

Fluid Spectrum Access – ERC grant 308267

Mid-term Scientific Report

Alexandre Proutiere

June 29, 2015

In this report, we highlight the scientific contributions of the first half of the ERC research project Fluid Spectrum Access. Most of the outcomes precisely correspond to the tasks planned in the application (see); but our work also yielded unexpected and interesting developments (see). The next four sections present the work done within each task (refer to the the ERC FSA application for a description of the work planned in each task). The last section provides a list of keynote / invited lectures / tutorials I gave on topics related to the project.

1 Task 1: Optimal Spectrum Exploration Algorithms

In this first task, we planned to investigate simple scenarios where a single transmitter explores the various available radio channels, with the objective of tracking the channel offering the best radio conditions. The main challenge in addressing this issue lies in the facts that (i) the quality of the various channels randomly varies over time in an unpredictable manner, and (ii) the transmitter gets information about the quality of the channel it is currently using only. In addition, it should be noted that the transmitter not only has to select a channel to transmit on, but also a coding rate and possibly a MIMO mode (if the transmitter has several antennas) depending on the quality of the selected channels. In summary, we aim at designing sequential channel and rate selection strategies maximizing throughput, i.e., the number of packets successfully transmitted over a given time horizon. This design problem can be naturally formalized as a bandit optimization problem where the set of decisions or *arms* covers the possible channel and rate pairs. The resulting bandit problem differs from the classical stochastic bandit problems in two ways. First, the various decisions are correlated in the sense that the function mapping the set of decisions to the corresponding average rewards, here the throughput, exhibits some structure that one should exploit to speed up the exploration and the selection of the best decision. For example, we know that on a given channel, the packet successful transmission probability is a decreasing function of the rate; furthermore, this function turns out to be unimodal. Then, in the bandit literature, the environment is assumed to be stationary: the packet successful transmission probabilities at the various (channel, rate) pairs do not evolve over time. In practice this is not the case, and the environment is typically non-stationary. In summary, the design of optimal sequential channel and rate selection strategies can be seen as a structured non-stationary bandit problem.

We spent considerable efforts revisiting and developing the theory of bandit optimization, and we then applied our results to the design of a provably optimal (channel, rate) selection protocols. Our approach significantly departs from existing techniques to design such protocols as all existing solutions rely on heuristics. The results obtained for bandit optimization have been presented at the three major machine learning conferences, namely NIPS, ICML, and COLT, whereas their applications to wireless systems are published in selective conferences and journals on communication systems (IEEE Infocom, IEEE JSAC, IEEE trans. on Networking). We describe below our contributions in more detail.

1.1 Bandit Optimization

Over the last decade, structured Multi-Armed Bandit (MAB) problems have attracted a lot of attention in the Machine Learning community, and in spite of the recent and abundant developments, our understanding of these problems remains very preliminary. For example, for the most common structures (e.g. linear, Lipschitz, convex, submodular), the solutions and algorithms devised to solve the problems have been driven by heuristics, and only exhibit weak performance guarantees. In contrast, the approach we proposed recently aims at first identifying (problem-specific) fundamental performance limits that no algorithm can exceed, and at then designing algorithms whose performance match these limits. To this aim, we have developed novel statistical tools, related to and extending results in hypothesis testing. Our versatile approach was applied to specific structures, including unimodal [6], Lipschitz [11], and combinatorial [4], but can readily be extended to any kind of structure. We also developed tools to handle non-stationary environments, and in particular proposed a generic method to extend algorithms designed to be optimal in stationary environments to the non-stationary setting [6].

Within this project, we also solved the so-called infinite-armed bandit problem [1]. Here the number of decisions is much larger than the time horizon, and all decisions cannot be explored. By solving, we mean that we identified fundamental performance limits satisfied by any sequential decision strategy, and devised an algorithm that matches these limits. This kind of bandit problem can be useful for spectrum exploration with delay-sensitive application (the transmitter cannot explore all resources before the packet it has to transmit becomes irrelevant).

Beyond the theoretical developments and their applications to the design of resource allocation schemes in wireless networks, we also applied our results on bandits in various fields, including online service (e.g., search engines) and quantum computing [5].

1.2 Optimal Rate Adaptation in 802.11 Systems

We first applied our results on structured bandits to the design of optimal rate sampling algorithms for 802.11 systems (WiFi). These algorithms constitute a crucial component of 802.11 systems, and yet the design of all existing algorithms (e.g. ARF) has been driven by heuristics. For the first time, we proposed a systematic formalism that leads to the best possible algorithm. Existing rate adaptation algorithms also turned out to perform poorly when the transmitter has several antennas and when a rate along with a MIMO mode have to be selected for the transmission of each packet. In contrast, our proposed algorithm, referred to as ORS (Optimal Rate

Sampling), also behaves optimally in this setting. We have verified the superiority of our algorithm through simulations and test-bed experiments. As an example, we present in Figure 1 the way our algorithm is able to track the best rate for transmission in 802.11g systems when the latter evolves over time (due for instance to the mobility of the transmitter). ORF clearly tracks the best rate, and outperforms state-of-the-art algorithms (here SampleRate). Refer to [8] for further experiments on systems with the latest standards (e.g. 802.11ac).

