RESEARCH CENTRE

Sophia Antipolis - Méditerranée

2020 ACTIVITY REPORT

Project-Team

NEO

Network Engineering and Operations

DOMAIN

Networks, Systems and Services, Distributed Computing

THEME

Networks and Telecommunications

Contents

Project-Team NEO

Creation of the Team: 2017 January 01, updated into Project-Team: 2017 December 01

Keywords

Computer sciences and digital sciences

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

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

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
- B6.2. – 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]
- Samir Medina Perlaza [Inria, Researcher]
- Giovanni Neglia [Inria, Researcher, HDR]

Post-Doctoral Fellows

- Andrei Bobu [Inria, until Nov. 2020]
- Kishor Yashavant Patil [Inria]
- Ramakrishnan Sambamoorthy [Inria, from Mar. 2020]
- Chuan Xu [Inria, until Feb. 2020]
- Sadaf Ul Zuhra [Inria, from Jul. 2020]

PhD Students

- Younes Ben Mazziane [Univ Côte d'Azur, from Oct. 2020]
- Olha Chuchuk [Organisation européenne pour la recherche nucléaire (CERN), from Sep. 2020]
- Mandar Datar [Inria]
- Maximilien Dreveton [Inria]
- Guilherme Iecker Ricardo [Institut Telecom]
- Mikhail Kamalov [Inria]
- Othmane Marfoq [Inria]
- Marie Masson [Orange Labs, CIFRE, until Oct. 2020]
- Tareq Si Salem [Univ Côte d'Azur]

Technical Staff

- Samaresh Bera [Inria, Engineer, until Apr. 2020]
- Ghilas Ferrat [Inria, Technician, until July 2020]
- Chuan Xu [Inria, Engineer, from Mar. 2020]

Interns and Apprentices

- Younes Ben Mazziane [Univ Côte d'Azur, from Mar. 2020 until Aug. 2020]
- Haoran Ding [Inria, from Mar. 2020 until Aug. 2020]
- Abdelkarim Hafid [Univ Côte d'Azur, from Mar. 2020 until Aug. 2020]
- Vidhya Kannan [Univ Côte d'Azur, until May 2020]
- Lucas Lopes Felipe [Inria, from Jan 2020 until Apr. 2020]
- Kaiyun Pan [Univ Côte d'Azur, from Mar. 2020 until Aug. 2020]

Administrative Assistant

- Laurie Vermeersch [Inria, until Nov 2020]

Visiting Scientists

- Vladimir Mazalov [Université d'état de Petrozavodsk - Russie, Feb. 2020]
- Maksim Mironov [Institut de physique et de technique Moscou - Russie, from Feb. 2020 until Mar. 2020]
- Philippe Nain [Inria, long-term visit from team DANTE]
- Seyedeh Haleh Seyed Dizaji [Université de Tabriz - Iran, from Feb. 2020 until Sep. 2020]

External Collaborator

- Patrick Brown [Orange]

2 Overall objectives

NEO is an Inria project-team whose members are located in Sophia Antipolis (S. Alouf, K. Avrachenkov, G. Neglia, S. Perlaza), in Avignon (E. Altman) at LIA (Lab. of Informatics of Avignon) and in Montpellier (A. Jean-Marie) at LIRMM (Lab. Informatics, Robotics and Microelectronics of Montpellier). 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

3.1 Stochastic Operations Research

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, and 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

- S. M. Perlaza has been appointed Visiting Research Collaborator in the EE Dept. at Princeton U., NJ, USA.
- One of the teaching projects of E. Altman received the best paper award at WINCOM [39] .

6 New software and platforms

6.1 New software

6.1.1 marmoteCore

Name: Markov Modeling Tools and Environments - the Core

Keywords: Modeling, Stochastic models, Markov model

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

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

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

Publications: hal-01651940, hal-01276456

Contacts: Alain Jean-Marie, Jean-Marc Vincent

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

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

7 New results

7.1 Stochastic Modeling

Participants Sara Alouf, Eitan Altman, Alain Jean-Marie, Samir M. Perlaza.

7.1.1 Random surfers and prefetching

Prefetching is a basic technique used to reduce the latency of diverse computer services. Deciding what to prefetch amounts to making a compromise between latency and the waste of resources (network bandwidth, storage, energy) if contents are mistakenly prefetched. We have pursued the analysis of the situation where the graph representing the different documents and their links is not completely known in advance. K. Keshava, under the supervision of S. Alouf and A. Jean-Marie, has studied a model where a tree, of depth one or two, is completed randomly after each movement by a uniform branching process.

He has identified the optimal prefetching policy for a finite-horizon situation, in the case where the prefetching budget is one or two documents at each round. This optimal policy appears to be nontrivial.

7.1.2 Data Injection Attacks

State estimation enables efficient, scalable, and secure operation of power systems. Monitoring and control processes are supported by supervisory control and data acquisition (SCADA) systems and more recently by advanced communication systems that acquire and transmit observations to a state estimator. This cyber-layer exposes the system to malicious attacks that exploit the vulnerabilities of the sensing and communication infrastructure solutions. One of the main threats faced by modern power systems are data injection attacks (DIAs) that alter the state estimate of the operator by compromising the system observations. In [41], S. M. Perlaza, X. Ye, I. Esnaola, and Robert F. Harrison (all from U. of Sheffield, UK) presented sparse attacks that minimize simultaneously the information obtained by the operator and the probability of detection. When the assumption on the sparsity is dropped, S. M. Perlaza, S. Ke and I. Esnaola have studied these attacks in [20]. An overview on data injection attacks is presented by the same authors in a book chapter [44].

7.1.3 Epidemic Surveillance with Non-Reliable Diagnostics

In the absence of a vaccination or effective medical treatment against the pandemics such as the SARS-CoV-2, the global population must cohabitate with these viruses. For succeeding in this task, different strategies to slow down the outbreak can be implemented, for example, encouraging social distancing, isolation of infected individuals, mobility restrictions, lockdowns, and contact tracing. The main objective is to guarantee that the number of infected individuals that develop critical forms of symptoms does not exceed the capacity of local health care systems. Nonetheless, most of the strategies to slow down the outbreak induce dramatic economical consequences, and thus, public health policies must be designed based on reliable predictions of the evolution of the pandemic to minimize undesired effects on the global economy. For doing so, estimating the values of variables such as the proportion of susceptible, infected and recovered individuals in the population, among other variables, is of paramount importance. In [46], S. M. Perlaza, E. Altman, I. Mounir (CHU de Nice), and F.Z. Najid (CHU d'Amiens) presented a formula for estimating the prevalence ratio of a disease in a population that is tested with imperfect tests. The formula is in terms of the fraction of positive test results and test parameters, i.e., probability of true positives (sensitivity) and the probability of true negatives (specificity).

7.1.4 The `marmoteCore` platform

The `marmoteCore` platform has been upgraded and moved to a new GitLab development environment, with the help of Inria's Service Experimentation Development. It is currently being ported from `C++` to `Python`. It is hosting the new RLGL method for Markov chain developed in the team. Its use for the analysis of protein-protein interaction networks, via Markov chains with restart, has been continued in collaboration with the project-team ABS.

7.2 Random Graph and Random Matrix Models

Participants Eitan Altman, Konstantin Avrachenkov, Andrei Bobu.

7.2.1 Random geometric graphs

Random geometric graphs have become now a popular object of research. Defined rather simply, these graphs describe real networks much better than classical Erdős–Rényi graphs due to their ability to produce tightly connected communities. The n vertices of a random geometric graph are points in d-dimensional Euclidean space, and two vertices are adjacent if they are close to each other. Many properties of these graphs have been revealed in the case when d is fixed. However, the case of growing dimension d is practically unexplored. This regime corresponds to a real-life situation when one has a data set of n observations with a significant number of features, a quite common case in data science today. In [13], K. Avrachenkov and A. Bobu study the clique structure of random geometric graphs when $n \to \infty$, and $d \to \infty$, and the average vertex degree grows significantly slower than n. They show that under these conditions, random geometric graphs do not contain cliques of size 4 a.s. if only $d >> \log^{1+\epsilon}(n)$. As for the cliques of size 3, we present new bounds on the expected number of triangles in the case $\log^2(n) << d << log^3(n)$ that improve previously known results. In addition, they provide new numerical results showing that the underlying geometry can be detected using the number of triangles even for small n.

