

AlgoGT 2013 Learning Algorithms and Dynamics in Distributed Systems

General Program and Book of Abstracts

Tuesday, July 2nd

10:45 - 11:00 **Opening**

11:00 - 12:30 **Technical Session:** Reinforcement Learning in Networks

Session Chair: Panayotis Mertikopoulos (CNRS)

- *Stéphane Sénécal*, The association problem in wireless networks: a Policy Gradient Reinforcement Learning approach
- *Marceau Coupechoux*, Self-Imitation in Cognitive Radio Networks
- *Yezekael Hayel*, Learning algorithms in queueing games

12:30 - 14:00 **Lunch Break**

14:00 - 15:30 **Technical Session:** Dynamical Systems for Learning and Decision-Making

Session Chair: Stéphane Sénécal (Orange Labs)

- *Rida Laraki*, Higher order game dynamics
- *Yannick Viossat*, Links between learning processes in games: no-regret dynamics and fictitious play
- *Ana Busic*, Density classification on infinite lattices and trees

15:30 - 15:50 **Coffee Break**

15:50 - 17:20 **Technical Session:** Information Filtering and Exploitation

Session Chair: Bruno Gaujal (Inria)

- *Nidhi Hegde*, Content curation games in social networks
- *Samson Lasaulce*, Distributed optimization under imperfect monitoring
- *Patrick Loiseau*, Classification games

17:20 - 19:45 **Discussions and social event**

20:00 - 23:00 **Drinks and dinner**

Wednesday, July 3rd

09:00 - 10:30 **Technical Session:** Dynamics and Algorithms for Pricing Problems

Session Chair: Johanne Cohen (PRISM - CNRS)

- *Veronica Belmega*, Hierarchical games and dynamics in Het-Nets pricing problems
- *Onno Zoeter*, Demand based pricing for parking. Theory and results from a first-of-its kind large scale experiment
- *Fabio Martignon*, Joint Pricing and Network Selection Game in Cognitive Radio Networks

10:30 - 10:50 **Coffee Break**

10:50 - 12:20 **Poster Session**

- | | |
|-----------------------------|------------------------------------|
| - <i>Marcella Tambuscio</i> | - <i>Alexandre Reiffers-Masson</i> |
| - <i>Raouia Masmoudi</i> | - <i>Ilaria Brunetti</i> |
| - <i>Lise Rodier</i> | - <i>Cengiz Hasan</i> |
| - <i>Mattia Minelli</i> | - <i>Josu Doncel</i> |
| - <i>François Mériaux</i> | - <i>Victor Ramiro</i> |
| - <i>Nicolas David</i> | |

12:20 - 14:00 **Lunch Break**

14:00 - 15:30 **Technical Session:** Congestion and Resource Allocation

Session Chair: Fabio Martignon (Paris-Sud University, LRI)

- *Cheng Wan*, Coalitions and delegation in congestion games
- *Giacomo Bacci*, An Energy-Efficient Perspective on Contention-Based Synchronization for 3G and 4G Communications
- *Francesco De Pellegrini*, Incentive Mechanisms based on Minority Games in Delay Tolerant Networks

15:30 - 15:45 **Closing and drinks**

Organizing committee:

Corinne Touati, Johanne Cohen, Panayotis Mertikopoulos, Marie-Anne Dauphin-Rizzi, Manon Dalban-Pilon.



Book of Abstracts

Session: Reinforcement Learning in Networks

(Tuesday, 11:00 - 12:30)

Stéphane Sénécal (Orange Labs) *The association problem in wireless networks: a Policy Gradient Reinforcement Learning approach.*

The purpose of these works is to develop a self-optimized association algorithm based on Policy Gradient Reinforcement Learning (PGRL), which is both scalable, stable and robust. The term robust means that performance degradation in the learning phase should be forbidden or limited to predefined thresholds. The algorithm is model-free (as opposed to Value Iteration) and robust (as opposed to Q-Learning). The association problem is modeled as a Markov Decision Process (MDP). The policy space is parameterized. The parameterized family of policies is then used as expert knowledge for the PGRL. The PGRL converges towards a local optimum and the average cost decreases monotonically during the learning process. The properties of the solution make it a good candidate for practical implementation. Furthermore, the robustness property allows to use the PGRL algorithm in an "always-on" learning mode. (Joint work with Richard Combes, Ilham El Bouloumi and Zwi Altman.)

Marceau Coupechoux (Telecom ParisTech)

Self-Imitation in Cognitive Radio Networks.

