

Delay and Cooperation in Nonstochastic Bandits

Nicolò Cesa-Bianchi

Università degli Studi di Milano, Italy

NICOLO.CESA-BIANCHI@UNIMI.IT

Claudio Gentile

University of Insubria, Italy

CLAUDIO.GENTILE@UNINSUBRIA.IT

Yishay Mansour

Microsoft Research and Tel-Aviv University, Israel

MANSOUR@TAU.AC.IL

Alberto Minora

University of Insubria, Italy

AMINORA@UNINSUBRIA.IT

Abstract

We study networks of communicating learning agents that cooperate to solve a common non-stochastic bandit problem. Agents use an underlying communication network to get messages about actions selected by other agents, and drop messages that took more than d hops to arrive, where d is a delay parameter. We introduce EXP3-COOP, a cooperative version of the EXP3 algorithm and prove that with K actions and N agents the average per-agent regret after T rounds is at most of order $\sqrt{(d+1 + \frac{K}{N}\alpha_{\leq d})(T \ln K)}$, where $\alpha_{\leq d}$ is the independence number of the d -th power of the communication graph G . We then show that for any connected graph, for $d = \sqrt{K}$ the regret bound is $K^{1/4}\sqrt{T}$, strictly better than the minimax regret \sqrt{KT} for noncooperating agents. More informed choices of d lead to bounds which are arbitrarily close to the full information minimax regret $\sqrt{T \ln K}$ when G is dense. When G has sparse components, we show that a variant of EXP3-COOP, allowing agents to choose their parameters according to their centrality in G , strictly improves the regret. Finally, as a by-product of our analysis, we provide the first characterization of the minimax regret for bandit learning with delay.

1. Introduction

Delayed feedback naturally arises in many sequential decision problems. For instance, a recommender system typically learns the utility of a recommendation by detecting the occurrence of certain events (e.g., a user conversion), which may happen with a variable delay after the recommendation was issued. Other examples are the communication delays experienced by interacting learning agents. Concretely, consider a network of geographically distributed ad servers using real-time bidding to sell their inventory. Each server sequentially learns how to set the auction parameters (e.g., reserve price) in order to maximize the network's overall revenue, and shares feedback information with other servers in order to speed up learning. However, the rate at which information is exchanged through the communication network is slower than the typical rate at which ads are

Extended abstract. Full version appears as Arxiv:1602.04741,v2.

served. This causes each learner to acquire feedback information from other servers with a delay that depends on the network’s structure.

Motivated by the ad network example, we consider networks of learning agents that cooperate to solve the same nonstochastic bandit problem, and study the impact of delay on the global performance of these agents. We introduce the EXP3-COOP algorithm, a distributed and cooperative version of the EXP3 algorithm of [Auer et al. \(2002\)](#). EXP3-COOP works within a distributed and synced model where each agent runs an instance of the same bandit algorithm (EXP3). All bandit instances are initialized in the same way irrespective to the agent’s location in the network (that is, agents have no preliminary knowledge of the network), and we assume the information about an agent’s actions is propagated through the network with a unit delay for each crossed edge. In each round t , each agent selects an action and incurs the corresponding loss (which is the same for all agents that pick that action in round t). Besides observing the loss of the selected action, each agent obtains the information previously broadcast by other agents with a delay equal to the shortest-path distance between the agents. Namely, at time t an agent learns what the agents at shortest-path distance s did at time $t - s$ for each $s = 1, \dots, d$, where d is a delay parameter. In this scenario, we aim at controlling the growth of the regret averaged over all agents (the so-called average welfare regret).

In the noncooperative case, when agents ignore the information received from other agents, the average welfare regret grows like \sqrt{KT} (the minimax rate for standard bandit setting), where K is the number of actions and T is the time horizon. We show that, using cooperation, N agents with communication graph G can achieve an average welfare regret of order $\sqrt{(d + 1 + \frac{K}{N}\alpha_{\leq d})(T \ln K)}$. Here $\alpha_{\leq d}$ denotes the independence number of the d -th power of G (i.e., the graph G augmented with all edges between any two pair of nodes at shortest-path distance less than or equal to d). When $d = \sqrt{K}$ this bound is at most $K^{1/4}\sqrt{T \ln K} + \sqrt{K}(\ln T)$ for any connected graph—see Remark 7 in Section 4.1—which is asymptotically better than \sqrt{KT} .

Networks of nonstochastic bandits were also investigated by [Awerbuch and Kleinberg \(2008\)](#) in a setting where the distribution over actions is shared among the agents without delay. [Awerbuch and Kleinberg \(2008\)](#) prove a bound on the average welfare regret of order $\sqrt{(1 + \frac{K}{N})T}$ ignoring polylog factors.¹ We recover the same bound as a special case of our bound when G is a clique and $d = 1$. In the clique case our bound is also similar to the bound $\sqrt{\frac{K}{N}(T \ln K)}$ achieved by [Seldin et al. \(2014\)](#) in a single-agent bandit setting where, at each time step, the agent can choose a subset of $N \leq K$ actions and observe their loss. In the case when $N = 1$ (single agent), our analysis can be applied to the nonstochastic bandit problem where the player observes the loss of each played action with a delay of d steps. In this case we improve on the previous result of $\sqrt{(d + 1)KT}$ by [Neu et al. \(2010, 2014\)](#), and give the first characterization (up to logarithmic factors) of the minimax regret, which is of order $\sqrt{(d + K)T}$.

In principle, the problem of delays in online learning could be tackled by simple reductions. Yet, these reductions give rise to suboptimal results. In the single agent setting, where the delay is constant and equal to d , one can use the technique of [Weinberger and Ordentlich \(2002\)](#) and run $d + 1$ in-

1. The rate proven in ([Awerbuch and Kleinberg, 2008](#), Theorem 2.1) has a worse dependence on T , but we believe this is due to the fact that their setting allows for dishonest agents and agent-specific loss vectors.

stances of an online algorithm for the nondelayed case, where each instance is used every $d+1$ steps. This delivers a suboptimal regret bound of $\sqrt{(d+1)KT}$. In the case of multiple delays, like in our multi-agent setting, one can repeat the same action for $d+1$ steps while accumulating information from the other agents, and then perform an update on scaled-up losses. The resulting (suboptimal) bound on the average welfare regret would be of the form $\sqrt{(d+1)(1 + \frac{K}{N}\alpha_{\leq d})(T \ln K)}$.

Rather than using reductions, the analysis of EXP3-COOP rests on quantifying the performance of suitable importance weighted estimates. In fact, in the single-agent setting with delay parameter d , using EXP3-COOP reduces to running the standard EXP3 algorithm performing an update as soon a new loss becomes available. This implies that at any round $t > d$, EXP3 selects an action without knowing the losses incurred during the last d rounds. The resulting regret is bounded by relating the standard analysis of EXP3 to a detailed quantification of the extent to which the distribution maintained by EXP3 can drift in d steps.

