Optimal Algorithms for Multiplayer Multi-Armed Bandits
(Supplementary Material)

Po-An Wang   **Alexandre Proutiere**   **Kaito Ariu**[*]   **Yassir Jedra**[*]   **Alessio Russo**[*]

KTH, Royal Institute of Technology
Stockholm, Sweden

Abstract

The paper addresses various Multiplayer Multi-Armed Bandit (MMAB) problems, where M decision-makers, or players, collaborate to maximize their cumulative reward. We first investigate the MMAB problem where players selecting the same arms experience a collision (and are aware of it) and do not collect any reward. For this problem, we present DPE1 (Decentralized Parsimonious Exploration), a decentralized algorithm that achieves the same asymptotic regret as that obtained by an optimal centralized algorithm. DPE1 is simpler than the state-of-the-art algorithm SIC-MMAB Boursier and Perchet (2019), and yet offers better performance guarantees. We then study the MMAB problem without collision, where players may select the same arm. Players sit on vertices of a graph, and in each round, they are able to send a message to their neighbours in the graph. We present DPE2, a simple and asymptotically optimal algorithm that outperforms the state-of-the-art algorithm DD-UCB Martínez-Rubio et al. (2019). Besides, under DPE2, the expected number of bits transmitted by the players in the graph is finite.

1 Introduction

In Multiplayer MAB, there are M independent decision-makers, or players. At every round: (i) each decision-maker selects an arm in the set $\mathcal{K} = \{1, \ldots, K\}$, (ii)

---

receives some feedback about this arm, and (iii) possibly communicates with neighbouring players. To simplify the presentation, we assume that in round t, when arm k is selected, the potential collected reward is a random variable (independent of the rewards of the other arms) $X_k(t)$ with Bernoulli distribution with mean μ_k. We further assume that the average rewards $\mu = (\mu_1, \ldots, \mu_K)$ are such that $\mu_1 > \mu_2 > \ldots > \mu_K$. MMAB problems have received a lot of attention recently. In this paper, we investigate the two most studied MMAB problems: MMAB *with* collisions, motivated by radio channel assignment problems in cognitive radios Jouini et al. (2009), and MMAB *without* collisions, motivated by sequential decisions in social networks Landgren et al. (2016).

1.1 Multiplayer MAB with collisions

In this model, when a player selects an arm, she collects the corresponding reward only if no other player has selected this same arm. More precisely, when in round t, the player selects k, she observes (1) whether her decision collides with those of other players, and (2) $X_k(t)$ in the absence of collision. This feedback scenario is referred to as *collision sensing* in Boursier and Perchet (2019). The different players are not communicating, and they only sense the presence of other players through experienced collisions. A policy π determines in each round which arm every decision-maker will select. We are interested in distributed policies where each decision-maker decides which arm to select independently. This choice depends on the available information to the decision-maker: the past observed collisions and rewards. We denote by $k_i^\pi(t)$ the arm selected by the decision-maker i in round t under the policy π.

Regret lower bound. The optimal expected reward that can be collected in each round is $\sum_{k=1}^M \mu_k$ (when the M best arms are played). Hence the regret up to

round T of a policy π is defined as:

$$R^\pi(T) = T \sum_{k=1}^{M} \mu_k - \sum_{t=1}^{T} \sum_{i=1}^{M} \mathbb{E}\big[\mu_{k_i^\pi(t)} \mathbf{1}_{\{k_i^\pi(t) \neq k_j^\pi(t), \forall j \neq i\}}\big].$$

As in the classical bandit literature Lai and Robbins (1985), we say that a policy π is *uniformly* good if it regret satisfies $R^\pi(T) = o(T^\alpha)$ for all $\alpha > 0$ for any possible μ. We know from Anantharam et al. (1987) that any uniformly good policy π, centralized or not, satisfies:

$$\liminf_{T \to \infty} \frac{R^\pi(T)}{\log(T)} \geq C(\mu) := \sum_{k > M} \frac{\mu_M - \mu_k}{\mathrm{kl}(\mu_k, \mu_M)}, \quad (1)$$

where $\mathrm{kl}(a, b)$ denotes the KL divergence between two Bernoulli distributions of respective means a and b. This result is a simple extension of the classical result derived by Lai and Robbins (1985). Anantharam et al. (1987) also presents a centralized policy achieving the above asymptotic regret lower bound.

