. Mertikopoulos and E. V. Belmega. ‘Derivative-Free Optimization over Multi-User MIMO Networks’. In: NetGCoop ’20: The 2020 International Conference on Network Games, Control and Optimization. Cargese, France, 2020. URL: <https://hal.archives-ouvertes.fr/hal-02881559>.
- [35] 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, 23rd Sept. 2020, pp. 1–8. URL: <https://hal.archives-ouvertes.fr/hal-03020330>.
- [36] M. Korczyński, Y. Nosyk, Q. Lone, M. Skwarek, B. Jonglez and A. Duda. ‘Inferring the Deployment of Inbound Source Address Validation Using DNS Resolvers’. In: ANRW ’20: Applied Networking Research Workshop. Madrid, Spain, 30th July 2020, pp. 9–11. DOI: [10.1145/3404868.3406668](https://doi.org/10.1145/3404868.3406668). URL: <https://hal.archives-ouvertes.fr/hal-02894078>.
- [37] M. Mendil, N. Gast and H.-J. Audéoud. ‘Collisions groupées lors du mécanisme d’évitement de collisions de CPL-G3’. In: Rencontres Francophones sur la Conception de Protocoles, l’Évaluation de Performance et l’Expérimentation des Réseaux de Communication. Lyon, France, 28th Sept. 2020. URL: <https://hal.archives-ouvertes.fr/hal-02879294>.
- [38] S. Thibault, L. Stanislav and A. Legrand. ‘Faithful Performance Prediction of a Dynamic Task-based Runtime System, an Opportunity for Task Graph Scheduling’. In: SIAM PP 2020 - SIAM Conference on Parallel Processing for Scientific Computing. Seattle, United States: <http://www.siam.org/meetings/pp20/>, 12th Feb. 2020. URL: <https://hal.inria.fr/hal-02943753>.

Scientific books

- [39] B. Pradelski and H. Nax. *Rational Choice*. 12th Oct. 2020. DOI: [10.1515/9783110673616](https://doi.org/10.1515/9783110673616). URL: <https://hal.archives-ouvertes.fr/hal-03100233>.

Scientific book chapters

- [40] 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>.

Doctoral dissertations and habilitation theses

- [41] B. De Moura Donassolo. ‘IoT Orchestration in the Fog’. Université Grenoble Alpes [2020-....], 4th Nov. 2020. URL: <https://tel.archives-ouvertes.fr/tel-03130144>.
- [42] N. Gast. ‘Refinements of Mean Field Approximation’. Université Grenoble Alpes, 30th Jan. 2020. URL: <https://tel.archives-ouvertes.fr/tel-02509756>.
- [43] B. Jonglez. ‘End-to-end mechanisms to improve latency in communication networks’. Université Grenoble Alpes [2020-....], 23rd Oct. 2020. URL: <https://tel.archives-ouvertes.fr/tel-03120529>.
- [44] S. Plassart. ‘Online optimization in dynamic real-time systems’. Université Grenoble Alpes [2020-....], 16th June 2020. URL: <https://tel.archives-ouvertes.fr/tel-02990646>.
- [45] D. Q. Vu. ‘Models and Solutions of Strategic Resource Allocation Problems: Approximate Equilibrium and Online Learning in Blotto Games’. Sorbonne Universites, UPMC University of Paris 6, 25th June 2020. URL: <https://tel.archives-ouvertes.fr/tel-02904074>.