7.2.2 Epidemic model and control

We have pursued our work on epidemiology in two application areas. The first has been to fight malware, e-viruses and e-worms which has been among our central research themes during the 10 last years, and developing tools to fight Covid 2019. Herd immunity, one of the most fundamental concepts in network epidemics, occurs when a large fraction of the population of devices is immune against a virus or malware. The few individuals who have not taken countermeasures against the threat are assumed to have very low chances of infection, as they are indirectly protected by the rest of the devices in the network. Although very fundamental, herd immunity does not account for strategic attackers scanning the network for vulnerable nodes. In face of such attackers, nodes who linger vulnerable in the network become easy targets, compromising cybersecurity. In [19], V. Rufino, D. Menasche and C. Lima from UFRJ, in collaboration with I. Cunha from UFMG, E. Altman, R. El-Azouzi and F. de Pellegrini from Avignon Univ and L. de Aguiar and A. Avritzer from Siemens, M. Grottke from Friedrich-Alexander U., propose an analytical model which allows us to capture the impact of countermeasures against attackers. Their model suggests that nodes should adopt countermeasures even when the remainder of the nodes has already decided to do so.

7.3 Data Analysis and Learning

Participants Konstantin Avrachenkov, Maximilien Dreveton, Othmane Marfoq, Giovanni Neglia, Chuan Xu.

7.3.1 Graph clustering

In [23] K. Avrachenkov together with M. Mironov (MIPT, Russia) consider a graph clustering problem with a given number of clusters and approximate desired sizes of the clusters. One possible motivation for such task could be the problem of databases or servers allocation within several given large computational clusters, where one wants related objects to share the same cluster in order to minimize latency and transaction costs. This task differs from the original community detection problem. To solve this task, the authors adopt some ideas from Glauber Dynamics and the Label Propagation Algorithm. At the same time they consider no additional information about node labels, so the task has the nature of unsupervised learning. They propose an algorithm for this problem, show that it works well for a large set of parameters of Stochastic Block Model (SBM) and theoretically show that its running time complexity for achieving almost exact recovery is of $O(n \cdot d \cdot \omega)$ for the mean-field SBM with d being the average degree and ω tending to infinity arbitrary slow. Other significant advantage of the proposed approach is its local nature, which means it can be efficiently distributed with no scheduling or synchronization.

7.3.2 Semi-supervised learning

Due to high interest in many applications, from social networks to blockchain to power grids, deep learning on non-Euclidean objects such as graphs and manifolds, coined Geometric Deep Learning (GDL), continues to gain an ever increasing interest. In [26] K. Avrachenkov together with Y. Chen (SMU, USA) and Y. Gel (U. Texas at Dallas, USA) propose a new Lévy Flights Graph Convolutional Networks (LFGCN) method for semi-supervised learning, which casts the Lévy Flights into random walks on graphs and, as a result, allows both to accurately account for the intrinsic graph topology and to substantially improve classification performance, especially for heterogeneous graphs. Furthermore, they propose a new preferential P-DropEdge method based on the Girvan-Newman argument. That is, in contrast to uniform removing of edges as in DropEdge, following the Girvan-Newman algorithm, we detect network periphery structures using information on edge betweenness and then remove edges according to their betweenness centrality. Their experimental results on semi-supervised node classification tasks demonstrate that the LFGCN coupled with P-DropEdge accelerates the training task, increases stability and further improves predictive accuracy of learned graph topology structure.

7.3.3 Speeding-up machine learning

The most popular framework for parallel training of machine learning models in a cluster is the (synchronous) parameter server (PS). This paradigm consists of n workers and a stateful parameter server, which waits for the responses of every worker's computation to proceed to the next iteration. Transient computation slowdowns or transmission delays can intolerably lengthen the time of each iteration. An efficient way to mitigate this problem is to let the PS wait only for the fastest $n-b$ updates, before generating the new parameters. The slowest b workers are called backup workers. In [40], C. Xu and G. Neglia, together with N. Sebastianelli (U. Côte d'Azur, France) show that the optimal number b of backup workers depends not only on the cluster configuration and workload, but also on the current stage of the training. They propose DBW, an algorithm that dynamically decides the number of backup workers during the training process to maximize the convergence speed at each iteration. Their experiments show that DBW 1) removes the necessity to tune b by preliminary time-consuming experiments, and 2) makes the training up to a factor 3 faster than the optimal static configuration.

The client-server architecture is also adopted in federated learning (that is distributed training at the scale of the Internet). The parameter server is often called orchestrator. This approach may be inefficient in cross-silo settings, as close-by data silos with high-speed access links may exchange information between them faster than with the orchestrator, and the orchestrator may become a communication bottleneck. In [33], O. Marfoq, C. Xu, and G. Neglia, together with R. Vidal (Accenture Labs, France) define the problem of topology design for cross-silo federated learning using the theory of max-plus linear systems to compute the system throughput—number of communication rounds per time unit. They also propose practical algorithms that, under the knowledge of measurable network characteristics, find a topology with the largest throughput or with provable throughput guarantees. In realistic Internet networks with 10 Gbps access links for silos, these algorithms speed up training by a factor 9 and 1.5 in comparison to the client-server architecture and to state-of-the-art MATCHA, respectively. Speedups are even larger with slower access links.

7.4 Game Theory

Participants Eitan Altman, Konstantin Avrachenkov, Mandar Datar, Ghilas Ferrat, Alain Jean-Marie, Samir M. Perlaza.

7.4.1 Discrete-time stochastic replicator dynamics

The population dynamics for the replicator equation are well studied in continuous time but there is less work that explicitly considers the evolution in discrete time. The discrete-time dynamics can often be justified indirectly, by establishing the relevant evolutionary dynamics for the corresponding continuous-time system, and then appealing to an appropriate approximation property. In [9] K. Avrachenkov together with A. Albrecht, P. Howlett and G. Verma (U. South Australia, Australia) study the discrete-time system directly, and establish basic stability results for the evolution of a population defined by a positive definite system matrix, where the population is disrupted by random perturbations to the genotype distribution either through migration or mutation, in each successive generation. One very interesting conclusion is that the replicator dynamics in discrete time is much more stable than in continuous time.

7.4.2 Two phases policies in dynamic opinion games

S. Dhamal (Chalmers U. of Technology, Sweden) and E. Altman, in collaboration with W. Ben-Ameur and T. Chahed from Institut Polytechnique de Paris propose in [15] a setting for two-phase opinion dynamic investment in social networks, where at each step, the final opinion of a node in the first phase acts as its initial biased opinion in the second phase. In this setting, we study the problem of competing camps aiming to maximize adoption of their respective opinions by strategically investing on nodes, where the effectiveness of a camp's investment on a node depends on the node's initial bias. We propose an extension of Friedkin-Johnsen model for our setting, and hence formulate the utility functions of the camps. We show the existence of Nash equilibria under reasonable assumptions, and that they can be computed in polynomial time. Our main conclusion is that, if nodes attribute high weightage to their initial biases, it is advantageous to have a high investment in the first phase, so as to exploit the manipulated biases in the second phase.

7.4.3 Routing games on loss networks

B. Toure and S. Paturel have followed in their third year of study in Centrale-Supéléc a 6-hours course for preparation to research. They published their research project [39] on routing games with losses in an International IEEE sponsored conference on wireless communications (WINCOM). Their work under the guidance of E. Altman received the best paper award of the conference. The loss criterion studied here is quite challenging as it is not additive and flow is not conserved. Further results on the structure of the equilibrium are obtained in [28] by E. Altman, M. Datar and G. Ferrat for the parallel link topology.

7.5 Applications in Telecommunications

Participants Eitan Altman, Konstantin Avrachenkov, Ghilas Ferrat, Marie Masson, Philippe Nain, Giovanni Neglia, Samir M. Perlaza, Guilherme Iecker Ricardo, Sadaf Ul Zuhra.

7.5.1 Elastic cloud caching services

Cloud operators now offer data storage that can be dynamically configured over short timescales (minutes). In [14] G. Neglia, D. Carra (U. of Verona, Italy), and P. Michiardi (EURECOM, France) continue their work on elastic resource provisioning in the cloud, focusing on in-memory key-value stores used as caches. The goal is to dynamically scale resources to the traffic pattern minimizing the overall cost, which includes not only the storage cost, but also the cost due to misses. In fact, a small variation of the cache miss ratio may have a significant impact on user perceived performance in modern web services, which in turn has an impact on the overall revenues for the content provider using such services. They propose and study a dynamic algorithm for TTL (Time To Live) caches, which is able to obtain close-to-minimal costs. On real-world traces collected from Akamai, they show that the TTL approach is able to track the optimal cache configuration and achieve significant cost savings specially in highly dynamic settings that are likely to require elastic cloud services.