In the multi-agent case, the importance weighted estimate of EXP3-COOP is designed in such a way that at each time $t > d$ the instance of the algorithm run by an agent v updates all actions that were played at time $t - d$ by agent v or by other agents not further away than d from v . Compared to the single agent case, here each agent can exploit the information circulated by the other agents. However, in order to compute the importance weighted estimates used locally by each agent, the probabilities maintained by the agents must be propagated together with the observed losses. Here, further concerns may show up, like the amount of communication, and the location of each agent within the network. In particular, when G has sparse components, we show that a variant of EXP3-COOP, allowing agents to choose their parameters according to their centrality within G , strictly improves on the regret of EXP3-COOP.

2. Additional Related Work

Many important ideas in delayed online learning, including the observation that the effect of delays can be limited by controlling the amount of change in the agent strategy, were introduced by Mesterharm (2005) —see also (Mesterharm, 2007, Chapter 8). A more recent investigation on delayed online learning is due to Neu et al. (2010, 2014), who analyzed exponential weights with delayed feedbacks. Further progress is made by Joulani et al. (2013), who also study delays in the general partial monitoring setting. Additional works (Joulani et al., 2016; Quanrud and Khashabi, 2015) prove regret bounds for the full-information case of the form $\sqrt{(D+T) \ln K}$, where D is the total delay experienced over the T rounds. In the stochastic case, bandit learning with delayed feedback was considered by Dudík et al. (2011); Joulani et al. (2013).

To the best of our knowledge, the first paper about nonstochastic cooperative bandit networks is (Awerbuch and Kleinberg, 2008). More papers analyze the stochastic setting, and the closest one to our work is perhaps (Szorenyi et al., 2013). In that paper, delayed loss estimates in a network of cooperating stochastic bandits are analyzed using a dynamic P2P random networks as communication model. A more recent paper is (Landgren et al., 2015), where the communication network is a fixed graph and a cooperative version of the UCB algorithm is introduced which uses a distributed consensus algorithm to estimate the mean rewards of the arms. The main result is an individual (per-agent) regret bound that depends on the network structure without taking delays into account.

Another interesting paper about cooperating bandits in a stochastic setting is (Kar et al., 2011). Similar to our model, agents sit on the nodes of a communication network. However, only one designated agent observes the rewards of actions he selects, whereas the others remain in the dark. This designated agent broadcasts his sampled actions through the networks to the other agents, who must learn their policies relying only on this indirect feedback. The paper shows that in any connected network this information is sufficient to achieve asymptotically optimal regret. Cooperative bandits with asymmetric feedback are also studied by Barrett and Stone (2011). In their model, an agent must teach the reward distribution to another agent while keeping the discounted regret under control. Tekin and van der Schaar (2015) investigate a stochastic contextual bandit model where each agent can either privately select an action or have another agent select an action on his behalf. In a related paper, Tekin et al. (2014) look at a stochastic bandit model with combinatorial actions in a distributed recommender system setting, and study incentives among agents who can now recommend items taken from other agents’ inventories. Another line of relevant work involves problems of decentralized bandit coordination. For example, Stranders et al. (2012) consider a bandit coordination problem where the the reward function is global and can be represented as a factor graph in which each agent controls a subset of the variables. A parallel thread of research concerns networks of bandits that compete for shared resources. A paradigmatic application domain is that of cognitive radio networks, in which a number of channels are shared among many users and any two or more users interfere whenever they simultaneously try to use the same channel. The resulting bandit problem is one of coordination in a competitive environment, because every time two or more agents select the same action at the same time step they both get a zero reward due to the interference —see (Rosenski et al., 2015) for recent work on stochastic competitive bandits and (Kleinberg et al., 2009) for a study of more general congestion games in a game-theoretic setting. Finally, there exists an extensive literature on the adaptation of gradient descent and related algorithms to distributed computing settings, where asynchronous processors naturally introduce delays —see, e.g., (Zinkevich et al., 2009; Agarwal and Duchi, 2011; Li et al., 2013; McMahan and Streeter, 2014; Quanrud and Khashabi, 2015; Liu et al., 2015; Duchi et al., 2015). However, none of these works considers bandit settings, which are an essential ingredient for our analysis.

3. Preliminaries

We now establish our notation, along with basic assumptions and preliminary facts related to our algorithms. Notation and setting here both refer to the single agent case. The cooperative setting with multiple agents (and notation thereof) will be introduced in Section 4. Proofs of all the results stated here can be found in (Cesa-Bianchi et al., 2016).

Let $A = \{1, \dots, K\}$ be the action set. A learning agent runs an exponentially-weighted algorithm with weights $w_t(i)$, and learning rate $\eta > 0$. Initially, $w_1(i) = 1$ for all $i \in A$. At each time step $t = 1, 2, \dots$, the agent draws action I_t with probability $\mathbb{P}(I_t = i) = p_t(i) = w_t(i)/W_t$, where $W_t = \sum_{j \in A} w_t(j)$. After observing the loss $\ell_t(I_t) \in [0, 1]$ associated with the chosen action I_t , and possibly some additional information, the agent computes, for each $i \in A$, nonnegative loss estimates $\hat{\ell}_t(i)$, and performs the exponential update

$$w_{t+1}(i) = p_t(i) \exp(-\eta \hat{\ell}_t(i)) \tag{1}$$

to these weights. The following two lemmas are general results that control the evolution of the probability distributions in the exponentially-weighted algorithm. As we said in the introduction, bounding the extent to which the distribution used by our algorithms can drift in d steps is key to controlling regret in a delayed setting. The first result bounds the *additive* change in the probability of any action, and it holds no matter how $\widehat{\ell}_t(i)$ is defined.

Lemma 1 *Under the update rule (1), for all $t \geq 1$ and for all $i \in A$,*

$$-\eta p_t(i) \widehat{\ell}_t(i) \leq p_{t+1}(i) - p_t(i) \leq \eta p_{t+1}(i) \sum_{j \in A} p_t(j) \widehat{\ell}_t(j)$$

holds deterministically with respect to the agent's randomization.

The second result delivers a *multiplicative* bound on the change in the probability of any action when the loss estimates $\widehat{\ell}_t(i)$ are of the following form:

$$\widehat{\ell}_t(i) = \begin{cases} \frac{\ell_{t-d}(i)}{q_{t-d}(i)} B_{t-d}(i) & \text{if } t > d, \\ 0 & \text{otherwise,} \end{cases} \quad (2)$$

where $d \geq 0$ is a delay parameter, $B_{t-d}(i) \in \{0, 1\}$, for $i \in A$, are indicator functions, and $q_{t-d}(i) \geq p_{t-d}(i)$ for all i and $t > d$. In all later sections, $B_{t-d}(i)$ will be instantiated to the indicator function of the event that action i has been played at time $t-d$ by some agent, and $q_{t-d}(i)$ will be the (conditional) probability of this event.

Lemma 2 *Let $\widehat{\ell}_t(i)$ be of the form (2) for each $t \geq 1$ and $i \in A$. If $\eta \leq \frac{1}{Ke(d+1)}$ in the update rule (1), then*

$$p_{t+1}(i) \leq \left(1 + \frac{1}{d}\right) p_t(i)$$

holds for all $t \geq 1$ and $i \in A$, deterministically with respect to the agent's randomization.