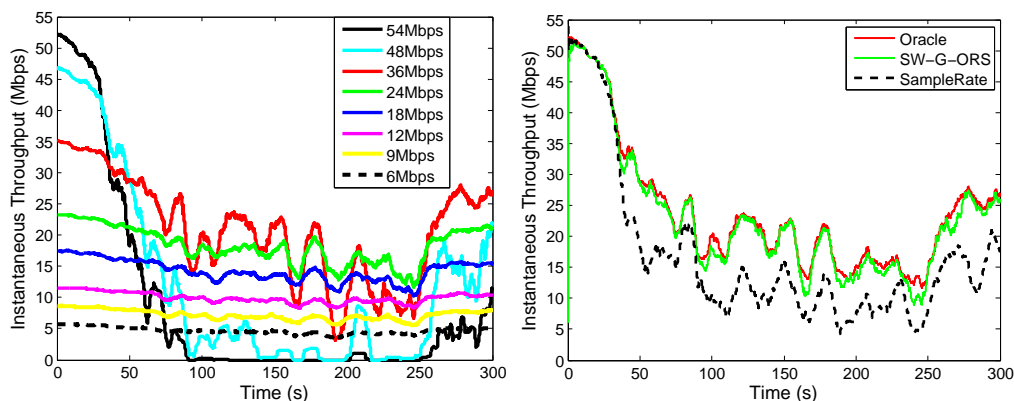


Figure 1: 802.11g test-bed traces – non-stationary environment. (Left) Throughputs obtained using the various rates vs. time. These curves would be available to an Oracle that hence could exactly track the best rate for transmission. In practice these curves are unknown. (Right) Throughput under our algorithm, ORS, SampleRate, and the Oracle algorithm.

Our results have been published at IEEE Infocom 2014 [9], and a journal version has been submitted to IEEE/ACM Transactions on Networking [8]. We are also currently trying to patent the algorithm, and influence the future 802.11 standards.

1.3 Optimal Rate and Channel Selection in Wireless Systems

The results and algorithms for rate adaptation have been extended to the case where the transmitter selects both a rate and a channel for the transmission of each packet. Again, we have designed optimal algorithms, significantly outperforming possible existing algorithms. This work is published in a IEEE JSAC special issue on Cognitive Radio Systems [7].

2 Task 2: Towards Decentralized Learning for Dynamic Spectrum Sharing

When multiple transmitters compete for the use of the various channels, we have to come up with schemes that learn the best allocation of transmitters to channels. This optimal allocation depends on the radio conditions for the various transmitters on the different channels, but also on the trade-off between efficiency and fairness we wish to achieve. Here we have chosen Proportional Fairness leading to a good and robust balance between total system throughput and

fairness. The challenges in identifying such allocations are (i) that the channel conditions for the various (transmitter, channel) pairs are initially unknown and may evolve over time, and (ii) that each transmitter needs to select a channel based on its very limited information about the entire system – each transmitter could well be aware of its current service rate (the rate achieved on the selected channel) only.

2.1 A Pure Learning Approach

To address (i), we can formalize the problem of spectrum allocation as a combinatorial bandit problem. Here the decision maker needs to select in each step a matching in the bi-partite graph consisting on the transmitters on one side, and of the channels on the other side. We have investigated generic combinatorial bandit problems in [10] and [4], and have derived fundamental performance limits, as well as algorithms that yield better performance guarantees than any other algorithm. Observe that we do assume here that the allocation of transmitters to channels is decided in a centralized manner. It remains unclear how to implement our algorithm in a decentralized manner.

2.2 Decentralized Optimization with Learning

To address (ii), we devise a simple randomised algorithm under which each transmitter periodically switches channels when it foresees a possible improvement in its throughput. In turn, under this simple scheme, the system converges to the Nash Equilibriums of a load balancing game. We have also shown that these equilibriums get closer to the ideal Proportional fair allocation when the number of transmitters increases. we are actually able to quantify at any time, how far is the channel allocation under our distributed scheme from the PF allocation. the gap is small and vanishes rapidly. Refer to [12], and [13] for more details.

3 Task 3: Handling Systems Heterogeneity

In this task, we planned to examine how heterogeneity could challenge the design of efficient and distributed radio resource sharing policies. So far, we have consider the case of heterogeneous transmit power. More precisely, we have looked at networks, where each transmitter has its own maximum transmit power, and wish to share resources (time slots or channels) in a fully distributed manner under the realistic SINR interference model. By fully distributed, we mean that transmitters are not allowed to exchange any signalling message, and the only feedback available at a given transmitter is the level of interference measured at the corresponding receiver on the various channels (just as in Foschini-Miljanic classical power control mechanisms). In this scenario, we aim at rate-optimal designing distributed radio resource sharing policies. Rate-optimality is defined as follows. The rate region is defined as the set of rates that can be simultaneously achieved on the various links using some (centralized) resource sharing policies. A policy is rate-optimal if it can achieve any rate vector within the rate region. The design of distributed rate-optimal schemes have been open since the 90's, and the design by Tassiulas-Ephremides of the max weight policy, a centralized rate-optimal scheme. Recently,

some researchers proposed a simple, distributed, and rate-optimal adaptive CSMA scheme that turns out to be rate-optimal under the simplistic assumption that interference can be modelled as a simple undirected graph. Finding such schemes under the realistic SINR interference model is much more challenging. Leveraging recent results in learning in games, we were able to devise such a scheme for the first time. This constitutes a major breakthrough. Refer to [2], [3] for more details.