7.5.2 Miss ratio curves computation

The Miss Ratio Curve (MRC) represents a fundamental tool for cache performance profiling. Approximate methods based on sampling provide a low-complexity solution for MRC construction. In [25], G. Neglia and D. Carra (U. of Verona, Italy) show that, in case of content with a large variance in popularity, the approximate MRC may be highly sensitive to the set of sampled content. They study in detail the impact of content popularity heterogeneity on the accuracy of the approximate MRC and observe that few, highly popular, items may cause large error at the head of the reconstructed MRC. From these observations, they design a new approach for building an approximate MRC, where they combine an exact portion of the MRC with an approximate one built from samples. This new algorithm computes MRC with an error up to 10 times smaller than state-of-the-art methods based on sampling, with similar computational and space overhead.

7.5.3 Bounding caching performance

Many cache eviction policies have been proposed and implemented to improve the hit probability. In [37], G. Neglia, P. Nain (Inria DANTE), N. Panigrahy and D. Towsley (UMass at Amherst, USA), propose a new method to compute an upper bound on hit probability for all non-anticipative caching policies, i.e., for policies that have no knowledge of future requests. At each object request arrival, they use hazard rate (HR) function based ordering to classify the request as a hit or not. Under some statistical assumptions, they prove that the proposed HR based ordering model computes the maximum achievable hit probability and serves as an upper bound for all non-anticipative caching policies. In simulation, they find this approach to be almost always tighter than Belady's bound.

7.5.4 Caching in dense cellular networks

In 5G and beyond network architectures, operators and content providers base their content distribution strategies on Heterogeneous Networks, where macro and small(er) cells are combined to offer better Quality of Service (QoS) to wireless users. On top of such networks, edge caching and Coordinated Multi-Point (CoMP) transmissions are used to further improve performance. The problem of optimally utilizing the cache space in dense and heterogeneous cell networks has been extensively studied under the name of "FemtoCaching". However, the literature usually assumes relatively simple physical layer (PHY) setups and known or stationary content popularity. In [43], G. I. Ricardo and G. Neglia, together with T. Spyropoulos (EURECOM, France), address these issues by proposing a class of fully distributed and dynamic caching algorithms that take advantage of CoMP capabilities towards minimizing PHY-aware metrics, such as end-to-end (E2E) delay. These policies outperform existing dynamic solutions that are PHY-unaware, under both synthetic and real (non-stationary) request processes, and converge to efficient centralized solutions, in static setups.

7.5.5 Simultaneous information and energy transmission

Battery dependency is a critical issue when communications systems are deployed in hard-to-reach locations, e.g., remote geographical areas, concrete structures, human bodies, or disaster/war zones. In this case, the lifetime of the electronic devices or even the whole communications system is determined by the battery life. An effective remedy is using energy harvesting technologies. Specifically, energy can be harvested from different ambient sources such as light, vibrations, heat, chemical reactions, physiological processes, or the radio frequency (RF) signals produced by other communications systems. This observation rises the idea of simultaneous information and energy transmission (SIET) via RF. S. U. Zuhra, S. M. Perlaza, and E. Altman have studied the fundamental limits on the rates at which information and energy can be simultaneously transmitted over an additive white Gaussian noise channel. The underlying assumption is that the number of channel input symbols (constellation size) is finite. The main results are mathematical expressions of the achievable and converse information-energy rates as a function of the constellation size, number of channel uses, decoding error probability, and energy-outage probability. As a by product, guidelines for optimal constellation design for SIET are obtained in terms of all real-system implementation parameters.

7.5.6 Minimum decoding error probability

Providing an upper bound on the minimum decoding error probability (DEP) in point-to-point memoryless channels at a fixed information rate and fixed transmission duration, e.g. number of channel uses, is central in the analysis of communications systems. Nonethless, only lower and upper bounds on the DEP are available in current literature. Moreover, such existing bounds are difficult to calculate as they involve dealing with the tails of cumulative distribution functions (c.d.f.) of n-dimensional random vectors. In [12], S. M. Perlaza, D. Anade (Inria, MARACAS), P. Mary (INSA de Rennes), and J.-M. Gorce (Inria, MARACAS) have presented an upper bound on the absolute difference between: (a) the cumulative distribution function (c.d.f.) of the sum of a finite number of independent and identically distributed random variables; and (b) a saddlepoint approximation of such c.d.f. This upper bound is general and particularly precise in the regime of large deviations. This result is used to study the DEP in [22].

7.5.7 5G cellular networks

Macro cell densification with Small Cells (SCs) is an effective solution to cope with traffic increase. To fully benefit from the additional SCs capacity, interference mitigation techniques are needed. Densification in 5G networks with Massive Multiple Input Multiple Output (M-MIMO) deployment needs to rethink interference mitigation to account for highly focused beams and MultiUser (MU) scheduling. M. Masson and E. Altman, in collaboration with Z. Altman (Orange Labs) present in [35] a low complexity collaborative Proportional Fair (PF) based scheduling that maximizes the throughput and improves fairness of the heterogeneous network. The solution is based on the calculation of a loss factor indicator that each SC provides to the macro cell at each scheduling period. These indicators allow the macro cell MU scheduler to efficiently select the set of users for scheduling, leading to a significant improvement in performance. Numerical results illustrate the interest of the collaborative solution.

F. de Pellegrini, F. Faticanti and D. Siracusa from U. Avignon and FBK in collaboration with E. Altman and M. Datar, study in [29] the tradeoff between running cost and processing delay in order to optimally orchestrate multiple fog applications. Fog applications process batches of objects' data along chains of containerised microservice modules, which can run either for free on a local fog server or run in cloud at a cost. Processor sharing techniques, in turn, affect the applications' processing delay on a local edge server depending on the number of application modules running on the same server. The fog orchestrator copes with local server congestion by offloading part of computation to the cloud trading off processing delay for a finite budget. Such problem can be described in a convex optimisation framework valid for a large class of processor sharing techniques. We show that the optimal solution is in threshold form and depends solely on the order induced by the marginal delays of N fog applications.

Slicing is emerging as a promising technique to support new differentiated services in 5G networks. It provides the necessary flexibility and scalability associated with future services. To maintain satisfactory services requirements and high profit for service providers, a slice may be designing according to the varying demands and resource availability. This paper develops a framework for resources allocation between slicing and business layer for multi-tenant slicing, e.g. virtual wireless operators, service providers and smart cities services. In [27], M. Datar, C. Touati and E. Altman (INRIA) in collaboration with F. de Pellegrini and R. El-Azouzi (LIA) propose a flexible mechanism based on a bidding scheme for slicing allocation, which achieve desirable fairness and efficiency among the network slices of the different tenants and their associated users. We then design a practical algorithms to realise the proposed desired solution. We also show through simulation the efficiency of our approach in terms of efficiency and fairness.

7.5.8 Energy saving protocols in 5G

Advanced Sleep Modes (ASM) correspond to a gradual deactivation of the base station's components according to the time needed by each of them to shut down then reactivate again. Each level of sleep has a different power consumption and imposes an extra delay on arriving traffic as it has to wait for the components to wake up and serve it. F. Ezzahra Salem, A. Gati, Z. Altman (Orange Labs) in collaboration with T. Chahed (Institut Polytechnique de Paris) and E. Altman, present in [38] a scalable management strategy of this feature based on Markov Decision Processes in order to derive the optimal policy allowing to choose the best sleep level according to the traffic load and to the tradeoff between delay and energy consumption while ensuring a low complexity. Our results show that this solution is very promising and allows to achieve high energy saving (up to 91%) if there is no constraint on the delay. Even with a high constraint, the energy reduction can reach up to 52% while the impact on the delay is negligible.

7.5.9 Fundamental Bounds on amount of transmitted information

Mobile phones rely on batteries to provide the power needed for transmission and for reception (up and downlink communications). Considering uplink, E. Altman, M. Datar and G. Ferrat analyse in [28] how the characteristics of the battery affect the amount of information that one can draw out from the terminal. We focus in particular on the impact of the charge in the battery on the internal resistance which grows as the battery depletes.