In this talk, we tackle the problem of resource allocation in cognitive radio networks and we propose a distributed learning mechanism based on self-imitation. The proposed scheme is fully distributed and does not require any information exchange between cognitive radios. The convergence analysis is based on the study of the stochastically stable states of a perturbed Markov chain. (Joint work with Stefano Iellamo and Lin Chen.)

Yezekael Hayel (University of Avignon)

Learning algorithms in queueing games.

Learning algorithms based on reinforcements have proved to be very efficient to propose decentralized protocols for optimization problems in which several decision makers interact and are in competition. Those mechanisms have many properties like convergence and speed if the non-cooperative game has some structural patterns. Interestingly, those algorithms are not used in games where the decision makers are competing into a queueing system. This talk will present some results using that kind of decentralized protocols where several decision makers interact through a queueing system. An application related to last-mile optimization problem in transportation networks will be presented.

Session: Dynamical Systems for Learning and Decision-Making

(Tuesday, 14:00 - 15:30)

Rida Laraki (CNRS - École Polytechnique)

Higher order game dynamics.

Continuous-time game dynamics are typically first order systems where payoffs determine the growth rate of the players' strategy shares. In this talk, we investigate what happens beyond first order by viewing payoffs as higher order forces of change, specifying e.g. the acceleration of the players' evolution instead of its velocity (a viewpoint which emerges naturally when it comes to aggregating empirical data of past instances of play). To that end, we derive a wide class of higher order game dynamics, generalizing first order imitative dynamics, and, in particular, the replicator dynamics. We show that strictly dominated strategies become extinct in n-th order payoff-monotonic dynamics n orders as fast as in the corresponding first order dynamics; furthermore, in stark contrast to first order, weakly dominated strategies also become extinct. All in all, higher order payoff-monotonic dynamics lead to the elimination of weakly dominated strategies, followed by the iterated deletion of strictly dominated strategies, thus providing a dynamic justification of the well-known epistemic rationalizability process of Dekel and Fudenberg (1990). Finally, we also establish a higher order analogue of the folk theorem and we show that convergence to strict equilibria in n-th order dynamics is n orders as fast as in first order.

Yannick Viossat (Université Paris-Dauphine)

Links between learning processes in games: no-regret dynamics and fictitious play.

No-regret dynamics and fictitious play are key learning processes in games, and may be of interest for distributed optimization. Though apparently quite different, these processes are connected. Indeed, wide classes of no-regret dynamics are perturbed versions of continuous fictitious play. We will explore the intuition behind these and their applications of these findings.

Ana Busic (Inria)

Density classification on infinite lattices and trees.

Consider an infinite graph with nodes initially labeled by independent Bernoulli random variables of parameter p . We address the density classification problem, that is, we want to design a deterministic or probabilistic dynamical system with a local and homogeneous updating rule that evolves on this graph and decides whether p is smaller or larger than $1/2$. Precisely, the trajectories should converge to the uniform configuration with only 0's if $p < 1/2$, and only 1's if $p > 1/2$. Two natural instantiations of dynamical systems are considered, one with synchronous updates of node labels (probabilistic cellular automaton), and one with asynchronous updates (finite range interacting particle system). The difficulty is twofold: first, it is impossible to centralize the information (nodes are indistinguishable); second, it is impossible to use classical counting techniques (labels contain only binary information). We present solutions to the problem on the regular grids of dimension d , for any $d > 1$, and on the regular infinite trees. For the bi-infinite line, we propose some candidates that we back up with numerical simulations. (Joint work with N. Fatès, J. Mairesse, and I. Marcovici.)



Session: Information Filtering and Exploitation

(Tuesday, 15:50 - 17:20)

Nidhi Hegde (Technicolor)*Content curation games in social networks.*

Social networks offer users new means of accessing information, essentially relying on "social filtering", i.e. propagation and filtering of information by social contacts. The sheer amount of data flowing in these networks, combined with the limited budget of attention of each user, makes it difficult to ensure that social filtering brings relevant content to the interested users. I will present two perspectives on filtering, or curating, of relevant content under a limited budget of attention. First, I will consider how users may self-organize their connections to receive content of interest to them. To this end I will introduce flow games, a simple abstraction that models network formation under selfish user dynamics, featuring user-specific interests and budget of attention. Second, I will consider curating of a user's personal stream when sourced from aggregators. Here I will use the framework of utility games to demonstrate the efficiency of incentive-based curation mechanisms.