Reports & preprints

- [46] O. Bilenne. *Dispatching to Parallel Servers: Solutions of Poisson's Equation for First-Policy Improvement*. 29th Aug. 2020. URL: <https://hal.archives-ouvertes.fr/hal-02925284>.
- [47] T. Cornebize and A. Legrand. *Simulation-based Optimization and Sensibility Analysis of MPI Applications: Variability Matters*. 15th Feb. 2021. URL: <https://hal.inria.fr/hal-03141988>.
- [48] B. Donassolo, A. Legrand, P. Mertikopoulos and I. Fajjari. *Online Reconfiguration of IoT Applications in the Fog: The Information-Coordination Trade-off*. 27th May 2020. URL: <https://hal.inria.fr/hal-02636987>.
- [49] N. Gast, B. Gaujal and C. Yan. *Exponential Convergence Rate for the Asymptotic Optimality of Whittle Index Policy*. 11th Dec. 2020. URL: <https://hal.inria.fr/hal-03041176>.
- [50] B. Gaujal, A. Girault and S. Plassart. *A Linear Time Algorithm Computing the Optimal Speeds Minimizing Energy Under Real-Time Constraints*. Inria Grenoble Rhône-Alpes, 10th Apr. 2020. URL: <https://hal.inria.fr/hal-02540230>.
- [51] C. Grasland and J.-M. Vincent. *Much ado about one (single) thing?: H2020 ODYCCEUS internal seminar*. 5th June 2020. URL: <https://halshs.archives-ouvertes.fr/halshs-02795093>.
- [52] N. Hallak, P. Mertikopoulos and V. Cevher. *Regret minimization in stochastic non-convex learning via a proximal-gradient approach*. 7th Dec. 2020. URL: <https://hal.archives-ouvertes.fr/hal-03043872>.
- [53] Y.-P. Hsieh, P. Mertikopoulos and V. Cevher. *The limits of min-max optimization algorithms: Convergence to spurious non-critical sets*. 7th Dec. 2020. URL: <https://hal.archives-ouvertes.fr/hal-03043862>.
- [54] L. Leandro Nesi, V. Garcia Pinto, M. Cogo Miletto and L. M. Schnorr. *StarVZ: Performance Analysis of Task-Based Parallel Applications*. 7th Oct. 2020. URL: <https://hal.inria.fr/hal-02960848>.
- [55] S. Zrigui, R. Y. De Camargo, A. Legrand and D. Trystram. *Improving the Performance of Batch Schedulers Using Online Job Runtime Classification*. 2020. URL: <https://hal.archives-ouvertes.fr/hal-03023222>.

Other scientific publications

- [56] M. Oliu-Barton and B. Pradelski. *A vaccination policy by zones*. Oct. 2020. URL: <https://hal.archives-ouvertes.fr/hal-03100256>.
- [57] M. Oliu-Barton, B. Pradelski and L. Attia. *policy report*. Apr. 2020. URL: <https://hal.archives-ouvertes.fr/hal-03100271>.
- [58] B. Pradelski and M. Oliu-Barton. *Focus mass testing: How to overcome low test accuracy*. Dec. 2020. URL: <https://hal.archives-ouvertes.fr/hal-03100250>.

11.2 Other

Softwares

- [59] [SW] S. Archipoff, C. Augonnet, O. Aumage, G. Beauchamp, B. Bramas, A. Buttari, A. Cassagne, J. Clet-Ortega, T. Cojean, N. Collin, V. Danjean, A. Denis, L. Eyraud-Dubois, N. Furmento, S. Henry, A. Hugo, M. Juhour, A. Juven, M. Keryell-Even, Y. Khorsi, T. Lambert, E. Leria, B. Lizé, M. Makni, S. Nakov, R. Namyst, L. Nesi Lucas, P. Joris, D. Pasqualinotto, S. Pitoiset, Q.-D. Nguyen, C. Roelandt, C. Sakka, C. Salingue, L. Mello Schnorr, M. Sergent, A. Simonet, L. Stanisic, B. Subervie, F. Tessier, S. Thibault, B. Videau, L. Villeveygoux and P.-A. Wacrenier, *StarPU* version 1.3.3, 20th Jan. 2020. HAL: [hal-02443512](https://hal.inria.fr/hal-02443512), URL: <https://hal.inria.fr/hal-02443512>, SWHID: [swh:1:dir:b6e19d99449a78805e7a55a341fbaba2bc431973;origin=https://hal.archives-ouvertes.fr/hal-02443512;visit=swh:1:snp:c21d3dfbd96e4fb502c534e59644dba14c542100;anchor=swh:1:rev:31be198773f103324593d26369f135fbde5b97f8;path=/](https://hal.archives-ouvertes.fr/hal-02443512).

