IN PARTNERSHIP WITH:
CNRS

Ecole normale supérieure de Paris

Activity Report 2018

Project-Team DYOGENE

Dynamics of Geometric Networks

IN COLLABORATION WITH: Département d'Informatique de l'Ecole Normale Supérieure

RESEARCH CENTER
Paris

THEME
Networks and Telecommunications

Table of contents

Project-Team DYOGENE

Creation of the Project-Team: 2013 July 01

Keywords:

Computer Science and Digital Science:

A1.2.4. - QoS, performance evaluation
A6.1.4. - Multiscale modeling
A6.2.3. - Probabilistic methods
A8.1. - Discrete mathematics, combinatorics
A8.2. - Optimization
A8.3. - Geometry, Topology
A8.6. - Information theory
A8.7. - Graph theory
A8.8. - Network science
A8.9. - Performance evaluation
A9.2. - Machine learning
A9.7. - AI algorithmics

Other Research Topics and Application Domains:

B4.3. - Renewable energy production
B6.2.2. - Radio technology
B6.3.4. - Social Networks

1. Team, Visitors, External Collaborators

Research Scientists

Bartlomiej Blaszczyszyn [Inria, Senior Researcher, HDR]
Marc Lelarge [Team leader, Inria, Senior Researcher, HDR]
François Baccelli [Inria, Senior Researcher, HDR]
Ana Busic [Inria, Researcher]
Christine Fricker [Inria, Researcher, from Mar 2018]
Laurent Massoulié [Inria, Senior Researcher, from Feb 2018, HDR]

Post-Doctoral Fellow

Cheng Wan [Université Paris-Sud and Inria, granted by PGMO, from Feb 2018 until Nov 2018]

PhD Students

Arnaud Cadas [PSL]
Alexis Galland [Ministère des Armées]
Md Umar Hashmi [PSL]
Hadrien Hendrikx [Inria, from Oct 2018]
Dalia-Georgiana Popescu [Bell Labs (Alcatel), until Oct 2018]
Alexandre Hollocou [Ministère de la Défense, until Nov 2018]
Quentin Le Gall [Inria]
Leo Miolane [Ecole polytechnique]
Edouard Pineau [Safran]
Sébastien Samain [Inria]
Ludovic Stephan [Ecole Normale Supérieure Paris, from Sep 2018]

Rémi Varloot [Inria, until Feb 2018]

Technical staff

Holger Keeler [Inria, granted by HUAWEI TECHNOLOGIES FRANCE]

Interns

Hugues Derogis [Inria, from Apr 2018 until Sep 2018]
Walid Ghanem [Inria, from Apr 2018 until Nov 2018]
Ludovic Stephan [Ecole Normale Supérieure Paris, from Apr 2018 until Jul 2018]

Administrative Assistants

Helene Bessin Rousseau [Inria, from Mar 2018]
Helene Milome [Inria]

Visiting Scientists

Akshay Goel [Feb 2018]
Sean Meyn [University of Florida, from Jul 2018 until Dec 2018]

External Collaborators

Anne Bouillard [Nokia, HDR]
Pierre Bremaud [EPFL]

2. Overall Objectives

2.1. Overall Objectives

The general scientific focus of DYOGENE is on the development of network mathematics. The following theories lie within our research interest: dynamical systems, queuing theory, optimization and control, information theory, stochastic processes, random graphs, stochastic geometry.

Our theoretical developments are motivated by and applied in the context of communication networks (Internet, wireless, mobile, cellular, peer-to-peer), social and economic networks, power grids.

We collaborate with many industrial partners. Our current industrial relations involve EDF, Google, Huawei, Microsoft, Nokia, Orange, Safran.

More specifically, the scientific focus of DYOGENE defined in 2013 was on geometric network dynamics arising in communications. By geometric networks we understand networks with a nontrivial, discrete or continuous, geometric definition of the existence of links between the nodes. In stochastic geometric networks, this definition leads to random graphs or stochastic geometric models.

A first type of geometric network dynamics is the one where the nodes or the links change over time according to an exogeneous dynamics (e.g. node motion and geometric definition of the links). We will refer to this as dynamics of geometric networks below. A second type is that where links and/or nodes are fixed but harbor local dynamical systems (in our case, stemming from e.g. information theory, queuing theory, social and economic sciences). This will be called dynamics on geometric networks. A third type is that where the dynamics of the network geometry and the local dynamics interplay. Our motivations for studying these systems stem from many fields of communications where they play a central role, and in particular: message passing algorithms; epidemic algorithms; wireless networks and information theory; device to device networking; distributed content delivery; social and economic networks, power grids.

3. Research Program

3.1. Initial research axes

The following research axes have been defined in 2013 when the project-team was created.

- Algorithms for network performance analysis, led by A. Bouillard and A. Busic.
- Stochastic geometry and information theory for wireless network, led by B. Blaszczyszyn and F. Baccelli.
- The cavity method for network algorithms, led by M. Lelarge.

Our scientific interests keep evolving. Research areas which received the most of our attention in 2017 are summarized in the following sections.

3.2. Distributed network control and smart-grids

Foundation of an entirely new science for distributed control of networks with applications to the stabilization of power grids subject to high volatility of renewable energy production is being developed A. Busic in collaboration with A. Bouillard and Sean Meyn [University of Florida].

3.3. Mathematics of wireless cellular networks

A comprehensive approach involving information theory, queueing and stochastic geometry to model and analyze the performance of large cellular networks, validated and implemented by Orange is being led by B. Blaszczyszyn in collaboration with F. Baccelli and M. K. Karray [Orange Labs]

3.4. High-dimensional statistical inference for social networks

Community detection and non-regular ramanujan graphs sole a conjecture on the optimality of non-backtracking spectral algorithm for community dectection in sparse stochastic block model graphs, as has been proved by M. Lelarge and L. Massoulié in collaboration with C. Bordenave [IMT Toulouse].

4. Application Domains

4.1. Physical communication networks

Internet, wireless, mobile, cellular networks.

4.2. Abstract networks

Social interactions, human communities, economic networks.

4.3. Power grids

Energy networks.

5. Highlights of the Year

5.1. Highlights of the Year

Publication of a monograph *Stochastic Geometry Analysis of Cellular Networks* by Cambridge University Press [30] that presents latest analytic techniques and results from stochastic geometry for modelling of heterogeneous cellular networks.

5.1.1. Awards

Our paper “ Optimal Algorithms for Non-Smooth Distributed Optimization in Networks” by K. Scaman, F. Bach, S. Bubeck, Y.T. Lee and L. Massoulié won a best paper award at the NeurIPS 2018 conference.

6. New Results

6.1. Energy Trade-offs for end-to-end Communications in Urban Vehicular Networks exploiting an Hyperfractal Model

In [33] presented this year at MSWIM DIVANet we show results on the trade-offs between the end-to-end communication delay and energy spent for completing a transmission in vehicular communications in urban settings. This study exploits our innovative model called " hyperfractal " that captures the self-similarity of the topology and vehicle locations in cities. We enrich the model by incorporating roadside infrastructure. We use analytical tools to derive theoretical bounds for the end-to-end communication hop count under two different energy constraints: either total accumulated energy, or maximum energy per node. More precisely, we prove that the hop count is bounded by $O(n(1{-}\alpha)/(dm{-}1))$ where $\alpha < 1$ and $m > 2$ is the precise hyperfractal dimension. This proves that for both constraints the energy decreases as we allow to chose among paths of larger length. In fact the asymptotic limit of the energy becomes significantly small when the number of nodes becomes asymptotically large. A lower bound on the network throughput capacity with constraints on path energy is also given. The results are confirmed through exhaustive simulations using different hyperfractal dimensions and path loss coefficients.

6.2. Broadcast Speedup in Vehicular Networks via Information Teleportation

In [32] presented this year at LCN our goal is to increase our understanding of the fundamental communication properties in urban vehicle-to-vehicle mobile networks by exploiting the self-similarity and hierarchical organization of modern cities. We use an innovative model called "hyperfractal" that captures the self-similarities of both the traffic and vehicle locations, and yet avoids the extremes of regularity and randomness. We use analytical tools to derive matching theoretical upper and lower bounds for the information propagation speed in an urban delay tolerant network (i.e., a network that is disconnected at all time, and thus uses a store-carry-and-forward routing model). We prove that the average broadcast time behaves as $n(1-\delta)$ (times a slowly varying function), where δ depends on the precise fractal dimension. Furthermore, we show that the broadcast speedup is due in part to an interesting self-similar phenomenon, that we denote as information teleportation. This phenomenon arises as a consequence of the topology of the vehicle traffic, and triggers an acceleration of the broadcast time. We show that our model fits real cities where open traffic data sets are available. The study presents simulations that confirm the validity of the bounds in multiple realistic settings, including scenarios with variable speed.