As we said in Section 1, the idea of controlling the drift of the probabilities in order to bound the effects of delayed feedback is not new. In particular, variants of Lemma 1 were already derived in the work of Neu et al. (2010, 2014). However, Lemma 2 appears to be new, and this is the key result to achieving our improvements.

4. The Cooperative Setting on a Communication Network

In our multi-agent bandit setting, there are N agents sitting on the vertices of a connected and undirected communication graph $G = (V, E)$, with $V = \{1, \dots, N\}$. The agents cooperate to solve the same instance of a nonstochastic bandit problem while limiting the communication among them. Let $N_s(v)$ be the set of nodes $v' \in V$ whose shortest-path distance $\text{dist}_G(v, v')$ from v in G is exactly s . At each time step $t = 1, 2, \dots$, each agent $v \in V$ draws an action $I_t(v)$ from the common action set A . Note that each action $i \in A$ delivers the same loss $\ell_t(i) \in [0, 1]$ to all agents v such that $I_t(v) = i$. At the end of round t , each agent v observes his own loss $\ell_t(I_t(v))$, and sends to his neighbors in G the message

$$m_t(v) = \langle t, v, I_t(v), \ell_t(I_t(v)), \mathbf{p}_t(v) \rangle$$

The cooperative bandit protocol

Parameters: Undirected communication graph $G = (V, E)$, hidden loss vectors $\ell_t = (\ell_t(1), \dots, \ell_t(K)) \in [0, 1]^K$ for $t \geq 1$, delay d .

For $t = 1, 2, \dots$

1. Each agent $v \in V$ plays action $I_t(v) \in A$ drawn according to distribution $\mathbf{p}_t(v)$;
2. Each agent $v \in V$ observes loss $\ell_t(I_t(v))$, sends to his neighbors the message $m_t(v)$, and receives from his neighbors messages $m_{t-s}(v')$;
3. Each agent $v \in V$ drops any message $m_{t-s}(v')$ received from some neighbor such that $s \geq d$, and forwards to the other neighbors the remaining messages.

Figure 1: The cooperative bandit protocol where all agents share the same delay parameter d .

where $\mathbf{p}_t(v) = (p_t(1, v), \dots, p_t(K, v))$ is the distribution of $I_t(v)$. Moreover, v also receives from his neighbors a variable number of messages $m_{t-s}(v')$. Each message $m_{t-s}(v')$ that v receives from a neighbor is used to update $\mathbf{p}_t(v)$ and then forwarded to the other neighbors only if $s < d$, otherwise it is dropped.² Here d is the maximum delay, a parameter of the communication protocol. Therefore, at the end of round t , each agent v receives one message $m_{t-s}(v')$ for each agent v' such that $\text{dist}_G(v, v') = s$, where $s \in \{1, \dots, d\}$. Graph G can thus be seen as a synchronous multi-hop communication network where messages are broadcast, each hop causing a delay of one time step. Our learning protocol is summarized in Figure 1, while Figure 2 contains a pictorial example.

Our model is similar to the LOCAL communication model in distributed computing (Linial, 1992; Suomela, 2013), where the output of a node depends only on the inputs of other nodes in a constant-size neighborhood of it, and the goal is to derive algorithms whose running time is independent of the network size. (The main difference is that the task here has no completion time, however, also in our model influence on a node is only through a constant-size neighborhood of it.)

One aspect deserving attention is that, apart from the common delay parameter d , the agents need not share further information. In particular, the agents need not know neither the topology of the graph G nor the total number of agents N . In Section 5, we show that our distributed algorithm can also be analyzed when each agent v uses a personalized delay $d(v)$, thus doing away with the need of a common delay parameter, and guaranteeing a generally better performance.

Further graph notation is needed at this point. Given G as above, let us denote by $G_{\leq d}$ the graph $(V, E_{\leq d})$ where $(u, v) \in E_{\leq d}$ if and only if the shortest-path distance between agents u and v in G is at most d (hence $G_{\leq 1} = G$). Graph $G_{\leq d}$ is sometimes called the d -th power of G . We also use G_0 to denote the graph (V, \emptyset) . Recall that an independent set of G is any subset $T \subseteq V$ such that no two $i, j \in T$ are connected by an edge in E . The largest size of an independent set is the *independence number* of G , denoted by $\alpha(G)$. Let d_G be the *diameter* of G (maximal length over all possible shortest paths between all pairs of nodes); then $G_{\leq d_G}$ is a clique, and one can easily see that $N = \alpha(G_0) > \alpha(G) \geq \alpha(G_{\leq 2}) \geq \dots \geq \alpha(G_{\leq d_G}) = 1$. We show in Section 4.1 that the

2. Dropping messages older than d rounds is clearly immaterial with respect to proving bandit regret bounds. We added this feature just to prove a point about the message complexity of the protocol. See Remark 8 in Section 5 for further discussion.




Figure 2: In this example, G is a line graph with $N = 6$ agents, and delay $d = 2$. At the end of time step t , agent 4 sends to his neighbors 3 and 5 message $m_t(4)$, receives from agent 3 messages $m_{t-1}(3)$, and $m_{t-2}(2)$, and from agent 5 messages $m_{t-1}(5)$ and $m_{t-2}(6)$. Finally, 4 forwards to 5 message $m_{t-1}(3)$ and forwards to 3 message $m_{t-1}(5)$. Any message older than $t - 1$ received by 4 at the end of round t will not be forwarded to his neighbors.

collective performance of our algorithms depends on $\alpha(G_{\leq d})$. If the graph G under consideration is directed (see Section 5), then $\alpha(G)$ is the independence number of the undirected graph obtained from G by disregarding edge orientation.

The adversary generating losses is oblivious: loss vectors $\ell_t = (\ell_t(1), \dots, \ell_t(K)) \in [0, 1]^K$ do not depend on the agents' internal randomization. The agents' goal is to control the *average welfare* regret R_T^{coop} , defined as

$$R_T^{\text{coop}} = \left(\frac{1}{N} \sum_{v \in V} \mathbb{E} \left[\sum_{t=1}^T \ell_t(I_t(v)) \right] - \min_{i \in A} \sum_{t=1}^T \ell_t(i) \right),$$

the expectation being with respect to the internal randomization of each agent's algorithm. In the sequel, we write $\mathbb{E}_t[\cdot]$ to denote the expectation w.r.t. the product distribution $\prod_{v \in V} \mathbf{p}_t(v)$, conditioned on $I_1(v), \dots, I_{t-1}(v)$, $v \in V$.