11.3 Cited publications

- [60] 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*. June 2003.
- [61] C. Baier, B. Haverkort, H. Hermanns and J.-P. Katoen. ‘Model-checking algorithms for continuous-time Markov chains’. In: *Software Engineering, IEEE Transactions on* 29.6 (2003). URL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1205180.
- [62] A. Basu, S. Fleming, J. Stanier, S. Naicken, I. Wakeman and V. K. Gurbani. ‘The State of Peer-to-peer Network Simulators’. In: *ACM Computing Survey*. 45.4 (Aug. 2013).
- [63] D. Becker, F. Wolf, W. Frings, M. Geimer, B. Wylie and B. Mohr. ‘Automatic Trace-Based Performance Analysis of Metacomputing Applications’. In: *Parallel and Distributed Processing Symposium, 2007. IPDPS 2007. IEEE International*. Mar. 2007. URL: <http://dx.doi.org/10.1109/IPDPS.2007.370238>.
- [64] P. Bedaride, A. Degomme, S. Genaud, A. Legrand, G. Markomanolis, M. Quinson, M. L. Stillwell, F. Suter and B. Videau. ‘Toward Better Simulation of MPI Applications on Ethernet/TCP Networks’. In: *PMBS13 - 4th International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems*. Denver, United States, Nov. 2013. URL: <https://hal.inria.fr/hal-00919507>.
- [65] G. Bianchi. ‘Performance analysis of the IEEE 802.11 distributed coordination function’. In: *Selected Areas in Communications, IEEE Journal on* 18.3 (2000).
- [66] L. Bobelin, A. Legrand, M. A. G. David, P. Navarro, M. Quinson, F. Suter and C. Thiery. ‘Scalable Multi-Purpose Network Representation for Large Scale Distributed System Simulation’. In: *CCGrid 2012 – The 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*. Ottawa, Canada, May 2012, p. 19. URL: <https://hal.inria.fr/hal-00650233>.
- [67] S. Böhm and C. Engelmann. ‘xSim: The Extreme-Scale Simulator’. In: *Proceedings of the International Conference on High Performance Computing and Simulation (HPCS) 2011*. Istanbul, Turkey: IEEE Computer Society, Los Alamos, CA, USA, July 2011.