6.3. Vehicle-to-Infrastructure Communications Design in Urban Hyperfractals

In [25] presented at SPAWC our goal is to increase the awareness about the communication opportunities that arise in urban vehicle networks when exploiting the self-similarity and hierarchical organization of modern cities. The work uses our innovative model called " hyperfractal " that captures the self-similarity of the urban vehicular networks as well as incorporating roadside infrastructure with its own self-similarity. We use analytical tools to provide achievable trade-offs in operating the roadside units under the constraint of minimum routing path delay while maintaining a reasonably balanced load. The models and results are supported by simulations with different city hyperfractal dimensions in two different routing scenarios: nearest neighbor routing with no collision and minimum delay routing model assuming slotted Aloha, signal to interference ratio (SIR) capture condition, power-path loss, Rayleigh fading.

6.4. Book on Stochastic Geometry Analysis of Cellular Networks

In 2018 we have published a monograph [30] in which we explain the very latest analytic techniques and results from stochastic geometry for modelling the signal-to-interference-plus-noise ratio (SINR) distribution in heterogeneous cellular networks. This book is supposed to help readers to understand the effects of combining different system deployment parameters on key performance indicators such as coverage and capacity, enabling the efficient allocation of simulation resources. In addition to covering results for network models based on the Poisson point process, this book presents recent results for when non-Poisson base station configurations appear Poisson, due to random propagation effects such as fading and shadowing, as well as non-Poisson models for base station configurations, with a focus on determinantal point processes and tractable approximation methods. Theoretical results are illustrated with practical Long-Term Evolution (LTE) applications and compared with real-world deployment results.

6.5. Gibbsian On-Line Distributed Content Caching Strategy for Cellular Networks

In [9], we develop Gibbs sampling based techniques for learning the optimal content placement in a cellular network. A collection of base stations are scattered on the space, each having a cell (possibly overlapping with other cells). Mobile users request for downloads from a finite set of contents according to some popularity distribution. Each base station can store only a strict subset of the contents at a time; if a requested content is not available at any serving base station, it has to be downloaded from the backhaul. Thus, there arises the problem of optimal content placement which can minimize the download rate from the backhaul, or equivalently maximize the cache hit rate. Using similar ideas as Gibbs sampling, we propose simple sequential content update rules that decide whether to store a content at a base station based on the knowledge of contents in neighbouring base stations. The update rule is shown to be asymptotically converging to the optimal content placement for all nodes. Next, we extend the algorithm to address the situation where content popularities and cell topology are initially unknown, but are estimated as new requests arrive to the base stations. Finally, improvement in cache hit rate is demonstrated numerically.

6.6. Location Aware Opportunistic Bandwidth Sharing between Static and Mobile Users with Stochastic Learning in Cellular Networks

In [7], we consider location-dependent opportunistic bandwidth sharing between static and mobile downlink users in a cellular network. Each cell has some fixed number of static users. Mobile users enter the cell, move inside the cell for some time and then leave the cell. In order to provide higher data rate to mobile users, we propose to provide higher bandwidth to the mobile users at favourable times and locations, and provide higher bandwidth to the static users in other times. We formulate the problem as a long run average reward Markov decision process (MDP) where the per-step reward is a linear combination of instantaneous data volumes received by static and mobile users, and find the optimal policy. The transition structure of this MDP is not known in general. To alleviate this issue, we propose a learning algorithm based on single timescale stochastic approximation. Also, noting that the unconstrained MDP can be used to solve a constrained problem, we provide a learning algorithm based on multi-timescale stochastic approximation. The results are extended to address the issue of fair bandwidth sharing between the two classes of users. Numerical results demonstrate performance improvement by our scheme, and also the trade-off between performance gain and fairness.

6.7. Performance analysis of cellular networks with opportunistic scheduling using queueing theory and stochastic geometry

In [38] submitted this year, combining stochastic geometric approach with some classical results from queuing theory,we propose a comprehensive framework for the performance study of large cellular networks featuring opportunistic scheduling. Rapid and verifiable with respect to real data, our approach is particularly useful for network dimensioning and long term economic planning. It is based on a detailed network model combining an information-theoretic representation of the link layer, a queuing-theoretic representation of the users' scheduler, and a stochastic-geometric representation of the signal propagation and the network cells. It allows one to evaluate principal characteristics of the individual cells, such as loads (defined as the fraction of time the cell is not empty), the mean number of served users in the steady state, and the user throughput. A simplified Gaussian approximate model is also proposed to facilitate study of the spatial distribution of these metrics across the network. The analysis of both models requires only simulations of the point process of base stations and the shadowing field to estimate the expectations of some stochastic-geometric functionals not admitting explicit expressions. A key observation of our approach , bridging spatial and temporal analysis, relates the SINR distribution of the typical user to the load of the typical cell of the network. The former is a static characteristic of the network related to its spectral efficiency while the latter characterizes the performance of the (generalized) processor sharing queue serving the dynamic population of users of this cell.

6.8. The Influence of Canyon Shadowing on Device-to-Device Connectivity in Urban Scenario

In [48] submitted this year, we use percolation theory to study the feasibility of large-scale connectivity of relay-augmented device-to-device (D2D) networks in an urban scenario, featuring a haphazard system of streets and canyon shadowing allowing only for line-of-sight (LOS) communications in a limited finite range. We use a homogeneous Poisson-Voronoi tessellation (PVT) model of streets with homogeneous Poisson users (devices) on its edges and independent Bernoulli relays on the vertices. Using this model, we demonstrated the existence of a minimal threshold for relays below which large-scale connectivity of the network is not possible, regardless of all other network parameters. Through simulations, we estimated this threshold to 71.3%. Moreover, if the mean street length is not larger than some threshold (predicted to 74.3% of the communication range; which might be the case in a typical urban scenario) then any (whatever small) density of users can be compensated by equipping more crossroads with relays. Above this latter threshold, good connectivity requires some minimal density of users, compensated by the relays in a way we make explicit. The existence of the above regimes brings interesting qualitative arguments to the discussion on the possible D2D deployment scenarios.

6.9. Determinantal thinning of point processes with network learning applications

In [39] submitted this year, a new type of dependent thinning for point processes in continuous space is proposed, which leverages the advantages of determinantal point processes defined on finite spaces and, as such, is particularly amenable to statistical, numerical, and simulation techniques. It gives a new point process that can serve as a network model exhibiting repulsion. The properties and functions of the new point process, such as moment measures, the Laplace functional, the void probabilities, as well as conditional (Palm) characteristics can be estimated accurately by simulating the underlying (non-thinned) point process, which can be taken, for example, to be Poisson. This is in contrast (and preference to) finite Gibbs point processes, which, instead of thinning, require weighting the Poisson realizations, involving usually intractable normalizing constants. Models based on determinantal point processes are also well suited for statistical (supervised) learning techniques, allowing the models to be fitted to observed network patterns with some particular geometric properties. We illustrate this approach by imitating with determinantal thinning the well-known Matérn II hard-core thinning, as well as a soft-core thinning depending on nearest-neighbour triangles. These two examples demonstrate how the proposed approach can lead to new, statistically optimized, probabilistic transmission scheduling schemes.

6.10. Analyzing LoRa long-range, low-power, wide-area networks using stochastic geometry

In [40] submitted this year, we present a simple, stochastic-geometric model of a wireless access network exploiting the LoRA (Long Range) protocol, which is a non-expensive technology allowing for long-range, single-hop connectivity for the Internet of Things. We assume a space-time Poisson model of packets transmitted by LoRA nodes to a fixed base station. Following previous studies of the impact of interference, we assume that a given packet is successfully received when no interfering packet arrives with similar power before the given packet payload phase. This is as a consequence of LoRa using different transmission rates for different link budgets (transmissions with smaller received powers use larger spreading factors) and LoRa intra-technology interference treatment. Using our model, we study the scaling of the packet reception probabilities per link budget as a function of the spatial density of nodes and their rate of transmissions. We consider both the parameter values recommended by the LoRa provider, as well as proposing LoRa tuning to improve the equality of performance for all link budgets. We also consider spatially non-homogeneous distributions of LoRa nodes. We show also how a fair comparison to non-slotted Aloha can be made within the same framework.

6.11. Statistical learning of geometric characteristics of wireless networks

In [41] to appear in Proc. INFOCOM 2019, motivated by the prediction of cell loads in cellular networks, we formulate the following new, fundamental problem of statistical learning of geometric marks of point processes: An unknown marking function, depending on the geometry of point patterns, produces characteristics (marks) of the points. One aims at learning this function from the examples of marked point patterns in order to predict the marks of new point patterns. To approximate (interpolate) the marking function, in our baseline approach, we build a statistical regression model of the marks with respect some local point distance representation. In a more advanced approach, we use a global data representation via the scattering moments of random measures, which build informative and stable to deformations data representation, already proven useful in image analysis and related application domains. In this case, the regression of the scattering moments of the marked point patterns with respect to the non-marked ones is combined with the numerical solution of the inverse problem, where the marks are recovered from the estimated scattering moments. Considering some simple, generic marks, often appearing in the modeling of wireless networks, such as the shot-noise values, nearest neighbour distance, and some characteristics of the Voronoi cells, we show that the scattering moments can capture similar geometry information as the baseline approach, and can reach even better performance, especially for non-local marking functions. Our results motivate further development of statistical learning tools for stochastic geometry and analysis of wireless networks, in particular to predict cell loads in cellular networks from the locations of base stations and traffic demand.