4.1. The Exp3-Coop algorithm

Our first algorithm, called EXP3-COOP (Cooperative Exp3) is described in Figure 3. The algorithm works in the learning protocol of Figure 1. Each agent $v \in V$ runs the exponentially-weighted algorithm (1), combined with a “delayed” importance-weighted loss estimate $\hat{\ell}_t(i, v)$ that incorporates the delayed information sent by the other agents. Specifically, denote by $N_{\leq d}(v) = \bigcup_{s \leq d} N_s(v)$ the set of nodes in G whose shortest-path distance from v is at most d , and note that, for all v , $\{v\} = N_{\leq 0}(v) \subseteq N_{\leq 1}(v) \subseteq N_{\leq 2}(v) \subseteq \dots$. If any of the agents in $N_{\leq d}(v)$ has played at time $t - d$ action i (that is, $B_{d,t-d}(i, v) = 1$ in Eq. in (3)), then the corresponding loss $\ell_{t-d}(i)$ is incorporated by v into $\hat{\ell}_t(i, v)$. The denominator $q_{d,t-d}(i, v)$ is simply, conditioned on the history, the probability of $B_{d,t-d}(i, v) = 1$, i.e., $q_{d,t-d}(i, v) = \mathbb{E}_t[B_{d,t-d}(i, v)]$. Observe that $\{v\} \subseteq N_{\leq d}(v)$ for all $d \geq 0$ implies $q_{d,t-d}(i, v) \geq p_{t-d}(i, v)$, as required by (2). It is also worth mentioning that, despite this is not strictly needed by our learning protocol, each agent v actually exploits the loss information gathered from playing action $I_t(v)$ only d time steps later. A relevant special case of this learning mode is when we only have a single bandit agent receiving delayed feedback (Section 6).

By their very definition, the loss estimates $\hat{\ell}_t(\cdot, \cdot)$ at time t are determined by the realizations of $I_s(\cdot)$, for $s = 1, \dots, t - d$. This implies that the numbers $p_t(\cdot, \cdot)$ defining $q_{d,t-d}(\cdot, \cdot)$, are determined by the realizations of $I_s(\cdot)$ for $s = 1, \dots, t - d - 1$ (because the probabilities $\mathbf{p}_t(v)$ at time t are

The Exp3-Coop Algorithm

Parameters: Undirected communication graph $G = (V, E)$; delay d ; learning rate η .

Init: Each agent $v \in V$ sets weights $w_1(i, v) = 1$ for all $i \in A$.

For $t = 1, 2, \dots$

1. Each agent $v \in V$ plays action $I_t(v) \in A$ drawn according to distribution $\mathbf{p}_t(v) = (p_t(1, v), \dots, p_t(K, v))$, where

$$p_t(i, v) = \frac{w_t(i, v)}{W_t(v)}, \quad i = 1, \dots, K, \quad \text{and} \quad W_t(v) = \sum_{j \in A} w_t(j, v);$$

2. Each agent $v \in V$ observes loss $\ell_t(I_t(v))$ and exchanges messages with his neighbors (Steps 2 and 3 of the protocol in Figure 1);
3. Each agent $v \in V$ performs the update $w_{t+1}(i, v) = p_t(i, v) \exp(-\eta \widehat{\ell}_t(i, v))$ for all $i \in A$, where

$$\widehat{\ell}_t(i, v) = \begin{cases} \frac{\ell_{t-d}(i)}{q_{d,t-d}(i, v)} B_{d,t-d}(i, v) & \text{if } t > d, \\ 0 & \text{otherwise,} \end{cases} \quad (3)$$

and $B_{d,t-d}(i, v) = \mathbb{I}\{\exists v' \in N_{\leq d}(v) : I_{t-d}(v') = i\}$ with

$$q_{d,t-d}(i, v) = 1 - \prod_{v' \in N_{\leq d}(v)} (1 - p_{t-d}(i, v')).$$

Figure 3: The Exp3-Coop algorithm where all agents share the same delay parameter d .

determined by the loss estimates up to time $t - 1$, see (1)). We have, for all $t > d$, $i \in A$, and $v \in V$,

$$\mathbb{E}_{t-d}[\widehat{\ell}_t(i, v)] = \ell_{t-d}(i). \quad (4)$$

Further, because of what we just said about $p_t(\cdot, \cdot)$ and $q_{d,t-d}(\cdot, \cdot)$ being determined by $I_1(\cdot), \dots, I_{t-d-1}(\cdot)$, we also have

$$\mathbb{E}_{t-d}[p_t(i, v) \widehat{\ell}_t(i, v)] = p_t(i, v) \ell_{t-d}(i), \quad \mathbb{E}_{t-d}[p_t(i, v) \widehat{\ell}_t(i, v)^2] = p_t(i, v) \frac{\ell_{t-d}(i)^2}{q_{d,t-d}(i, v)}. \quad (5)$$

The following theorem quantifies the behavior of EXP3-COOP in terms of a free parameter γ in the learning rate, the tuning of which will be addressed in the subsequent Theorem 4.

Theorem 3 *The regret of EXP3-COOP run over a network $G = (V, E)$ of N agents, each using delay d and learning rate $\eta = \gamma / (Ke(d + 1))$, for $\gamma \in (0, 1]$, satisfies*

$$R_T^{\text{coop}} \leq 2d + \frac{Ke(d + 1) \ln K}{\gamma} + \gamma \left(\frac{\alpha(G_{\leq d})}{2(1 - e^{-1})(d + 1)N} + \frac{3}{Ke} \right) T.$$

With this bound handy, we might be tempted to optimize for γ . However, this is not a legal learning rate setting in a distributed scenario, for the optimized value of γ would depend on the global

quantities N and $\alpha(G_{\leq d})$. Thus, instead of this global tuning, we let each agent set its own learning rate γ through a “doubling trick” played locally. The doubling trick³ works as follows. For each $v \in V$, we let $\gamma_r(v) = Ke(d+1)\sqrt{(\ln K)/2^r}$ for each $r = r_0, r_0+1, \dots$, where $r_0 = \lceil \log_2 \ln K + 2 \log_2(Ke(d+1)) \rceil$ is chosen in such a way that $\gamma_r(v) \leq 1$ for all $r \geq r_0$. Let T_r be the random set of consecutive time steps where the same $\gamma_r(v)$ was used. Whenever the local algorithm is running with $\gamma_r(v)$ and detects $\sum_{s \in T_r} Q_s(v) > 2^r$, then we restart the algorithm with $\gamma(v) = \gamma_{r+1}(v)$. We have the following result.

