RESEARCH CENTRE

Inria Centre
at Université Côte d'Azur

2023 ACTIVITY REPORT

Project-Team NEO

Network Engineering and Operations

DOMAIN

Networks, Systems and Services, Distributed Computing

THEME

Networks and Telecommunications

Inria

Contents

Project-Team NEO 1

1 Team members, visitors, external collaborators 3

2 Overall objectives 4

3 Research program 4

4 Application domains 5
4.1 Network Science . . . 5
4.2 Network Engineering . . . 5

5 Highlights of the year 6
5.1 Awards . . . 6
5.2 Foreign Appointments . . . 6

6 New software, platforms, open data 6
6.1 New software . . . 6
6.1.1 marmote . . . 6

7 New results 7
7.1 Game Theory and Applications . . . 7
7.1.1 Games with Commitments and Noisy Observations . . . 7
7.1.2 Pricing Strategies for Digital Renting Platforms . . . 7
7.1.3 Constrained games in energy markets . . . 8
7.2 Stochastic Modeling . . . 8
7.2.1 Retrial Queues . . . 8
7.2.2 Spatial Queues with Nearest Neighbor Shifts . . . 8
7.2.3 Performance Models for Similary Caching . . . 8
7.2.4 Branching processes . . . 9
7.3 Data Analysis and Learning . . . 9
7.3.1 Empirical Risk Minimization (ERM) . . . 9
7.3.2 Characterization of the Generalization Error . . . 10
7.3.3 Online Learning . . . 11
7.3.4 Reinforcement Learning . . . 11
7.3.5 Federated Learning . . . 11
7.3.6 Privacy and Fairness . . . 12
7.3.7 Data Injection Attacks in Power Systems . . . 12
7.3.8 Graph Clustering . . . 12
7.4 Applications in Telecommunications . . . 13
7.4.1 Optimal Rate Memory Tradeoff in Multi-Access Coded Caching . . . 13
7.4.2 Deployment of Microservice-based Applications . . . 13
7.4.3 Learning Optimal Edge Processing with Offloading and Energy Harvesting . . . 13
7.4.4 Pricing and resource allocation . . . 14
7.5 Applications to Environmental Issues . . . 14

8 Bilateral contracts and grants with industry 15
8.1 Bilateral contracts with industry . . . 15
8.1.1 Accenture contract on the topic “Distributed Machine Learning for IoT applications” (Dec 2019 – November 2023) . . . 15
8.1.2 Accenture “Plan de Relance” (PLR) contract on the topic “Energy-Aware Federated Learning” (Oct 2022 – September 2024) . . . 15
8.1.3 Cifre contract with NSP-SmartProfile “Automated and responsible recommendation systems for digital marketing” (August 2022 – July 2025) . . . 15

8.1.4 Cifre contract with Orange Labs “Analytical modeling of large-scale wireless networks integrating RIS” (March 2023 – March 2026) . . . 16
8.1.5 Cifre contract with SAP “Privacy and fairness for ML” (December 2021 – December 2024) . . . 16
8.1.6 QITI 3IA Start-It-Up contract on the topic “Reinforcement Learning for Conversational Recommender Systems (RLCRS)” (April 2023 – March 2024) . . . 17

9 Partnerships and cooperations **17**
9.1 International initiatives . . . 17
9.1.1 Inria associate team not involved in an IIL or an international program . . . 17
9.1.2 STIC/MATH/CLIMAT AmSud projects . . . 17
9.2 International research visitors . . . 18
9.2.1 Visits of international scientists . . . 18
9.2.2 Visits to international teams . . . 24
9.3 European initiatives . . . 27
9.3.1 Horizon Europe . . . 27
9.3.2 H2020 projects . . . 28
9.4 National initiatives . . . 28

10 Dissemination **31**
10.1 Promoting scientific activities . . . 31
10.1.1 Scientific events: organisation . . . 31
10.1.2 Scientific events: selection . . . 32
10.1.3 Journal . . . 33
10.1.4 Invited talks . . . 34
10.1.5 Invited papers . . . 35
10.1.6 Leadership within the scientific community . . . 35
10.1.7 Research administration . . . 36
10.2 Teaching - Supervision - Juries . . . 36
10.2.1 Teaching . . . 36
10.2.2 Supervision . . . 37
10.2.3 Juries . . . 37

11 Scientific production **38**
11.1 Major publications . . . 38
11.2 Publications of the year . . . 39

Project-Team NEO

Creation of the Project-Team: 2017 December 01

Keywords

Computer sciences and digital sciences

A1.1.11. – Quantum architectures
A1.2.4. – QoS, performance evaluation
A1.2.5. – Internet of things
A1.2.6. – Sensor networks
A1.5. – Complex systems
A1.5.1. – Systems of systems
A1.5.2. – Communicating systems
A3.3.3. – Big data analysis
A3.4. – Machine learning and statistics
A3.5. – Social networks
A3.5.2. – Recommendation systems
A4.1. – Threat analysis
A5.9. – Signal processing
A6.1.1. – Continuous Modeling (PDE, ODE)
A6.1.2. – Stochastic Modeling
A6.2.2. – Numerical probability
A6.2.3. – Probabilistic methods
A6.2.6. – Optimization
A6.4.1. – Deterministic control
A6.4.2. – Stochastic control
A6.4.6. – Optimal control
A7.1. – Algorithms
A7.1.1. – Distributed algorithms
A7.1.2. – Parallel algorithms
A7.1.4. – Quantum algorithms
A8.1. – Discrete mathematics, combinatorics
A8.2.1. – Operations research
A8.6. – Information theory
A8.8. – Network science
A8.9. – Performance evaluation
A8.11. – Game Theory
A9.2. – Machine learning

A9.6. – Decision support

A9.9. – Distributed AI, Multi-agent

Other research topics and application domains

B2.3. – Epidemiology

B2.5.1. – Sensorimotor disabilities

B3.1. – Sustainable development

B3.1.1. – Resource management

B4.3.4. – Solar Energy

B4.4. – Energy delivery

B4.4.1. – Smart grids

B4.5.1. – Green computing

B6. – IT and telecom

B6.2. – Network technologies

B6.2.1. – Wired technologies

B6.2.2. – Radio technology

B6.3.3. – Network Management

B6.3.4. – Social Networks

B6.4. – Internet of things

B6.6. – Embedded systems

B8.1. – Smart building/home

B9.2.1. – Music, sound

B9.5.1. – Computer science

B9.5.2. – Mathematics

B9.6.3. – Economy, Finance

B9.6.4. – Management science

B9.6.5. – Sociology

1 Team members, visitors, external collaborators

Research Scientists

- Alain Jean-Marie [Team leader, INRIA, Senior Researcher]
- Sara Alouf [INRIA, Researcher, HDR]
- Eitan Altman [INRIA, Senior Researcher, HDR]
- Konstantin Avrachenkov [INRIA, Senior Researcher, HDR]
- Philippe Nain [INRIA, Emeritus, HDR]
- Giovanni Neglia [INRIA, Senior Researcher, HDR]
- Samir M. Perlaza [INRIA, Researcher, HDR]

Post-Doctoral Fellows

- Emmanouil Marios Athanasakos [INRIA, Post-Doctoral Fellow, from Feb 2023]
- Vinay Kumar Bindiganavile Ramadas [Université Côte d'Azur (until Jan 2023), INRIA (from Feb 2023), Post-Doctoral Fellow]
- Jake Clarkson [INRIA, Post-Doctoral Fellow, until Feb 2023]
- Mandar Datar [INRIA, Post-Doctoral Fellow, until Feb 2023]
- Francescomaria Faticanti [INRIA, Post-Doctoral Fellow, until Aug 2023]
- Diego Goldsztajn [INRIA, Post-Doctoral Fellow, from Dec 2023]
- Vijith Kumar Kizhakke Purakkal [INRIA, Post-Doctoral Fellow, until Oct 2023]
- Ashok Krishnan Komalan Sindhu [INRIA, Post-Doctoral Fellow]
- Charlotte Rodriguez [INRIA, Post-Doctoral Fellow]
- Sadaf Ul Zuhra [INRIA, until Jan 2023]

PhD Students

- Younes Ben Mazziane [Université Côte d'Azur]
- Olha Chuchuk [CERN, until Aug 2023]
- Francisco Daunas [The University of Sheffield, UK]
- Ibtihal El Mimouni [NSP SmartProfile, CIFRE]
- Louis Hauseux [Université Côte d'Azur, from Oct 2023]
- Caelin Kaplan [SAP, CIFRE]
- Othmane Marfoq [INRIA, until Nov 2023]
- Angelo Rodio [INRIA]
- Julian Alfonso Santos Bustos [ORANGE, CIFRE, from Sep 2023]
- Xiuzhen Ye [The University of Sheffield, UK, until Mar 2023]
- Xufeng Zhang [INRIA, from Dec 2023]
- Xinying Zou [INRIA]

Technical Staff

- Louis Hauseux [INRIA, Engineer, from May 2023 until Sep 2023]
- Hariprasad Manjunath Hegde [INRIA, Engineer, from Apr 2023]
- Othmane Marfoq [INRIA, Engineer, from Dec 2023]

Interns and Apprentices

- Mihir Raghvendra Deshpande [INRIA, Intern, from May 2023 until Jul 2023]
- Francesco Diana [INRIA, Intern, from Mar 2023 until Aug 2023]
- Cosimo Giani [INRIA, Intern, until Mar 2023]
- Mohamed Salah Jebali [INRIA, Intern, from Mar 2023 until Aug 2023]
- Tejas Sanjaykumar Pagare [INRIA, Intern, from May 2023 until Jul 2023]
- Xufeng Zhang [INRIA, Intern, from Mar 2023 until Aug 2023]

Administrative Assistant

- Jane Desplanques [INRIA]

Visiting Scientists

- Lucas Gamertsfelder [Macquarie University, Australia, until Jan 2023]
- Jacopo Talpini [UNIV MILAN BICOCCA, from Nov 2023]
- Kavitha Veeraruna [IIT BOMBAY, from Sep 2023]
- Lotte Weedage [UNIV TWENTE, from Apr 2023 until Jul 2023]

2 Overall objectives

NEO is an Inria project-team whose members are located in Sophia Antipolis (S. Alouf, K. Avrachenkov, G. Neglia, and S. M. Perlaza), in Avignon (E. Altman) at LIA (Lab. of Informatics of Avignon) and in Montpellier (A. Jean-Marie). E. Altman is also with the LINCS (Lab. for Information, Networking and Communication Sciences). S. M. Perlaza is also with the ECE department at Princeton Univ., N.J. USA; and the Mathematics Department of the Univ. de la Polynésie française (Laboratoire GAATI), Faaa, Tahiti.

The team is positioned at the intersection of Operations Research and Network Science. By using the tools of Stochastic Operations Research, we model situations arising in several application domains, involving networking in one way or the other. The aim is to understand the rules and the effects in order to influence and control them so as to engineer the creation and the evolution of complex networks.

3 Research program

The problems studied in NEO involve generally optimization, dynamic systems or randomness, and often all at the same time. The techniques we use to tackle these problems are those of Stochastic Operations Research, Applied Probabilities and Information Theory.

Stochastic Operations Research is a collection of modeling, optimization and numerical computation techniques, aimed at assessing the behavior of man-made systems driven by random phenomena, and at helping to make decisions in such a context.

The discipline is based on applied probability and focuses on effective computations and algorithms. Its core theory is that of Markov chains over discrete state spaces. This family of stochastic processes has,

at the same time, a very large modeling capability and the potential of efficient solutions. By "solution" is meant the calculation of some *performance metric*, usually the distribution of some random variable of interest, or its average, variance, etc. This solution is obtained either through exact "analytic" formulas, or numerically through linear algebra methods. Even when not analytically or numerically tractable, Markovian models are always amenable to "Monte-Carlo" simulations with which the metrics can be statistically measured.

An example of this is the success of classical Queueing Theory, with its numerous analytical formulas. Another important derived theory is that of the Markov Decision Processes, which allows to formalize *optimal* decision problems in a random environment. This theory allows to characterize the optimal decisions, and provides algorithms for calculating them.

Strong trends of Operations Research are: a) an increasing importance of multi-criteria multi-agent optimization, and the correlated introduction of Game Theory in the standard methodology; b) an increasing concern of (deterministic) Operations Research with randomness and risk, and the consequent introduction of topics like Chance Constrained Programming and Stochastic Optimization. Data analysis is also more and more present in Operations Research: techniques from statistics, like filtering and estimation, or Artificial Intelligence like clustering, are coupled with modeling in Machine Learning techniques like Q-Learning.

4 Application domains

4.1 Network Science

Network Science is a multidisciplinary body of knowledge, principally concerned with the emergence of global properties in a network of individual agents. These global properties emerge from "local" properties of the network, namely, the way agents interact with each other. The central model of "networks" is the graph (of Graph Theory/Operations Research). Nodes represent the different entities managing information and taking decisions, whereas, links represent the fact that entities interact, or not. Links are usually equipped with a "weight" that measures the intensity of such interaction. Adding evolution rules to this quite elementary representation leads to dynamic network models, the properties of which Network Science tries to analyze.

A classical example of properties sought in networks is the famous "six degrees of separation" (or "small world") property: how and why does it happen so frequently? Another ubiquitous property of real-life networks is the Zipf or "scale-free" distribution for degrees. Some of these properties, when properly exploited, lead to successful business opportunities: just consider the PageRank algorithm of Google, which miraculously connects the relevance of some Web information with the relevance of the other information that points to it.

4.2 Network Engineering

In its primary acceptation, Network Science involves little or no engineering: phenomena are assumed to be "natural" and emerge without external interventions. However, the idea comes fast to intervene in order to modify the outcome of the phenomena. This is where NEO is positioned. Beyond the mostly descriptive approach of Network Science, we aim at using the techniques of Operations Research so as to engineer complex networks.