6.12. Ressource allocation in bike sharing systems

Vehicle sharing systems are becoming an urban mode of transportation, and launched in many cities, as Velib' and Autolib' in Paris. Managing such systems is quite difficult. One of the major issues is the availability of the resources: vehicles or free slots. These systems became a hot topic in Operation Research and the importance of stochasticity on the system behavior leads us to propose mathematical stochastic models. The aim is to understand the system behavior and how to manage these systems in order to improve the allocation of both resources to users.

To improve BSS (bike-sharing systems), two types of policies can be deployed: incentives to the users to choose a better station, called *natural* or *green* regulation, or redistribution by trucks, called *active* regulation. In a simple mathematical model, we proved the efficiency of the 2-choice incentive policy for BSS (bike-sharing systems). The drawback of the model is that it ignores the geometry of the system, where the choice is only local. The purpose of this first work is to deal with this policy in real systems.

We use data trip data obtained from JCDecaux and reports on station status collected as open data, to test local choice policy. Indeed we designed and tested a new policy relying on a local small change in user behaviors, by adapting their trips to resource availability around their departure and arrival stations, based on 2-choice policy. Results show that, even with a small user collaboration, the proposed method increases significantly the global balance of the bike sharing system and therefore the user satisfaction. This is done using trip data sets and detecting spatial outliers, stations having a behavior significantly different from their spatial neighbors, in a context where neighbors are heavily correlated. For that we proposed an improved version of the well-known

Moran scatterplot method, using a robust distance metric called Gower similarity. Using this new version of Moran scatterplot, we show that, for the occupancy data set obtained by modifiying trips, the number of spatial outliers drastically decreases. We generalize this study with W. Ghanem and L. Massoulié testing incentive and redistribution policies on a simulator, where the tradeoff between the number of frustrated trips and the penalty for the users can be measured. We propose new versions of these policies including prediction.

6.13. Analyzing the choice of the least loaded queue between two neighboring queues

A model of N queues, with a local choice policy, is studied. Each one-server queue has a Poissonian arrival of customers. When a customer arrives at a queue, he joins the least loaded queue between this queue and the next one, ties solved at random. Service times have exponential distribution. The system is stable if the arrival-to-service rate ratio, also called load, is less than one. When the load tends to zero, we derive the first terms of the expansion in this parameter for the stationary probabilities that a queue has few customers. Then we provide explicit asymptotics, as the load tends to zero, for the stationary probabilities of the queue length. We used the analyticity of the stationary probabilities as a function of the load. It shows the behavior difference between this local choice policy and the 2-choice policy (*supermarket model*).

6.14. Optimal Content Replication and Request Matching in Large Caching Systems

We consider models of content delivery networks in which the servers are constrained by two main resources: memory and bandwidth. In such systems, the throughput crucially depends on how contents are replicated across servers and how the requests of specific contents are matched to servers storing those contents. In this paper, we first formulate the problem of computing the optimal replication policy which if combined with the optimal matching policy maximizes the throughput of the caching system in the stationary regime. It is shown that computing the optimal replication policy for a given system is an NP-hard problem. A greedy replication scheme is proposed and it is shown that the scheme provides a constant factor approximation guarantee. We then propose a simple randomized matching scheme which avoids the problem of interruption in service of the ongoing requests due to re-assignment or repacking of the existing requests in the optimal matching policy. The dynamics of the caching system is analyzed under the combination of proposed replication and matching schemes. We study a limiting regime, where the number of servers and the arrival rates of the contents are scaled proportionally, and show that the proposed policies achieve asymptotic optimality. Extensive simulation results are presented to evaluate the performance of different policies and study the behavior of the caching system under different service time distributions of the requests.

6.15. Statistical thresholds for Tensor PCA

This is a joint work with Aukosh Jagannath and Patrick Lopatto. We study the statistical limits of testing and estimation for a rank one deformation of a Gaussian random tensor. We compute the sharp thresholds for hypothesis testing and estimation by maximum likelihood and show that they are the same. Furthermore, we find that the maximum likelihood estimator achieves the maximal correlation with the planted vector among measurable estimators above the estimation threshold. In this setting, the maximum likelihood estimator exhibits a discontinuous BBP-type transition: below the critical threshold the estimator is orthogonal to the planted vector, but above the critical threshold, it achieves positive correlation which is uniformly bounded away from zero.

6.16. The distribution of the Lasso: Uniform control over sparse balls and adaptive parameter tuning

This is a joint work with Andrea Montanari. The Lasso is a popular regression method for high-dimensional problems in which the number of parameters $\theta_1, \cdots, \theta_N$, is larger than the number n of samples: $N > n$. A useful heuristics relates the statistical properties of the Lasso estimator to that of a simple soft-thresholding denoiser, in a denoising problem in which the parameters $(\theta_i)_{i \leq N}$ are observed in Gaussian noise, with a carefully tuned variance. Earlier work confirmed this picture in the limit $n, N \to \infty$, pointwise in the parameters θ, and in the value of the regularization parameter.

Here, we consider a standard random design model and prove exponential concentration of its empirical distribution around the prediction provided by the Gaussian denoising model. Crucially, our results are uniform with respect to θ belonging to ℓ_q balls, $q \in [0, 1]$, and with respect to the regularization parameter. This allows to derive sharp results for the performances of various data-driven procedures to tune the regularization.

Our proofs make use of Gaussian comparison inequalities, and in particular of a version of Gordon's minimax theorem developed by Thrampoulidis, Oymak, and Hassibi, which controls the optimum value of the Lasso optimization problem. Crucially, we prove a stability property of the minimizer in Wasserstein distance, that allows to characterize properties of the minimizer itself.

6.17. Phase transitions in spiked matrix estimation: information-theoretic analysis

We study here the so-called spiked Wigner and Wishart models, where one observes a low-rank matrix perturbed by some Gaussian noise. These models encompass many classical statistical tasks such as sparse PCA, submatrix localization, community detection or Gaussian mixture clustering. The goal of these notes is to present in a unified manner recent results (as well as new developments) on the information-theoretic limits of these spiked matrix/tensor models. We compute the minimal mean squared error for the estimation of the low-rank signal and compare it to the performance of spectral estimators and message passing algorithms. Phase transition phenomena are observed: depending on the noise level it is either impossible, easy (i.e. using polynomial-time estimators) or hard (information-theoretically possible, but no efficient algorithm is known to succeed) to recover the signal.

6.18. Accelerated decentralized optimization with local updates for smooth and strongly convex objectives

We study the problem of minimizing a sum of smooth and strongly convex functions split over the nodes of a network in a decentralized fashion. We propose the algorithm $ESDACD$, a decentralized accelerated algorithm that only requires local synchrony. Its rate depends on the condition number κ of the local functions as well as the network topology and delays. Under mild assumptions on the topology of the graph, $ESDACD$ takes a time $O((\tau_{\max} + \Delta_{\max})\sqrt{\kappa/\gamma}\ln(\epsilon^{-1}))$ to reach a precision ϵ where γ is the spectral gap of the graph, $\tau_{\max}$ the maximum communication delay and $\Delta_{\max}$ the maximum computation time. Therefore, it matches the rate of $SSDA$, which is optimal when $\tau_{\max} = \Omega(\Delta_{\max})$. Applying $ESDACD$ to quadratic local functions leads to an accelerated randomized gossip algorithm of rate $O(\sqrt{\theta_{\text{gossip}}/n})$ where θ_{gossip} is the rate of the standard randomized gossip. To the best of our knowledge, it is the first asynchronous gossip algorithm with a provably improved rate of convergence of the second moment of the error. We illustrate these results with experiments in idealized settings.

6.19. Group synchronization on grids

Group synchronization requires to estimate unknown elements $(\theta_v)_{v \in V}$ of a compact group $\mathbb{G}$ associated to the vertices of a graph $G = (V, E)$, using noisy observations of the group differences associated to the edges. This model is relevant to a variety of applications ranging from structure from motion in computer vision to graph localization and positioning, to certain families of community detection problems.

We focus on the case in which the graph G is the d-dimensional grid. Since the unknowns θ_v are only determined up to a global action of the group, we consider the following weak recovery question. Can we determine the group difference $\theta_u^{-1}\theta_v$ between far apart vertices u, v better than by random guessing? We prove that weak recovery is possible (provided the noise is small enough) for $d \geq 3$ and, for certain finite groups, for $d \geq 2$. Vice-versa, for some continuous groups, we prove that weak recovery is impossible for $d = 2$. Finally, for strong enough noise, weak recovery is always impossible.