7.6 Applications in Social Networks

Participants Eitan Altman, Konstantin Avrachenkov, Mikhail Kamalov, Giovanni Neglia, Kishor Yashavant Patil.

7.6.1 Deep learning framework for graphlet counting

Graphlet counting is a widely-explored problem in network analysis and has been successfully applied to a variety of applications in many domains, most notably bioinformatics, social science and infrastructure network studies. Efficiently computing graphlet counts remains challenging due to the combinatorial explosion, where a naive enumeration algorithm needs $O(Nk)$ time for k-node graphlets in a network of size N. Recently, many works introduced carefully designed combinatorial approximations and sampling methods with encouraging results. However, the existing methods ignore the fact that graphlet counts and the graph structural information are correlated. They always consider a graph as a new input and repeat the tedious counting procedure on a regular basis even if it is similar or exactly isomorphic to previously studied graphs. This provides an opportunity to speed up the graphlet count estimation procedure by exploiting this correlation via learning methods. In [18] K. Avrachenkov together with X. Liu, Y.-Z. Chen and J.C.S. Lui (Chinese U. of Hong Kong) raise a novel Graphlet Count Learning (GCL) problem: given a set of historical graphs with known graphlet counts, how to learn to estimate/predict graphlet count for unseen graphs coming from the same (or similar) underlying distribution. They develop a deep learning framework which contains two *convolutional neural network* (CNN) models and a series of data *preprocessing techniques* to solve the GCL problem. Extensive experiments are conducted on three types of synthetic random graphs and three types of real world graphs for all 3,4,5-node graphlets to demonstrate the accuracy, efficiency and generalizability of their framework.

7.6.2 Generative PageRank framework for semi-supervised learning on citation graphs

Nowadays, Semi-Supervised Learning (SSL) on citation graph data sets is a rapidly growing area of research. However, the recently proposed graph-based SSL algorithms use a default adjacency matrix with binary weights on edges (citations), that causes a loss of the nodes (papers) similarity information. In [31] M. Kamalov and K. Avrachenkov propose a framework focused on embedding PageRank SSL in a generative model. This framework allows one to do joint training of nodes latent space representation and label spreading through the reweighted adjacency matrix by node similarities in the latent space. They explain that a generative model can improve accuracy and reduce the number of iteration steps for PageRank SSL. Moreover, we show that our framework outperforms the best graph-based SSL algorithms on four public citation graph data sets and improves the interpretability of classification results.

7.6.3 Web page crawling

For providing quick and accurate results, a search engine maintains a local snapshot of the entire web. And, to keep this local cache fresh, it employs a crawler for tracking changes across various web pages. However, finite bandwidth availability and server restrictions impose some constraints on the crawling frequency. Consequently, the ideal crawling rates are the ones that maximise the freshness of the local cache and also respect the above constraints. Azar et al. 2018 recently proposed a tractable algorithm to solve this optimisation problem. However, they assume the knowledge of the exact page change rates, which is unrealistic in practice. K. Avrachenkov, K. Patil and G. Thoppe (IISc Bangalore, India) address this issue in [24]. Specifically, they provide two novel schemes for online estimation of page change rates. Both schemes only need partial information about the page change process, i.e., the schemes only need to know if the page has changed or not since the last crawled instance. For both these schemes, the authors prove convergence and, also, derive their convergence rates. Finally, they provide some numerical experiments to compare the performance of the estimators they proposed with the existing ones (e.g., MLE).

7.7 Applications to Energy and Environmental Issues

Participants Sara Alouf, Alain Jean-Marie.

7.7.1 Short-time-scale modeling of solar irradiance

S. Alouf and A. Jean-Marie have proposed in [10] an efficient forecasting method for solar irradiance. The model is a stochastic process at the minute scale, whose parameters depend on both the period within the day and a category of day (sunny, cloudy, changing...). The categories and the distributions of interest are identified for a given location by clustering observed data. Then, based on the weather forecast for next day, the proper day category is selected and random trajectories can be generated. Experiments show that this new model outperforms several previous proposals when the solar input is used in a datacenter model, consisting of a storage, and an energy consumption corresponding to a real workload.

8 Bilateral contracts and grants with industry

8.1 Bilateral contracts with industry

NEO members are involved in the

- Inria-Nokia Bell Labs joint laboratory: the joint laboratory consists of five ADRs (Action de Recherche/ Research Action) in its third phase (starting October 2017). NEO members participate in two ADRs: "Distributed Learning and Control for Network Analysis" (see §8.1.1) and "Rethinking the network: virtualizing network functions, from middleboxes to application" (see §8.1.2);
- Inria-QWANT joint laboratory "Smart search is privacy" (see §8.1.3);
- Inria-Orange Labs joint laboratory (see §8.1.4).

NEO has contracts with Accenture (see §8.1.5), Azursoft (see §8.1.6), MyDataModels (see §8.1.7) and Payback Network (see §8.1.8).

8.1.1 ADR Nokia on the topic "Distributed Learning and Control for Network Analysis" (October 2017 – September 2021)

Participants Eitan Altman, Konstantin Avrachenkov, Mandar Datar, Maximilien Dreveton.

- Contractor: Nokia Bell Labs (http://www.bell-labs.com)
- Collaborator: Gérard Burnside

Over the last few years, research in computer science has shifted focus to machine learning methods for the analysis of increasingly large amounts of user data. As the research community has sought to optimize the methods for sparse data and high-dimensional data, more recently new problems have emerged, particularly from a networking perspective that had remained in the periphery.

The technical program of this ADR consists of three parts: Distributed machine learning, Multiobjective optimisation as a lexicographic problem, and Use cases / Applications. We address the challenges related to the first part by developing distributed optimization tools that reduce communication overhead, improve the rate of convergence and are scalable. Graph-theoretic tools including spectral analysis, graph partitioning and clustering will be developed. Further, stochastic approximation methods and D-iterations or their combinations will be applied in designing fast online unsupervised, supervised and semi-supervised learning methods.

8.1.2 ADR Nokia on the topic "Rethinking the network: virtualizing network functions, from middleboxes to application" (October 2017 – September 2021)

Participants Sara Alouf, Giovanni Neglia.

- Contractor: Nokia Bell Labs (http://www.bell-labs.com)
- Collaborators: Fabio Pianese, Massimo Gallo

A growing number of network infrastructures are being presently considered for a software-based replacement: these range from fixed and wireless access functions to carrier-grade middle boxes and server functionalities. On the one hand, performance requirements of such applications call for an increased level of software optimization and hardware acceleration. On the other hand, customization and modularity at all layers of the protocol stack are required to support such a wide range of functions. In this scope the ADR focuses on two specific research axes: (1) the design, implementation and evaluation of a modular NFV architecture, and (2) the modelling and management of applications as virtualized network functions. Our interest is in low-latency machine learning prediction services and in particular how the quality of the predictions can be traded off with latency.

8.1.3 Qwant contract on "Asynchronous on-line computation of centrality measures" (December 2017 – May 2020)

Participants Konstantin Avrachenkov, Patrick Brown.

- Contractor: Qwant
- Collaborators: Sylvain Peyronnet, Thomas Aynaud

We shall study asynchronously distributed methods for network centrality computation. The asynchronous distributed methods are very useful because they allow efficient and flexible use of computational resources on the one hand (e.g., using a cluster or a cloud) and on the other hand they allow quick local update of centrality measures without the need to recompute them from scratch.

8.1.4 Orange CIFRE on the topic "Self-organizing features in the virtual 5G radio access network" (November 2017 – October 2020)

Participants Eitan Altman, Marie Masson.

- Contractor: Orange Labs (https://www.orange.com/en/Infographics/Orange-and-Research/Orange-and-Research)
- Collaborator: Zwi Altman

The considerable extent of the complexity of 5G networks and their operation is in contrast with the increasing demands in terms of simplicity and efficiency. This antagonism highlights the critical importance of network management. Self-Organizing Networks (SON), which cover self-configuration, self-optimization and self-repair, play a central role for 5G Radio Access Network (RAN).