Samson Lasaulce (CNRS)*Distributed optimization under imperfect monitoring.*

The goal of this presentation is to present a result which has been derived by using information-theoretic tools and allows one to know to what extent two decision makers can coordinate when a given imperfect monitoring structure is available. Numerical results for the case of power allocation in wireless channels will be presented.

Patrick Loiseau (EURECOM)*Classification games.*

Classification games are games between a player (the defender) who performs a classification task and a player (the attacker) who controls part of the training dataset. Such games have many potential applications in areas where machine learning algorithms are used in a partially adversarial environment (e.g., spam filtering). As a simple example, consider a toy model where a defender classifies an attacker between two possible types: a dangerous spy or a benign spammer. Attackers hit one of two servers ("mail" and "file") at each discrete time instant and the classification is based on the number of file server hits observed in a fixed time window. The spammer naively hits both servers with a commonly known distribution (more concentrated on his main target, the mail server). The game is then between the spy and the defender. The defender chooses a classification threshold on the number of file server hits, to balance the cost of missed spy detection and the cost of false alarm. The spy chooses the number of hits on the file server (his main target) to balance the benefit of file server hits and the cost of being classified as spy (and therefore more heavily tracked!). In this talk, we give results on the structure of Nash equilibria of classification games in mixed strategies for a family of different nonzero-sum payoff functions and we discuss on-going work. (Joint work with Lemonnia Dritsoula and John Musacchio (UC Santa Cruz))

Session: Dynamics and Algorithms for Pricing Problems

(Wednesday, 09:00 - 10:30)

Veronica Belmega (ETIS / ENSEA - UCP - CNRS)*Hierarchical games and dynamics in HetNets pricing problems.*

The strategic deployment of small and femto-cells, overlaid on existing wireless communication infrastructure is envisaged as a key technology enabling network providers to offer a wide variety of services to their customers. Several challenges arise due to the selfish behavior of the customers wishing to optimize their chosen service's quality-price tradeoff. Hence, the behavior and choices of the customers must be anticipated by the providers in order to select the most profitable pricing policy. A game-theoretic model of the complex interactions between customers and providers, i.e., the Stackelberg formulation in which the service providers act as leaders and the customers as followers, will be investigated. (Joint work with Luca Rose, Walid Saad and Merouane Debbah.)

Onno Zoeter (Xerox Research Centre Europe)*Demand based pricing for parking. Theory and results from a first-of-its kind large scale experiment.*

The LA ExpressPark project is a large scale project in downtown LA that puts demand based pricing principles to the test. In LA more than 6000 on-street parking spaces have been equipped with sensors. This data can be used to understand the underlying patterns in parking. This is a first for such an important part of daily life. In addition, the data can be used to guide the pricing of parking, to try to encourage parkers to avoid peak hours and peak locations. In a sense such a system delivers upon the promise made by William Vickrey (Nobel laureate in 1996) in the 1950s. He argued eloquently that publically owned utilities, even though they might be paid for and owned by the general public, should not be put available to use for free to ensure most efficient use. For parking he proposed an ex-post parking meter. In this talk I will present an ex-ante payment mechanism, a posted price mechanism, and a practical way to learn the optimal prices iteratively. The final method balances the need to stay close to the patterns in demand, and the need to be simple enough to be memorized by drivers. It also improves over the most basic versions of demand based pricing as used in similar projects. Our methods have been in use in LA since June 2012. If time permits I will also share some of our initial results. (Joint work with Eduardo Cardenas, Stephane Clinchant, Damien Cramet, Chris Dance, Honglu Du, James Glasnapp, Philippe Rerole, and Frank Torres.)

Fabio Martignon (Paris-Sud University, LRI)*Joint Pricing and Network Selection Game in Cognitive Radio Networks .*

We address the joint pricing and network selection problem in cognitive radio networks, considering both the point of view of Primary/Secondary operators and network users. The problem is formulated as a Stackelberg game where first the Primary and Secondary operators set the network subscription price to maximize their revenue. Then, users perform the network selection process, deciding



whether to pay more for a guaranteed service, or use a cheaper, best-effort secondary network, where congestion and low throughput may be experienced. We derive optimal stable price and network selection settings, investigating the efficiency of the equilibria reached by operators and users through the determination of the Price of Anarchy. Furthermore, we study network users dynamics using a population game model, and determine its convergence properties under replicator dynamics.