6.20. An Impossibility Result for Reconstruction in a Degree-Corrected Planted-Partition Model

We consider the Degree-Corrected Stochastic Block Model (DC-SBM): a random graph on n nodes, having i.i.d. weights $(\phi_u)_{u=1}^n$ (possibly heavy-tailed), partitioned into $q \geq 2$ asymptotically equal-sized clusters. The model parameters are two constants $a, b > 0$ and the finite second moment of the weights $\Phi^{(2)}$. Vertices u and v are connected by an edge with probability $(\phi_u \phi_v / n)a$ when they are in the same class and with probability $(\phi_u \phi_v / n)b$ otherwise. We prove that it is information-theoretically impossible to estimate the clusters in a way positively correlated with the true community structure when $(a-b)2\Phi^{(2)} \leq q(a+b)$. As by-products of our proof we obtain (1) a precise coupling result for local neighbourhoods in DC-SBM's, that we use in a follow up paper [Gulikers et al., 2017] to establish a law of large numbers for local-functionals and (2) that long-range interactions are weak in (power-law) DC-SBM's.

6.21. On the capacity of information processing systems

We propose and analyze a family of information processing systems, where a finite set of experts or servers are employed to extract information about a stream of incoming jobs. Each job is associated with a hidden label drawn from some prior distribution. An inspection by an expert produces a noisy outcome that depends both on the job's hidden label and the type of the expert, and occupies the expert for a finite time duration. A decision maker's task is to dynamically assign inspections so that the resulting outcomes can be used to accurately recover the labels of all jobs, while keeping the system stable. Among our chief motivations are applications in crowd-sourcing, diagnostics, and experiment designs, where one wishes to efficiently learn the nature of a large number of items, using a finite pool of computational resources or human agents. We focus on the capacity of such an information processing system. Given a level of accuracy guarantee, we ask how many experts are needed in order to stabilize the system, and through what inspection architecture. Our main result provides an adaptive inspection policy that is asymptotically optimal in the following sense: the ratio between the required number of experts under our policy and the theoretical optimal converges to one, as the probability of error in label recovery tends to zero.

6.22. Optimal Algorithms for Non-Smooth Distributed Optimization in Networks

In this work, we consider the distributed optimization of non-smooth convex functions using a network of computing units. We investigate this problem under two regularity assumptions: (1) the Lipschitz continuity of the global objective function, and (2) the Lipschitz continuity of local individual functions. Under the local regularity assumption, we provide the first optimal first-order decentralized algorithm called multi-step primal-dual (MSPD) and its corresponding optimal convergence rate. A notable aspect of this result is that, for non-smooth functions, while the dominant term of the error is in $O(1/\sqrt{t})$, the structure of the communication network only impacts a second-order term in $O(1/t)$, where t is time. In other words, the error due to limits in communication resources decreases at a fast rate even in the case of non-strongly-convex objective functions. Under the global regularity assumption, we provide a simple yet efficient algorithm called distributed randomized smoothing (DRS) based on a local smoothing of the objective function, and show that DRS is within a $d^{1/4}$ multiplicative factor of the optimal convergence rate, where d is the underlying dimension.

6.23. Zap Meets Momentum: Stochastic Approximation Algorithms with Optimal Convergence Rate

There are two well known Stochastic Approximation techniques that are known to have optimal rate of convergence (measured in terms of asymptotic variance): the Ruppert-Polyak averaging technique, and stochastic Newton-Raphson (SNR)(a matrix gain algorithm that resembles the deterministic Newton-Raphson method). The Zap algorithms, introduced by Devraj and Meyn in 2017, are a version of SNR designed to behave more closely like their deterministic cousin. It is found that estimates from the Zap Q-learning algorithm converge remarkably quickly, but the per-iteration complexity can be high. In [43], we introduce a new class of stochastic approximation algorithms based on matrix momentum. For a special choice of the matrix momentum and gain sequences, it is found in simulations that the parameter estimates obtained from the algorithm couple with those obtained from the more complex stochastic Newton-Raphson algorithm. Conditions under which coupling is guaranteed are established for a class of linear recursions. Optimal finite-n error bounds are also obtained.

6.24. Ergodic theory for controlled Markov chains with stationary inputs

Consider a stochastic process $\mathbf{X}$ on a finite state space $X = \{1, \cdots, d\}$. It is conditionally Markov, given a real-valued 'input process' ζ. This is assumed to be small, which is modeled through the scaling, $\zeta_t = \varepsilon\zeta_t^1,\ 0 \le \varepsilon \le 1$, where ζ^1 is a bounded stationary process. The following conclusions are obtained, subject to smoothness assumptions on the controlled transition matrix and a mixing condition on ζ:

- A stationary version of the process is constructed, that is coupled with a stationary version of the Markov chain $\mathbf{X}^\bullet$ obtained with $\zeta \equiv 0$. The triple $(\mathbf{X}, \mathbf{X}^\bullet, \zeta)$ is a jointly stationary process satisfying $\mathsf{P}\{X(t) \neq X^\bullet(t)\} = O(\varepsilon)$. Moreover, a second-order Taylor-series approximation is obtained:

$$\mathsf{P}\{X(t) = i\} = \mathsf{P}\{X^\bullet(t) = i\} + \varepsilon^2 \varrho(i) + o(\varepsilon^2), \quad 1 \le i \le d,$$

  with an explicit formula for the vector $\varrho \in \mathfrak{R}^d$.
- For any $m \ge 1$ and any function $f : \{1, \cdots, d\} \times \mathfrak{R} \to \mathfrak{R}^m$, the stationary stochastic process $Y(t) = f(X(t), \zeta(t))$ has a power spectral density S_f that admits a second order Taylor series expansion: A function $\mathrm{S}^{(2)}{}_f : [-\pi, \pi] \to C^{m \times m}$ is constructed such that

$$\mathrm{S}_f(\theta) = \mathrm{S}_f^\bullet(\theta) + \varepsilon^2 \mathrm{S}^{(2)}{}_f(\theta) + o(\varepsilon^2), \quad \theta \in [-\pi, \pi].$$

  An explicit formula for the function $\mathrm{S}^{(2)}{}_f$ is obtained, based in part on the bounds in (i).

The results are illustrated using a version of the timing channel of Anantharam and Verdu.

6.25. Ordinary Differential Equation Methods for Markov Decision Processes and Application to Kullback–Leibler Control Cost

A new approach to computation of optimal policies for MDP (Markov decision process) models is introduced in [5], published in SICON this year. The main idea is to solve not one, but an entire family of MDPs, parameterized by a scalar ζ that appears in the one-step reward function. For an MDP with d states, the family of relative value functions $\{h_\zeta^* : \zeta \in \mathbb{R}\}$ is the solution to an ODE, $\frac{d}{d\zeta} h_\zeta^* = \mathcal{V}(h_\zeta^*)$, where the vector field $\mathcal{V} : R^d \to R^d$ has a simple form, based on a matrix inverse. Two general applications are presented: Brockett's quadratic-cost MDP model, and a generalization of the "linearly solvable" MDP framework of Todorov in which the one-step reward function is defined by Kullback–Leibler divergence with respect to nominal dynamics. This technique was introduced by Todorov in 2007, where it was shown under general conditions that the solution to the average-reward optimality equations reduce to a simple eigenvector problem. Since then many authors have sought to apply this technique to control problems and models of bounded rationality in economics. A crucial assumption is that the input process is essentially unconstrained. For example, if the nominal dynamics include randomness from nature (eg, the impact of wind on a moving vehicle), then the optimal control solution does not respect the exogenous nature of this disturbance. In [16] we introduce a technique to solve a more general class of action-constrained MDPs.

6.26. Distributed control design for balancing the grid using flexible loads

Inexpensive energy from the wind and the sun comes with unwanted volatility, such as ramps with the setting sun or a gust of wind. Controllable generators manage supply-demand balance of power today, but this is becoming increasingly costly with increasing penetration of renewable energy. It has been argued since the 1980s that consumers should be put in the loop: "demand response" will help to create needed supply-demand balance. However, consumers use power for a reason and expect that the quality of service (QoS) they receive will lie within reasonable bounds. Moreover, the behavior of some consumers is unpredictable, while the grid operator requires predictable controllable resources to maintain reliability.

The goal of the book chapter [31] is to describe an emerging science for demand dispatch that will create virtual energy storage from flexible loads. By design, the grid-level services from flexible loads will be as controllable and predictable as a generator or fleet of batteries. Strict bounds on QoS will be maintained in all cases. The potential economic impact of these new resources is enormous. California plans to spend billions of dollars on batteries that will provide only a small fraction of the balancing services that can be obtained using demand dispatch. The potential impact on society is enormous: a sustainable energy future is possible with the right mix of infrastructure and control systems.