To quote two examples: controlling the spread of diseases through a "network" of people is of primarily interest for mankind. Similarly, controlling the spread of information or reputation through a social network is of great interest in the Internet. Precisely, given the impact of web visibility on business income, it is tempting (and quite common) to manipulate the graph of the web by adding links so as to drive the PageRank algorithm to a desired outcome.

Another interesting example is the engineering of community structures. Recently, thousands of papers have been written on the topic of community *detection* problem. In most of the works, the researchers propose methods, most of the time, heuristics, for detecting communities or dense subgraphs inside a large network. Much less effort has been put in the understanding of community formation process and even much less effort has been dedicated to the question of how one can influence the

process of community formation, e.g. in order to increase overlap among communities and reverse the fragmentation of the society.

Our ambition for the medium term is to reach an understanding of the behavior of complex networks that will make us capable of influencing or producing a certain property in a given network. For this purpose, we will develop families of models to capture the essential structure, dynamics, and uncertainty of complex networks. The "solution" of these models will provide the correspondence between metrics of interest and model parameters, thus opening the way to the synthesis of effective control techniques.

In the process of tackling real, very large size networks, we increasingly deal with large graph data analysis and the development of decision techniques with low algorithmic complexity, apt at providing answers from large datasets in reasonable time.

5 Highlights of the year

5.1 Awards

K. Avrachenkov has been invited by Australian Mathematical Science Institute (AMSI) and Australia and New Zealand Industrial and Applied Mathematics (ANZIAM) society to a AMSI-ANZIAM lecture tour to give a keynote at the ANZIAM Conference and a series of lectures at Australian Universities, February 5-23, 2023.

G. Neglia was recognized TPC (Technical Program Committee) distinguished member for IEEE International Conference on Computer Communications (INFOCOM) and top reviewer for the Thirty-seventh Annual Conference on Neural Information Processing Systems (NeurIPS)

5.2 Foreign Appointments

S. Perlaza was re-appointed "Visiting Research Collaborator" in the Department of Electrical and Computer Engineering at **Princeton University** for the academic year 2023-2024. He was also re-appointed "Associate Researcher" in the Laboratory of Algebraic Geometry and Applications to Information Theory (GAATI) at the **Université de la Polynésie Française** for the academic year 2023-2024.

6 New software, platforms, open data

6.1 New software

6.1.1 marmote

Name: MARkovian MOdeling: The Environment

Keyword: Markov model

Functional Description: marmote is a C++ library 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. The current release features the library marmoteMDP for modeling Markov Decision Processes and solving them.

This software was started within the ANR MARMOTE project: ANR-12-MONU-00019 under the name marmoteCore. Within the marmote project, the code conforms the latest C++ standards and the library is available on multiple platforms via a conda distribution.

Release Contributions: This version improves the code portability and its conformity to recent C++ standards. It contains the first release of marmoteMDP, a library for manipulating Markov Decision Processes, built on top of marmote's core. It also provides for the first time a complete documentation for installing and using the library, as well as numerous examples.

News of the Year: The first public release of marmote occured in May 2023 with version 0.2.0. The release featured, for the first time, installation procedures on multiple architectures, a documentation and numerous examples. It also featured the library marmoteMDP, built on top of marmote's core, devoted to the modeling and solution of Markov Decision Processes.

URL: https://marmote.gitlabpages.inria.fr/marmote/

Publications: hal-03770430v1, hal-03781620v1, hal-02395100v1, hal-04176076v1

Contact: Alain Jean-Marie

Participants: Alain Jean-Marie, Patrick Brown, Emmanuel Hyon

Partner: Université Paris Nanterre

7 New results

7.1 Game Theory and Applications

Participants: Eitan Altman, Emmanouil Athanasakos, Konstantin Avrachenkov, Mandar Datar, Alain Jean-Marie, Ashok Krishnan Komalan Sindhu, Samir M. Perlaza.

7.1.1 Games with Commitments and Noisy Observations

In [54], S. Perlaza, A. Jean-Marie and K. Sun studied classical zero-sum games under the following assumptions: (1) One of the players (the leader) commits to choose its actions by sampling a given probability measure (strategy); (2) The leader announces its action, which is observed by its opponent (the follower) through a binary channel; and (3) the follower chooses its strategy based on the knowledge of the leader's strategy and the noisy observation of the leader's action. Under these conditions, the equilibrium is shown to always exist. Interestingly, even subject to noise, observing the actions of the leader is shown to be either beneficial or immaterial for the follower. More specifically, the payoff at the equilibrium of this game is upper bounded by the payoff at the Stackelberg equilibrium (SE) in pure strategies; and lower bounded by the payoff at the Nash equilibrium, which is equivalent to the SE in mixed strategies. Finally, necessary and sufficient conditions for observing the payoff at equilibrium to be equal to its lower bound are presented. Sufficient conditions for the payoff at equilibrium to be equal to its upper bound are also presented. This work was presented as a conference in [39].

S. Perlaza, A. Jean-Marie, and E. Athanasakos are currently extending these results to more general classes of channels. In particular, Gaussian channels.

7.1.2 Pricing Strategies for Digital Renting Platforms

In [51], E. Altman, S. Perlaza, and A. Krishnan K.S. considered different pricing models for a platform-based rental system, such as Airbnb. A linear model is assumed for the demand response to price, and existence and uniqueness conditions for Nash equilibria are obtained. The Stackelberg equilibrium prices for the game are also obtained, and an iterative scheme is provided, which converges to the Nash equilibrium. Different cooperative pricing schemes are studied, and splitting of revenues based on the Shapley value is discussed. It is shown that a division of revenue based on the Shapley value gives a revenue to the platform proportional to its control of the market. The demand response function is modified to include user response to quality of service. It is shown that when the cost to provide quality of service is low, both renter and the platform will agree to maximize the quality of service. However, if this cost is high, they may not always be able to agree on what quality of service to provide. This work was presented as a conference in [31].

7.1.3 Constrained games in energy markets

In [40] R. Taisant (INOCS), M. Datar, H. Le Cadre (INOCS), and E. Altman consider a peer-to-peer electricity market modeled as a network game, where End Users (EUs) minimize their cost by computing their demand and generation while satisfying a set of local and coupling constraints. The nominal demand of EUs constitutes sensitive information, that EUs might want to keep private. The authors of [40] prove that the network game admits a unique Variational Equilibrium, which depends on the private information of all the EUs. A data aggregator is introduced, which aims to learn the EUs' private information. The EUs might have incentives to report biased and noisy readings to preserve their privacy, which creates shifts in their strategies. Relying on performative prediction, the authors define a decision-dependent game G^{stoch} to couple the network game with a data market. Two variants of the Repeated Stochastic Gradient Method (RSGM) are proposed to compute the Performatively Stable Equilibrium solution of G^{stoch}, that outperform RSGM with respect to efficiency gap minimization, privacy preservation, and convergence rates in numerical simulations. This work has been extended in [52] by H. Le Cadre and M. Guckert (INOCS), M. Datar and E. Altman (also LIA).

7.2 Stochastic Modeling

Participants: Sara Alouf, Eitan Altman, Konstantin Avrachenkov, Younes Ben Mazziane, Vinay Kumar Bindiganavile Ramadas, Giovanni Neglia.

7.2.1 Retrial Queues

In [13] K. Avrachenkov in collaboration with E. Morozov and R. Nekrasova (Karelian Institute of Applied Mathematical Research, Russia) establish stability criterion for a two-class retrial system with Poisson inputs, general class-dependent service times and class-dependent constant retrial rates. They also characterize an interesting phenomenon of partial stability when one orbit is tight but the other orbit goes to infinity in probability. All theoretical results are illustrated by numerical experiments.

7.2.2 Spatial Queues with Nearest Neighbor Shifts

Motivated primarily by electric vehicles (EV) queueing at charging stations, B.R. Vinay Kumar studies multiple server queues on an Euclidean space. He considers N servers that are distributed uniformly in $[0,1]^d$. Customers or EV users arrive at the servers according to Poisson processes of intensity λ. However, they probabilistically decide whether to join the queue they arrived at, or move to one of the nearest neighbors. The strategy followed by the customers affects the load on the servers in the long run. In [41, 55], B.R. Vinay Kumar is interested in characterizing the fraction of servers that bear a larger load as compared to when the users do not follow any strategy, i.e., they join the queue at which they arrive. These are called *overloaded servers*. He evaluates the expected fraction of overloaded servers in the system for the one dimensional case ($d = 1$) when the users follow probabilistic nearest neighbor shift strategies.

7.2.3 Performance Models for Similary Caching

Similarity caching allows requests for an item to be served by a similar item. Applications include recommendation systems, multimedia retrieval, and machine learning. Recently, many similarity caching policies have been proposed, like SIM-LRU (Similarity Least Recently Used) and its generalization RND-LRU (Random Least Recently Used), but the performance analysis of their hit ratio is still wanting. Y. Ben Mazziane, S. Alouf, and G. Neglia, together with D. S. Menasche (Federal Univ. of Rio de Janeiro, Brazil) are pursuing their effort to estimate the hit ratio of the similarity caching policy RND-LRU. They extend the popular time-to-live approximation in classic caching to similarity caching. They introduce the RND-TTL (Random Time-to-Live) approximation and the RND-TTL cache model and tune the model's parameters in such a way as to mimic the behavior of RND-LRU. The parameter tuning involves solving a fixed point system of equations for which they provide an algorithm for numerical resolution and sufficient conditions for its convergence.

7.2.4 Branching processes

In [17], R. Dhounchak and V. Kavitha (IIT Mumbai) in collaboration with E. Altman consider the inherent timeline structure of the appearance of content in online social networks (OSNs) while studying content propagation. They model the propagation of a post/content of interest by a multi-type branching process. The latter allows one to predict the emergence of global macro properties (e.g., the spread of a post in the network) from the laws and parameters that determine local interactions. The local interactions largely depend upon the timeline (an inverse stack capable of holding many posts and one dedicated to each user) structure and the number of friends (i.e., connections) of users, etc. They explore the use of multi-type branching processes to analyze the viral properties of the post, e.g., to derive the expected number of shares, the probability of virality of the content, etc. In OSNs, the new posts push down the existing contents in timelines, which can greatly influence content propagation; their analysis considers this influence. They find that one leads to draw incorrect conclusions when the timeline (TL) structure is ignored. One cannot capture some interesting paradigm shifts/phase transitions; for example, virality chances are not monotone with network activity parameter, as shown by analysis including TL influence.

7.3 Data Analysis and Learning

Participants: Eitan Altman, Sara Alouf, Konstantin Avrachenkov, Younes Ben Mazziane, Francisco Daunas, Francescomaria Faticanti, Alain Jean-Marie, Caelin Kaplan, Othmane Marfoq, Giovanni Neglia, Samir M. Perlaza, Angelo Rodio, Tareq Si Salem, Xinying Zou.

7.3.1 Empirical Risk Minimization (ERM)

Classical problems such as classification, pattern recognition, regression, and density estimation can be posed as special cases of the ERM problem. Unfortunately, ERM is prone to training data memorization, a phenomenon also known as overfitting. For this reason, ERM is often regularized in order to provide generalization guarantees. That is, to identify models using available training datasets that induce low empirical risk with respect to unseen datasets. At NEO special attention is payed to the study of the statistical properties of the solutions to ERM problems subject to particular regularizations. The main feature of this research effort is that contrary to the main stream in the community, this analysis is made for fixed training datasets, which provides new and insightful mathematical tools for the analysis of generalization capabilities of machine learning algorithms.

Regularization by f-Divergences In [50], S. Perlaza and F. Daunas, together with I. Esnaola (Univ. of Sheffield) and H.V. Poor (Princeton Univ.) present the solution to the empirical risk minimization with f-divergence regularization, under mild conditions on the function f. Under such conditions, the optimal measure is shown to be unique and to always exist. The solution is presented as a closed-form expression of the Radon-Nikodym derivative of the optimal probability measure with respect to the reference measure. Examples for particular choices of the function f are presented. For some choices, existing results are obtained as special cases of the main result. These include the unique solutions to the empirical risk minimization with relative entropy regularization, further studied within NEO in [53] and [27].

Regularization by Relative Entropy in [53], S. Perlaza and A. Jean-Marie together with G. Bisson (Univ. de la Polynésie française), I. Esnaola (Univ. of Sheffield), and S. Rini (National Chiao Tung Univ.) have continued the study of ERM problem with relative entropy regularization (ERM-RER) under the assumption that the reference measure is a σ-finite measure, and not necessarily a probability measure. Under this assumption, which leads to a generalization of the ERM-RER problem allowing a larger degree of flexibility for incorporating prior knowledge, numerous relevant properties are stated. Among these properties, the solution to this problem, if it exists, is shown to be a unique probability measure, often mutually absolutely continuous with the reference measure. Such a solution exhibits a probably-approximately-correct guarantee for the ERM problem independently of whether the latter possesses a solution. For

a fixed dataset, the empirical risk is shown to be a sub-Gaussian random variable when the models are sampled from the solution to the ERM-RER problem. The generalization capabilities of the solution to the ERM- RER problem (the Gibbs algorithm) are studied via the sensitivity of the expected empirical risk to deviations from such a solution towards alternative probability measures. Finally, an interesting connection between sensitivity, generalization error, and lautum information is established.