4 Tasks 4 and 5: Simulation Platform and Real Experiments

We have developed a simulation platform, and a test-bed to test the various algorithms proposed in the project. In particular we have implemented the bandit algorithms described in Section 1 in a 802.11g and 802.11n test-bed [9], [8]. The development has been made by a Korean student, Donggyu Yun, who unfortunately went back to KAIST in Korea. We are looking for new developers to extend his work, and continue our implementation efforts. We are also starting collaborating with researchers in James Gross' group at KTH to implement our proposed algorithms in their Software-Defined-Radio (SDR) platform.

5 Dissemination

We list below invited keynotes, tutorials, and invited course given by myself on topics related to the project.

- Keynote speaker at ISWCS (<http://www.iswcs2015.org/index.php/keynote-speakers>), August 2015
- Keynote speaker at ITAs (<http://itas2015.iitp.ru/en/index.html>), September 2015
- Keynote speaker at the Lunteren Conference (<http://www.lnmb.nl/conferences/2015>), January 2015
- Invited speaker at LCCC workshop "Dynamics and Control of Networks", Lund, Sweden, October 2014
- Invited course on Network Optimization, Lipari School of Computer Science, July 2014
- Keynote speaker at WiOpt (<http://www.hanalab.org/wiopt2014>), May 2014
- Invited course on Bandit Optimization and Distributed Optimization at KAIST, Korea, February 2014
- Distinguished lecture on Bandit Optimization in the LCCC seminar, Lund, Sweden, October 2013
- Lecture on Bandit Optimization in the Advanced Networks Colloquia series, UMD, USA, October 2013
- Invited course at the program on Stochastic Processes in Engineering, IIT Bombay, India, March 2013

References

- [1] T. Bonald and A. Proutiere. Two-target algorithms for infinite-armed bandits with bernoulli rewards. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States.*, pages 2184–2192, 2013.
- [2] P. Chaporkar and A. Proutiere. Optimal distributed scheduling in wireless networks under SINR interference model. In *2013 51st Annual Allerton Conference on Communication, Control, and Computing, Allerton Park & Retreat Center, Monticello, IL, USA, October 2-4, 2013*, pages 1372–1379, 2013.
- [3] P. Chaporkar, A. Proutiere, and S. Magureanu. Optimal distributed scheduling in wireless networks under SINR interference model. *IEEE Transactions on Networking*, 2015 to appear.
- [4] R. Combes, M. Lelarge, A. Proutiere, and M. S. Talebi. Stochastic and adversarial combinatorial bandits. *CoRR*, abs/1502.03475, 2015.
- [5] R. Combes, S. Magureanu, A. Proutiere, and C. Laroche. Learning to rank: Regret lower bounds and efficient algorithms. In *Proceedings of the 2015 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems, Portland, OR, USA, June 15-19, 2015*, pages 231–244, 2015.
- [6] R. Combes and A. Proutiere. Unimodal bandits: Regret lower bounds and optimal algorithms. In *Proc. of ICML*, 2014.
- [7] R. Combes and A. Proutiere. Dynamic rate and channel selection in cognitive radio systems. *IEEE Journal on Selected Areas in Communications*, 33(5):910–921, 2015.
- [8] R. Combes, A. Proutiere, D. Yun, J. Ok, and Y. Yi. Optimal rate sampling in 802.11 systems. <http://arxiv.org/abs/1307.7309>, 2013.
- [9] R. Combes, A. Proutiere, D. Yun, J. Ok, and Y. Yi. Optimal rate sampling in 802.11 systems. In *Proc. of IEEE INFOCOM*, 2014.
- [10] M. Lelarge, A. Proutiere, and M. S. Talebi. Spectrum bandit optimization. In *2013 IEEE Information Theory Workshop, ITW 2013, Sevilla, Spain, September 9-13, 2013*, pages 1–5, 2013.
- [11] S. Magureanu, R. Combes, and A. Proutiere. Lipschitz bandits: Regret lower bound and optimal algorithms. In *Proceedings of The 27th Conference on Learning Theory, COLT 2014, Barcelona, Spain, June 13-15, 2014*, pages 975–999, 2014.
- [12] S. Yun and A. Proutiere. Distributed load balancing in heterogenous systems. In *48th Annual Conference on Information Sciences and Systems, CISS 2014, Princeton, NJ, USA, March 19-21, 2014*, pages 1–6, 2014.

- [13] S. Yun and A. Proutiere. Distributed proportional fair load balancing in heterogenous systems. In *Proceedings of the 2015 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems, Portland, OR, USA, June 15-19, 2015*, pages 17–30, 2015.