In [17], presented at IEEE CDC 2018, a natural notion of *energy capacity* is proposed for the special case of thermostatically controlled loads (TCLs). It is shown that this quantity is closely approximated by thermal energy capacity, which is a component of the "leaky battery model" introduced in prior work. Simulation experiments in a distributed control setting show that these energy limits, and accompanying power capacity limits, are reliable indicators of online capacity, even for a heterogeneous population of loads. A feedforward/feedback control scheme is proposed for a large collection of heterogeneous loads. At the local level, control loops are used to create cooperative responses from each load in a given class of homogeneous loads. This simplifies control of the aggregate based on two pieces of information: aggregate power consumption from each class of loads and the *state of charge* surrogate that is a part of the leaky battery model. This information is required at a slow time-scale (say, 5 minute sampling).

In [18], we study the problem of coordination of a collection of on/off thermostatically controlled loads (TCLs) to act as a "virtual battery". Virtual Energy Storage (VES) is provided by the collection by either consuming more (charging) or less (discharging) power than the baseline. VES can be an inexpensive alternative to batteries when a large share of the electricity comes from volatile sources such as solar and wind. Almost all prior work has assumed that the outside weather - which significantly effects a TCLs behavior - is constant. We combine the above distributed load control design with a grid level MPC (model predictive control) that uses predictions of disturbances (weather) over a planning horizon. Additionally, irrespective of the choice of control architecture, there is a fundamental limit to the power and energy capacity of the collection of TCLs. We partially address this issue by scaling the reference signal by a function of the outside air temperature.

6.27. Estimation and control of quality of service in demand dispatch

Flexibility of energy consumption can be harnessed for the purposes of grid-level ancillary services. In particular, through distributed control of a collection of loads, a balancing authority regulation signal can be tracked accurately, while ensuring that the quality of service (QoS) for each load is acceptable on average. Subject to distributed control approaches advocated in recent research, the histogram of QoS is approximately Gaussian, and consequently, each load will eventually receive poor service. In [11], published this year in IEEE Transactions on Smart Grid, statistical techniques are developed to estimate the mean and variance of QoS as a function of the power spectral density of the regulation signal. It is also shown that additional local control can eliminate risk. The histogram of QoS is truncated through this local control, so that strict bounds on service quality are guaranteed. While there is a tradeoff between the grid-level tracking performance (capacity and accuracy) and the bounds imposed on QoS, it is found that the loss of capacity is minor in typical cases.

The previous designs for distributed control of TCLs ensure that the indoor temperature remains within a pre-specified bound, but other QoS metrics, especially the frequency of turning on and off was not limited. In [19], presented at ACM BuildSys 2018, we propose a more advanced control architecture that reduces the cycling rate of TCLs. We show through simulations that the proposed controller is able to reduce the cycling of individual TCLs compared to the previous designs with little loss in tracking of the grid-supplied reference signal.

6.28. Optimal control of energy storage

Energy storage revenue estimation is essential for analyzing financial feasibility of investment in batteries. In [22], we quantify the cycles of operation considering depth-of-discharge (DoD) of operational cycles and provide an algorithm to calculate equivalent 100% DoD cycles. This facilitates in comparing cycles of different DoDs. The battery life is frequently defined as a combination of cycle and calendar life. We propose a battery capacity degradation model based on the cycle and the calendar life and operational cycles. Using equivalent 100% DoD cycles and revenue generated, we calculate the dollars per cycle revenue of storage performing electricity price based arbitrage and ancillary services for load balancing in real time. Using PJM's (a regional transmission organization in the United States) real data we calculate short term and long term financial potential for the year of 2017. We observe that participating in ancillary services is significantly more beneficial for storage owners compared to participating in energy arbitrage.

Battery life is often described a combination of cycle life and calendar life. In [21], we propose a mechanism to limit the number of cycles of operation over a time horizon in an optimal arbitrage algorithm proposed in our previous work. The cycles of operation have to be tuned based on price volatility to maximize the battery life and arbitrage gains.

In [23], we analyze the effect of real time electricity price (RTP) on the amount of ancillary services required for load balancing in presence of responsive users, information asymmetry and forecast errors in demand and renewable energy sources (RES) generation. We consider a RTP that is determined by the forecasted generation and ramping cost. A community choice aggregator manages the load of all the consumers by setting the price. The consumer's objective is to minimize their overall cost of consumption. Ancillary services are called upon to balance the load in real time. With zero RES in the power network and a high degree of load flexibility, the proposed RTP flattens and the volatility in demand vanishes. However, in presence of RES the volatility in price and demand is reduced up to an extent and ancillary services are required for load balancing. The amount of ancillary services required increases with forecast errors. We also propose a real time algorithm that approximates the optimal consumer behavior under the complete information setting. Extensive numerical simulations are provided using real data from Pecan Street and Elia Belgium.

6.29. Dynamic matching models

The model of First Come First Served infinite bipartite matching was introduced in Caldentey, Kaplan and Weiss, 2009. In this model, there is a sequence of items that are chosen i.i.d. from a finite set $\mathcal{C}$ and an independent sequence of items that are chosen i.i.d. from a finite set $\mathcal{S}$, and a bipartite compatibility graph G between $\mathcal{C}$ and $\mathcal{S}$. Items of the two sequences are matched according to the compatibility graph, and the matching is FCFS, meaning that each item in the one sequence is matched to the earliest compatible unmatched item in the other sequence. In Adan and Weiss, 2012, a Markov chain associated with the matching was analyzed, a condition for stability was derived, and a product form stationary distribution was obtained. In [2], we present several new results that unveil the fundamental structure of the model. First, we provide a pathwise Loynes' type construction which enables to prove the existence of a unique matching for the model defined over all the integers. Second, we prove that the model is dynamically reversible: we define an exchange transformation in which we interchange the positions of each matched pair, and show that the items in the resulting permuted sequences are again independent and i.i.d., and the matching between them is FCFS in reversed time. Third, we obtain product form stationary distributions of several new Markov chains associated with the model. As a by-product, we compute useful performance measures, for instance the link lengths between matched items.

In [51], we propose an explicit construction of the stationary state of Extended Bipartite Matching (EBM) models, as defined in (Busic et. al., 2013). We use a Loynes-type backwards scheme similar in flavor to that in (Moyal et al., 2017), allowing to show the existence and uniqueness of a bi-infinite perfect matching under various conditions, for a large class of matching policies and of bipartite matching structures. The key algebraic element of our construction is the sub-additivity of a suitable stochastic recursive representation of the model, satisfied under most usual matching policies. By doing so, we also derive stability conditions for the system under general stationary ergodic assumptions, subsuming the classical markovian settings.

In [42], we consider holding costs for the items that are waiting to be matched. We model this problem as an MDP (Markov decision process) and study the discounted cost and the average cost case. We first consider a model with two types of supply and two types of demand items with an N matching graph. For linear cost function, we prove that an optimal matching policy gives priority to the end edges of the matching graph and is of threshold type for the diagonal edge. In addition, for the average cost problem, we compute the optimal threshold value. According to our preliminary numerical experiments, threshold-type policies performs also very well for more general bipartite graphs.

7. Bilateral Contracts and Grants with Industry

7.1. CRE with Huawei

18-month contract titled “Mathematical Modeling of 5G Ultra Dense Wireless Networks” between Inria represented by B. Blaszczyszyn (PI) and F. Baccelli, and Huawei comes to an end in December 2018. It aimed at investigating obstacle-based shadowing fields in the spatial models of cellular networks and efficient scheduling policies. Paul Keeler was hired by Inria as a research engineer thanks to this contract. The publication [39] is one of the deliverable of this contract.

7.2. CIFRE with Nokia

Contract with Nokia started in 2015 for the co-advising by B. Blaszczyszyn of a PhD student of Nokia, Dalia-Georgiana Herculea came to an end in December 2018. Dalia-Georgiana Herculea has successfully defended her PhD Thesis in November 2018.

7.3. CIFRE with Orange

Contract with Orange started in 2017 and continued in 2018 for the co-advising by B. Blaszczyszyn of a PhD student of Orange, Quentin Le Gall.

8. Partnerships and Cooperations

8.1. Regional Initiatives

8.1.1. Laboratory of Information, Networking and Communication Sciences (LINCS)

Dyogene participates in LINCS https://www.lincs.fr/, a research centre co-founded by Inria, Institut Mines-Télécom, UPMC and Alcatel-Lucent Bell Labs (currently Nokia Bell Labs) dedicated to research and innovation in the domains of future information and communication networks, systems and services. S. Meyn [University of Florida] was invited professor by LINCS and ENS from July to December 2018.

8.1.2. PGMO

Dyogene participates to the PGMO (Gaspard Monge Program for Optimization, operations research, and their interactions with data science) via the project a 2 year project “Distributed control of flexible loads” funded through the ICODE/IROE call. This is a collaborative project between University Paris-Sud (PI: Gilles Stoltz) and Inria (PI: Ana Busic). Post-doc Cheng Wan was financed by this project from Feb-Nov 2018.