State-of-the-art algorithm. In a recent paper Boursier and Perchet (2019), the authors develop SIC-MMAB, an algorithm that uses collisions as a communication tool, and whose regret satisfies:

$$
\begin{aligned}
R^{\mathrm{SIC}}(T) \leq & c_1 \sum_{k > M} \min \left\{ \frac{\log T}{\mu_M - \mu_k}, \sqrt{T \log T} \right\} \\
& + c_2 K M \log T \\
& + c_3 K M^3 \log^2 \left(\min \left\{ \frac{\log T}{(\mu_M - \mu_{M+1})^2}, T \right\} \right),
\end{aligned}
$$

for some constants $c_1, c_2, c_3 > 0$. The regret of SIC-MMAB is logarithmically increasing with the time horizon, but does not match the regret lower bound (1). In addition, SIC-MMAB needs to know the time horizon in advance. More importantly, it requires involved communication phases (players need to exchange their estimates of the arms' mean rewards), and in turn, the number of collisions used to communicate grows large with T.

Our contributions. We present DPE1 (Decentralized Parsimonious Exploration), a simple policy that achieves the asymptotic fundamental regret limit (1). The policy relies on the observation that in a MAB problem where the decision-maker selects M arms in each round (a model referred to as MAB with multiple plays Anantharam et al. (1987)), an optimal algorithm consists in playing the $(M-1)$ best empirical arms and exploring using the remaining arm according to an optimal index policy, such as KL-UCB Garivier and Cappé (2011). This observation that such *parsimonious* exploration suffices was already made and exploited in Combes et al. (2015) for the design of learning-to-rank

algorithms. It is powerful in the design of a decentralized MMAB algorithm: indeed, it implies that the exploration can be only performed by a single player, the so-called *leader*. The leader maintains the set of the M best empirical arms based on the rewards she received so far for the various arms. The other players, referred to as the *followers*, just need to play these best empirical arms greedily. To this aim, the leader just needs to inform the followers when the set of the M best empirical arms changes – and it can be done using collisions as proposed in Boursier and Perchet (2019).

Our finite-time analysis of the regret of DPE1 reveals that: for all $T \geq 3$ and any $0 < \delta < \delta_0 = \min_{1 \leq k \leq K-1} \frac{\mu_k - \mu_{k+1}}{2}$:

$$
\begin{aligned}
R^{\mathrm{DPE1}}(T) \leq & \sum_{k > M} \frac{\mu_M - \mu_k}{\mathrm{kl}(\mu_k + \delta, \mu_M - \delta)} f(T) \\
& + K^2 M^2 \left[\frac{1}{K - M} + 284 K^{1/2} M(7 + \delta^{-2}) \right],
\end{aligned}
$$

where $f(T) = \log(T) + 4 \log \log(T)$. In particular, by letting first T tend to ∞ and then δ tend to 0, the above result implies that DPE1 is asymptotically optimal: its regret matches the regret lower bound (1). DPE1 achieves the regret of the best possible centralized algorithm. In addition, DPE1 is simpler than SIC-MMAB, since the leader just needs to communicate the indexes of the best empirical arms, when the latter change. In fact, the expected number of collisions used for communication – equivalently the number of communication bits (one may see a collision as communicating one bit) is finite (it is upper bounded by $K^2 M^2 \left[\frac{1}{(K-M)} + 242 K^{1/2}(7 + \delta^{-2}) \right]$).

1.2 Multiplayer MAB without collision

In the absence of collisions, different players can select the same arm. When a player selects arm k in round t, she collects the reward $X_k(t)$[1]. In this model, the players are the vertices of a *communication* graph $G = (V, E)$. At the end of each round, a player can communicate to her neighbors in G. In recent papers Landgren et al. (2016, 2018); Martínez-Rubio et al. (2019), players are assumed to be able to communicate real numbers to their neighbors in each round. We study a more realistic setting where players can only send a finite number of bits per round. As for the model with collision, we are interested in distributed arm selection policies. Under such a policy π, a player i selects in round t arm $k_i^\pi(t)$ and design messages to be

---

[1]When two arms select the same arm, we assume that they receive the same random rewards for simplicity. However, they could well collect stochastically independent rewards; the analysis of the average regret would not be affected.

sent to neighbors depending on her past observations (the collected rewards and the messages received from her neighbors).

Regret lower bound. The maximum expected reward that can be collected in one round is $M\mu_1$. Hence the regret up to round T of a policy π is defined as:

$$R^\pi(T) = MT\mu_1 - \sum_{t=1}^{T}\sum_{i=1}^{M}\mathbb{E}[\mu_{k_i^\pi(t)}].$$

The regret of any uniformly good policy (centralized or not) should satisfy:

$$\lim_{T\to\infty}\inf\frac{R^\pi(T)}{\log(T)} \geq C_1(\mu) := \sum_{k>1}\frac{\mu_1 - \mu_k}{\mathrm{kl}(\mu_k, \mu_1)}. \quad (2)$$

Indeed, the asymptotic regret above corresponds to the best possible regret of a single player.