This CIFRE thesis aims at innovating in the field of managing 5G RAN, with a special focus on the features of the SON-5G. Three objectives are identified: a) develop self-organizing features (SON in 5G-RAN), b) develop cognitive managing mechanisms for the SON-5G features developed, and c) demonstrate how do the self-organizing mechanisms fit in the virtual RAN.

8.1.5 Accenture contract on the topic “Distributed Machine Learning for IoT applications” (Dec 2019 – May 2020)

Participant Giovanni Neglia.

- Contractor: Accenture Labs (https://www.accenture.com/fr-fr/accenture-lab-sophia-antipolis)
- 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. This 6-month contract prepares a CIFRE.

8.1.6 AzurSoft contract on the topic “Proof of concept on automatic detection of false alarms” (May 2019 – April 2020)

Participants Konstantin Avrachenkov, Andrei Bobu.

- Contractor: AzurSoft (https://www.azursoft.com/)
- Collaborators: Marc Vaillant, Beatrice Escuyer

Intrusion detection or telesurveillance systems generates signals from sensors that allow to raise alarm and start a checking procedure for a potential intrusion or anomaly. Typically, one telesurveillance system surveys many sites and is challenged by a stream of false alarms. In this project, we aim to reduce the rate of false alarms by using various supervised and semi-supervised learning methods.

8.1.7 MyDataModels contract on the topic “Semi supervised variational autoencoders for versatile data” (June 2019 – May 2022)

Participants Konstantin Avrachenkov, Mikhail Kamalov.

- Contractor: MyDataModels (https://www.mydatamodels.com/)
- Collaborators: Denis Bastiment, Carlo Fanara

Variational autoencoders are highly flexible machine learning techniques for learning latent dimension representation. This model is applicable for denoising data as well as for classification purposes. In this thesis we plan to add semi-supervision component to the variational autoencoder techniques. We plan to develop methods which are universally applicable to versatile data such as categorical data, images, texts, etc. Initially starting from static data we aim to extend the methods to time-varying data such as audio, video, time-series, etc. The proposed algorithms can be integrated into the internal engine of MyDataModels company and tested on use cases of MyDataModels.

8.1.8 Consulting contract with Payback Network (November 2019 - January 2020)

Participant Giovanni Neglia.

- Contractor: Payback Network
- Collaborators: Tanguy Racinet, Anne Legencre

Consulting with the startup Payback Network on differential privacy techniques.

9 Partnerships and cooperations

9.1 International initiatives

9.1.1 Participation in other international programs

The Embassy of France in the United States, via the progamme **"make our planet great again"**, has funded an initiative led by S. M. Perlaza and A. Tajer (RPI, USA) for addressing foundational questions pertinent to two emerging wireless communication technologies: (i) energy harvesting (EH) systems, and (ii) ultra low-latency systems for critical missions. This project explores two strongly symbiotic research directions for establishing the fundamental limits of (i) data transmission and, (ii) simultaneous energy and data transmission, in mission critical systems empowered by EH. The expected results have applications in, e.g., disaster relief, medical instruments, cyber-physical systems, and the Internet of things. This program was launched by the President of France Emmanuel Macron in June 2017 and handled by the Embassy of France in the United States of America. The aim of such fellowships is to reinforce the international engagements of the 2015 Paris Agreement on Climate Change by fostering collaborations between scholars in both US and France.

9.2 International research visitors

9.2.1 Visits of international scientists

Professors/Researchers

Vladimir Mazalov, Date : 17-28 February, Petrozavodsk State U. and Karelian Institute for Applied Mathematics (Russia)

Postdoc/PhD Students

Maksim Mironov, Date: 17 February - 13 May, PhD student at MFTI Moscow (Russia). Visit interrupted on 13 March due to lockdown.

Haleh Dizaji, Date: 1 February - 30 September, PhD student at U. of Tabriz (Iran)

Dadja Anade, Date: 6 - 18 July, PhD student at Inria MARACAS.

9.2.2 Internships

Kausthub Keshava, Date: 15 June-15 Dec, Institution: IISER Mohali (India), Supervisors: S. Alouf and A. Jean-Marie. Internship done remotely due to travel restrictions.

9.2.3 Visits to international teams

Research stays abroad

- K. Avrachenkov visited Leiden U. (The Netherlands), 4-7 February 2020, and Aalto U. (Finland), 1-8 March 2020.
- A. Bobu visited U. of Twente (The Netherlands), 24-28 February 2020.
- M. Dreveton visited Aalto U. (Finland), 1-8 March 2020 and 16 October 2020 to 1 November 2020, and visited KTH (Sweden), 2-8 November.
- A. Jean-Marie visited the U. Nacional de Rosario (Argentina), 26 February 2020 to 4 March 2020.

9.3 European initiatives

9.3.1 FP7 & H2020 Projects

TESTBED2

Participants Samir Perlaza, Eitan Altman, Sadaf Ul Zuhra.

Project Acronym: TESTBED2

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

Coordinator: University of Durham, UK.

Duration: February 2020 – January 2024

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 decarbonised 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

PIA ANSWER

Participants Konstantin Avrachenkov, Abhishek Bose, Kishor Yashavant Patil.

Project Acronym: ANSWER

Project Title: Advanced aNd Secured Web Experience and seaRch

Coordinator: QWANT

Duration: 15 November 2017 – 31 December 2020

Others Partners: Inria Project-Teams WIMMICS, INDES, COFFEE

Abstract: ANSWER is a joint project between QWANT and Inria, funded by the French Government's initiative PIA "Programme d'Investissement d'Avenir".

The aim of the ANSWER project is to develop the new version of the search engine http://www.qwant.com by introducing radical innovations in terms of search criteria as well as indexed content and security. This initiative is a part of the Big Data Big Digital Challenges field, since a Web search engine deals with large volumes of heterogeneous and dynamic data.

Of the five characteristics of big data, the ANSWER project will focus more particularly on the aspects of Velocity in terms of near real-time processing of results, and Variety for the integration of new indicators (emotions, sociality, etc.) and meta-data. The Volume, Value and Veracity aspects will necessarily be addressed jointly with these first ones and will also be the subject of locks, especially on the topics of crawling and indexing.

This registration of the search engine in the Big Data domain will only be reinforced by developments in the Web such as the Web of data, and generally by the current trend to integrate the Web of increasingly diverse, rich and complex resources.

ANR MAESTRO5G

Participant Eitan Altman.

Project Acronym: MAESTRO5G

Project Title: MAnagEment of Slices in The Radio access Of 5G networks

Coordinator: Orange Labs

Duration: February 2019 – January 2022

Others Partners: Nokia Bell Labs, U. Avignon, Inria Project-Team AGORA, Sorbonne U., Telecom SudParis, CentraleSupélec.

Abstract: The project develops enablers for implementing and managing slices in the 5G radio access network, not only for the purpose of serving heterogeneous services, but also for dynamic sharing of infrastructure between operators. MAESTRO-5G develops a framework for resource allocation between slices and a business layer for multi-tenant slicing. It provides an orchestration framework based on Software Define Networking that manages resources and virtual functions for slices. A hardware demonstrator brings the slicing concept to reality and showcases the project's innovations.

10 Dissemination

10.1 Promoting scientific activities

10.1.1 Scientific events: organisation

General chair, scientific chair

- E. Altman is chairman of the steering committee of the workshop NetGCoop - Networking Games, Control and Optimisation;

- E. Altman is member of the steering committee of the conferences WiOpt and Valuetools;
- G. Neglia is member of the steering committee of the Workshop on Intelligent Things and Services (InThingS) since November 2020.
- S. M. Perlaza was one of the general chairs of the "Workshop on Resource Allocation, Cooperation and Competition in Wireless Networks (RAWNET)", June 19, 2020, Volvos, Greece.
- G. Neglia is one of the co-chairs of the 1st International Workshop on Intelligent Things and Services (InThingS 2020, online).