8.2. National Initiatives

8.2.1. GdR GeoSto

Members of Dyogene participate in Research Group GeoSto (Groupement de recherche, GdR 3477) http://gdr-geostoch.math.cnrs.fr/ on Stochastic Geometry led by and David Coupier [Université de Valenciennes].

This is a collaboration framework for all French research teams working in the domain of spatial stochastic modeling, both on theory development and in applications. This year DYOGENE has co-organized yearly conference of the GdR *Stochastic Geometry Days 2018* 14–18 mai 2018 Paris (France); https://geosto-2018.sciencesconf.org/.

8.2.2. GdR RO

Members of Dyogene participate in GdR-RO (Recherche Opérationelle; GdR CNRS 3002), http://gdrro.lip6.fr/, working group COSMOS (Stochastic optimization and control, modeling and simulation), lead by A. Busic and E. Hyon (LIP 6); http://gdrro.lip6.fr/?q=node/78

8.2.3. ANR JCJC PARI

Probabilistic Approach for Renewable Energy Integration: Virtual Storage from Flexible Loads. The project started in January 2017. PI — A. Bušić. This project is motivated by current and projected needs of a power grid with significant renewable energy integration. Renewable energy sources such as wind and solar have a high degree of unpredictability and time variation, which makes balancing demand and supply challenging. There is an increased need for ancillary services to smooth the volatility of renewable power. In the absence of large, expensive batteries, we may have to increase our inventory of responsive fossil-fuel generators, negating the environmental benefits of renewable energy. The proposed approach addresses this challenge by harnessing the inherent flexibility in demand of many types of loads. The objective of the project is to develop decentralized control for automated demand dispatch, that can be used by grid operators as ancillary service to regulate demand-supply balance at low cost. We call the resource obtained from these techniques virtual energy storage (VES). Our goal is to create the necessary ancillary services for the grid that are environmentally friendly, that have low cost and that do not impact the quality of service (QoS) for the consumers. Besides respecting the needs of the loads, the aim of the project is to design local control solutions that require minimal communications from the loads to the centralized entity. This is possible through a systems architecture that includes the following elements: i) local control at each load based on local measurements combined with a grid-level signal; ii) frequency decomposition of the regulation signal based on QoS and physical constraints for each class of loads.

8.3. International Initiatives

8.3.1. Inria International Partners

8.3.1.1. IFCAM Project “Geometric statistics of stationary point processes”

B. Błaszczyszyn and Yogeshwaran D. from Indian Statistical Institute (ISI), Bangalore, have got in 2018 the approval from Indo-French Centre for Applied Mathematics (IFCAM), for their joint project on “Geometric statistics of stationary point processes" for the period 2018–2021. B. Błaszczyszyn was visiting ISI Bangalore for two weeks in November–December 2018.

8.3.1.2. Informal International Partners

- University of Florida: collaborations with Prof Sean Meyn (ECE), Associate Prof Prabir Barooah (MAE), and the PhD students: A. Devraj (ECE), A. Coffman (MAE), N. Cammardella (ECE), J. Mathias (ECE).

8.4. International Research Visitors

8.4.1. Visits of International Scientists

- D. Yogeshwaran [Indian Statistical Institute, Bangalore, India]
- S. Meyn [University of Florida, USA] was invited Prof at ENS and LINCS, July - December 2018

8.4.1.1. Internships

- Master Probabilités et Modèles aléatoires UPMC, Walid Ghanem, *Hydrodynamic limit of a network with moving servers*, 04-07/2018, encadrant Christine Fricker.
- Master MASH (Mathématiques appliquées aux sciences humaines) ENS-Paris Dauphine University, *Using customer oriented policies based on probabilistic methods to enhance the Bike Sharing System Velib'*, 08-011/2018, encadrants Christine Fricker et Laurent Massoulié.
- Akshay Goel [Kyushu University, Fukuoka, Japan] Mars 2018,
- Tokuyama Kiichi [Tokyo Tech, Tokyo, Japan], April 2018,

8.4.2. Visits to International Teams

8.4.2.1. Research Stays Abroad

- B. Błaszczyszyn was visiting Yogeshwaran D. at the Indian Statistical Institute Bangalore for two weeks in November–December 2018 (IFCAM project).
- A. Busic was a long-term participant (March-Mai 2018) of the Real-Time Decision Making program, Simons Institute, UC Berkeley, USA; https://simons.berkeley.edu/programs/realtime2018

9. Dissemination

9.1. Promoting Scientific Activities

9.1.1. Scientific Events Organisation

9.1.1.1. Member of the Organizing Committees

- B. Błaszczyszyn was in the Organizing Commitee of *Stochastic Geometry Days 2018* 14–18 mai 2018 Paris (France); https://geosto-2018.sciencesconf.org/.
- A. Busic was in the Organizing Commitee of *ALEA Days 2018* 12–16 March 2018 at CIRM, Luminy (France); https://conferences.cirm-math.fr/1776.html

9.1.2. Scientific Events Selection

9.1.2.1. Member of the Editorial Boards

- M. Lelarge: IEEE's Transactions on Network Science and Engineering, Bernoulli Journal, Queueing Systems.

9.1.3. Invited Talks

- Analysis of Algorithms 2018, Uppsala (Sweden), June 2018.
- European Conference on Queueing Theory (ECQT) 2018, Jerusalem, July 2018.
- Rencontres de Probabilités de Rouen 2018, September 2018, invited talk.
- Societal Networks, RTDM program, Simons Institute, UC Berkeley, USA, March 2018, invited talk; video: https://simons.berkeley.edu/talks/ana-busic-3-29-18.
- Advances in Modelling and Control for Power Systems of the Future (CAESARS 2018), EDF Palaiseau, September 2018, invited talk.

9.2. Teaching - Supervision - Juries

9.2.1. Teaching

- Licence: B. Blaszczyszyn (Cours) **Théorie de l'information et du codage** 24 heqTD, L3, ENS, France.
- Licence: A. Busic (Cours) and S. Samain (TD) Structures et algorithmes aléatoires 60heqTD, L3, ENS, France.
- Master: B. Blaszczyszyn (Cours) **Processus ponctuels, graphes aléatoires et géeométrie stochastique** 39heqTD, M2 Probabilités et Modèles Aléatoires, UPMC, France
- Master: A. Busic (Cours) and L. Stephan (TD) **Modéles et algorithmes de réseaux** 60heqTD, M1, ENS, France.
- Master: A. Busic (Cours) **Fondements de la modélisation des réseaux** 18 heqTD, M2 MPRI, France.
- Master: M. Lelarge (Cours) **Deep Learning Do it Yourself**, M1, ENS
- Doctorat: A. Busic (Cours) **Markov chains and exact sampling**, 7heqTD, Ecole thématique CNRS, MathExp 2018 "Mathématiques expérimentales: méthodes et pratiques", 21 mai-1 juin 2018, Saint Flour, France. https://mathexp2018.sciencesconf.org/

9.2.2. Supervision

- PhD: Dalia-Georgiana Herculea "Les Hyperfractales pour la Modelisation des Reseaux sans Fil" since October 2016, defence 21 November 2018; PhD CIFRE co-advised by B. Błaszczyszyn, and Ph. Jacquet.
- PhD: Alexandre Hollocou, defense December 19 2018, Nouvelles approches pour le partitionnement de graphes, co-advised by M. Lelarge and T. Bonald (Telecom ParisTech)
- PhD in progress: Léo Miolane, since 2016, High dimensional statistics, advised by M. Lelarge
- PhD in progress: Alexis Galland, since 2017, Deep Learning on Graphs, advised by M. Lelarge
- PhD in progress: Quentin Le Gall "Crowd networking : modélisation de la connectivité D2D" since October 2017; PhD CIFRE co-advised by B. Błaszczyszyn and E. Cali (Orange).
- PhD in progress: Antoine Brochard "Signal processing for point processes and statistical learning for telecommunications", since September 2018; PhD CIFRE co-advised by B. Błaszczyszyn and Georgios Paschos (Huawei).
- PhD in progress: Md Umar Hashmi, Decentralized control for renewable integration in smartgrids, since from December 2015, advised by A. Busic.
- PhD in progress: Sébastien Samain, Monte Carlo methods for performance evaluation and reinforcement learning, since November 2016, advised by A. Busic.
- PhD in progress: Arnaud Cadas, Dynamic matching models, since October 2017, supervised by A. Busic.

9.2.3. Juries

- B. Błaszczyszyn, member of the PhD defense jury of Dalia-Herculea Popescu; 21 November 2018.
- A. Busic, member of the PhD defense jury of J. Horta (Télécom ParisTech, France); 16 Avril 2018.
- C. Fricker, PhD defense of Farah Slim (Orange- Bretagne Loire University), 13/03/2018.