- [68] L. Bortolussi and J. Hillston. ‘Model checking single agent behaviours by fluid approximation’. In: *Information and Computation* 242 (2015). URL: <http://dx.doi.org/10.1016/j.ic.2015.03.002>.
- [69] L. Bortolussi and R. Lanciani. ‘Fluid Model Checking of Timed Properties’. In: *Formal Modeling and Analysis of Timed Systems*. Springer International Publishing, 2015.
- [70] L. Bortolussi and R. Lanciani. ‘Model Checking Markov Population Models by Central Limit Approximation’. In: *Quantitative Evaluation of Systems*. Lecture Notes in Computer Science 8054. Springer Berlin Heidelberg, 2013. (Visited on 11/06/2015).
- [71] H. Brunst, D. Hackenberg, G. Juckeland and H. Rohling. ‘Comprehensive Performance Tracking with Vampir 7’. In: *Tools for High Performance Computing 2009*. Ed. by M. S. Müller, M. M. Resch, A. Schulz and W. E. Nagel. Springer Berlin Heidelberg, 2010. URL: http://dx.doi.org/10.1007/978-3-642-11261-4_2.
- [72] A. Busic, B. Gaujal, G. Gorgo and J.-M. Vincent. ‘PSI2 : Envelope Perfect Sampling of Non Monotone Systems’. In: *QEST 2010 - International Conference on Quantitative Evaluation of Systems*. Williamsburg, VA, United States: IEEE, Sept. 2010, pp. 83–84. URL: <https://hal.inria.fr/hal-00788884>.
- [73] A. Busic, B. Gaujal and F. Perronnin. ‘Perfect Sampling of Networks with Finite and Infinite Capacity Queues’. In: *19th International Conference on Analytical and Stochastic Modelling Techniques and Applications (ASMTA) 2012*. Ed. by K. Al-Begain, D. Fiems and J.-M. Vincent. Vol. 7314. Lecture Notes in Computer Science. Grenoble, France: Springer, 2012, pp. 136–149. DOI: 10.1007/978-3-642-30782-9_10. URL: <https://hal.inria.fr/hal-00788003>.
- [74] H. Casanova, A. Giersch, A. Legrand, M. Quinson and F. Suter. ‘Versatile, Scalable, and Accurate Simulation of Distributed Applications and Platforms’. In: *Journal of Parallel and Distributed Computing* 74.10 (June 2014), pp. 2899–2917. DOI: 10.1016/j.jpdc.2014.06.008. URL: <https://hal.inria.fr/hal-01017319>.
- [75] A. Chaintreau, J.-Y. Le Boudec and N. Ristanovic. ‘The Age of Gossip: Spatial Mean Field Regime’. In: *SIGMETRICS Perform. Eval. Rev.* 37.1 (June 2009). DOI: 10.1145/2492101.1555363. URL: <http://doi.acm.org/10.1145/2492101.1555363>.
- [76] K. Coulomb, M. Faverge, J. Jazeix, O. Lagrasse, J. Marcouelle, P. Noisette, A. Redondy and C. Vuchener. *Visual trace explorer (ViTE)*. 2009.
- [77] J. Doncel, N. Gast and B. Gaujal. ‘Mean-Field Games with Explicit Interactions’. Feb. 2016. URL: <https://hal.inria.fr/hal-01277098>.
- [78] S. Durand, B. Gaujal, F. Perronnin and J.-M. Vincent. ‘A perfect sampling algorithm of random walks with forbidden arcs’. In: *QEST 2014 - 11th International Conference on Quantitative Evaluation of Systems*. Vol. 8657. Florence, Italy: Springer, Sept. 2014, pp. 178–193. DOI: 10.1007/978-3-319-10696-0_15. URL: <https://hal.inria.fr/hal-01069975>.
- [79] C. Fricker and N. Gast. ‘Incentives and redistribution in homogeneous bike-sharing systems with stations of finite capacity’. In: *EURO Journal on Transportation and Logistics* (June 2014), p. 31. DOI: 10.1007/s13676-014-0053-5. URL: <https://hal.inria.fr/hal-01086009>.
- [80] C. Fricker, N. Gast and H. Mohamed. ‘Mean field analysis for inhomogeneous bike sharing systems’. In: *AofA*. Montreal, Canada, July 2012. URL: <https://hal.inria.fr/hal-01086055>.
- [81] D. Fudenberg and D. K. Levine. *The Theory of Learning in Games*. Vol. 2. Economic learning and social evolution. Cambridge, MA: MIT Press, 1998.
- [82] R. M. Fujimoto. ‘Parallel Discrete Event Simulation’. In: *Commun. ACM* 33.10 (Oct. 1990). DOI: 10.1145/84537.84545. URL: <http://doi.acm.org/10.1145/84537.84545>.
- [83] N. Gast and B. Gaujal. ‘Markov chains with discontinuous drifts have differential inclusion limits’. In: *Performance Evaluation* 69.12 (2012), pp. 623–642. DOI: 10.1016/j.peva.2012.07.003. URL: <https://hal.inria.fr/hal-00787999>.
- [84] N. Gast, B. Gaujal and J.-Y. Le Boudec. ‘Mean field for Markov Decision Processes: from Discrete to Continuous Optimization’. In: *IEEE Transactions on Automatic Control* 57.9 (2012), pp. 2266–2280. DOI: 10.1109/TAC.2012.2186176. URL: <https://hal.inria.fr/hal-00787996>.