Impact of Relative Entropy Asymmetry In [49], S. Perlaza and F. Daunas, together with I. Esnaola (Univ. of Sheffield) and H.V. Poor (Princeton Univ.) study the effect of the relative entropy asymmetry in the empirical risk minimization with relative entropy regularization (ERM-RER) problem. A novel regularization is introduced, coined Type-II regularization, that allows for solutions to the ERM-RER problem with a support that extends outside the support of the reference measure. The solution to the new ERM-RER Type-II problem is analytically characterized in terms of the Radon-Nikodym derivative of the reference measure with respect to the solution. The analysis of the solution unveils the following properties of relative entropy when it acts as a regularizer in the ERM-RER problem: i) relative entropy forces the support of the Type-II solution to collapse into the support of the reference measure, which introduces a strong inductive bias that dominates the evidence provided by the training data; ii) Type-II regularization is equivalent to classical relative entropy regularization with an appropriate transformation of the empirical risk function. Closed-form expressions of the expected empirical risk as a function of the regularization parameters are provided. This work was presented as a conference in [27].

7.3.2 Characterization of the Generalization Error

The expected generalization error (GE) is a central performance metric for the analysis of generalization capabilities of machine learning algorithms. In a nutshell, the GE characterizes the ability of a learning algorithm to correctly find patterns in datasets that are not available during the training stage. Specifically, it is defined for a fixed training dataset and a specific model instance, as the difference between the population risk induced by the model and the empirical risk with respect to the training dataset. At NEO our research focuses on the search of closed-form expressions of the GE for specific algorithms under the assumption that datasets follow a specific probability distribution that is consistent with the training dataset.

The Worst-Case Data-Generating Probability Distributions In [56], S. Perlaza, E. Altman, and X. Zou, together with I. Esnaola (Univ. of Sheffield) have introduced the worst-case probability measure over the data as a tool for characterizing the generalization capabilities of machine learning algorithms. More specifically, the worst-case probability measure is a Gibbs probability measure and the unique solution to the maximization of the expected loss under a relative entropy constraint with respect to a reference probability measure. Fundamental generalization metrics, such as the sensitivity of the expected loss, the sensitivity of the empirical risk, and the generalization gap are shown to have closed-form expressions involving the worst-case data-generating probability measure. Existing results for the Gibbs algorithm, such as characterizing the generalization gap as a sum of mutual information and lautum information, up to a constant factor, are recovered. A novel parallel is established between the worst-case data-generating probability measure and the Gibbs algorithm. Specifically, the Gibbs probability measure is identified as a fundamental commonality of the model space and the data space for machine learning algorithms.

Data Aggreagation and Validation of Gibbs Algorithms In [35], S. Perlaza and A. Jean-Marie together with G. Bisson (Univ. de la Polynésie française), I. Esnaola (Univ. of Sheffield), and H. V. Poor (Princeton Univ.) analytically characterized the dependence on training data of the Gibbs algorithm (GA). By adopting the expected empirical risk as the performance metric, the sensitivity of the GA is obtained in closed form. In this case, sensitivity is the performance difference with respect to an arbitrary alternative algorithm. This description enables the development of explicit expressions involving the training errors and test errors of GAs trained with different datasets. Using these tools, dataset aggregation is studied and different figures of merit to evaluate the generalization capabilities of GAs are introduced. For particular sizes of such datasets and parameters of the GAs, a connection between Jeffrey's divergence, training and test errors is established.

7.3.3 Online Learning

Online learning algorithms have been successfully used to design caching policies with regret guarantees. Existing algorithms assume that the cache knows the exact request sequence, but this may not be feasible in high load and/or memory-constrained scenarios, where the cache may have access only to sampled requests or to approximate requests' counters. In [26], Y. Ben Mazziane, F. Faticanti, G. Neglia, and S. Alouf propose the Noisy-Follow-the-Perturbed-Leader (NFPL) algorithm, a variant of the classic Follow-the-Perturbed-Leader (FPL) when request estimates are noisy, and they show that the proposed solution has sublinear regret under specific conditions on the requests estimator. The experimental evaluation compares the proposed solution against classic caching policies and validates the proposed approach under both synthetic and real request traces.

In [28], F. Faticanti and G. Neglia study a batched version of optimistic online algorithms for caching. The analysis consists of updating the cache state less frequently with respect to traditional caching algorithms that update the cache state after receiving every single new request. This new approach proposes to cumulate a batch of requests before updating the cache given the high computational complexity of online algorithms based on Follow-The-Regularized-Leader (FTRL). No-regret results are showed in this new setting and experimental results show that the batched versions of the online algorithms outperform traditional caching policies on both synthetic and real traces.

The abstract of the paper "Enabling Long-term Fairness in Dynamic Resource Allocation " has been published in the Proceedings of the 2023 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems [38]. The corresponding research activity is described in NEO activity report for 2022.

7.3.4 Reinforcement Learning

In [34] K. Avrachenkov in collaboration with T. Pagare and V. Borkar (IIT Bombay, India) extend the provably convergent Full Gradient DQN (Deep Q-Network) algorithm for discounted reward Markov decision processes from Avrachenkov et al. (2021) to average reward problems. They experimentally compare widely used RVI (Relative Value Iteration) Q-learning with recently proposed Differential Q-learning in the neural function approximation setting with Full Gradient DQN and DQN. They also extend this to learn Whittle indices for Markovian restless multi-armed bandits and observe a better convergence rate of the proposed Full Gradient variant across different tasks.

K. Avrachenkov in collaboration with V. Borkar and J. Nair (IIT Bombay, India) have organized and edited a special volume of *Dynamic Games and Applications* journal on the topic "Multi-Agent Dynamic Decision Making and Learning" [44]. This field interests multiple communities such as dynamic games, control theory and machine learning, especially, reinforcement learning. 14 papers have been accepted and published in this special volume.

7.3.5 Federated Learning

In [36], A. Rodio, F. Faticanti, O. Marfoq, and G. Neglia in collaboration with E. Leonardi (Polytechnic Univ. of Turin, Italy) provide the first convergence result for Federated Learning algorithms under heterogeneous and correlated client availability. Their analysis shows the negative impact of correlation on the algorithms' convergence rate and highlights a trade-off between optimization error (related to convergence speed) and bias error (indicative of model quality). Their proposed Correlation-Aware FL (CA-Fed) algorithm effectively balances convergence speed and model quality by adjusting client aggregation weights and selectively excluding highly correlated, low-availability clients. Throughout simulations, CA-Fed achieves higher time-average accuracy and reduced standard deviation compared to existing methods on both synthetic and real datasets.

Subsequently, in [20], A. Rodio, F. Faticanti, O. Marfoq, and G. Neglia in collaboration with E. Leonardi (Polytechnic Univ. of Turin, Italy) further enhance their CA-Fed algorithm by introducing a new hyper-parameter to optimize the trade-off between convergence speed and model quality. A sensitivity analysis confirms this parameter's impact on convergence. In contrast to [36], which relied on prior knowledge of clients' availability and correlation, they propose a Bayesian estimator with a beta prior, requiring only a limited amount of observations to outperform existing methods. Additionally, they address a gap in [36]

by evaluating CA-Fed in spatially correlated scenarios, where the availability patterns are correlated among clients.

In [37], A. Rodio and G. Neglia in collaboration with F. Busacca and S. Palazzo (Univ. of Catania, Italy), S. Mangione and I. Tinnirello (Univ. of Palermo, Italy), and F. Restuccia (Northeastern Univ., USA) address training Federated Learning algorithms in wireless networks with packet losses. Contrary to conventional approaches that focus on mitigating packet losses through retransmission or error correction, they show that FL algorithms can effectively learn in asymmetric lossy channels while maintaining the same computational and communication efficiency. Their proposed algorithm, UPGA-PL (Unbiased Pseudo-Gradient Aggregation with Packet Losses), employs a pseudo-gradient step rather than model averaging, and adjusts the aggregation weights for heterogeneous packet losses. In experimental evaluation, UPGA-PL outperforms existing methods in lossy environments and matches Federated Learning algorithms in lossless scenarios within a limited number of communication rounds.

Most work on federated learning assumes that clients operate on static datasets collected before training starts. This approach may be inefficient because 1) it ignores new samples clients collect during training, and 2) it may require a potentially long preparatory phase for clients to collect enough data. Moreover, learning on static datasets may be simply impossible in scenarios with small aggregate storage across devices. It is, therefore, necessary to design federated algorithms able to learn from data streams. In [33] O. Marfoq and G. Neglia, together with L. Kameni and R. Vidal from (Accenture Labs, France) formulate and study the problem of federated learning for data streams. They propose a general FL algorithm to learn from data streams through an opportune weighted empirical risk minimization. Their theoretical analysis provides insights to configure such an algorithm, and they evaluate its performance on a wide range of machine learning tasks.

7.3.6 Privacy and Fairness

In [42], C. Kaplan in collaboration with A.S. de Oliveira (SAP Labs France), Khawla Mallat (SAP Labs France), and Tanmay Chakraborty (SAP Labs France, Eurecom) investigate the impact of differential privacy on fairness notions for tabular data. They empirically analyze how different fairness notions, belonging to distinct classes of statistical fairness criteria, are impacted when one selects a model architecture suitable for DP-SGD (differentially private stochastic gradient descent), optimized for utility. Using standard datasets from ML fairness literature, they show that by selecting the optimal model architecture for DP-SGD, the differences across groups concerning the relevant fairness metrics more often decrease or are negatively impacted, compared to the non-private baseline, for which the optimal model architecture has also been selected to maximize utility. These findings challenge the understanding that differential privacy will necessarily exacerbate unfairness in deep learning models trained on biased datasets.

7.3.7 Data Injection Attacks in Power Systems

Since 2020, S. Perlaza in collaboration with X. Ye, I. Esnaola, and R. Harrison (Univ. of Sheffield) have studied sparse stealth attack constructions that minimize the mutual information between the state variables and the observations. In [24], the attack construction is formulated as the design of a multivariate Gaussian distribution that aims to minimize the mutual information while limiting the Kullback-Leibler divergence between the distribution of the observations under attack and the distribution of the observations without attack. The sparsity constraint is incorporated as a support constraint of the attack distribution. Two heuristic greedy algorithms for the attack construction are proposed. The first algorithm assumes that the attack vector consists of independent entries, and therefore, requires no communication between different attacked locations. The second algorithm considers correlation between the attack vector entries and achieves a better disruption to stealth tradeoff at the cost of requiring communication between different locations. Numerical evaluations show that it is feasible to construct stealth attacks that generate significant disruption with a low number of compromised sensors.

7.3.8 Graph Clustering

In [25] K. Avrachenkov and B.R. Vinay Kumar in collaboration with K. Alaluusua and L. Leskelä (Aalto University, Finland) consider the community recovery problem on a multilayer variant of the hypergraph

stochastic block model (HSBM). Each layer is associated with an independent realization of a d-uniform HSBM on N vertices. Given the similarity matrix containing the aggregated number of hyperedges incident to each pair of vertices, the goal is to obtain a partition of the N vertices into disjoint communities. In this work, they investigate a semidefinite programming (SDP) approach and obtain information-theoretic conditions on the model parameters that guarantee exact recovery both in the assortative and the disassortative cases.

In [30, 58] K. Avrachenkov and L. Hauseux in collaboration with J. Zerubia (Ayana team) propose an original density estimator built from a cloud of points $X \in R^d$. To do this, they consider geometric graphs $G(X,r)$ on the cloud. These graphs depend on a radius r. By varying the radius, they see the emergence of large components around certain critical radii, which is the phenomenon of continuum percolation. Percolation allows us to have both a local view of the data (through local constraints on the radius r) and a global one (the emergence of macro-structures). With this tool, they address the problem of galaxy filament extraction. The density estimator gives us a relevant graph on galaxies. With an algorithm sharing the ideas of the Fréchet mean, they extract a subgraph from this graph, the galaxy filaments.

7.4 Applications in Telecommunications

Participants: Eitan Altman, Mandar Datar, Francescomaria Faticanti, Philippe Nain, Vijith Kumar Kizhakke Purakkal, Samir M. Perlaza.

7.4.1 Optimal Rate Memory Tradeoff in Multi-Access Coded Caching

In [32] Vijith Kumar K. P. together with B. Kumar Rai and T. Jacob (IIT Guwahati, India) consider the (N,K,L) multi-access caching network where K users and K caches are connected to a server with N files, each of size F bits, through a shared error-free broadcast channel. Each user has access to L nearby caches, each of size MF bits, in a cyclic wrap-around manner. Even after several previous attempts, the exact characterization of the optimal rate memory tradeoff is still an open problem except in the case where $L = K-1$ and $L = 1$ with large cache $M \in [N/L \cdot (K-1)/K, N/L]$. This paper determines the optimal rate memory tradeoff for the cache network with $L = K-2$ and $M \in [N/(K-2) \cdot (K-1)/K, N/(K-2)]$. This is done by proposing a new caching scheme that operates at the memory rate pair $(N/(K-2) \cdot (K-1)/K, 1/K)$ and deriving a set of lower bounds to demonstrate the optimality of the scheme.

7.4.2 Deployment of Microservice-based Applications

In [18] F. Faticanti in collaboration with M. Savi (Univ. Milano-Bicocca, Italy), F. De Pellegrini (Univ. Avignon) and D. Siracusa (Fondazione Bruno Kessler, Italy) propose a new approach for the deployment of microservice-based applications in a Federated Fog Computing scenario under locality constraints for a subset of microservices of the application. The approach is based on a Breadth-First-Search (BFS) visit of the search space for the deployment of applications where the main objective is to minimize the deployment cost towards external fog domains with respect to the main provider. Experiments show that the proposed approach outperforms traditional deployment methods based on Depth-First-Search visits of the search space.

7.4.3 Learning Optimal Edge Processing with Offloading and Energy Harvesting

Modern portable devices can execute increasingly sophisticated AI models on sensed data. The complexity of such processing tasks is data-dependent and has relevant energy cost. In [29], A. Fox, F. De Pellegrini (Univ. Avignon) and E. Altman develop an Age of Information Markovian model for a system where multiple battery-operated devices perform data processing and energy harvesting in parallel. Part of their computational burden is offloaded to an edge server which polls devices at a given rate. The structural properties of the optimal policy for a single device-server system are derived. They permit to derive a new model-free reinforcement learning method specialized for monotone policies, namely Ordered Q-Learning, providing a fast procedure to learn the optimal policy. The method is oblivious to

the devices' battery capacities, the cost and the value of data batch processing and to the dynamics of the energy harvesting process. Finally, the polling strategy of the server is optimized by combining such policy improvement techniques with stochastic approximation methods. Extensive numerical results provide insight into the system properties and demonstrate that the proposed learning algorithms outperform existing baselines.