9.3. Popularization

9.3.1. Internal or external Inria responsibilities

- B. Błaszczyszyn is an ENS adjunct professor since September 2018.
- B. Błaszczyszyn and A. Busic are members of the ENS Computer Science Department Board (Conseil du Laboratoire).
- A. Busic is member of the Committee for the Technological Development, Inria Paris

9.3.2. Education

C. Fricker: member of the jury of *Agrégation de Mathématiques*.

10. Bibliography

Publications of the year

Articles in International Peer-Reviewed Journals

[1] E. ABBE, L. MASSOULIÉ, A. MONTANARI, A. SLY, N. SRIVASTAVA. *Group synchronization on grids*, in "Mathematical Statistics and Learning", September 2018, https://arxiv.org/abs/1706.08561 , https://hal.archives-ouvertes.fr/hal-01940467

[2] I. ADAN, A. BUŠIC, J. MAIRESSE, G. WEISS. *Reversibility and further properties of FCFS infinite bipartite matching*, in "Mathematics of Operations Research", May 2018, vol. 43, n^{o} 2, pp. 347-692, https://arxiv.org/abs/1507.05939 - 27 pages, 12 figures [*DOI :* 10.1287/MOOR.2017.0874], https://hal.inria.fr/hal-01273897

[3] J. BARRE, M. LELARGE, J. MITSCHE. *On rigidity, orientability, and cores of random graphs with sliders*, in "Random Structures & Algorithms", May 2018, vol. 52, n^{o} 3, pp. 419-453, https://hal.archives-ouvertes.fr/hal-01963872

[4] C. BORDENAVE, M. LELARGE, L. MASSOULIÉ. *Nonbacktracking spectrum of random graphs: Community detection and nonregular Ramanujan graphs*, in "Annals of Probability", January 2018, vol. 46, n^{o} 1, pp. 1-71, https://hal.archives-ouvertes.fr/hal-01940486

[5] A. BUŠIC, S. MEYN. *Ordinary Differential Equation Methods for Markov Decision Processes and Application to Kullback–Leibler Control Cost*, in "SIAM Journal on Control and Optimization", February 2018, vol. 56, n^{o} 1, pp. 343-366, https://arxiv.org/abs/1605.04591 [*DOI :* 10.1137/16M1100204], https://hal.archives-ouvertes.fr/hal-01968522

[6] F. CALTAGIRONE, M. LELARGE, L. MIOLANE. *Recovering Asymmetric Communities in the Stochastic Block Model*, in "IEEE Transactions on Network Science and Engineering", July 2018, vol. 5, n^{o} 3, pp. 237-246, https://hal.archives-ouvertes.fr/hal-01963866

[7] A. CHATTOPADHYAY, B. BLASZCZYSZYN, E. ALTMAN. *Location Aware Opportunistic Bandwidth Sharing between Static and Mobile Users with Stochastic Learning in Cellular Networks*, in "IEEE Transactions on Mobile Computing", 2018, https://arxiv.org/abs/1608.04260 - 16 Pages, 1 Figure, 1 Table, https://hal.inria.fr/hal-01401007

[8] A. CHATTOPADHYAY, B. BLASZCZYSZYN, E. ALTMAN. *Two-tier cellular networks for throughput maximization of static and mobile users*, in "IEEE Transactions on Wireless Communications", 2018, https://arxiv.org/abs/1605.07341 - title of the previous version: ”Cell planning for mobility management in heterogeneous cellular networks” [*DOI :* 10.1109/TWC.2018.2887386], https://hal.inria.fr/hal-01331936

[9] A. CHATTOPADHYAY, B. BLASZCZYSZYN, H. P. KEELER. *Gibbsian On-Line Distributed Content Caching Strategy for Cellular Networks*, in "IEEE Transactions on Wireless Communications", February 2018, vol. 17, n^{o} 2, pp. 969 - 981, https://arxiv.org/abs/1610.02318 [*DOI :* 10.1109/TWC.2017.2772911], https://hal.inria.fr/hal-01401010

[10] Y. CHEN, A. BUŠIC, S. MEYN. *Ergodic Theory for Controlled Markov Chains with Stationary Inputs*, in "The Annals of Applied Probability : an official journal of the institute of mathematical statistics", 2018, vol. 28, n^{o} 1, pp. 79-111, https://hal.archives-ouvertes.fr/hal-01672476

[11] Y. CHEN, A. BUŠIC, S. MEYN. *Estimation and Control of Quality of Service in Demand Dispatch*, in "IEEE Transactions on Smart Grid", 2018, vol. 9, n^{o} 5, pp. 5348 - 5356, https://hal.archives-ouvertes.fr/hal-01672458

[12] L. GULIKERS, M. LELARGE, L. MASSOULIÉ. *An impossibility result for reconstruction in the degree-corrected stochastic block model*, in "The Annals of Applied Probability : an official journal of the institute of mathematical statistics", October 2018, vol. 28, n^{o} 5, pp. 3002-3027, https://arxiv.org/abs/1511.00546 , https://hal.archives-ouvertes.fr/hal-01940494

[13] E. KAUFMANN, T. BONALD, M. LELARGE. *A spectral algorithm with additive clustering for the recovery of overlapping communities in networks*, in "Theoretical Computer Science", September 2018, vol. 742, pp. 3-26, https://hal.archives-ouvertes.fr/hal-01963868

[14] L. MASSOULIÉ, K. XU. *On the Capacity of Information Processing Systems*, in "Operations Research", April 2018, vol. 66, n^{o} 2, pp. 568-586, https://hal.archives-ouvertes.fr/hal-01940447

International Conferences with Proceedings

[15] A. BOUILLARD, F. MATHIEU, P. SEHIER, T. DEISS. *Dimensionnement du fronthaul 5G : c'est simple comme un coût de file*, in "CORES 2018 - Rencontres Francophones sur la Conception de Protocoles, l'Évaluation de Performance et l'Expérimentation des Réseaux de Communication", Roscoff, France, May 2018, pp. 1-4, https://hal.inria.fr/hal-01787187

[16] A. BUŠIC, S. MEYN. *Action-Constrained Markov Decision Processes With Kullback-Leibler Cost*, in "Conference On Learning Theory (COLT)", Stockholm, Sweden, July 2018, https://arxiv.org/abs/1807.10244 , https://hal.archives-ouvertes.fr/hal-01968536

[17] N. CAMMARDELLA, J. MATHIAS, M. KIENER, A. BUŠIC, S. MEYN. *Balancing California's Grid Without Batteries*, in "57th IEEE Conference on Decision and Control (CDC 2018)", Miami, United States, December 2018, https://hal.archives-ouvertes.fr/hal-01968606

[18] A. R. COFFMAN, A. BUŠIC, P. BAROOAH. *A Study of Virtual Energy Storage From Thermostatically Controlled Loads Under Time-Varying Weather Conditions*, in "International High Performance Buildings Conference", Purdue, United States, July 2018, https://hal.archives-ouvertes.fr/hal-01968587

[19] A. R. COFFMAN, A. BUŠIC, P. BAROOAH. *Virtual energy storage from TCLs using QoS preserving local randomized control*, in "5th Conference on Systems for Built Environments (BuildSys '18)", Shenzen, China, November 2018, pp. 93-102, https://hal.archives-ouvertes.fr/hal-01968602

[20] P. S. DESTER, C. FRICKER, H. MOHAMED. *Stationary Distribution Analysis of a Queueing Model with Local Choice*, in "29th International Conference on Probabilistic, Combinatorial and Asymptotic Methods for the Analysis of Algorithms (AofA 2018)", Uppsala, Sweden, J. A. FILL, M. D. WARD (editors), LIPIcs-Leibniz International Proceedings in Informatics, June 2018, vol. 110, n^{o} 22 [*DOI :* 10.4230/LIPICS.AOFA.2018.22], https://hal.inria.fr/hal-01666326

[21] M. U. HASHMI, A. BUŠIC. *Limiting Energy Storage Cycles of Operation*, in "10th IEEE Green Technologies Conference (GreenTech 2018)", Austin, TX, United States, April 2018 [*DOI :* 10.1109/GREENTECH.2018.00022], https://hal.archives-ouvertes.fr/hal-01806487

[22] M. U. HASHMI, W. LABIDI, A. BUŠIC, S.-E. ELAYOUBI, T. CHAHED. *Long-Term Revenue Estimation for Battery Performing Arbitrage and Ancillary Services*, in "SmartGridComm 2018 - IEEE International Conference Communications, Control, and Computing Technologies for Smart Grids", Aalborg, Denmark, October 2018, pp. 1-7, https://hal.archives-ouvertes.fr/hal-01867341