Theorem 4 *The regret of EXP3-COOP run over a network $G = (V, E)$ of N agents, each using delay d , and an individual learning rate $\eta(v) = \gamma(v)/(Ke(d+1))$, where $\gamma(v) \in (0, 1]$ is adaptively selected by each agent through the above doubling trick, satisfies, when T grows large,⁴*

$$R_T^{\text{coop}} = \mathcal{O} \left(\sqrt{(\ln K) \left(d + 1 + \frac{K}{N} \alpha(G_{\leq d}) \right) T + d \log T} \right).$$

Remark 5 *Theorem 4 shows a natural trade-off between delay and information. To make it clear, suppose $N \approx K$. In this case, the regret bound becomes of order $\sqrt{(d + \alpha(G_{\leq d}))T \ln K + d \ln T}$. Now, if d is as big as the diameter d_G of G , then $\alpha(G_{\leq d}) = 1$. This means that at every time step all $N \approx K$ agents observe (with some delay) the losses of each other’s actions. This is very much reminiscent of a full information scenario, and in fact our bound becomes of order $\sqrt{(d_G + 1)T \ln K + d_G \ln T}$, which is close to the full information minimax rate $\sqrt{(d+1)T \ln K}$ when feedback has a constant delay d (Weinberger and Ordentlich, 2002). When G is sparse (i.e., d_G is likely to be large, say $d_G \approx N$), then agents have no advantage in taking $d = d_G$ since $d_G \approx N \approx K$. In this case, agents may even give up cooperation (choosing $d = 0$ in Figure 3), and fall back on the standard bandit bound $\sqrt{TK \ln K}$, which corresponds to running EXP3-COOP on the edgeless graph G_0 . (No doubling trick is needed in this case, hence no extra $\log T$ term appears.)*

Remark 6 *When $d = d_G$, each neighborhood $N_{\leq d}(v)$ used in the loss estimate (3) is equal to V , hence all agents receive the same feedback. Because they all start off from the same initial weights, the agents end up computing the same updates. This in turn implies that: (1) the individual regret incurred by each agent is the same as the average welfare regret R_T^{coop} ; (2) the messages exchanged by the agents (see Figure 1) may be shortened by dropping the distribution part $\mathbf{p}_{t-s}(v')$.*

Remark 7 *An interesting question is whether the agents can come up with a reasonable choice for the value of d even when they lack any information whatsoever about the global structure of G . A partial answer to this question follows. It is easy to show that the choice $d = \sqrt{K}$ in Theorem 4 yields a bound on the average welfare regret of the form $K^{1/4}\sqrt{T \ln K} + \sqrt{K}(\ln T)$ for all G (and irrespective to the value of $N = |V|$), provided G is connected. This holds because, for any connected graph G , the independence number $\alpha(G_{\leq d})$ is always bounded by⁵ $\lceil 2N/(d+2) \rceil$. To*

3. There has been some recent work on adaptive learning rate tuning applied to nonstochastic bandit algorithms (Kocák et al., 2014; Neu, 2015). One might wonder whether the same techniques may apply here as well. Unfortunately, the specific form of our update (1) makes this adaptation nontrivial, and this is why we resorted to a more traditional “doubling trick”.

4. The big-oh notation here hides additive terms that are independent of T and do depend polynomially on the other parameters.

5. Because it holds for a worst-case (connected) G , this upper bound on $\alpha(G_{\leq d})$ can be made tighter when specific graph topologies are considered.

see why this latter statement is true, observe that the neighborhood $N_{\leq d/2}(v)$ of any node v in $G_{\leq d/2}$ contains at least $d/2 + 1$ nodes (including v), and any pair of nodes $v', v'' \in N_{\leq d/2}(v)$ are adjacent in $G_{\leq d}$. Therefore, no independent set of $G_{\leq d}$ can have size bigger than $\lceil 2N/(d+2) \rceil$. A more detailed bound is contained, e.g., in (Firby and Haviland, 1997).

5. Extensions: Cooperation with Individual Parameters

In this section, we analyze a modification of EXP3-COOP that allows each agent v in the network to use a delay parameter $d(v)$ different from that of the other agents. We then show how such individual delays may improve the average welfare regret of the agents. In the previous setting, where all agents use the same delay parameter d , messages have an implicit time-to-live equal to d . In this setting, however, agents may not have a detailed knowledge of the delay parameters used by the other agents. For this reason we allow an agent v to generate messages with a time-to-live $tll(v)$ possibly different from the delay parameter $d(v)$. Note that the role of the two parameters $d(v)$ and $tll(v)$ is inherently different. Whereas $d(v)$ rules the extent to which v uses the messages received from the other agents, $tll(v)$ limits the number of times a message from v is forwarded to the other agents, thereby limiting the message complexity of the algorithm. In order to accommodate this additional parameter, we are required to modify the cooperative bandit protocol of Figure 1. As in Section 4, we have an undirected communication network $G = (V, E)$ over the agents. However, in this new protocol the message that at the end of round t each agent v sends to his neighbors in G has the format

$$m_t(v) = \langle t, v, tll(v), I_t(v), \ell_t(I_t(v)), \mathbf{p}_t(v) \rangle$$

where $tll(v)$ is the time-to-live parameter of agent v . Each message $m_{t-s}(v')$, which v receives from a neighbor, first has its time-to-live decremented by one. If the resulting value is positive, the message is forwarded to the other neighbors, otherwise it is dropped. Moreover, v uses this message to update $\mathbf{p}_t(v)$ only if $s \leq d(v)$. Hence, at time t an agent v uses the message sent at time $t - s$ by v' if and only if $\text{dist}_G(v', v) = s$ with $s \leq \min\{d(v), tll(v')\}$, where $\text{dist}_G(v, v')$ is the shortest-path distance from v' to v in G .

Based on the collection $\mathcal{P} = \{d(v), tll(v)\}_{v \in V}$ of individual parameters, we define the directed graph $G_{\mathcal{P}} = (V, E_{\mathcal{P}})$ as follows: arc $(v', v) \in E_{\mathcal{P}}$ if and only if $\text{dist}_G(v, v') \leq \min\{d(v), tll(v')\}$. The in-neighborhood $N_{\mathcal{P}}^-(v)$ of v thus contains the set of all $v' \in V$ whose distance from v is not larger than $\min\{d(v), tll(v')\}$. Notice that, with this definition, $v \in N_{\mathcal{P}}^-(v)$, so that $(V, E_{\mathcal{P}})$ includes all self-loops (v, v) . Figure 4(a) illustrates these concepts through a simple pictorial example.

Remark 8 *It is important to remark that the communication structure encoded by \mathcal{P} is an exogenous parameter of the regret minimization problem, and so our algorithms cannot trade it off against regret. In addition to that, the parameterization $\mathcal{P} = \{d(v), tll(v)\}_{v \in V}$ defines a simple and static communication graph which makes it relatively easy to express the tradeoff between regret and communication. This would not be possible if we had each individual node v decide whether to forward a message based, say, on its own local delay parameter $d(v)$. To see why, consider the situation where nodes v and v' are along the route of a message that is reaching v before v' . The decision of v to drop the message may clash with the willingness of v' to receive it, and this may clearly*




Figure 4: **(a)** In this example, the communication network G is an undirected line graph with $N = 6$ agents, whose edges are depicted in black. Close to each node $v = 1, \dots, 6$ is the individual delay $d(v)$ (in blue), and the individual time-to-leave $tll(v)$ (in red). The arcs (a.k.a., directed edges) of the induced directed graph $G_{\mathcal{P}}$ are also depicted in blue. Self-loops are not depicted. For instance, we have $N_{\mathcal{P}}^-(5) = \{4, 5, 6\}$ and $N_{\mathcal{P}}^-(3) = \{3\}$. **(b)** A communication network having a dense (red nodes) and a sparse (black nodes) region. The black region has $N^{1/2}$ agents, the red one has $N - N^{1/2}$ agents. **(c)** A star graph with long rays. The center v (in red) sets a small $d(v)$ and a large $tll(v)$. The peripheral nodes v' (in green) set a large $d(v')$ and a small $tll(v')$.

happen when $d(v) < d(v')$. The structure of the communication graph resulting from this individual behavior of the nodes would be rather complicated. On the contrary, the time-to-live-based parametrization, which is commonly used in communication networks to control communication complexity, does not have this issue.