7.4.4 Pricing and resource allocation

We have been extending our research on games in networks where users compete over resources, to the case in which competition arises also between the users and the supplier of services and further to the case when the suppliers compete between each other through the prices of their services. This bundelling of services called Network slicing (NS) is a key technology in 5G that enables the customization and efficient sharing of network resources to support the diverse requirements of next-generation services.

In [15], M. Datar, E. Altman and H. Le Cadre (INOCS) consider a marketplace in the context of 5G network slicing, where Application Service Providers (ASP), i.e., slice tenants, providing heterogeneous services, are in competition for the access to the virtualized network resource owned by a Network Slice Provider (NSP), who relies on network slicing. They model the interactions between the end users (followers) and the ASPs (leaders) as a Stackelberg game. They prove that the competition between the ASPs results in a multi-resource Tullock rent-seeking game. To determine resource pricing and allocation, They devise two innovative market mechanisms.

In [16], S. Dhamal, W. Ben-Ameur, and T. Chahed (Télécom SudParis), E. Altman, A. Sunny (ITT Palakkad), and S. Poojary (Qualcomm India) study a distributed computing setting wherein a central entity seeks power from computational providers by offering a certain reward in return. The computational providers are classified into long-term stakeholders that invest a constant amount of power over time and players that can strategize on their computational investment. In this paper, they model and analyze a stochastic game where players arrive and depart over time. They prove that, in Markov perfect equilibrium, only players with cost parameters in a relatively low range which collectively satisfy a certain constraint in a given state, invest. They infer that players need not have knowledge about the system state and other players' parameters, if the total power that is being received by the central entity is communicated to the players as part of the system's protocol. If players are homogeneous and the system consists of a reasonably large number of players, They observe that the total power received by the central entity is proportional to the offered reward and does not vary significantly despite the players' arrivals and departures, thus resulting in a robust and reliable system. They study by simulations and mean field approximation, how the players' utilities are influenced by the system parameters.

In [48], M. Datar (currently at Orange Innovation), N. Modina (CNAM), R. El Azouzi (Univ. Avignon), and E. Altman propose an allocation scheme for network slicing based on the Fisher-market model and the Trading-post mechanism. The scheme aims to achieve efficient resource utilization while ensuring multi-level fairness, dynamic load conditions, and the protection of service level agreements (SLAs) for slice tenants. In the proposed scheme, each service provider (SP) is allocated a budget representing its infrastructure share or purchasing power in the market.

7.5 Applications to Environmental Issues

Participants: Alain Jean-Marie.

The nature of fishing activities is such that marine habitats can be deteriorated when employing destructive fishing gear. This makes even more complex the determination of sustainable fishing policies and has led some authors to propose dynamic models which take into account this habitat degradation. In [19], A. Jean-Marie and M. Tidball (INRAE) analyze in detail one of these models, an extension of the single-species Gordon–Schaefer model to two state interrelated variables: stock of fish and habitat. The model assumes that stock and carrying capacity are positively linked, and that the fishing activity has a direct and negative impact on the carrying capacity. The authors extend and characterize Clark's most rapid approach optimal solution to this case.

8 Bilateral contracts and grants with industry

8.1 Bilateral contracts with industry

NEO has contracts with Accenture (see §8.1.1 and §8.1.2), NSP SmartProfile (see §8.1.3), Orange Labs (see §8.1.4), QITI (see §8.1.6), and SAP (see §8.1.5).

8.1.1 Accenture contract on the topic "Distributed Machine Learning for IoT applications" (Dec 2019 – November 2023)

Participants: Othmane Marfoq, Giovanni Neglia.

- Contractor: Accenture Labs
- Collaborators: Laetitia Kameni, Richard Vidal

IoT applications will become one of the main sources to train data-greedy machine learning models. Until now, IoT applications were mostly about collecting data from the physical world and sending them to the Cloud. Google's federated learning already enables mobile phones, or other devices with limited computing capabilities, to collaboratively learn a machine learning model while keeping all training data locally, decoupling the ability to do machine learning from the need to store the data in the cloud. While Google envisions only users' devices, it is possible that part of the computation is executed at other intermediate elements in the network. This new paradigm is sometimes referred to as Edge Computing or Fog Computing. Model training as well as serving (provide machine learning predictions) are going to be distributed between IoT devices, cloud services, and other intermediate computing elements like servers close to base stations as envisaged by the Multi-Access Edge Computing framework. The goal of this project is to propose distributed learning schemes for the IoT scenario, taking into account in particular its communication constraints. O. Marfoq is funded by this project. A first 12-month pre-PhD contract has been followed by a PhD grant.

8.1.2 Accenture "Plan de Relance" (PLR) contract on the topic "Energy-Aware Federated Learning" (Oct 2022 – September 2024)

Participants: Giovanni Neglia, Charlotte Rodriguez.

- Contractor: Accenture Labs
- Collaborators: Laura Degioanni, Laetitia Kameni, Richard Vidal

Deep neural networks have enabled impressive accuracy improvements across many machine learning tasks. Often the highest scores are obtained by the most computationally-hungry models. As a result, training a state-of-the-art model now requires substantial computational resources which demand considerable energy, along with the associated economic and environmental costs. Research and development of new models multiply these costs by thousands of times due to the need to try different model architectures and different hyper-parameters. In this project, we investigate a more algorithmic/system-level approach to reduce energy consumption for distributed ML training over the Internet. The postdoc of C. Rodriguez is funded by this project.

8.1.3 Cifre contract with NSP-SmartProfile "Automated and responsible recommendation systems for digital marketing" (August 2022 – July 2025)

Participants: Konstantin Avrachenkov, Ibtihal El Mimouni.

- Contractor: NSP-SmartProfile
- Collaborators: Hervé Baile, Julien Musso

SmartProfile is a marketing platform that allows to collect, to enhance and to analyze marketing data. Digital marketing campaigns continue to expand across all digital channels and media. The 'mass marketing' strategies implemented by most companies show limits in terms of performance and acceptance by clients, as well as in terms of their impact on the environment. In opposite to these practices, we believe that current technologies, particularly in terms of Artificial Intelligence (AI), should make marketing interactions more efficient and virtuous. Through this research project, we want to create an alternative solution to mass marketing by switching to an intelligent, automated and eco-responsible system, which will support the heterogeneity of data and the diversity of sectors, and whose purpose is to recommend the best content by determining the most relevant target and taking into account the communication constraints. This contract complements the Cifre thesis of Ibtihal El Mimouni.

8.1.4 Cifre contract with Orange Labs "Analytical modeling of large-scale wireless networks integrating RIS" (March 2023 – March 2026)

Participants: Eitan Altman, Julian Alfonso Santos Bustos.

- Contractor: Orange Labs
- Collaborators: Jean-Marc Kelif

A Reconfigurable Intelligent Surface (RIS) is a programmable surface structure that allows one to control the reflection of electromagnetic (EM) waves by changing the electric and magnetic properties of the surface. In the absence of RIS, short wavelentghs signals as in 5G, are subject to a huge attenuation when there is no direct line of sight channel. Within our collaboration we shall evaluate and optimize the position of RIS.

This contract complements the Cifre thesis of J. Santos.

8.1.5 Cifre contract with SAP "Privacy and fairness for ML" (December 2021 – December 2024)

Participants: Caelin Kaplan, Giovanni Neglia.

- Contractor: SAP Labs France
- Collaborators: Anderson Santana de Oliveira

There are increasing concerns among scholars and the public about bias, discrimination, and fairness in AI and machine learning. Decision support systems may present biases, leading to unfair treatment of some categories of individuals, for instance, systematically assigning high risk of recidivism in a criminal offense analysis system. Essentially, the analysis of whether an algorithm's output is fair (e.g. does not disadvantages a group with respect to others) depends on substantial contextual information that often requires human intervention. There are though several metrics for fairness that have been developed, which may rely on collecting additional sensitive attributes (e.g., ethnicity) before imposing strong privacy guarantees to be used in any situation. It is known that differential privacy has the effect of hiding outliers from the data analysis, perhaps compounding existing bias in certain situations. This project encompasses the search for a mitigating strategy. This contract complements the Cifre thesis of C. Kaplan.

8.1.6 QITI 3IA Start-It-Up contract on the topic "Reinforcement Learning for Conversational Recommender Systems (RLCRS)" (April 2023 – March 2024)

Participants: Konstantin Avrachenkov, Hariprasad Manjunath Hegde.

- Contractor: QITI
- Collaborators: Christophe Bremard, Guillaume Dion

Qiti is a start-up created in Nice in 2021 which among other things develops a Conversational Recommender Systems (CRS) for insurance holders and insurers. The CRS should reduce the load on the workers of the insurers and to simplify the process of insurance establishment and modification. The goal of the present cooperation is to test and to improve various Reinforcement Learning schemes for CRS. The post-doc of H. Manjunath is funded by this contract.

9 Partnerships and cooperations

9.1 International initiatives

9.1.1 Inria associate team not involved in an IIL or an international program

Cefipra-EA LION

Participants: Eitan Altman, Konstantin Avrachenkov, Samir M. Perlaza.

Title: Learning In Operations and Networks

Duration: 2022 ->

Coordinator: Kavitha Veeraruna (vkavitha@iitb.ac.in)

Partners:

- Indian Institute of Technology Bombay Bombay (India) the collaboration work is cofinanced by CEFIPRA (Indo-French Centre for the Promotion of Advanced Research)

Inria contact: Eitan Altman

Summary: Artificial Intelligence has affected all walks of life. We propose to study its application in various domains like 1) Learning and Control in Healthcare: Our aim is to use novel AI methodologies, to predict the results of possible actions of involved decision-makers, using the available data. 2) Dual Learning Algorithms in wireless networks: We aim to develop learning algorithms for beam alignment in 5G Wireless networks to maintain high rates. We propose to use AoI as a metric. 3) Distributed and reinforcement learning: We will develop and analyze Deep Q Network (DQN) based learning algorithms and analyze their performance.

9.1.2 STIC/MATH/CLIMAT AmSud projects

MATH-AmSud GSA

Participants: Konstantin Avrachenkov, Alain Jean-Marie, Hariprasad Manjunath Hegde.

Title: Graph Spectra and its Applications

Program: MATH-AmSud

Duration: January 1, 2023 – December 31, 2024

Local supervisor: Konstantin Avrachenkov

Partners:

- V. Trevisan (South America side coordinator), C. Hoppen, L.E. Allem, A.M. França (Brazil)
- A. Pastine (Argentina)
- L. Medina (Chile)
- A.M. França (Brazil)

Inria contact: Konstantin Avrachenkov

Summary: The present project proposes the establishment of a network of collaboration among Argentina, Chile, Brazil, and France, using the strength of 5 mathematics groups of 5 different institutions. The research topic of the proposal is Algebraic Graph Theory, an important and modern area of discrete mathematics. The proposal is structured in such a way that the training of highly qualified human resources and research activities are intertwined, this will ensure the generation of new knowledge in a relevant scientific area and leave permanent ties of collaboration between the different research groups beyond the completion of the project.

9.2 International research visitors

9.2.1 Visits of international scientists

Kavitha Veeraruna

Status: Professor

Institution of origin: IIT Bombay

Country: India

Dates: September 4 to December 9, 2023

Context of the visit: collaboration on stochastic processes, queueing theory and game theory

Mobility program/type of mobility: sabbatical

Gholamali Aminian

Status: Research Associate

Institution of origin: Alain Turing Institute

Country: United Kingdom

Dates: October 25 to October 26, 2023

Context of the visit: Seminar: Generalization error via measure-valued calculus

Mobility program/type of mobility: lecture

Nirupam Gupta

Status: Postdoc

Institution of origin: EPFL

Country: Switzerland

Dates: December 13 to December 14, 2023

Context of the visit: Seminar: Machine Learning in Untrusted Environment

Mobility program/type of mobility: lecture

Nelly Litvak

Status: Full Professor

Institution of origin: Eindhoven University of Technology

Country: The Netherlands

Dates: December 4 to December 6, 2023

Context of the visit: Morgenstern Colloquium: Projection methods for community detection in complex networks

Mobility program/type of mobility: lecture

Nahuel Soprano Loto

Status: Postdoc

Institution of origin: LAAS-CNRS

Country: France

Dates: December 18 to December 19, 2023

Context of the visit: Seminar: Stochastic matching and online matching algorithms

Mobility program/type of mobility: lecture

Uri Yechiali

Status: Professor Emeritus

Institution of origin: Tel Aviv University

Country: Israel

Dates: April 13 to April 14, 2023

Context of the visit: Seminar: Queues with service resetting

Mobility program/type of mobility: lecture

Other international visits to the team: research stay

Kalle Alaluusua

Status: PhD student

Institution of origin: Aalto University

Country: Finland

Dates: October 9 to November 7, 2023

Context of the visit: collaboration on community detection, hypergraphs, Bayesian statistics and related problems

Mobility program/type of mobility: research stay

Luiz Allem

Status: Full Professor

Institution of origin: Universidade Federal do Rio Grande do Sul (Porto Alegre)

Country: Brazil

Dates: June 2 to June 9, 2023

Context of the visit: MATH AmSUD exchange program (project GSA)

Mobility program/type of mobility: research stay

Vivek Borkar

Status: Professor

Institution of origin: IIT Bombay

Country: India

Dates: July 3 to July 15, 2023

Context of the visit: collaboration on inference and optimization on dynamic graphs