[23] M. U. HASHMI, D. MUTHIRAYAN, A. BUŠIC. *Effect of Real-Time Electricity Pricing on Ancillary Service Requirements*, in "Ninth International Conference on Future Energy Systems (ACM e-Energy 2018). Workshops", Karlsruhe, Germany, June 2018, pp. 550-555 [*DOI :* 10.1145/3208903.3214350], https://hal.archives-ouvertes.fr/hal-01968604

[24] P. JACQUOT, C. WAN. *Routing Game on Parallel Networks: the Convergence of Atomic to Nonatomic*, in "CDC 2018 - IEEE 57th Conference on Decision and Control", Miami, United States, Proceedings of the 57th IEEE Conference on Decision and Control, IEEE, December 2018, vol. 1, https://hal.archives-ouvertes.fr/hal-01762547

[25] D. POPESCU, P. JACQUET. *Vehicle-to-Infrastructure Communications Design in Urban Hyperfractals*, in "19th IEEE International Workshop on Signal Processing Advances in Wireless Communications, SPAWC 2018", Kalamata, Greece, June 2018, https://hal.inria.fr/hal-01787602

Conferences without Proceedings

[26] T. BONALD, B. CHARPENTIER, A. GALLAND, A. HOLLOCOU. *Hierarchical Graph Clustering using Node Pair Sampling*, in "MLG 2018 - 14th International Workshop on Mining and Learning with Graphs", London, United Kingdom, August 2018, https://hal.archives-ouvertes.fr/hal-01887669

[27] T. BONALD, A. HOLLOCOU, M. LELARGE. *Weighted Spectral Embedding of Graphs*, in "56th Annual Allerton Conference on Communication, Control, and Computing", Urbana-Champaign, United States, October 2018, https://hal.archives-ouvertes.fr/hal-01887680

[28] A. MUKHOPADHYAY, N. HEGDE, M. LELARGE. *Optimal Content Replication and Request Matching in Large Caching Systems*, in "INFOCOM 2018 - IEEE International Conference on Computer Communications", Honolulu, United States, April 2018, https://arxiv.org/abs/1801.02889 , https://hal.archives-ouvertes.fr/hal-01963864

[29] K. SCAMAN, F. BACH, S. BUBECK, Y. T. LEE, L. MASSOULIÉ. *Optimal Algorithms for Non-Smooth Distributed Optimization in Networks*, in "Advances In Neural Information Processing systems", Montreal, Canada, December 2018, https://arxiv.org/abs/1806.00291 - 17 pages, https://hal.archives-ouvertes.fr/hal-01957013

Scientific Books (or Scientific Book chapters)

[30] B. BLASZCZYSZYN, M. HAENGGI, H. P. KEELER, S. MUKHERJEE. *Stochastic Geometry Analysis of Cellular Networks*, Cambridge University Press, March 2018 [*DOI :* 10.1017/9781316677339], https://hal.inria.fr/hal-01958627

[31] Y. CHEN, M. U. HASHMI, J. MATHIAS, A. BUŠIC, S. MEYN. *Distributed Control Design for Balancing the Grid Using Flexible Loads*, in "IMA Volume on the Control of Energy Markets and Grids", S. MEYN, T. SAMAD, I. HISKENS, J. STOUSTRUP (editors), Springer, 2018, vol. 162, pp. 383-411, https://hal.archives-ouvertes.fr/hal-01656726

Scientific Popularization

[32] P. JACQUET, D. POPESCU, B. MANS. *Broadcast Speedup in Vehicular Networks via Information Teleportation*, in "The 43rd Annual IEEE Conference on Local Computer Networks (LCN 2018)", Chicago, United States, October 2018, https://hal.inria.fr/hal-01963430

[33] P. JACQUET, D. POPESCU, B. MANS. *Energy Trade-offs for end-to-end Communications in Urban Vehicular Networks exploiting an Hyperfractal Model*, in "MSWIM DIVANet", Montreal, Canada, IEEE, October 2018, https://hal.inria.fr/hal-01674685

Patents and standards

[34] F. MATHIEU, A. BOUILLARD. *A Method for Scoring Objects, a Method of Ranking an Object, a Software Product, and an Apparatus*, February 2018, n^{o} 18305203.4, https://hal.archives-ouvertes.fr/hal-01916427

[35] S. MEYN, A. BUŠIC. *Using loads with discrete finite states of power to provide ancillary services for a power grid*, July 2018, n^{o} 15740328, https://hal.archives-ouvertes.fr/hal-01968598

Other Publications

[36] F. BACCELLI, M.-O. HAJI-MIRSADEGHI, A. KHEZELI. *On the Dimension of Unimodular Discrete Spaces, Part I: Definitions and Basic Properties*, January 2019, https://arxiv.org/abs/1807.02980 - working paper or preprint, https://hal.inria.fr/hal-01976265

[37] F. BACCELLI, M.-O. HAJI-MIRSADEGHI, A. KHEZELI. *On the Dimension of Unimodular Discrete Spaces, Part II: Relations with Growth Rate*, January 2019, https://arxiv.org/abs/1808.02551 - working paper or preprint, https://hal.inria.fr/hal-01976257

[38] B. BLASZCZYSZYN, M. K. KARRAY. *Performance analysis of cellular networks with opportunistic scheduling using queueing theory and stochastic geometry*, June 2018, working paper or preprint, https://hal.inria.fr/hal-01824986

[39] B. BLASZCZYSZYN, H. P. KEELER. *Determinantal thinning of point processes with network learning applications*, December 2018, https://arxiv.org/abs/1810.08672 - 8 pages, 6 figures. All code available online: https://github.com/hpaulkeeler/DetPoisson_MATLAB [*DOI :* 10.08672], https://hal.inria.fr/hal-01958924

[40] B. BLASZCZYSZYN, P. MUHLETHALER. *Analyzing LoRa long-range, low-power, wide-area networks using stochastic geometry*, December 2018, https://arxiv.org/abs/1811.01886 - 8 pages, https://hal.inria.fr/hal-01958939

[41] A. BROCHARD, B. BLASZCZYSZYN, S. MALLAT, S. ZHANG. *Statistical learning of geometric characteristics of wireless networks*, 2018, https://arxiv.org/abs/1812.08265 - working paper or preprint, https://hal.inria.fr/hal-01962191

[42] A. CADAS, A. BUŠIC, J. DONCEL. *Optimal Control of Dynamic Bipartite Matching Models*, 2018, https://arxiv.org/abs/1810.08541 - working paper or preprint [*DOI :* 10.08541], https://hal.archives-ouvertes.fr/hal-01968549

[43] A. M. DEVRAJ, A. BUŠIC, S. MEYN. *Zap Meets Momentum: Stochastic Approximation Algorithms with Optimal Convergence Rate*, September 2018, https://arxiv.org/abs/1809.06277 - working paper or preprint, https://hal.archives-ouvertes.fr/hal-01968558

[44] M. U. HASHMI. *Load Flexibility for Price based Demand Response*, November 2018, working paper or preprint, https://hal.archives-ouvertes.fr/hal-01911254

[45] H. HENDRIKX, F. BACH, L. MASSOULIÉ. *Accelerated decentralized optimization with local updates for smooth and strongly convex objectives*, October 2018, working paper or preprint, https://hal.inria.fr/hal-01893568

[46] A. JAGANNATH, P. LOPATTO, L. MIOLANE. *Statistical thresholds for Tensor PCA*, December 2018, https://arxiv.org/abs/1812.03403 - working paper or preprint, https://hal.archives-ouvertes.fr/hal-01960964

[47] H. P. KEELER, B. JAHNEL, O. MAYE, D. ASCHENBACH, M. BRZOZOWSKI. *Disruptive events in high-density cellular networks*, December 2018, working paper or preprint, https://hal.archives-ouvertes.fr/hal-01960597

[48] Q. LE GALL, B. BLASZCZYSZYN, E. CALI, T. EN-NAJJARY. *The Influence of Canyon Shadowing on Device-to-Device Connectivity in Urban Scenario*, October 2018, https://arxiv.org/abs/1810.03498 - 8 pages, 7 figures, 2 tables [*DOI :* 10.03498], https://hal.inria.fr/hal-01891024

[49] L. MIOLANE. *Phase transitions in spiked matrix estimation: information-theoretic analysis*, December 2018, https://arxiv.org/abs/1806.04343 - These notes present in a unified manner recent results (as well as new developments) on the information-theoretic limits in spiked matrix/tensor estimation, https://hal.archives-ouvertes.fr/hal-01960925

[50] L. MIOLANE, A. MONTANARI. *The distribution of the Lasso: Uniform control over sparse balls and adaptive parameter tuning*, December 2018, https://arxiv.org/abs/1811.01212 - 68 pages, 2 figures, https://hal.archives-ouvertes.fr/hal-01960945

[51] P. MOYAL, A. BUŠIC, J. MAIRESSE. *Loynes construction for the extended bipartite matching*, March 2018, https://arxiv.org/abs/1803.02788 - working paper or preprint, https://hal.archives-ouvertes.fr/hal-01968565