Poster Session**(Wednesday, 10:50 - 12:20)**

Marcella Tambuscio (Università di Pisa - BCAM), *Worst-case (in)efficiency of a congested duopoly on equilibria and Edgeworth cycles*

Raouia Masmoudi (ETIS/ ENSEA, University of Cergy-Pontoise), *Modified Iterative Water-filling Algorithm for the Power Minimization Problem under QoS and Cognitive Radio Interference Constraints*

Lise Rodier (PRISM), *SLA learning from past failures, a Multi-Armed Bandit approach*

Mattia Minelli (ENST), *Simulated Annealing Algorithm for Optimal Relay Placement in Cellular Networks*

François Mériaux (Laboratoire des Signaux et Systèmes), *Achievability of efficient satisfaction equilibria in self-configuring networks*

Nicolas David (LIA), *Simulations of learning methods for noncooperative queueing systems*

Alexandre Reiffers-Masson (Inria/LIA), *Pricing Agreement between Service and Content Provider: A Net Neutrality Issue*

Ilaria Brunetti (Inria Sophia Antipolis), *Revisiting Evolutionary Games*

Cengiz Hasan (Inria), *Matching Games in Networking Problems*

Josu Doncel (LAAS-CNRS), *On the Efficiency of Non-Cooperative Load Balancing*

Victor Ramiro (ISAE / University of Toulouse), *On the limits of DTN monitoring*

Session: Congestion and Resource Allocation**(Wednesday, 14:00 - 15:30)**

Cheng Wan (UPMC - Paris 6, IMJ)

Coalitions and delegation in congestion games.

In a network congestion game, a nonatomic player holds an infinitesimal flow while a splittable (resp. integer splittable) atomic player holds a flow which can be divided arbitrarily (resp. into several parts of integer weight) and sent by different paths. The first part of the talk is about the impact of the coalitions in composite congestion games where nonatomic and atomic splittable players coexist. Two-terminal parallel-path network is considered. Once the existence and the uniqueness of composite equilibrium are established, several results on the connection between the composition of the player set and the equilibrium costs will be presented. The second part of the talk concerns the behavior of delegation in integer splittable congestion games. A new class of games will be introduced: delegation games, where the players delegate their tasks to several independent players. Delegation equilibrium and consistent delegation equilibrium will be defined and their existence will be shown.

Giacomo Bacci (University of Pisa and Princeton University) *An Energy-Efficient Perspective on Contention-Based Synchronization for 3G and 4G Communications.*

Since the early days of wireless communications, radio resource management has emerged as a key issue in network design. In the era of wireless computing, the need for an efficient allocation of the resources (namely, power, bandwidth, and computational complexity) has become more and more challenging. A smart management of the available resources should include not only the data detection phase, that has received a remarkable interest in the last decade, but also the initial phase of network association, during which the terminals exchange some information with the base station of infrastructured networks, and eventually get access to the system. In the first part of this talk, we will focus on CDMA-based communications, adopting the analytical tools of game theory to derive an iterative and distributed algorithm to trade off satisfactory QoS and prolonging battery life of uplink users. We will then extend this approach to contention-based synchronization for 4G communications, encompassing both OFDMA and SC-FDMA technologies. In this context, we formulate a game-theoretic framework to identify optimization criteria for the system design and to derive a low-complexity and scalable resource allocation procedure based on estimated parameters.

Francesco De Pellegrini (CREATE-NET)

Incentive Mechanisms based on Minority Games in Delay Tolerant Networks.

In this talk we describe an incentive mechanism for heterogeneous Delay Tolerant Networks (DTNs). The proposed mechanism tackles a core problem of such systems: how to induce coordination of DTN relays in order to achieve a target performance figure, e.g., delivery probability or end-to-end delay, under a given constraint in term of network resources, e.g., number of active nodes or energy consumption. Also, we account for the realistic case when the cost for taking part in the forwarding process varies with the devices' technology or the users' habits. Finally, the scheme is truly applicable to DTNs since it works with no need for end-to-end connectivity and it does not require stringent assumptions on the mobility pattern. In this context, one can introduce a basic coordination mechanism leveraging the notion of a Minority Game. In this game, relays compete to be in the population minority and their utility is defined in combination with a rewarding mechanism. The rewards in turn configure as a control by which the network operator controls the desired operating point for the DTN. To this aim, the equilibria of the game in the case of heterogeneous DTNs can be determined. Finally, a learning algorithm based on stochastic approximations provably drives the system to the equilibrium without requiring perfect state information at relay nodes or at the source node and without using end-to-end communications to implement the rewarding scheme.