10.1.2 Scientific events: selection

Member of the conference program committees

- ACM Sigmetrics 2021 (Beijing, China) (S. Alouf);
- ACM Sigmetrics 2020 (Boston, USA) (S. Alouf);
- EAI International Conference on Performance Evaluation Methodologies and Tools (VALUETOOLS 2020, Tsukuba (online), Japan) (A. Jean-Marie, K. Avrachenkov);
- IEEE 28th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2020, Nice (online), France) (A. Jean-Marie);
- IEEE International Conference on Computer Communications (INFOCOM 2021, Virtual Conference) (G. Neglia);
- AAAI 2021 Workshop Towards Robust, Secure and Efficient Machine Learning (RSEML 2021, Virtual Conference) (G. Neglia);
- IEEE Intl. Conf. on Network Protocols (IEEE ICNP 2020, Madrid, Spain, online) (K. Avrachenkov);
- SIAM Intl. Conf. on Data Mining (SDM20, Cincinnati, Ohio, U.S., online) (K. Avrachenkov);
- The 9th Intl. Conf. on Complex Networks and their Applications 2020 (Madrid, Spain, online) (K. Avrachenkov);
- The 18th Intl. Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt 2020) (Greece, online) (K. Avrachenkov);
- The 17th Workshop on Algorithms and Models for the Web Graph (WAW 2020) (Warsaw, Poland, online) (K. Avrachenkov);
- The 11th Conference on Decision and Game Theory for Security (GameSec 2020), (Maryland, College Park, U.S.) (K. Avrachenkov);
- 22nd Workshop on MAthematical performance Modeling and Analysis (MAMA 2020, Boston (online), USA) (A. Jean-Marie);
- 21st Conf. of the Société Française de Recherche Opérationnelle et d'Aide à la Décision (ROADEF 2020, Montpellier, France) (A. Jean-Marie);
- IEEE International Conference on Communications (ICC) (S. M. Perlaza);
- IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC) (S. M. Perlaza);
- International Conference on Computing, Networking and Communications (ICNC) (S. M. Perlaza);
- IEEE Wireless Communications and Networking Conference (WCNC) (S. M. Perlaza).

10.1.3 Journal

- Guest editor for *IEEE Transactions on Network Science and Engineering* special issue on "Communication-Efficient Distributed Machine Learning" (G. Neglia);
- Guest editor for *IEEE Internet of Things Journal* special Issue on "Artificial Intelligence Powered Edge Computing for the Internet of Things" (S. M. Perlaza).

Member of the editorial boards

- *Wiley Transactions on Emerging Telecommunications Technologies (ETT)* (S. Alouf, until Aug 2020);
- *Frontiers in Communications and Networks* (S. M. Perlaza, Associate Editor);
- *ACM Transactions on Modeling and Perforunimance Evaluation of Computing Systems (ACM ToMPECS)* (K. Avrachenkov);
- *Elsevier Computer Communications (COMCOM)* (G. Neglia);
- *IEEE Network Magazine* (K. Avrachenkov);
- *Elsevier International Journal of Performance Evaluation* (K. Avrachenkov);
- *IEEE Transactions on Mobile Computing* (G. Neglia);
- *IEEE Transactions on Communications* (S. Perlaza, Editor);
- *IET Smart Grids* (S. Perlaza, Editor);
- *Probability in the Engineering and Informational Sciences* (K. Avrachenkov);
- *Stochastic Models* (K. Avrachenkov);
- *Springer: Dynamic Games and Applications* (E. Altman);
- *AIMS: Journal of Dynamic Games* (E. Altman);
- *IEEE/ACM Transaction of Networking* (E. Altman, Editor at large).

Reviewer - reviewing activities NEO members are reviewers for *Environmental Modeling & Assessment*, *IEEE/ACM Transactions on Networking*, *IEEE Transactions on Mobile Computing*, *IEEE Transactions on Network and Service Management*, *IEEE Transactions on Information Theory*, *IEEE Transactions on Wireless Communications*, *IEEE Transactions on Communications*, *IEEE Journal on Selected Topics in Signal Processing*, *IEEE Journal on Selected Areas in Communications*, and many other journals.

10.1.4 Invited talks

- G. Neglia gave a keynote talk on "Machine Learning Training: Research Challenges and Opportunities for the Networking Community" at the Italian Networking Workshop, 29 January 2020, Cavalese, Italy.
- S. M. Perlaza is invited to deliver the talk "An upper bound on the error induced by saddlepoint approximations: Applications to information theory" in the School of Electrical, Computer, and Systems Engineering at Rensselear Polytechnic Institute (RPI), NY, USA. (postponed until further notice)
- K. Avrachenkov gave a keynote talk on "Distributed optimization of caching devices with geographic constraints" at the 2nd Intl. Workshop on Stochastic Modeling and Applied Research of Technology (SMARTY 2020, online) Petrozavodsk, Karelia, Russia on August 16-20, 2020.
- E. Altman was Invited speaker on 10 July 2020 to the workshop in honour of Prof. Anurag Kumar upon his retirement. Virtual presence through visioconf. He was also invited (teleportation) to U. of Pennsylvania's ESE Department, presented two talks on 2020 Fall Seminar, October 20th : 1. On audience retention of youTube, and 2. The evolution of languages in twitter.

10.1.5 Leadership within the scientific community

S. Alouf

- is an elected member at the Board of Directors of ACM SIGMETRICS;
- is a member of the Equality and Diversity committee of ACM SIGMETRICS;
- is a member of the Conference Advisory committee of ACM SIGMETRICS.

E. Altman

- is a fellow member of IEEE (Class of 2010);
- is a fellow member of EAI (Class of 2019);
- is an elected member of IFIP WG 7.3 on "Computer System Modeling";
- is a member of WG 6.3 of IFIP on Performance of Communication Systems.
- Best paper award at the conference WINCOM 2020.

K. Avrachenkov

- is a member of the scientific committee for Labex UCN@Sophia;
- is a member of Conseil Scientifique & Pédagogique EUR DS4H Université Côte d'Azur.

A. Jean-Marie

- was a member of the Steering Committee of the GDR RO, a national research initiative on Operations Research sponsored by the CNRS, until Dec. 2020;
- is an elected member of IFIP WG 7.3 on "Computer System Modeling".

G. Neglia

- is member of the scientific animation committee for the IDEX UCA$^{\mathrm{JEDI}}$ research program "Social Interactions and Complex Dynamics" since 2017.

S. M. Perlaza

- is a Senior Member of the IEEE (class of 2015);
- is a Marie Sklodowska-Curie Action (MSCA) Fellow. Individual Fellowship, (Class of 2015);
- is a Member of the On-Line Committee of the IEEE Information Theory Society.

10.1.6 Research administration

S. Alouf

- was in the recruitment committee for a *Maître de Conférences* position at U. Savoie Mont Blanc;
- 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.

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);

- is member of the managing sub-committee of the Project-Team Committee of the Inria Sophia Antipolis – Méditerranée research center since December 2017;
- is Head of project-team NEO since January 2017.

G. Neglia

- has been the scientific delegate for European partnerships for Inria Sophia Antipolis – Méditerranée from 2014 to 2020;
- is member of the Inria COST GTRI (International Relations Working Group of Inria's Scientific and Technological Orientation Council) since 2016.

10.2 Teaching - Supervision - Juries

Teaching Note: UCA is the Univ Côte d'Azur.

- Licence : M. Dreveton, "Tools for biology – Statistics", 24H, L1 Life Sciences, UCA, France;
- Master : E. Altman, "Routing Games", 6H, M2, Centrale-Supelec, France;
- Master : K. Avrachenkov, "Statistical Analysis of Graphs", 6H, M2 Data Science, UCA, France;
- Master : M. Dreveton, "Statistical Analysis of Graphs", 26H, M2 Data Science, UCA, France;
- Master : G. Neglia, "Distributed Optimization and Games", 31.5H, M2 Ubinet, UCA, France;
- Master : G. Neglia, "Machine Learning: Theory and Algorithms", 31.5H, M2 Ubinet, UCA, France;
- Master: S. M. Perlaza, "Selected Topics in Information Theory", 30H, M2 Foundations of Computer Science, ENS de Lyon, France.