Mobility program/type of mobility: research stay

Damiano Carra

Status: Associate Professor

Institution of origin: University of Verona

Country: Italy

Dates: September 11 to September 15, 2023

Context of the visit: collaboration on no-regret algorithms for online caching

Mobility program/type of mobility: research stay

Pavel Chebotarev

Status: Researcher

Institution of origin: Technion

Country: Israel

Dates: September 25 to September 29, 2023

Context of the visit: collaboration on centrality measures and their comparative and clustering analysis

Mobility program/type of mobility: research stay

Maximilen Dreveton

Status: Post-Doc

Institution of origin: EPFL

Country: Switzerland

Dates: April 3 to April 7, 2023

Context of the visit: collaboration on community recovery in non-binary and temporal stochastic block models

Mobility program/type of mobility: research stay

Lucas Gamertsfelder

Status: PhD student

Institution of origin: Macquarie University

Country: Australia

Dates: November 1, 2022 to February 28, 2023

Context of the visit: collaboration on partially observable Markov decision processes

Mobility program/type of mobility: research stay

Iñaki Esnaola

Status: Senior Lecturer

Institution of origin: The University of Sheffield

Country: United Kingdom

Dates: September 18 to September 22, 2023

Context of the visit: collaboration on generalizing learning algorithms using tools from the theory of information, measurement theory, and game theory

Mobility program/type of mobility: research stay

Carlos Hoppen

Status: Associate Professor

Institution of origin: Universidade Federal do Rio Grande do Sul (Porto Alegre)

Country: Brazil

Dates: June 2 to June 9, 2023

Context of the visit: MATH AmSUD exchange program (project GSA)

Mobility program/type of mobility: research stay

Lasse Leskelä

Status: Professor

Institution of origin: Aalto University

Country: Finland

Dates: April 3 to April 7, and October 19 to October 25, 2023

Context of the visit: collaboration on Community recovery in non-binary and temporal stochastic block models

Mobility program/type of mobility: research stay

Floske Spieksma

Status: Professor

Institution of origin: Leiden University

Country: The Netherland

Dates: October 22 to October 26, 2023

Context of the visit: collaboration on deviation matrix for Markov processes and its application to retrial queues

Mobility program/type of mobility: research stay

Jacopo Talpini

Status: PhD student

Institution of origin: University of Milano - Bicocca

Country: Italy

Dates: November 1, 2023 to April 30, 2024

Context of the visit: collaboration on federated learning

Mobility program/type of mobility: research stay

Konstantinos Varsos

Status: ERCIM Fellow

Institution of origin: CWI

Country: The Netherlands

Dates: May 15 to May 26, 2023

Context of the visit: collaboration on game theory and incomplete information

Mobility program/type of mobility: research stay

Lotte Weedage

Status: PhD student

Institution of origin: University of Twente

Country: The Netherlands

Dates: April 17 to July 10, 2023

Context of the visit: collaboration on clustering problems in wireless networks

Mobility program/type of mobility: research stay

Other international visits to the team: internship

Mihir Deshpande

Status: intern (bachelor)

Institution of origin: IIT Kanpur

Country: India

Dates: May 10 to July 20, 2023

Context of the visit: working on distributed control with delayed sharing patterns

Mobility program/type of mobility: internship

Francesco Diana

Status: intern (master/eng)

Institution of origin: University of Bari/ Univeristé Côte d'Azur

Country: Italy/ France

Dates: March 1 to August 31, 2023

Context of the visit: working on cooperative inference

Mobility program/type of mobility: internship

Cosimo Giani

Status: intern (master/eng)

Institution of origin: University of Florence

Country: Italy

Dates: October 1, 2022 to March 30, 2023

Context of the visit: working on federated learning for inference delivery networks

Mobility program/type of mobility: internship

Tejas Pagare

Status: intern (bachelor)

Institution of origin: IIT Mumbai

Country: India

Dates: May 9 to July 14, 2023

Context of the visit: working on reinforcement Learning and its applications in graph optimization problems

Mobility program/type of mobility: internship

Mohamed Salah Jebali

Status: intern (master/eng)

Institution of origin: University of Florence

Country: Italy

Dates: March 1 to August 31, 2023

Context of the visit: working on federated learning

Mobility program/type of mobility: internship

Xufeng Zhang

Status: intern (master/eng)

Institution of origin: University of Rome La Sapienza

Country: Italy

Dates: March 1 to August 31, 2023

Context of the visit: working on online learning

Mobility program/type of mobility: internship

9.2.2 Visits to international teams

Research stays abroad

Konstantin Avrachenkov

Visited institution: University of South Australia, RMIT, Australian Bureau of Statistics, University of Newcastle, The University of Queensland

Country: Australia

Dates: February 3 to February 25, 2023

Context of the visit: AMSI lecture tour

Mobility program/type of mobility: research stay, lectures

Visited institution: University Johns Hopkins, Washington DC

Country: United States

Dates: June 17 to June 18, 2023

Context of the visit: collaboration with Prof. Ali Eshragh

Mobility program/type of mobility: research stay

Visited institution: Aalto University

Country: Finland

Dates: August 6 to August 12, 2023

Context of the visit: collaboration with Prof. L. Leskela

Mobility program/type of mobility: research stay

Visited institution: Umea University

Country: Sweden

Dates: October 1 to October 10 , 2023

Context of the visit: collaboration with Prof. L. Freidovich

Mobility program/type of mobility: research stay

Vinay Kumar Bindiganavile Ramadas

Visited institution: Aalto University

Country: Finland

Dates: August 6 to September 3, 2023

Context of the visit: Aalto Science Institute Visitor Fellowship

Mobility program/type of mobility: research stay

Francescomaria Faticanti

Visited institution: University of Milano-Bicocca

Country: Italy

Dates: January 19 to January 21, 2023

Context of the visit: collaboration with Professor Marco Savi

Mobility program/type of mobility: research stay

Visited institution: Fondazione Bruno Kessler, Trento

Country: Italy

Dates: Avril 27 to May 1, 2023

Context of the visit: collaboration with Dr. Domenico Siracusa

Mobility program/type of mobility: research stay, lecture

Visited institution: Vienna university of Technology

Country: Austria

Dates: April 3 to April 11, and August 13 to August 21, 2023

Context of the visit: collaboration with Prof. Ivona Brandi

Mobility program/type of mobility: research stay

Othmane Marfoq

Visited institution: University of Texas at Austin

Country: United States

Dates: January 6 to June 27, 2023

Context of the visit: collaboration on personalized federated learning

Mobility program/type of mobility: research stay

Hariprasad Manjunath Hegde

Visited institution: University Federal Rio Grande de Sul, Porto Alegre

Country: Brazil

Dates: November 18 to December 2, 2023

Context of the visit: MATH AmSUD exchange program (project GSA)

Mobility program/type of mobility: research stay

Giovanni Neglia

Visited institution: University of Florence

Country: Italy

Dates: March 23 to March 24, 2023

Context of the visit: collaboration with Fabrizio Argenti

Mobility program/type of mobility: research stay

Samir Perlaza

Visited institution: CWI

Country: The Netherlands

Dates: September 13 - 19, 2023

Context of the visit: Kick-off international research collaboration Inria-CWI

Mobility program/type of mobility: Research Stay

Visited institution: Universidad del Cauca

Country: Colombia

Dates: March 16 - 27, 2023

Context of the visit: Collaboration with Prof. Victor Quintero

Mobility program/type of mobility: Research Stay

9.3 European initiatives

9.3.1 Horizon Europe

dAIEDGE

Participants: Sara Alouf, Alain Jean-Marie, Giovanni Neglia.

Project Title: A network of excellence for distributed, trustworthy, efficient and scalable AI at the Edge

Program: HORIZON-CL4-2022-HUMAN-02-02

Coordinator: DFKI

Duration: September 1, 2023 – August 31, 2026

Local supervisor: Giovanni Neglia

Other Partners: see website

Summary: The dAIEDGE Network of Excellence seeks to strengthen and support the development of the dynamic European edge AI ecosystem under the umbrella of the European AI Lighthouse and to sustain the advanced research and innovation of distributed AI at the edge as essential digital, enabling, and emerging technology in an extensive range of industrial sectors.

9.3.2 H2020 projects

TESTBED2

Participant: Samir M. Perlaza.

Project Title: Testing and Evaluating Sophisticated information and communication Technologies for enaBling scalablE smart griD Deployment

Program: H2020-MSCA-RISE

Coordinator: University of Durham, UK

Duration: February 2020 – June 2025

Local supervisor: Samir M. Perlaza

Others Partners: The University of Durham (UDUR); University of Tuebingen (EKUT); Heriot-Watt University (HWU); University of Klagenfurt (AAU); University of Northumbria at Newcastle (UNN); DotX Control Solutions (DotX); BEIA Consult International (BEIA); DEPSys (DEPS); Hellenic Telecommunications Organization S.A (OTE); Princeton University (PU); University of California, Santa Barbara (UC); University of Nebraska–Lincoln (UNL); Institute of Electrical Engineering of the Chinese Academy of Sciences (CAS); China Electric Power Research Institute (EPRI); Southeast University (SEU); and Jinan University (JNU)

Abstract: TESTBED2 is a major interdisciplinary project that combines wisdoms in three academic disciplines - Electronic & Electrical Engineering, Computing Sciences and Macroeconomics, for developing new techniques to improve the scalability of smart grid services, particularly considering the joint evolution of decarbonized power, heat and transport systems. Moreover, new experimental testbeds will be created to evaluate scalable smart grid solutions. Overall, the main objective of this project is to coordinate the action of 12 Universities and 5 enterprises (3 SMEs and 2 large enterprises) with complementary expertise to develop and test various promising strategies for ensuring the scalability of smart grid services, thereby facilitating successful deployment and full roll-out of smart grid technologies.

9.4 National initiatives

NF-FOUNDS PC9 PEPR 5G

Participants: Eitan Altman, Samir M. Perlaza.

Project Acronym: PEPR

Project Title: Networks of the Future - Foundations of Future Communications Networks

Program: ANR-22-PEFT-0010

Coordinator: CEA (Dmitri Kténas), CNRS (Serge Verdeyme), IMT (Daniel Koffman)

Duration: 2023 - 2030

Other Partners: EURECOM

Abstract: The 5G network and the networks of the future represent a key issue for French and European industry, society and digital sovereignty. This is why the French government has decided to launch a dedicated national strategy. One of this strategy's priority ambitions is to produce significant public research efforts so the national scientific community contributes fully to making progress that

clearly responds to the challenges of 5G and the networks of the future. In this context, the CNRS, the CEA and the Institut Mines-Télécom (IMT) are co-leading the '5G' acceleration PEPR to support upstream research into the development of advanced technologies for 5G and the networks of the future. NEO is involved in the theme "Networks and Telecommunications" and more specifically in the 'targeted projet 9 (PC9)" Foundations of Future Communications Networks (FOUNDS).

ANR PARFAIT

Participants: Eitan Altman, Ashok Krishnan Komalperlan Sindhu, Girik Maskara, Samir Medina Perlaza, Xinying Zou.

Project Acronym: PARFAIT

Project Title: Planning And leaRning For AI-Edge compuTing

Coordinator: Avignon Univ.

Duration: October 2021 – September 2025

Other Partners: Conservatoire National des Arts et Métiers (CNAM), Univ. Savoie Mont Blanc (USMB)

Abstract: The PARFAIT project develops theoretical foundations for distributed and scalable resource allocation schemes on edge computing infrastructures tailored for AI-based processing tasks. Algorithmic solutions will be developed based on the theory of constrained, delayed, and distributed Markov decision processes to account for edge service orchestration actions and quantify the effect of orchestration policies. Furthermore, using both game and team formulations, the project will pave the way for a theory of decentralized orchestration, a missing building block necessary to match the application quest for data proximity and the synchronization problems that arise when multiple edge orchestrators cooperate under local or partial system view. Finally, to achieve efficient online edge service orchestration, such solutions will be empowered with reinforcement learning techniques to define a suit of orchestration algorithms able to at once adapt over time to the applications' load and cope with the uncertain information available from AI-based applications' footprints.

BPI/PLR/PIA 5G Events Lab

Participants: Eitan Altman, Mandar Datar, Ashok Krishnan Komalan Sindhu, Vijith Kumar Kizhakke Purakkal, Giovanni Neglia, Samir M. Perlaza.

Project Acronym: 5G Events Labs

Project Title: 5G Events Labs

Coordinator: Orange

Duration: November 2020 – November 2023

Other Partners: CEA Saclay, Ericsson.

Abstract: The project aims to boost economic activity in the event, cultural and sports sectors, around major Olympic sites in France where Orange and its partners offer increased 5G coverage, technological platforms and appropriate support allowing companies to appropriate these technologies and incubate innovations in the field of services to spectators and organizers.

Inria Challenge FedMalin

Participants: Othmane Marfoq, Giovanni Neglia, Angelo Rodio.

Project Acronym: FedMalin

Project Title: FEDerated MAchine Learning over the INternet

Coordinator: Giovanni Neglia and Aurélien Bellet (Inria team MAGNET)

Duration: November 2022 - November 2026

Abstract: In many use-cases of Machine Learning (ML), data is naturally decentralized: medical data is collected and stored by different hospitals, crowdsensed data is generated by personal devices, etc. Federated Learning (FL) has recently emerged as a novel paradigm where a set of entities with local datasets collaboratively train ML models while keeping their data decentralized.

FedMalin is a research project that spans 10 Inria research teams and aims to push FL research and concrete use-cases through a multidisciplinary consortium involving expertise in ML, distributed systems, privacy and security, networks, and medicine. We propose to address a number of challenges that arise when FL is deployed over the Internet, including privacy and fairness, energy consumption, personalization, and location/time dependencies. FedMalin will also contribute to the development of open-source tools for FL experimentation and real-world deployments, and use them for concrete applications in medicine and crowdsensing. The FedMalin Inria Challenge is supported by Groupe La Poste, sponsor of the Inria Foundation.