Figure 5 contains our algorithm (called EXP3-COOP2) for this setting. EXP3-COOP2 is a strict generalization of EXP3-COOP, and so is its analysis. The main difference between the two algorithms is that EXP3-COOP2 deals with directed graphs. This fact prevents us from using the same techniques of Section 4.1 in order to control the regret. Intuitively, adding orientations to the edges reduces the information available to the agents and thus increases the variance of their loss estimates. Thus, in order to control this variance, we need a lower bound⁶ on the probabilities $p_t(i, v)$. From Figure 5, one can easily see that

$$1 = \sum_{i \in A} \frac{w_t(i, v)}{W_t(v)} \leq \tilde{P}_t(v) \leq \sum_{i \in A} \left(\frac{w_t(i, v)}{W_t(v)} + \frac{\delta}{K} \right) = 1 + \delta \quad (6)$$

implying the lower bound $p_t(i, v) \geq \frac{\delta}{K(1+\delta)}$, holding for all i, t , and v .

The following theorem is the main result of this section.

Theorem 9 *The regret of EXP3-COOP2 run over a network $G = (V, E)$ of N agents, each agent v using individual delay $d(v)$, individual time-to-leave $tll(v)$, exploration parameter $\delta = 1/T$, and*

6. We find it convenient to derive this lower bound without mixing with the uniform distribution over A —see, e.g., (Auer et al., 2002)—but in a slightly different manner. This facilitates our delayed feedback analysis.

The Exp3-Coop2 Algorithm

Parameters: Undirected graph $G = (V, E)$; learning rate η ; exploration parameter $\delta > 0$.

Init: Each $v \in V$ sets weights $w_1(i, v) = 1$, for all $i \in A$, delay $d(v)$, and time-to-live $tll(v)$.

For $t = 1, 2, \dots$

1. Each agent $v \in V$ plays action $I_t(v) \in A$ drawn according to distribution $\mathbf{p}_t(v) = (p_t(1, v), \dots, p_t(K, v))$, where

$$p_t(i, v) = \frac{\tilde{p}_t(i, v)}{\tilde{P}_t(v)}, \quad \tilde{P}_t(v) = \sum_{j \in A} \tilde{p}_t(j, v),$$

and

$$\tilde{p}_t(i, v) = \max \left\{ \frac{w_t(i, v)}{W_t(v)}, \frac{\delta}{K} \right\}, \quad W_t(v) = \sum_{j \in A} w_t(j, v);$$

2. Each agent $v \in V$ observes loss $\ell_t(I_t(v))$ and exchanges messages with his neighbors (see main text for an explanation);
3. Each agent $v \in V$ performs the update $w_{t+1}(i, v) = p_t(i, v) \exp(-\eta \hat{\ell}_t(i, v))$ for all $i \in A$, where

$$\hat{\ell}_t(i, v) = \begin{cases} \frac{\ell_{t-d(v)}(i)}{q_{\mathcal{P}, t-d(v)}(i, v)} B_{\mathcal{P}, t-d(v)}(i, v) & \text{if } t > d(v), \\ 0 & \text{otherwise,} \end{cases}$$

and $B_{\mathcal{P}, t-d(v)}(i, v) = \mathbb{I}\{\exists v' \in N_{\mathcal{P}}^-(v) : I_{t-d(v)}(v') = i\}$, with

$$q_{\mathcal{P}, t-d(v)}(i, v) = 1 - \prod_{v' \in N_{\mathcal{P}}^-(v)} (1 - p_{t-d(v)}(i, v')).$$

Figure 5: The Exp3-Coop2 algorithm with individual delay and time-to-live parameters.

learning rate η such that $\eta \rightarrow 0$ as $T \rightarrow \infty$ satisfies, when T grows large,

$$R_T^{\text{coop}} = \mathcal{O} \left(\frac{\ln K}{\eta} + \eta \left(\bar{d}_V + \frac{K}{N} \alpha(G_{\mathcal{P}}) \ln(TNK) \right) T \right), \quad \text{where } \bar{d}_V = \frac{1}{N} \sum_{v \in V} d(v).$$

Using a doubling trick in much the same way we used it to prove Theorem 4, we can prove the following result.

Corollary 10 *The regret of EXP3-COOP2 run over a network $G = (V, E)$ of N agents, each agent v using individual delay $d(v)$, individual time-to-leave $tll(v)$, exploration parameter $\delta = 1/T$, and individual learning rate $\eta(v)$ adaptively selected by each agent through a doubling trick, satisfies, when T grows large*

$$R_T^{\text{coop}} = \mathcal{O} \left(\sqrt{(\ln K) \left(\bar{d}_V + 1 + \frac{K}{N} \alpha(G_{\mathcal{P}}) \ln(TNK) \right) T} + \bar{d}_V (\ln T + \ln \ln(TNK)) \right).$$

To illustrate the advantage of having individual delays as opposed to sharing the same delay value, it suffices to consider a communication network including regions of different density. Concretely, consider the graph in Figure 4(b) with a large densely connected region (red agents) and a small sparsely connected region (black agents). In this example, the black agents prefer a large value of their individual delay so as to receive more information from nearby agents, but this comes at the price of a larger bias for their estimators $\widehat{\ell}_t(i, v)$. On the contrary, information from nearby agents is readily available to the red agents, so that they do not gain any regret improvement from a large delay parameter. A similar argument applies here to the individual time-to-live values: red agents v will set a small $tll(v)$ to reduce communication. Black agents v' may decide to set $tll(v')$ depending on their intention to reach the red nodes. But because the red agents have set a small $d(v)$, any effort made by v' trying to reach them would be a communication waste. Hence, it is reasonable for a black agent v' to set a moderately large value for $tll(v')$, but perhaps not so large as to reach the red agents. One can read this off the bounds in both Theorem 9 and Corollary 10, as explained next. Suppose for simplicity that $K \approx N$ so that, disregarding log factors, these bounds depend on parameters \mathcal{P} only through the quantity $H = \bar{d}_V + \alpha(G_{\mathcal{P}})$. Now, in the case of a common delay parameter d (Section 4.1), it is not hard to see that the best setting for d in order to minimize H is of the form $d = N^{1/4}$, resulting in $H = \Theta(N^{1/4})$. On the other hand, the best setting for the individual delays is $d(v) = 1$ when v is red, and $d(v) = \sqrt{N}$ when v is black, resulting in $H = \Theta(1)$.