Supervision

- PhD in progress: Dadja Anade, "Communications in the non-asymptotic block-length regime", INSA Lyon, 1. Oct. 2018, co-advisor: S. M. Perlaza.
- PhD in progress: Younes Ben Mazziane, "Online learning for Caching at the Edge", UCA, 1 Oct. 2020, advisors: Sara Alouf and Giovanni Neglia.
- PhD in progress: Olha Chuchuk, "Optimization of data access at CERN and in the World Large Hadron Collider Computing Grid (WLCG)", UCA, 1 Sept. 2020, advisors: Giovanni Neglia.
- PhD in progress: Mandar Datar, "Singular perturbation approach for machine learning in multiobjective optimisation", Univ Avignon, 1 May 2018, advisor: Eitan Altman.
- PhD in progress: Maximilien Dreveton, "Statistical Physics Methods for Distributed Machine Learning", UCA, 1 Oct. 2018, advisor: Konstantin Avrachenkov.
- PhD in progress: Guilherme Iecker Ricardo, "Caching for wireless networks", UCA, 1 Sept. 2018, advisors: Giovanni Neglia and Pietro Elia (EURECOM).
- PhD in progress: Mikhail Kamalov, "Semi-supervised variational autoencoders versatile data", UCA, 1 June 2019, advisor: Konstantin Avrachenkov.
- PhD in progress: Othmane Marfoq, "Distributed machine learning for IoT applications", UCA, 1 December 2020, advisor: Giovanni Neglia.
- PhD in progress: Marie Masson, "Fonctionnalités auto-organisantes dans le réseau d'accès radio 5G virtuels", UCA, 1 Dec. 2017, advisors: Eitan Altman and Zwi Altman (Orange).
- PhD in progress: Tareq Si Salem, "Federated Learning", UCA, 1 October 2019, advisor: Giovanni Neglia.
- PhD in progress: Xiuzhen Ye, "Data Injection Attacks in Smart Grids", 1 Oct. 2019, U. of Sheffield, co-advisor: S. M. Perlaza.

Juries NEO members participated in the Ph.D. committees of (in alphabetical order):

- Dario Bega, "Deep Learning solutions for next generation slicing-aware mobile networks", U. Carlos III de Madrid, 17 April 2020 (G. Neglia, jury member);
- Santiago Duran, "Resource allocation with observable and unobservable environments", U. Paul Sabatier, 27 February 2020 (K. Avrachenkov, reviewer);
- Dwi Ertinin, "Structural properties of single server queueing systems: Efficient methods via lumping and dynamic programming", U. Leiden, 5 February 2020 (K. Avrachenkov, reviewer);
- Alexis Galland, "Deep learning techniques for graph embedding at different scales", ENS, 17 December 2020 (K. Avrachenkov, jury president);
- Nicola Piovesan, "Network Resource Allocation Policies with Energy Transfer Capabilities", U. Politècnica de Catalunya, 5 June 2020 (G. Neglia, jury member);
- Stephan Plassart, "Online energy optimisation for real-time systems", U. Grenoble Alpes, 16 June (S. Alouf, jury member);
- Samuel Unicomb, "Threshold driven contagion on complex networks", U. Lyon, 14 January (K. Avrachenkov, jury member).

11 Scientific production

11.1 Major publications

[1] 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.

[2] K. Avrachenkov, P. Chebotarev and A. Mishenin. 'Semi-supervised learning with regularized Laplacian'. In: *Optimization Methods and Software* 32.2 (Jan. 2017), pp. 222–236. DOI: 10.1080/10556788.2016.1193176. URL: https://hal.inria.fr/hal-01671800.

[3] 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.

[4] 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.

[5] 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.

[6] 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.

[7] 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.

[8] 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.

11.2 Publications of the year

International journals

[9] A. Albrecht, K. Avrachenkov, P. Howlett and G. Verma. 'Evolutionary dynamics in discrete time for the perturbed positive definite replicator equation'. In: *The ANZIAM Journal* 62.2 (Apr. 2020), pp. 148–184. URL: https://hal.inria.fr/hal-02892186.

[10] S. Alouf and A. Jean-Marie. 'Short-Scale Stochastic Solar Energy Models: A Datacenter Use Case'. In: *Mathematics.* Queue and Stochastic Models for Operations Research 8.12 (Dec. 2020), pp. 1–26. DOI: 10.3390/math8122127. URL: https://hal.inria.fr/hal-03046453.

[11] E. Altman, I. Mounir, F.-Z. Najid and S. M. Perlaza. 'On the true number of COVID-19 infections: Effect of Sensitivity, Specificity and Number of Tests on Prevalence Ratio Estimation'. In: *International Journal of Environmental Research and Public Health* (1st July 2020). URL: https://hal.archives-ouvertes.fr/hal-03128312.

[12] D. Anade, J.-M. Gorce, P. Mary and S. Perlaza. 'An Upper Bound on the Error Induced by Saddlepoint Approximations - Applications to Information Theory'. In: *Entropy.* Wireless Networks: Information Theoretic Perspectives 22.6 (20th June 2020), p. 690. DOI: 10.3390/exx010005. URL: https://hal.archives-ouvertes.fr/hal-02884541.

[13] K. Avrachenkov and A. V. Bobu. 'Cliques in high-dimensional random geometric graphs'. In: *Applied Network Science* 5 (23rd Nov. 2020). DOI: 10.1007/s41109-020-00335-6. URL: https://hal.inria.fr/hal-03023124.

[14] D. Carra, G. Neglia and P. Michiardi. 'Elastic Provisioning of Cloud Caches: A Cost-Aware TTL Approach'. In: *IEEE/ACM Transactions on Networking* 28.3 (June 2020), pp. 1283–1296. DOI: 10.1109/TNET.2020.2980105. URL: https://hal.inria.fr/hal-03078608.

[15] S. Dhamal, W. Ben-Ameur, T. Chahed and E. Altman. 'A two phase investment game for competitive opinion dynamics in social networks'. In: *Information processing & management* 57.2 (Mar. 2020), 102064:1–20. DOI: 10.1016/j.ipm.2019.102064. URL: https://hal.inria.fr/hal-01924838.

[16] K. Fredj, A. Jean-Marie, G. Martín-Herrán and M. Tidball. 'Effects of transaction costs and discount rate on the banking decision of emission permits trading'. In: *Advances in Economics and Business* 8.1 (2020), pp. 63–71. DOI: 10.13189/aeb.2020.080106. URL: https://hal.inria.fr/hal-02372292.

[17] E. Hyon and A. Jean-Marie. 'Optimal control of admission in service in a queue with impatience and setup costs'. In: *Performance Evaluation* 144 (Dec. 2020), p. 102134. DOI: 10.1016/j.peva.2020.102134. URL: https://hal.inria.fr/hal-02961019.

[18] X. Liu, Y.-Z. J. Chen, J. C. S. Lui and K. Avrachenkov. 'Learning to count: A deep learning framework for graphlet count estimation'. In: *Network Science* (Sept. 2020), p. 30. DOI: 10.1017/nws.2020.35. URL: https://hal.inria.fr/hal-02942321.

[19] V. Rufino, D. Menasché, C. Lima, Í. Cunha, L. P. De Aguiar, E. Altman, R. El-Azouzi, F. De Pellegrini, A. Avritzer and M. Grottke. 'Beyond Herd Immunity Against Strategic Attackers'. In: *IEEE Access* (2020). DOI: 10.1109/ACCESS.2017.DOI. URL: https://hal.inria.fr/hal-02522148.

[20] K. Sun, I. Esnaola, S. Perlaza and H. Vincent Poor. 'Stealth Attacks on the Smart Grid'. In: *IEEE Transactions on Smart Grid* (1st Mar. 2020). URL: https://hal.inria.fr/hal-01857366.

International peer-reviewed conferences

[21] E. Altman, M. Datar and G. Ferrat. 'Maximizing amount of transferred traffic for battery powered mobiles'. In: Netgcoop 2020 - International Conference on NETwork Games, Control and Optimisation. Cargese, France, 22nd Sept. 2021. URL: https://hal.inria.fr/hal-02931313.

[22] D. Anade, J.-M. Gorce, P. Mary and S. Perlaza. 'On the saddlepoint approximation of the dependence testing bound in memoryless channels'. In: IEEE International Conference on Communications. Proc. of the International Conference on Communications (ICC). Dublin, Ireland, 7th June 2020, pp. 1–5. URL: https://hal.inria.fr/hal-02457361.

[23] K. Avrachenkov and M. Mironov. 'Cluster-size constrained network partitioning'. In: The 25th International Conference on Pattern Recognition (ICPR 2020). Milano, Italy, 10th Jan. 2021. URL: https://hal.inria.fr/hal-02972577.

[24] K. Avrachenkov, K. Patil and G. Thoppe. 'Change Rate Estimation and Optimal Freshness in Web Page Crawling'. In: ValueTools 2020: The 13th EAI International Conference on Performance Evaluation Methodologies and Tools. Tsukuba, Japan, 18th May 2020. URL: https://hal.inria.fr/hal-03123809.