Exploratory action Inria MAMMALS

Participants: Francescomaria Faticanti, Giovanni Neglia.

Project Acronym: MAMMALS

Project Title: Memory-augmented Models for low-latency Machine-learning Serving

Coordinator: Giovanni Neglia

Duration: November 2020 - June 2023

Other Partners: Univ. Turin, Polytechnic Turin, Univ. Verona, Univ. of Massachusetts, Northeastern Univ.

Abstract: A machine learning model is often trained for inference's purposes. Inference does not involve complex iterative algorithms and is therefore generally presumed to be easy. Nevertheless, it presents fundamental challenges that are likely to become dominant as machine learning adoption increases and machine learning systems are ubiquitously deployed and need to make timely and safe decisions in unpredictable environments. MAMMALS aims to provide low-latency inferences by running—close to the end user—simple machine learning models that can also take advantage of a (small) local data store of examples. The focus is on algorithms to learn online what to store locally to improve inference quality and achieve domain adaptation.

Exploratory action Inria IDEM

Participants: Emmanouil Athanasakos, Alain Jean-Marie, Samir M. Perlaza, Ke Sun.

Project Acronym: IDEM

Project Title: Information and Decision Making

Coordinators: Samir M. Perlaza, Alain Jean-Marie

Duration: November 2021 - November 2023

Other Partners: Univ. of Sheffield, the National Chiao Tung Univ. in Taiwan, Princeton Univ.

Abstract: IDEM aims to characterize the interplay between data acquisition and information processing in decentralized decision making by bringing together tools from information theory and game theory. This characterization is central in the comprehension of problems including decentralized optimization and Machine Learning subject to local information constraints.

10 Dissemination

Participants: Sara Alouf, Eitan Altman, Konstantin Avrachenkov, Younes Ben Mazziane, Vinay Kumar Bindiganavile Ramadas, Jake Clarkson, Louis Hauseux, Alain Jean-Marie, Caelin Kaplan, Philippe Nain, Giovanni Neglia, Samir M. Perlaza, Angelo Rodio, Charlotte Rodriguez.

10.1 Promoting scientific activities

10.1.1 Scientific events: organisation

Steering committees

- **E. Altman** is
  - Chairperson of the steering committee and co-founder of "Workshop on Networking Games Control and Optimization (NetGcoop)";
  - Member of the steering committee and founder of "Workshop on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt)";
  - Member of the steering committee of "International Conference on Performance Evaluation Methodologies and Tools (ValueTools)".

General chair, scientific chair

- **S. Alouf** was a Chair of the 2023 ACM SIGMETRICS Student Research Competition, Orlando, FL, United States, 20-22 June 2023 [43];
- **S. Perlaza** was a Chair of the “Workshop on Resource Allocation, Cooperation and Competition in Wireless Networks (RAWNET)”, with Samson Lasaulce, Anne Savard, Vincent Tan, and Chao Zhang. Held jointly with WiOpt 2023. Singapore, 24 August 2023.

Member of the organizing committees

- **S. Alouf** was a Student Activities Chair of ACM SIGMETRICS, Orlando, FL, United States, 20-22 June 2023;
- **K. Avrachenkov** was a Lightning Chair of the 12th International Conference on Complex Networks and their Applications, Menton, France, 28-30 November 2023;
- **K. Avrachenkov** was an Award Chair of the 3rd French Regional Conference on Complex Systems (FRCCS), Le Havre, France, 31 May - 2 June 2023;
- **S. Perlaza** was a Publicity Chair of the "International Symposium on Modeling and Optimization in Mobile, Ad hoc, and Wireless Networks" (WiOpt2023). Singapore, 24-27 August 2023.

10.1.2 Scientific events: selection

Chair of conference program committees

- **S. Alouf** was a TPC Chair of the "35th International Teletraffic Congress (ITC 35)", Turin, Italy, 2-4 October 2023.
- **K. Avrachenkov** was a TPC Co-Chair at ACM SIGMETRICS 2023 (Orlando, Florida, USA) [45, 46].
- **S. Perlaza** was a Technical Symposium Chair (Privacy and Security) at the "IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids" (SmartGridComm). Glasgow, Scotland, 31 October - 3 November 2023.

Member of the conference program committees

- 24th International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing (ACM MobiHoc 2023), 23-26 October 2023, Washington DC, United States (S. Alouf);
- ACM SIGMETRICS/ IFIP Performance 2024, Summer and Fall TPCs, 10-14 June 2024, Venice, Italy (K. Avrachenkov);
- ACM SIGMETRICS 2023, Winter TPCs, 19-22 June 2023, Orlando, Florida, United States (S. Alouf, K. Avrachenkov);
- The 27th International Conference on Analytical & Stochastic Modelling Techniques & Applications / 19th European Performance Engineering Workshop (ASMTA/EPEW 2023), 20-23 June 2023, Florence, Italy (K. Avrachenkov, A. Jean-Marie);
- 18th Workshop on Algorithms and Models for the Web Graph (WAW), 23-26 May 2023, Fields Institute, Toronto, Canada (K. Avrachenkov);
- 14th Conference on Decision and Game Theory for Security (GameSec), 18-20 October 2023, Avignon, France (K. Avrachenkov);
- SIAM Conference on Data Mining (SDM), 27-29 April 2023, Minneapolis, United States (K. Avrachenkov);
- 16th International Conference on Performance Evaluation Methodologies and Tools (ValueTools), 6-7 September 2023, Crete, Greece (K. Avrachenkov);
- 12th International Conference on Complex Networks and their Applications, 28-30 November 2023, Menton, France (K. Avrachenkov);
- 3rd French Regional Conference on Complex Systems (FRCCS), 31 May - 2 June 2023, Le Havre, France (K. Avrachenkov);

- IEEE International Symposium on Information Theory (ISIT), 25–30 June 2023, Taipei, Taiwan (S. Perlaza);
- IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids" (SmartGridComm) 31 October - 3 November 2023, Glasgow, Scotland, UK (S. Perlaza);
- IEEE International Conference in Communications (ICC), Wireless Communications Symposium and Selected Areas in Communications (Machine Learning in Communications and Networking (MLCN); and Integrated Sensing and Communications (ISC)). 28 May - 1 June 2023, Rome, Italy (S. Perlaza);
- IEEE Wireless Communications and Networking Conference (WCNC), 26–29 March 2023, Glasgow, Scotland, UK (S. Perlaza);
- International ITG Workshop on Smart Antennas (WSA) and Conference on Systems, Communications, and Coding (CSCC) 27 February - 3 March 2023, Braunschweig, Germany (S. Perlaza);
- 21st International Symposium on Modeling and Optimization in Mobile, Ad hoc, and Wireless Networks (WiOpt 2023), 24-27 August 2023, Singapore (S. Alouf, K. Avrachenkov);
- 31st International Symposium on the Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2023) 21-23 October 2023, Stony Brook, United States (A. Jean-Marie);
- 35th International Teletraffic Congress (ITC 35), 3-5 October 2023, Torino, Italy (A. Jean-Marie);
- 24th Conf. of the Société Française de Recherche Opérationnelle et d'Aide à la Décision (ROADEF 2023), 20-23 February 2023, Rennes, France (A. Jean-Marie);
- 25th Workshop on MAthematical performance Modeling and Analysis (MAMA 2023), 19 June 2023, Orlando, Florida, United States (A. Jean-Marie);
- 20th IEEE International Conference on Sensing, Communication, and Networking (IEEE SECON 2023), 11–14 September 2023, Madrid, Spain (F. Faticanti);
- IEEE Workshop on Edge Network Softwarization (ENS 2023), 23 June 2023, Milan, Italy (F. Faticanti);
- IEEE International Conference on Computer Communications (INFOCOM 2024), 20–23 May 2024, Vancouver, Canada (G. Neglia);
- The First International Workshop on the Integration between Distributed Machine Learning and the Internet of Things (AIoT 2023) 23 October 2023, Washington DC, United States (G. Neglia).

Reviewer (outside program committee duties)

- 21st International Symposium on Modeling and Optimization in Mobile, Ad hoc, and Wireless Networks (WiOpt 2023), 24-27 August 2023, Singapore (A. Jean-Marie);
- 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2023), 4-10 June 2023, Rhodes Island, Greece (G. Neglia).

10.1.3 Journal

Member of the editorial boards

- ACM Transactions on Modeling and Performance Evaluation of Computing Systems (ACM ToMPECS) (K. Avrachenkov, since 2016);
- AIMS Journal of Dynamics and Games (E. Altman, since 2015);
- Birkhauser Journal on Dynamic Games and Applications (E. Altman, since 2012);

- Elsevier Computer Communications (S. Alouf, since 2021; G. Neglia, since 2014);
- Elsevier Performance Evaluation (K. Avrachenkov, since 2009; P. Nain, since 2018);
- IEEE/ACM Transactions on Networking (E. Altman, since 2013);
- IEEE Network Magazine (K. Avrachenkov, since 2020);
- Springer Iran Journal of Computer Science (E. Altman, advisory board member);
- Stochastic Models (K. Avrachenkov, since 2019).

K. Avrachenkov in collaboration with V. Borkar and J. Nair (IIT Bombay, India) have organized and edited a special volume of *Dynamic Games and Applications* journal on the topic "Multi-Agent Dynamic Decision Making and Learning" [44]. See Section 7.3.4.

Reviewer - reviewing activities NEO members regularly perform reviews for journals such as IEEE/ACM Transactions on Networking, IEEE Transactions on Information Theory, IEEE Transactions on Wireless Communications, IEEE Transactions on Communications, IEEE Transactions on Network and Service Management, IEEE Transactions on Network Science and Engineering, Performance Evaluation, Elsevier Computer Communications, Elsevier Computer Networks.

10.1.4 Invited talks

- **K. Avrachenkov** delivered a
  - plenary talk "Graph Clustering Problem: Beyond Binary Interactions" at the 59th ANZIAM Conference, Cairns, Australia, held on 5-9 February 2023;
  - Specialist Lecture "Singularly Perturbed Markovian Models: From Queues to Web Ranking and Reinforcement Learning" on Monday 13 February, at University of South Australia;
  - Specialist Lecture "Random-walk Based Sampling in Social Networks" on Wednesday 15 February, at RMIT;
  - Specialist Lecture "Random Graph Models, Network Centralities and Graph Clustering" on Friday 17 February, at Australian Bureau of Satistics;
  - Specialist Lecture "Reinforcement Learning for Restless Bandits" on Monday 20 February, at University of Newcastle;
  - Public Lecture "Aesthetics and ubiquitous applications of Markov chains" on Wednesday 22 February, at The University of Queensland;
  - an invited presentation "Reinforcement Learning Methods for Weakly Coupled MDPs" at the Workshop on restless bandits, index policies and applications in reinforcement learning, Grenoble, on 21 November.
- **J. Clarkson** delivered an invited talk "On the Price of Information in Queueing" at 21st INFORMS Applied Probability Society Conference, Nancy, on June 28-30.
- **C. Kaplan** delivered the following invited talks:
  - "A Cautionary Tale: On the Role of Reference Data in Empirical Privacy Defenses" at SAP's Machine Learning Conference, March 22 - March 23, Remote.
  - "Membership Inference Attacks in Machine Learning" for SAP's explAIn Series, December 7, Remote.
- **V. Kumar B.R.** delivered the following invited talks
  - "Community detection on block models with geometric kernels" at 21st INFORMS Applied Probability Society Conference, Nancy, on June 28-30.

  - "Community detection on Block Models with Geometric Kernels", in the CNI Networks Seminar Series at IISc, Bangalore, on December 5.

- **G. Neglia** delivered the following invited talks:
  - "Personalized Federated Learning" at Stellantis'Data Science Community meetup, online, 22/9/2023.
  - "Personalized Federated Learning" at IEEE Systems, Man, and Cybernetics (SMC) Society Italy Chapter, 5/5/2023.

- **S. Perlaza** delivered the following invited talks:
  - "On the Validation of Gibbs Algorithms: Training Datasets, Test Datasets and their Aggregation" at the Ecole Normale Supérieure (ENS) de Lyon. Laboratoire de l'Informatique du Parallélisme (LIPS). Lyon, France, June 21, 2023. Host: Prof. Francesco Bronzino.
  - "An Upper Bound on the Error Induced by Saddlepoint Approximations—Applications to Wireless Communications" at the Workshop on "Performance Guarantees in Wireless Networks". Laboratory for Information, Networking and Communication Sciences (LINCS). Palaiseau, France, March 8, 2023. Host: Prof. Francois Baccelli and Prof. Jean-Marie Gorce.
  - "Pricing Models for Digital Renting Platforms" at Booking.com, Amsterdam, Netherlands, October 18, 2023. Host: Dr. Christina Katsimerou.

- **C. Rodriguez** delivered a talk on "How can we evaluate the energy consumption of Machine Learning?" at Journées de Recherche en Apprentissage Frugal, Grenoble, 13-14 December, 2023.

10.1.5 Invited papers

S. Alouf was invited to write a technical perspective on a paper appearing in the Research Highlights section of the Communications of the ACM [57].

10.1.6 Leadership within the scientific community

- **S. Alouf**
  - is an elected member at the Board of Directors of ACM SIGMETRICS (July 2019 - June 2023);
  - is a member of the Equality and Diversity committee of ACM SIGMETRICS (November 2019 - June 2023);
  - is a member of the Conference Advisory committee of ACM SIGMETRICS (November 2019 - June 2023).

- **E. Altman**
  - is Fellow Member of IEEE;
  - is Member of WG 7.3 of IFIP on Computer System Modeling;
  - is Member of WG 6.3 of IFIP on Performance of Communications Systems.