- [85] N. Gast, J.-Y. Le Boudec and D.-C. Tomozei. ‘Impact of Demand-Response on the Efficiency and Prices in Real-Time Electricity Markets’. In: *ACM e-Energy 2014*. Cambridge, United Kingdom, June 2014. DOI: [10.1145/2602044.2602052](https://doi.org/10.1145/2602044.2602052). URL: <https://hal.inria.fr/hal-01086036>.
- [86] N. Gast and B. Van Houdt. ‘Transient and Steady-state Regime of a Family of List-based Cache Replacement Algorithms’. In: *ACM SIGMETRICS 2015*. Portland, United States, June 2015. DOI: [10.1145/2745844.2745850](https://doi.org/10.1145/2745844.2745850). URL: <https://hal.inria.fr/hal-01143838>.
- [87] J. Gonzalez, J. Gimenez and J. Labarta. ‘Automatic detection of parallel applications computation phases’. In: *Parallel and Distributed Processing Symposium, International 0* (2009). URL: <http://doi.ieeecomputersociety.org/10.1109/IPDPS.2009.5161027>.
- [88] M. Heath and J. Etheridge. ‘Visualizing the performance of parallel programs’. In: *IEEE software* 8.5 (1991).
- [89] T. Hoefler, T. Schneider and A. Lumsdaine. ‘LogGOPSim - Simulating Large-Scale Applications in the LogGOPS Model’. In: *Proc. of the ACM Workshop on Large-Scale System and Application Performance*. June 2010.
- [90] L. Hu, J.-Y. Le Boudec and M. Vojnović. ‘Optimal channel choice for collaborative ad-hoc dissemination’. In: *INFOCOM, 2010 Proceedings IEEE*. IEEE, 2010.
- [91] 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).
- [92] J. C. de Kergommeaux, B. Stein and P. Bernard. ‘Paje, an interactive visualization tool for tuning multi-threaded parallel applications’. In: *Parallel Computing* 10.26 (2000), pp. 1253–1274.
- [93] T. G. Kurtz. *Approximation of population processes*. Vol. 36. SIAM, 1981.
- [94] Y.-B. Lin and E. D. Lazowska. ‘A Time-division Algorithm for Parallel Simulation’. In: *ACM Trans. Model. Comput. Simul.* 1.1 (Jan. 1991). DOI: [10.1145/102810.214307](https://doi.org/10.1145/102810.214307). URL: <http://doi.acm.org/10.1145/102810.214307>.
- [95] G. Llorc, J. González, H. Servat, J. Giménez and J. Labarta. ‘On-line Detection of Large-scale Parallel Application’s Structure’. In: *24th IEEE International Parallel and Distributed Processing Symposium (IPDPS’2010)*. 2010.
- [96] L. Mello Schnorr and A. Legrand. ‘Visualizing More Performance Data Than What Fits on Your Screen’. In: *Tools for High Performance Computing 2012*. Ed. by A. Cheptsov, S. Brinkmann, J. Gracia, M. M. Resch and W. E. Nagel. Springer Berlin Heidelberg, 2013, pp. 149–162. DOI: [10.1007/978-3-642-37349-7_10](https://doi.org/10.1007/978-3-642-37349-7_10). URL: <https://hal.inria.fr/hal-00842761>.
- [97] S. Meyn, P. Barooah, A. Busic and J. Ehren. ‘Ancillary service to the grid from deferrable loads: the case for intelligent pool pumps in Florida’. In: *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*. IEEE, 2013.
- [98] M. Mitzenmacher. ‘The power of two choices in randomized load balancing’. In: *Parallel and Distributed Systems, IEEE Transactions on* 12.10 (2001).
- [99] K. Mohror, K. Karavanic and A. Snively. ‘Scalable Event Trace Visualization’. In: *Euro-Par 2009 – Parallel Processing Workshops*. Ed. by H.-X. Lin, M. Alexander, M. Forsell, A. Knüpfer, R. Prodan, L. Sousa and A. Streit. Vol. 6043. Lecture Notes in Computer Science. Springer Berlin / Heidelberg, 2010. URL: http://dx.doi.org/10.1007/978-3-642-14122-5_27.
- [100] W. Nagel, A. Arnold, M. Weber, H. Hoppe and K. Solchenbach. ‘VAMPIR: Visualization and Analysis of MPI Resources’. In: *Supercomputer* 12.1 (1996).
- [101] 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. Transputer and Occam Engineering. IOS Press, 1995.
- [102] J. Propp and D. Wilson. ‘Coupling from the past: a user’s guide’. In: *DIMACS Series on Discrete Mathematics and Theoretical Computer Science* 41 (1998). Microsurveys in discrete probability.
- [103] M. L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.