The time-to-live parameters $tll(v)$ affect the regret bound only through $\alpha(G_{\mathcal{P}})$, but they clearly play the additional role of bounding the message complexity of the algorithm. In our example of Figure 4(b), we essentially have $d(v) \approx tll(v)$ for all v . A typical scenario where agents may have $d(v) \neq tll(v)$ is illustrated in Figure 4(c). In this case, we have star-like graph where a central agent is connected through long rays to all other agents. The center v prefers to set a small $d(v)$, since it has a large degree, but also a large $tll(v)$ in order to reach the green peripheral nodes. The green nodes v' are reasonably doing the opposite: a large $d(v')$ in order to gather information from other nodes, but also a smaller time-to-live than the center, for the information transmitted by v' is comparatively less valuable to the whole network than the one transmitted by the center.

Agents can set their individual parameters in a topology-dependent manner using any algorithm for assessing the centrality of nodes in a distributed fashion —e.g., (Wehmuth and Ziviani, 2013), and references therein. This can be done at the beginning in a number of rounds which only depends on the network topology (but not on T). Hence, this initial phase would affect the regret bound only by an additive constant.

6. Delayed Losses (for a Single Agent)

EXP3-COOP can be specialized to the setting where a single agent is facing a bandit problem in which the loss of the chosen action is observed with a fixed delay d . In this setting, at the end of each round t the agent incurs loss $\ell_t(I_t)$ and observes $\ell_{t-d}(I_{t-d})$, if $t > d$, and nothing otherwise. The regret is defined in the usual way,

$$R_T = \mathbb{E} \left[\sum_{t=1}^T \ell_t(I_t) \right] - \min_{i=1, \dots, K} \sum_{t=1}^T \ell_t(i).$$

This problem was studied by [Weinberger and Ordentlich \(2002\)](#) in the full information case, for which they proved that $\sqrt{(d+1)T \ln K}$ is the optimal order for the minimax regret. The result was extended to the bandit case by [Neu et al. \(2010, 2014\)](#)—see also [Joulani et al. \(2013\)](#)—whose techniques can be used to obtain a regret bound of order $\sqrt{(d+1)KT}$. Yet, no matching lower bound was available for the bandit case.

As a matter of fact, the upper bound $\sqrt{(d+1)KT}$ for the bandit case is easily obtained: just run in parallel $d+1$ instances of the minimax optimal bandit algorithm for the standard (no delay) setting, achieving $R_T \leq \sqrt{KT}$ (ignoring constant factors). At each time step $t = (d+1)r + s$ (for $r = 0, 1, \dots$ and $s = 0, \dots, d$), use instance $s+1$ for the current play. Hence, the no-delay bound applies to every instance and, assuming $d+1$ divides T , we immediately obtain $R_T \leq \sum_{s=1}^{d+1} \sqrt{K \frac{T}{d+1}} \leq \sqrt{(d+1)KT}$, again, ignoring constant factors.

Next, we show that the machinery we developed in [Section 4.1](#) delivers an improved upper bound on the regret for the bandit problem with delayed losses, and then we complement this result by providing a lower bound matching the upper bound up to log factors, thereby characterizing (up to log factors) the minimax regret for this problem.

Corollary 11 *In the nonstochastic bandit setting with $K \geq 2$ actions and delay $d \geq 0$, where at the end of each round t the predictor has access to the losses $\ell_1(I_1), \dots, \ell_s(I_s) \in [0, 1]^K$ for $s = \max\{1, t-d\}$, the minimax regret is of order $\sqrt{(K+d)T}$, ignoring logarithmic factors.*

7. Conclusions and Ongoing Research

We have investigated a cooperative and nonstochastic bandit scenario where cooperation comes at the price of delayed information. We have proven average welfare regret bounds that exhibit a natural tradeoff between amount cooperation and delay, the tradeoff being ruled by the underlying communication network topology. As a by-product of our analysis, we have also provided the first characterization to date of the regret of learning with (constant) delayed feedback in an adversarial bandit setting. There are a number of possible extensions which we are currently considering:

1. So far our analysis only delivers average welfare regret bounds. It would be interesting to show simultaneous regret bounds that hold for each agent individually. We conjecture that the individual regret bound of an agent v is of the form $\sqrt{(\ln K) \left(d + \frac{K}{|N_{\leq d}(v)|} \right) T}$, where $|N_{\leq d}(v)|$ is the degree of v in $G_{\leq d}$ (plus one). Such bound would in fact imply, e.g., the one in [Theorem 4](#). A possible line of attack to solve this problem could be the use of graph sparsity along the lines of ([Pan et al., 2015](#); [Duchi et al., 2013](#); [Mania et al., 2015](#); [McMahan and Streeter, 2014](#)).
2. It would be nice to characterize the average welfare regret by complementing our upper bounds with suitable *lower* bounds: Is the upper bound of [Theorem 4](#) optimal in the communication model considered here?
3. The two algorithms we designed do not use the loss information in the most effective way, for they both postpone the update step by d ([Figure 3](#)) or $d(v)$ ([Figure 5](#)) time steps. In fact, we do have generalized versions of both algorithms where all losses $\ell_{t-s}(i)$ coming from agents

at distance s from any given agent v are indeed used at time t by agent v i.e., as soon as these losses become available to v . The resulting regret bounds mix delays and independence numbers of graphs at different levels of delay. (Details will be given in the full version of this paper.) More ambitiously, it is natural to think of ways to adaptively tune our algorithms so as to automatically determine the best delay parameter d . For instance, disregarding message complexity, is there a way for each agent to adaptively tune d locally so to minimize the bound in Theorem 4?

4. Our messages $m_t(v)$ contain both action/loss information and distribution information. Is it possible to drop the distribution information and still achieve average welfare regret bounds similar to those in Theorems 3 and 4?
5. Even for the single-agent setting, we do not know whether regret bounds of the form $\sqrt{(D + T) \ln K}$, where D is the total delay experienced over the T rounds, could be proven—see (Joulani et al., 2016; Quanrud and Khashabi, 2015) for similar results in the full-information setting. In general, the study of learning on a communication network with time-varying delays, and its impact on the regret rates, is a topic which is certainly worth of attention.

Acknowledgments

We thank the anonymous reviewers for their careful reading, and for their thoughtful suggestions that greatly improved the presentation of this paper. Yishay Mansour is supported in part by the Israeli Centers of Research Excellence (I-CORE) program, (Center No. 4/11), by a grant from the Israel Science Foundation (ISF), by a grant from United States-Israel Binational Science Foundation (BSF) and by a grant from the Len Blavatnik and the Blavatnik Family Foundation.