- **K. Avrachenkov**
  - was the 2023 AMSI-ANZIAM Lecturer;
  - is a member of Conseil Scientifique & Pédagogique EUR DS4H Univ. Côte d'Azur.

- **S. Perlaza**
  - is a member of the Digital Presence Committee of the IEEE Information Theory Society.
  - is a workpackage leader of the PEPR - Réseaux du Futur – A project funded by the French National Agency for Research (ANR) via the project n°ANR-22-PEFT-0010 of the France 2030 program
  - is the organizer of the PC9 Seminar on Wireless Communications, a national online seminar part of the PEPR – Réseaux du Futur.

10.1.7 Research administration

- **S. Alouf**
  - was member of the competitive exam jury for a researcher position (CRCN/ISFP) at Inria Sophia Antipolis Méditerranée in 2023.
  - is a member of the Colloquium Jacques Morgenstern Committee at Inria Sophia Antipolis Méditerranée, since March 2023
  - is a member of NICE, the Invited Researchers Committee of Inria Sophia Antipolis Méditerranée, since June 2020;
  - is member of CLF, the training committee of Inria Sophia Antipolis Méditerranée, since November 2014;
  - is vice-head of project-team Neo since January 2017.
- **K. Avrachenkov** was in charge of the Activity Report process for Inria Sophia Antipolis Méditerranée.
- **A. Jean-Marie**
  - is the scientific coordinator of Inria activities in Montpellier (since 2008); as part of this duty, he represents Inria at: the Scientific Council of the Doctoral School "Sciences and Agrosciences" of the Univ. of Avignon; at the Regional Conference of Research Organisms (CODOR); at the board of the Labex NUMEV.
  - is Head of project-team NEO since January 2017
- **P. Nain** chaired the Inria ad-hoc commission in charge of evaluating applications for the double 2023-2024 bonus campaign (RIPEC-C3). [He also chaired this commission for the 2022 campaign.]
- **G. Neglia** is
  - an elected member of Inria evaluation committee, since September 2023;
  - a member of the steering commitee of UniCA Graduate School of Digital Systems for Humans (DS4H) since September 2022.
- **S. Perlaza**
  - is a representant (suppléant, collège A) before the Comité de Centre Inria d'Université Côte d'Azur
  - is a Member of the Bureau de l'Académie Réseaux, Information et Société Numérique, Académie d'Excellence d'Univ Côte d'Azur.
  - is a Member of the Département Disciplinaire Informatique de l'Univ Côte d'Azur.

10.2 Teaching - Supervision - Juries

10.2.1 Teaching

PhD

- G. Neglia, "Communication Aspect of Federated Learning", 1.25h for GDR RSD Summer School on Distributed Learning, Lyon, France 19-20/9/2023.

Master

- L. Hauseux, K. Avrachenkov, "Statistical Analysis of Networks", 24H, M2 Data Science and Artificial Intelligence;
- G. Neglia, O. Marfoq, "Machine Learning: Theory and Algorithms", 21H, M2 Ubinet, UniCA, France;
- G. Neglia and A. Rodio together with C. Xu (COATIteam) "Federated Learning & Data Privacy", 30H, M2 Data Science and Artificial Intelligence;
- F. Faticanti, "Resource allocation in Cloud and Fog Computing", 2H, M2 University of Trento, Italy.

Bachelor

- Y. Ben Mazziane, “Base de Programmation en Python”, 26h, BAT3 S1, Polytech Nice Sophia, UniCA, France.

10.2.2 Supervision

PhD defended

- Othmane Marfoq, "Distributed machine learning for IoT applications", UniCA, since 1 December 2020, defended 7 December 2023, advisor: Giovanni Neglia.
- Xiuzhen Ye, "Sparsity and Coordination Constraints on Stealth Data Injection Attacks", defended March 28, 2023, Univerisity of Sheffield, UK, external-advisor: S. Perlaza.

PhD in progress

- Younes Ben Mazziane, "Online learning for Caching at the Edge", UniCA, 1 Oct. 2020, advisors: Sara Alouf and Giovanni Neglia.
- Olha Chuchuk, "Optimization of data access at CERN and in the World Large Hadron Collider Computing Grid (WLCG)", UniCA, 1 Sept. 2020, advisor: Giovanni Neglia.
- Francisco Daunas, "Data Injection Attacks in Machine Learning Systems", Univ. of Sheffield, 1 Oct. 2020, co-advisor: S. M. Perlaza.
- Ibtihal El Mimouni, "Systèmes de recommandation automatisés et responsables pour le marketing digital", UniCA, Thèse Cifre avec NSP SmartProfile, 1 Oct. 2022, advisor: Konstantin Avrachenkov.
- Louis Hauseux, "Classifiers on Random Graphs with applications to Social Networks and Image Processing", UniCA, 1 Oct. 2023, advisors: Konstantin Avrachenkov and Josiane Zerubia.
- Caelin Kaplan, "Privacy and fairness for machine learning", UniCA, 1 July 2021, advisors: Giovanni Neglia and Alain Jean-Marie.
- Angelo Rodio, "Sustainable distributed machine learning", UniCA, 1 April 2021, advisors: Giovanni Neglia and Alain Jean-Marie.
- Julian Santos, "Modélisation analytique de réseaux sans fils grande échelle intégrant les RIS - Optimisation de l’allocation dynamique des ressources", UniCA, 1 Sep. 2023, advisor: Eitan Altman.
- Xufeng Zhang, "Incentives for Federated Learning", UniCA, 1 Dec. 2023, advisors: Sara Alouf and Giovanni Neglia.
- Xinying Zou, "Classical and non-classical information patterns in distributed control with delayed sharing information", UniCA, 1 Dec. 2022, advisors: Eitan Altman and Samir M. Perlaza.

10.2.3 Juries

PhD

- Anthony Bardou, “Online Learning for the Black-Box Optimization of Wireless Networks”, Ecole normale supérieure de Lyon, 7 September 2023 (G. Neglia, reviewer);
- Marina Costantini “Optimization methods over networks for popular content delivery and distributed machine learning”, U. Sorbonne, 26 June 2023 (G. Neglia, jury member);
- David Jia, “Topological Interference Management in Clustered Ad-hoc Networks”, U. Paris Saclay, 9 October 2023 (S. Perlaza, reviewer);
- Kimang Khun, “Algorithms for Markovian bandits : Indexability and Learning”, U. Grenoble Alpes, 4 March 2023 (K. Avrachenkov, reviewer);

- Arthur Da Cunha, “Pruning random structures”, UniCA, 13 September 2023 (K. Avrachenkov, jury president);
- Mohamad Mestoukirdi, “Reliable and Communication-Efficient Federated Learning for Future Intelligent Edge Networks”, U. Sorbonne, 4 December 2023 (G. Neglia, jury member);
- Matei Moldoveanu, “In-Network Learning: Distributed Training and Inference in Communication Networks”, U. Gustave Eiffel, 30 January 2023 (S. Perlaza, reviewer);
- Pierre Popineau, “Study of the dynamics of spatial point processes in wireless communication networks”, ENS - U. PSL, 14 June 2023 (S. Alouf, jury member);
- Raoul Raftopoulos “Integrating Deep Reinforcement Learning in 6G Edge Environments: Towards Intelligent Network Optimization”, U. Catania, Italy, November 2023, (G. Neglia, reviewer);
- Louis-Sébastien Rebuffi, “Reinforcement Learning Algorithms for Controlled Queueing Systems”, U. Grenoble, 11 December 2023 (A. Jean-Marie, jury member);
- Ilia Shilov, “Game-Theoretic Approaches for Peer-to-Peer Energy Trading”, U. PSL, 20 September 2023 (A. Jean-Marie, jury member);
- Robin Vacus, “Algorithmic Perspectives to Collective Natural Phenomena”, U. Paris Cité, 18 December 2023 (A. Jean-Marie, jury president);
- Paul Youssef, “Online Learning at the Edge”, U. Grenoble Alpes, 13 March 2023 (G. Neglia, reviewer).

11 Scientific production

11.1 Major publications

[1] D. Anade, J.-M. Gorce, P. Mary and S. M. Perlaza. ‘Saddlepoint Approximations of Cumulative Distribution Functions of Sums of Random Vectors’. In: ISIT 2021 - IEEE International Symposium on Information Theory. Melbourne / Virtual, Australia: IEEE, 12th July 2021, pp. 1–6. URL: https://hal.inria.fr/hal-03226009.

[2] K. Avrachenkov, A. Bobu and M. Dreveton. ‘Higher-Order Spectral Clustering for Geometric Graphs’. In: *Journal of Fourier Analysis and Applications* 27 (15th Mar. 2021). DOI: 10.1007/s00041-021-09825-2. URL: https://hal.inria.fr/hal-03169834.

[3] K. Avrachenkov and V. S. Borkar. ‘Whittle Index Policy for Crawling Ephemeral Content’. In: *IEEE Transactions on Control of Network Systems* 5.1 (Mar. 2018), pp. 446–455. DOI: 10.1109/TCNS.2016.2619066. URL: https://hal.inria.fr/hal-01937994.

[4] K. Avrachenkov, A. Y. Kondratev, V. V. Mazalov and D. Rubanov. ‘Network partitioning algorithms as cooperative games’. In: *Computational Social Networks* 5.11 (Oct. 2018). DOI: 10.1186/s40649-018-0059-5. URL: https://hal.inria.fr/hal-01935419.

[5] K. Avrachenkov, A. Piunovskiy and Y. Zhang. ‘Hitting Times in Markov Chains with Restart and their Application to Network Centrality’. In: *Methodology and Computing in Applied Probability* 20.4 (Dec. 2018), pp. 1173–1188. DOI: 10.1007/s11009-017-9600-5. URL: https://hal.inria.fr/hal-01937983.

[6] V. Bucarey López, E. Della Vecchia, A. Jean-Marie and F. Ordoñez. ‘Stationary Strong Stackelberg Equilibrium in Discounted Stochastic Games’. In: *IEEE Transactions on Automatic Control* (2023). URL: https://hal.inria.fr/hal-03934114.

[7] E. Leonardi and G. Neglia. ‘Implicit Coordination of Caches in Small Cell Networks under Unknown Popularity Profiles’. In: *IEEE Journal on Selected Areas in Communications* 36.6 (June 2018), pp. 1276–1285. DOI: 10.1109/JSAC.2018.2844982. URL: https://hal.inria.fr/hal-01956307.

[8] A. R. Masson, Y. Hayel and E. Altman. 'Posting behaviour Dynamics and Active Filtering for Content Diversity in Social Networks'. In: *IEEE transactions on Signal and Information Processing over Networks* 3.2 (2017), pp. 376–387. DOI: 10.1109/TSIPN.2017.2696738. URL: https://hal.inria.fr/hal-01536172.

[9] K. P. Naveen, E. Altman and A. Kumar. 'Competitive Selection of Ephemeral Relays in Wireless Networks'. In: *IEEE Journal on Selected Areas in Communications* 35 (2017), pp. 586–600. DOI: 10.1109/JSAC.2017.2659579. URL: https://hal.inria.fr/hal-01536123.

[10] G. Neglia, D. Carra, M. Feng, V. Janardhan, P. Michiardi and D. Tsigkari. 'Access-Time-Aware Cache Algorithms'. In: *ACM Transactions on Modeling and Performance Evaluation of Computing Systems* 2.4 (Dec. 2017), pp. 1–29. DOI: 10.1145/3149001. URL: https://hal.inria.fr/hal-01956285.

[11] G. Neglia, D. Carra and P. Michiardi. 'Cache Policies for Linear Utility Maximization'. In: *IEEE/ACM Transactions on Networking* 26.1 (Feb. 2018), pp. 302–313. DOI: 10.1109/TNET.2017.2783623. URL: https://hal.inria.fr/hal-01956319.

[12] A. Tajer, S. M. Perlaza and H. V. Poor. *Advanced Data Analytics for Power Systems*. Cambridge University Press, 1st Jan. 2021. URL: https://hal.archives-ouvertes.fr/hal-03128425.

11.2 Publications of the year

International journals

[13] K. Avrachenkov, E. Morozov and R. Nekrasova. 'Stability analysis of two-class retrial systems with constant retrial rates and general service times'. In: *Performance Evaluation* 159 (Jan. 2023), p. 102330. DOI: 10.1016/j.peva.2022.102330. URL: https://inria.hal.science/hal-04372687.

[14] V. Bucarey López, E. Della Vecchia, A. Jean-Marie and F. Ordoñez. 'Stationary Strong Stackelberg Equilibrium in Discounted Stochastic Games'. In: *IEEE Transactions on Automatic Control* 68.9 (2023), pp. 5271–5286. DOI: 10.1109/TAC.2022.3220512. URL: https://inria.hal.science/hal-03934114.

[15] M. Datar, E. Altman and H. Le Cadre. 'Strategic Resource Pricing and Allocation in a 5G Network Slicing Stackelberg Game'. In: *IEEE Transactions on Network and Service Management* 20.1 (2023), pp. 1932–4537. DOI: 10.1109/TNSM.2022.3216588. URL: https://inria.hal.science/hal-03824540.

[16] S. Dhamal, W. Ben-Ameur, T. Chahed, E. Altman, A. Sunny and S. Poojary. 'A game theoretic framework for distributed computing with dynamic set of agents'. In: *Annals of Operations Research* (23rd Feb. 2023). DOI: 10.1007/s10479-023-05231-7. URL: https://hal.science/hal-04302615.

[17] R. Dhounchak, V. Kavitha and E. Altman. 'Viral Marketing Branching Processes'. In: *Computer Communications* 198 (15th Jan. 2023), pp. 140–156. DOI: 10.1016/j.comcom.2022.11.015. URL: https://inria.hal.science/hal-03858831.