[25] D. Carra and G. Neglia. 'Efficient Miss Ratio Curve Computation for Heterogeneous Content Popularity'. In: ATC'20 - USENIX Annual Technical Conference. online conference, Unknown Region, 15th July 2020. URL: https://hal.inria.fr/hal-03078642.

[26] Y. Chen, Y. R. Gel and K. Avrachenkov. 'LFGCN: Levitating over Graphs with Levy Flights'. In: IEEE ICDM 2020 - International Conference on Data Mining. Sorrento, Italy, 17th Nov. 2020. URL: https://hal.inria.fr/hal-03123684.

[27] M. Datar, E. Altman, F. De Pellegrini, R. E. Azouzi and C. Touati. 'A Mechanism for Price Differentiation and Slicing in Wireless Networks'. In: WiOpt 2020 - 18th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks. Volos, Greece, June 2020, pp. 1–8. URL: https://hal.inria.fr/hal-02931452.

[28] M. Datar, E. Altman and G. Ferrat. 'Routing into parallel collision channels'. In: MAMA 2020 - Workshop in conjunction with ACM/Sigmetrics 2020. Boston / Virtual, United States, 8th June 2020. URL: https://hal.inria.fr/hal-02931301.

[29] F. De Pellegrini, F. Faticanti, M. Datar, E. Altman and D. Siracusa. 'Optimal Blind and Adaptive Fog Orchestration under Local Processor Sharing'. In: RAWNET 2020 - 15th Workshop on Resource Allocation, Cooperation and Competition in Wireless Networks in conjunction with WiOPT 2020 - 18th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks, Volos, Greece, 19th June 2020. URL: https://hal.inria.fr/hal-02931451.

[30] M. Garetto, E. Leonardi and G. Neglia. 'Similarity Caching: Theory and Algorithms'. In: IEEE INFOCOM 2020 - IEEE Conference on Computer Communications. Beijing, China, 27th Apr. 2020. URL: https://hal.inria.fr/hal-02411268.

[31] M. Kamalov and K. Avrachenkov. 'GenPR: Generative PageRank Framework for Semi-supervised Learning on Citation Graphs'. In: INL 2020 - 9th Conference on Artificial Intelligence and Natural Language. Helsinki, Finland, 30th Sept. 2020, pp. 158–165. DOI: 10.1007/978-3-030-59082-6_12. URL: https://hal.inria.fr/hal-02977308.

[32] V. Kavitha and E. Altman. 'Controlling Packet Drops to Improve Freshness of information'. In: Netgcoop 2020 - International Conference on NETwork Games, Control and Optimisation. Cargese, France, 22nd Sept. 2021. URL: https://hal.inria.fr/hal-02931314.

[33] O. Marfoq, C. Xu, G. Neglia and R. Vidal. 'Throughput-Optimal Topology Design for Cross-Silo Federated Learning'. In: NeurIPS 2020 - 34th Conference on Neural Information Processing Systems. Vancouver / Online, Canada, 6th Dec. 2020. URL: https://hal.inria.fr/hal-03007834.

[34] M. Masson, Z. Altman and E. Altman. 'Coordinated scheduling based on automatic neighbor beam relation'. In: NETGCOOP 2020. Cargèse, France, 2021. URL: https://hal.inria.fr/hal-02566350.

[35] M. Masson, Z. Altman and E. Altman. 'Multi-User collaborative scheduling in 5G massive MIMO heterogeneous networks'. In: IFIP Networking 2020. Paris / Online, France, 22nd June 2020. URL: https://hal.inria.fr/hal-02566253.

[36] G. Neglia, C. Xu, D. Towsley and G. Calbi. 'Decentralized gradient methods: does topology matter?' In: AISTATS 2020 - 23rd International Conference on Artificial Intelligence and Statistics. Palermo /Online, Italy, 26th Aug. 2020. URL: https://hal.inria.fr/hal-02430485.

[37] N. K. Panigrahy, P. Nain, G. Neglia and D. Towsley. 'A New Upper Bound on Cache Hit Probability for Non-anticipative Caching Policies'. In: Performance 2020 - 38th International Symposium on Computer Performance, Modeling, Measurements and Evaluation. Milan / Virtual, Italy, 2nd Nov. 2020, pp. 1–6. URL: https://hal.inria.fr/hal-02987388.

[38] F. E. Salem, T. Chahed, E. Altman, A. Gati and Z. Altman. 'Scalable Markov decision process model for advanced sleep modes management in 5G networks'. In: *VALUETOOLS '20: Proceedings of the 13th EAI International Conference on Performance Evaluation Methodologies and Tools*. VALUETOOLS 2020 - 13th EAI International Conference on Performance Evaluation Methodologies and Tools. Tsukuba, Japan, 2020, pp. 136–141. DOI: 10.1145/3388831.3388852. URL: https://hal.archives-ouvertes.fr/hal-03107237.

[39] ***Best Paper***
B. Toure, S. Paturel and E. Altman. 'Congestion load balancing game with losses'. In: WINCOM'20 - 8th International Conference on Wireless Networks and Mobile Communications. Reims, France, 27th Oct. 2020. URL: https://hal.inria.fr/hal-02931311.

[40] C. Xu, G. Neglia and N. Sebastianelli. 'Dynamic Backup Workers for Parallel Machine Learning'. In: IFIP Networking 2020. Paris / Online, France, 22nd June 2020. URL: https://hal.inria.fr/hal-03044393.

[41] X. Ye, I. Esnaola, S. M. Perlaza and R. F. Harrison. 'Information Theoretic Data Injection Attacks with Sparsity Constraints'. In: International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm 2020). Proc. of the International Conference on Communications, Control, and Computing Technologies for the Smart Grids (SmartGridComm). Virtual Conference, United States, 11th Nov. 2020. URL: https://hal.archives-ouvertes.fr/hal-03128257.

Conferences without proceedings

[42] N. Benammar, P. Chrétienne, E. Hyon and A. Jean-Marie. 'Rolling Horizon approaches for a stochastic planning problem'. In: ROADEF 2020 - 21ème congrès annuel de la société Française de Recherche Opérationnelle et d'Aide à la Décision. Montpellier, France, 19th Feb. 2020. URL: https://hal.archives-ouvertes.fr/hal-02967919.

[43] G. IECKER RICARDO, G. Neglia and T. Spyropoulos. 'Caching Policies for Delay Minimization in Small Cell Networks with Joint Transmissions'. In: ICC 2020 - IEEE International Conference on Communications (ICC). Dublin, Ireland, 7th June 2020, pp. 1–6. DOI: 10.1109/ICC40277.2020.9149237. URL: https://hal.archives-ouvertes.fr/hal-03051751.

Scientific book chapters

[44] I. Esnaola, S. M. Perlaza and K. Sun. 'Data-Injection Attacks'. In: *Advanced Data Analytics for Power Systems*. 28th Jan. 2021. URL: https://hal.archives-ouvertes.fr/hal-03129791.

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

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

Reports & preprints

[46] E. Altman, I. Mounir, F.-Z. Najid and S. Perlaza. *On the true number of COVID-19 infections: Effect of Sensitivity, Specificity and Number of Tests on Prevalence Ratio Estimation*. INRIA Sophia Antipolis - Méditerranée, 27th May 2020. URL: https://hal.inria.fr/hal-02633844.

[47] D. Anade, J.-M. Gorce, P. Mary and S. Perlaza. *An upper bound on the error induced by saddlepoint approximations - Applications to information theory*. INRIA Grenoble - Rhône-Alpes, Apr. 2020, pp. 1–55. URL: https://hal.inria.fr/hal-02557887.

[48] D. Anade, J.-M. Gorce, P. Mary and S. M. Perlaza. *Saddlepoint Approximations of Cumulative Distribution Functions of Sums of Random Vectors*. INRIA Grenoble - Rhone-Alpes, Feb. 2021, pp. 1–33. URL: https://hal.inria.fr/hal-03143508.

[49] B. Jiang, P. Nain and D. Towsley. *Covert Cycle Stealing in a Single FIFO Server (extended version)*. 6th Feb. 2021. URL: https://hal.inria.fr/hal-02497557.

[50] C. Xu, G. Neglia and N. Sebastianelli. *Dynamic Backup Workers for Parallel Machine Learning*. 7th Dec. 2020. URL: https://hal.inria.fr/hal-03044199.