- [104] 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., Proceedings of the*. 1993.
- [105] W. H. Sandholm and M. Staudigl. ‘A Sample Path Large Deviation Principle for a Class of Population Processes’. In: *arXiv preprint arXiv:1511.07897* (2015).
- [106] W. H. Sandholm. *Population Games and Evolutionary Dynamics*. Economic learning and social evolution. Cambridge, MA: MIT Press, 2010.
- [107] H. Servat, G. Llort, J. Giménez, K. Huck and J. Labarta. ‘Folding: detailed analysis with coarse sampling’. In: *Tools for High Performance Computing 2011*. Springer Berlin Heidelberg, 2012.
- [108] H. Servat, G. Llort, J. Gonzalez, J. Gimenez and J. Labarta. ‘Identifying code phases using piece-wise linear regressions’. In: *Parallel and Distributed Processing Symposium, 2014 IEEE 28th International*. IEEE. 2014.
- [109] B. Shneiderman. ‘The eyes have it: A task by data type taxonomy for information visualizations’. In: *Visual Languages, 1996. Proceedings., IEEE Symposium on*. IEEE. 1996.
- [110] M. Tikir, M. Laurenzano, L. Carrington and A. Snavey. ‘PSINS: An Open Source Event Tracer and Execution Simulator for MPI Applications’. In: *Proc. of the 15th International Euro-Par Conference on Parallel Processing*. LNCS 5704. Springer, Aug. 2009.
- [111] B. Van Houdt. ‘A Mean Field Model for a Class of Garbage Collection Algorithms in Flash-based Solid State Drives’. In: *Proceedings of the ACM SIGMETRICS*. SIGMETRICS ’13. Pittsburgh, PA, USA: ACM, 2013. DOI: [10.1145/2465529.2465543](https://doi.org/10.1145/2465529.2465543). URL: <http://doi.acm.org/10.1145/2465529.2465543>.
- [112] P. Velho, L. Mello Schnorr, H. Casanova and A. Legrand. ‘On the Validity of Flow-level TCP Network Models for Grid and Cloud Simulations’. In: *ACM Transactions on Modeling and Computer Simulation* 23.4 (Oct. 2013). URL: <https://hal.inria.fr/hal-00872476>.
- [113] J. J. Wilke, K. Sargsyan, J. P. Kenny, B. Debusschere, H. N. Najm and G. Hendry. ‘Euro-Par 2013 Parallel Processing: 19th International Conference, Aachen, Germany, August 26-30, 2013. Proceedings’. In: Berlin, Heidelberg: Springer Berlin Heidelberg, 2013. Chap. Validation and Uncertainty Assessment of Extreme-Scale HPC Simulation through Bayesian Inference.
- [114] F. Wolf and B. Mohr. ‘Automatic performance analysis of hybrid MPI/OpenMP applications’. In: *Journal of Systems Architecture* 49.10-11 (2003).
- [115] T. Yang, P. G. Mehta and S. P. Meyn. ‘A mean-field control-oriented approach to particle filtering’. In: *American Control Conference (ACC), 2011*. IEEE. 2011.
- [116] L. Ying. ‘On the Rate of Convergence of Mean-Field Models: Stein’s Method Meets the Perturbation Theory’. In: *arXiv preprint arXiv:1510.00761* (2015).
- [117] 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). URL: <http://dx.doi.org/10.1177/109434209901300310>.
- [118] G. Zheng, G. Kakulapati and L. Kalé. ‘BigSim: A Parallel Simulator for Performance Prediction of Extremely Large Parallel Machines’. In: *Proc. of the 18th International Parallel and Distributed Processing Symposium (IPDPS)*. Apr. 2004.