[18] F. Faticanti, M. Savi, F. De Pellegrini and D. Siracusa. 'Locality-aware deployment of application microservices for multi-domain fog computing'. In: *Computer Communications* 203 (Apr. 2023), pp. 180–191. DOI: 10.1016/j.comcom.2023.02.012. URL: https://inria.hal.science/hal-04367135.

[19] A. Jean-Marie and M. Tidball. 'Dynamic fishing with endogenous habitat damage'. In: *Dynamic Games and Applications* (19th July 2023). DOI: 10.1007/s13235-023-00511-0. URL: https://hal.inrae.fr/hal-04136790.

[20] A. Rodio, F. Faticanti, O. Marfoq, G. Neglia and E. Leonardi. 'Federated Learning Under Heterogeneous and Correlated Client Availability'. In: *IEEE/ACM Transactions on Networking* (2023), pp. 1–10. DOI: 10.1109/TNET.2023.3324257. URL: https://hal.science/hal-04364293.

[21] T. S. Salem, G. Castellano, G. Neglia, F. Pianese and A. Araldo. 'Toward Inference Delivery Networks: Distributing Machine Learning With Optimality Guarantees'. In: *IEEE/ACM Transactions on Networking* (31st Aug. 2023), p. 1. DOI: 10.1109/TNET.2023.3305922. URL: https://inria.hal.science/hal-04385717.

[22] T. Si Salem, G. Neglia and S. Ioannidis. 'No-regret Caching via Online Mirror Descent'. In: *ACM Transactions on Modeling and Performance Evaluation of Computing Systems* 8.4 (11th Aug. 2023), pp. 1–32. DOI: 10.1145/3605209. URL: https://hal.science/hal-04181387.

[23] G. Vardoyan, P. Nain, S. Guha and D. Towsley. 'On the Capacity Region of Bipartite and Tripartite Entanglement Switching'. In: *ACM Transactions on Modeling and Performance Evaluation of Computing Systems* 8.1-2 (June 2023), pp. 1–18. DOI: 10.1145/3571809. URL: https://inria.hal.science/hal-04018593.

[24] X. Ye, I. Esnaola, S. M. Perlaza and R. F. Harrison. 'Stealth Data Injection Attacks with Sparsity Constraints'. In: *IEEE Transactions on Smart Grid* 14.4 (1st July 2023), pp. 3201–3209. DOI: 10.1109/TSG.2023.3238913. URL: https://hal.science/hal-03516567.

International peer-reviewed conferences

[25] K. Alaluusua, K. Avrachenkov, B. R. Vinay Kumar and L. Leskelä. 'Multilayer Hypergraph Clustering Using the Aggregate Similarity Matrix'. In: *Lecture Notes in Computer Science*. WAW 2023 - Algorithms and Models for the Web Graph. Vol. LNCS-13894. Algorithms and Models for the Web Graph 18th International Workshop, WAW 2023, Toronto, ON, Canada, May 23–26, 2023, Proceedings. Toronto, Canada: Springer Nature Switzerland, 17th Mar. 2023, pp. 83–98. DOI: 10.1007/978-3-031-32296-9_6. URL: https://inria.hal.science/hal-04372200.

[26] Y. Ben Mazziane, F. Faticanti, G. Neglia and S. Alouf. 'No-Regret Caching with Noisy Request Estimates'. In: IEEE VCC 2023 - IEEE Virtual Conference on Communications. New York (ONLINE), United States, 28th Nov. 2023. URL: https://hal.science/hal-04318435.

[27] F. Daunas, I. Esnaola, S. M. Perlaza and H. V. Poor. 'Analysis of the Relative Entropy Asymmetry in the Regularization of Empirical Risk Minimization'. In: *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*. (ISIT 2023 - IEEE International Symposium on Information Theory. Taipei, Taiwan, 24th June 2023. DOI: 10.1109/ISIT54713.2023.10206876. URL: https://hal.science/hal-04097637.

[28] F. Faticanti and G. Neglia. 'Optimistic Online Caching for Batched Requests'. In: *IEEE Xplore*. ICC 2023 - IEEE International Conference on Communications. ICC 2023 - IEEE International Conference on Communications. Rome, France: IEEE, 28th May 2023, pp. 6243–6248. DOI: 10.1109/ICC45041.2023.10278692. URL: https://inria.hal.science/hal-04367129.

[29] A. Fox, F. De Pellegrini and E. Altman. 'Learning Optimal Edge Processing with Offloading and Energy Harvesting'. In: MSWiM '23: Int'l ACM Conference on Modeling Analysis and Simulation of Wireless and Mobile Systems. Montreal Quebec Canada, Canada: ACM, 30th Oct. 2023, pp. 83–92. DOI: 10.1145/3616388.3617516. URL: https://hal.science/hal-04022507.

[30] L. Hauseux, K. Avrachenkov and J. Zerubia. 'Graph Based Approach for Galaxy Filament Extraction'. In: *Studies in Computational Intelligence*. Complex Networks 2023 - The 12th International Conference on Complex Networks and their Applications. Proceedings of The Twelfth International Conference on Complex Networks and Their Applications. Menton, France, 28th Nov. 2023. URL: https://inria.hal.science/hal-04231772.

[31] A. Krishnan K. S., S. M. Perlaza and E. Altman. 'Pricing Models for Digital Renting Platforms'. In: ASMTA 2023 - International Conference on Analytical & Stochastic Modelling Techniques & Applications. Florence, France, 20th June 2023. DOI: 10.1007/978-3-031-43185-2_10. URL: https://hal.science/hal-04136929.

[32] V. Kumar K. P., B. Kumar Rai and T. Jacob. 'The Optimal Rate Memory Tradeoff in Multi-Access Coded Caching: Large Cache Size'. In: 2023 IEEE Information Theory Workshop (ITW). Saint-Malo, France: IEEE, 23rd Apr. 2023, pp. 165–169. DOI: 10.1109/ITW55543.2023.10161659. URL: https://inria.hal.science/hal-04388612.

[33] O. Marfoq, G. Neglia, L. Kameni and R. Vidal. 'Federated Learning for Data Streams'. In: *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*. The 26th International Conference on Artificial Intelligence and Statistics. Valencia (Espagne), Spain, 25th Apr. 2023. URL: https://inria.hal.science/hal-04385635.

[34] T. Pagare, V. S. Borkar and K. Avrachenkov. 'Full Gradient Deep Reinforcement Learning for Average-Reward Criterion'. In: *Proceedings of Machine Learning Research (PMLR)*. L4DC - The 5th Annual Learning for Dynamics and Control Conference. Vol. 211. Learning for Dynamics and Control Conference, 15-16 June 2023, Philadelphia, PA. Philadelphia, United States, 14th June 2023, pp. 235–247. URL: https://inria.hal.science/hal-04372096.

[35] S. M. Perlaza, I. Esnaola, G. Bisson and H. V. Poor. 'On the Validation of Gibbs Algorithms: Training Datasets, Test Datasets and their Aggregation'. In: ISIT 2023 - IEEE International Symposium on Information Theory. Taipei, Taiwan, 30th June 2023. DOI: 10.1109/ISIT54713.2023.10206506. URL: https://hal.science/hal-04096054.

[36] A. Rodio, F. Faticanti, O. Marfoq, G. Neglia and E. Leonardi. 'Federated Learning under Heterogeneous and Correlated Client Availability'. In: IEEE INFOCOM 2023 - IEEE Conference on Computer Communications. IEEE Annual Joint Conference: INFOCOM, IEEE Computer and Communications Societies. New York City, United States: IEEE, 17th May 2023, pp. 1–10. DOI: 10.1109/INFOCOM53939.2023.10228876. URL: https://hal.science/hal-04364285.

[37] A. Rodio, G. Neglia, F. Busacca, S. Mangione, S. Palazzo, F. Restuccia and I. Tinnirello. 'Federated Learning with Packet Losses'. In: *IEEE Xplore*. WPMC 2023 - 26th International Symposium on Wireless Personal Multimedia Communications. 2022 25th International Symposium on Wireless Personal Multimedia Communications (WPMC). Tampa, United States: IEEE, 19th Nov. 2023, pp. 1–6. DOI: 10.1109/WPMC59531.2023.10338845. URL: https://hal.science/hal-04364289.

[38] T. Si Salem, G. Iosifidis and G. Neglia. 'Enabling Long-term Fairness in Dynamic Resource Allocation'. In: *ACM digital library*. 2023 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems. SIGMETRICS '23: ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems. Orlando (FL), United States, 27th June 2023. DOI: 10.1145/3578338.3593541. URL: https://inria.hal.science/hal-04386354.

[39] K. Sun, S. M. Perlaza and A. Jean-Marie. '2×2 Zero-Sum Games with Commitments and Noisy Observations'. In: ISIT 2023 - IEEE International Symposium on Information Theory. Taipei, Taiwan, 25th June 2023. DOI: 10.1109/ISIT54713.2023.10206806. URL: https://inria.hal.science/hal-04091712.

[40] R. Taisant, M. Datar, H. Le Cadre and E. Altman. 'Learning Market Equilibria Using Performative Prediction: Balancing Efficiency and Privacy'. In: ECC 2023 - European Control Conference. 2023 European Control Conference (ECC). Bucharest, Romania, June 2023. DOI: 10.23919/ECC57647.2023.10178247. URL: https://inria.hal.science/hal-03816949.

[41] B. R. Vinay Kumar. 'Spatial Queues with Nearest Neighbour Shifts'. In: ITC-35 - 35th International Teletraffic Congress. Torino, Italy, 3rd Oct. 2023. URL: https://inria.hal.science/hal-04386596.

Conferences without proceedings

[42] A. Santana de Oliveira, C. Kaplan, K. Mallat and T. Chakraborty. 'An Empirical Analysis of Fairness Notions under Differential Privacy *'. In: The Fourth AAAI Workshop on Privacy-Preserving Artificial Intelligence. Washington D.C., United States, 6th Feb. 2023. URL: https://inria.hal.science/hal-04387685.

Edition (books, proceedings, special issue of a journal)

[43] *ACM SIGMETRICS 2023 Student Research Competition* 51.3 (Dec. 2023). DOI: 10.1145/3639830.3639832. URL: https://inria.hal.science/hal-04264945.

[44] *Special Issue : Multi-Agent Dynamic Decision Making and Learning* 13 (31st Jan. 2023). DOI: 10.1007/s13235-023-00493-z. URL: https://inria.hal.science/hal-04372613.

[45] *POMACS V7, N1, March 2023 Editorial* 7.1 (2nd Mar. 2023). DOI: 10.1145/3579311. URL: https://inria.hal.science/hal-04391378.

[46] *POMACS V7, N2, June 2023 Editorial* 7.2 (22nd May 2023). DOI: 10.1145/3589972. URL: https://inria.hal.science/hal-04391361.

Reports & preprints

[47] F. Cabo, M. Tidball and A. Jean-Marie. *Positional and conformist effects in public good provision. Strategic interaction and inertia*. 27th June 2023. URL: https://hal.inrae.fr/hal-04147447.

[48] M. Datar, N. Modina, R. El Azouzi and E. Altman. *Fisher Market Model based Resource Allocation for 5G Network Slicing*. 2023. URL: https://inria.hal.science/hal-04173931.

[49] F. Daunas, I. Esnaola, S. M. Perlaza and H. V. Poor. *Empirical Risk Minimization with Relative Entropy Regularization Type-II*. RR-9508. INRIA, Centre Inria d'Université Côte d'Azur, Sophia Antipolis, 31st May 2023. URL: https://hal.science/hal-04110899.

[50] F. Daunas, I. Esnaola, S. M. Perlaza and H. V. Poor. *Empirical Risk Minimization with f-Divergence Regularization in Statistical Learning*. RR-9521. Inria, 25th Oct. 2023. URL: https://hal.science/hal-04258765.

[51] A. Krishnan K. S., S. M. Perlaza and E. Altman. *Pricing for Platforms: Games, Equilibria and Cooperation*. RR-9510. Inria, 21st June 2023. URL: https://hal.science/hal-04132372.

[52] H. Le Cadre, M. Datar, M. Guckert and E. Altman. *Learning Market Equilibria Preserving Statistical Privacy Using Performative Prediction*. 23rd June 2023. URL: https://inria.hal.science/hal-04343535.

[53] S. M. Perlaza, G. Bisson, I. Esnaola, A. Jean-Marie and S. Rini. *Empirical Risk Minimization with Relative Entropy Regularization*. 29th Nov. 2023. URL: https://hal.science/hal-03849748.

[54] K. Sun, S. M. Perlaza and A. Jean-Marie. *2 × 2 Zero-Sum Games with Commitments and Noisy Observations*. RR-9505. Inria - Sophia Antipolis, 30th May 2023. URL: https://hal.science/hal-03838009.

[55] B. R. Vinay Kumar. *Spatial Queues with Nearest Neighbour Shifts*. 8th Sept. 2023. URL: https://hal.science/hal-04201357.

[56] X. Zou, S. M. Perlaza, I. Esnaola and E. Altman. *The Worst-Case Data-Generating Probability Measure*. RR-9515. INRIA, 21st Aug. 2023. URL: https://inria.hal.science/hal-04181971.

Other scientific publications

[57] S. Alouf. 'Technical Perspective: Can We Uncover Private Backbone Infrastructures?' In: *Communications of the ACM* 66.8 (Aug. 2023), p. 94. DOI: 10.1145/3604621. URL: https://inria.hal.science/hal-04177294.

[58] L. Hauseux, K. Avrachenkov and J. Zerubia. 'Un nouvel estimateur des niveaux de densité utilisant les graphes et complexes simpliciaux.Application à la détection des clusters de galaxies.' In: XVIe Journées de géostatistiques (Fontainebleau, 7-8 septembre 2023) organisé par les Mines de Paris - PSL. Fontainebleau, France, 7th Sept. 2023. URL: https://inria.hal.science/hal-04222280.