References

- Alekh Agarwal and John C Duchi. Distributed delayed stochastic optimization. In J. Shawe-Taylor, R.S. Zemel, P.L. Bartlett, F. Pereira, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 873–881. Curran Associates, Inc., 2011.
- Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multi-armed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002.
- Baruch Awerbuch and Robert Kleinberg. Competitive collaborative learning. *Journal of Computer and System Sciences*, 74(8):1271–1288, 2008.
- Samuel Barrett and Peter Stone. Ad hoc teamwork modeled with multi-armed bandits: An extension to discounted infinite rewards. In *Proceedings of 2011 AAMAS Workshop on Adaptive and Learning Agents*, pages 9–14, 2011.
- Nicolo’ Cesa-Bianchi, Claudio Gentile, Yishay Mansour, and Alberto Minora. Delay and cooperation in nonstochastic bandits. *arXiv preprint, arXiv:1602.04741v2*, 2016.

- John Duchi, Michael I Jordan, and Brendan McMahan. Estimation, optimization, and parallelism when data is sparse. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 2832–2840. Curran Associates, Inc., 2013.
- John C Duchi, Sorathan Chaturapruek, and Christopher Ré. Asynchronous stochastic convex optimization. *arXiv preprint arXiv:1508.00882*, 2015.
- Miroslav Dudík, Daniel J. Hsu, Satyen Kale, Nikos Karampatziakis, John Langford, Lev Reyzin, and Tong Zhang. Efficient optimal learning for contextual bandits. In *UAI 2011, Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence, Barcelona, Spain, July 14-17, 2011*, pages 169–178, 2011.
- P. Firby and J. Haviland. Independence and average distance in graphs. *Discrete Applied Mathematics*, 75:27–37, 1997.
- Pooria Joulani, András György, and Csaba Szepesvári. Online learning under delayed feedback. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pages 1453–1461, 2013.
- Pooria Joulani, András György, and Csaba Szepesvári. Delay-tolerant online convex optimization: Unified analysis and adaptive-gradient algorithms. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, pages 1744–1750, 2016.
- Soumya Kar, H Vincent Poor, and Shuguang Cui. Bandit problems in networks: Asymptotically efficient distributed allocation rules. In *50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)*, pages 1771–1778. IEEE, 2011.
- Robert Kleinberg, Georgios Piliouras, and Éva Tardos. Multiplicative updates outperform generic no-regret learning in congestion games. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 533–542. ACM, 2009.
- Tomáš Kocák, Gergely Neu, Michal Valko, and Remi Munos. Efficient learning by implicit exploration in bandit problems with side observations. In *Advances in Neural Information Processing Systems 27*, pages 613–621. 2014.
- Peter Landgren, Vaibhav Srivastava, and Naomi Ehrlich Leonard. On distributed cooperative decision-making in multiarmed bandits. *arXiv preprint arXiv:1512.06888*, 2015.
- Mu Li, David G. Andersen, and Alexander Smola. Distributed delayed proximal gradient methods. In *NIPS Workshop on Optimization for Machine Learning*, 2013.
- Nathan Linial. Locality in distributed graph algorithms. *SIAM J. Comput.*, 21(1):193–201, 1992.
- Ji Liu, Stephen J Wright, Christopher Ré, Victor Bittorf, and Srikrishna Sridhar. An asynchronous parallel stochastic coordinate descent algorithm. *The Journal of Machine Learning Research*, 16(1):285–322, 2015.

- Horia Mania, Xinghao Pan, Dimitris Papailiopoulos, Benjamin Recht, Kannan Ramchandran, and Michael I Jordan. Perturbed iterate analysis for asynchronous stochastic optimization. *arXiv preprint arXiv:1507.06970*, 2015.
- Brendan McMahan and Matthew Streeter. Delay-tolerant algorithms for asynchronous distributed online learning. In Z. Ghahramani, M. Welling, C. Cortes, N.D. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2915–2923. Curran Associates, Inc., 2014.
- Chris Mesterharm. On-line learning with delayed label feedback. In *Algorithmic Learning Theory*, pages 399–413. Springer, 2005.
- Chris Mesterharm. *Improving Online Learning*. PhD thesis, Department of Computer Science, Rutgers University, 2007.
- Gergely Neu. Explore no more: Improved high-probability regret bounds for non-stochastic bandits. In *Advances in Neural Information Processing Systems 28 (NIPS)*, 2015.
- Gergely Neu, Andras Antos, András György, and Csaba Szepesvári. Online Markov decision processes under bandit feedback. In *Advances in Neural Information Processing Systems 23*, pages 1804–1812. Curran Associates, Inc., 2010.
- Gergely Neu, Andras Gyorgy, Csaba Szepesvari, and Andras Antos. Online markov decision processes under bandit feedback. *Automatic Control, IEEE Transactions on*, 59(3):676–691, 2014.
- Xinghao Pan, Dimitris Papailiopoulos, Samet Oymak, Benjamin Recht, Kannan Ramchandran, and Michael I Jordan. Parallel correlation clustering on big graphs. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 82–90. Curran Associates, Inc., 2015.
- Kent Quanrud and Daniel Khashabi. Online learning with adversarial delays. In C. Cortes, N.D. Lawrence, D.D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 1270–1278. Curran Associates, Inc., 2015.
- Jonathan Rosenski, Ohad Shamir, and Liran Szlak. Multi-player bandits – a musical chairs approach. *CoRR*, abs/1512.02866, 2015.
- Yevgeny Seldin, Peter Bartlett, Koby Crammer, and Yasin Abbasi-Yadkori. Prediction with limited advice and multiarmed bandits with paid observations. In *Proceedings of The 31st International Conference on Machine Learning*, pages 280–287, 2014.
- Ruben Stranders, Long Tran-Thanh, Francesco M Delle Fave, Alex Rogers, and Nicholas R Jennings. Dcops and bandits: Exploration and exploitation in decentralised coordination. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pages 289–296. International Foundation for Autonomous Agents and Multiagent Systems, 2012.
- Jukka Suomela. Survey of local algorithms. *ACM Computing Surveys*, 45(2):24, 2013.

- Balazs Szorenyi, Róbert Busa-Fekete, István Hegedüs, Róbert Ormándi, Márk Jelasity, and Balázs Kégl. Gossip-based distributed stochastic bandit algorithms. In *30th International Conference on Machine Learning (ICML 2013)*, volume 28, pages 19–27. ACM Press, 2013.
- Cem Tekin and Mihaela van der Schaar. Distributed online learning via cooperative contextual bandits. *IEEE Transactions on Signal Processing*, 63(14):3700–3714, 2015.
- Cem Tekin, Simpson Z. Zhang, and Mihaela van der Schaar. Distributed online learning in social recommender systems. *J. Sel. Topics Signal Processing*, 8(4):638–652, 2014.
- Klaus Wehmuth and Artur Ziviani. Daccor: Distributed assessment of the closeness centrality ranking in complex networks. *Computer Networks*, 57(13):2536–2548, 2013.
- Marcelo J Weinberger and Erik Ordentlich. On delayed prediction of individual sequences. *IEEE Transactions on Information Theory*, 48(7):1959–1976, 2002.
- Martin Zinkevich, John Langford, and Alex J. Smola. Slow learners are fast. In Y. Bengio, D. Schuurmans, J.D. Lafferty, C.K.I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems 22*, pages 2331–2339. Curran Associates, Inc., 2009.