State-of-the-art algorithm. In Martínez-Rubio et al. (2019), the authors present Distributed Delayed UCB (DD-UCB), an algorithm that combines UCB and a consensus algorithm. The latter is meant so that all players share similar estimates of the mean rewards of the arms, and requires that each player sends a few real numbers to her neighbours in each round. DD-UCB enjoys the following finite-time regret guarantee[2]:

$$R^{\mathrm{DDUCB}}(T) \leq c_1 \sum_{k>1}\frac{\log(MT)}{\mu_1 - \mu_k}$$
$$+ c_2 M\log(M)\sum_{k>1}(\mu_1 - \mu_k),$$

for some constants $c_1, c_2 > 0$. DD-UCB suffers from the same issues as SIC-MMAB for the MMAB with collisions: its regret does not match the regret lower bound (2), and it requires players to communicate a lot.

Our contributions. We present DPE2, an algorithm based on the same *parsimonious exploration* principle as for DPE1. The algorithm starts by electing a leader among the players. After this election, the leader is the only player exploring arms, again using KL-UCB indexes. The other players, the followers, just play the best empirical arm announced by the leader. We show that the regret of DPE2 satisfies: for all $T \geq 3$ and any $0 < \delta < \delta_0$:

$$R^{\mathrm{DPE2}}(T) \leq \sum_{k>1}\frac{\mu_1 - \mu_k}{\mathrm{kl}(\mu_k + \delta, \mu_1 - \delta)}f(T)$$
$$+ 9DKM(29 + K\delta^{-2}),$$

---

[2]This upper bound is derived for subgaussian rewards, so it is valid for Bernoulli rewards as well.

where D is the diameter of the graph G. Hence, DPE2 achieves the regret of the best possible centralized algorithm. In addition, under DPE2, the expected number of bits used for communication is finite (it is upper bounded by $4DM^2\log_2(M) + 8KM^2D\log_2(K)(29 + K\delta^{-2})$).

2 Multiplayer MAB with collisions

This section is devoted to problems where players experience a collision when they select the same arm. We present the DPE1 algorithm, and analyze its performance.

2.1 The DPE1 algorithm

We provide a detailed description of DPE1 and explain its advantages over SIC-MMAB. DPE1 starts with an initialization phase whose objective is to assign different ranks in $\{1, \ldots, M\}$ to players. This rank assignment will be used to avoid collisions. After the initialization phase, DPE1 alternates between exploitation and exploration as usual. However, exploration is conducted by one player only (the leader).

2.1.1 Initialization phase

The first phase consists of coordinating the players. After this phase, a single player becomes the *leader*; this player is ranked first and is aware of this rank. The other players are *followers* and get to know their respective ranks $2, \ldots, M$. All players learn, in passing, the number of players M. After this phase, they can coordinate and avoid collisions except if they need collisions to communicate. SIC-MMAB also starts with such an initialization phase; this phase has by design a fixed duration $T_0 = \lceil K\log(T) \rceil$, which implies in particular that its cost in terms of expected regret is $KM\log(T)$. In contrast, the initialization phase in DPE1 has a random duration: it lasts until all its objectives are achieved. The expected duration of the DPE1 initialization phase is finite, and hence just generates a constant expected regret.

DPE1 initialization phase consists of two sub-phases:

A. Orthogonalization. This first sub-phase aims at assigning in a distributed manner M different arms within $\{1, \ldots, K - 1\}$ to the various players. In this sub-phase, the players maintain an internal state with values in $\{0, 1\ldots, K - 1\}$: when the state is '0', it means that the player is not satisfied, and still needs to find a free arm. When the state is 'k' with $k \neq 0$, it means that the player manages to select arm k without collision, and she will keep this state until the end of the sub-phase. The sub-phase consists of a sequence of

blocks of $K+1$ rounds: in the first round of a block, players with state different than '0' select the arm corresponding to their state, and players with '0' state randomly select an arm in $\{0, 1 \ldots, K-1\}$. The K remaining rounds of the block are used to communicate the outcomes of the first round. This communication is done by selecting arm K and by observing collisions. More precisely, if a player is in state $k \neq 0$, then she selects arm k except in the k-th round where she selects K. If a player is in state '0', she selects arm K in the K rounds. Note that as long as there is a player in state '0', collisions are experienced by all players in the K last rounds of the block. Hence, all the players know that all players are satisfied when no collision is experienced in a block. When such a block occurs for the first time, the sub-phase terminates, and all players are aware of this termination. We prove (see Appendix) that the expected duration of the orthogonalization phase does not exceed $\frac{M(K-1)(K+1)}{K-M}$ rounds.

B. Rank assignment. After the orthogonalization sub-phase, all the players have different states in $\{1, \ldots, K-1\}$. The rank assignment sub-phase consists of $2K-2$ consecutive rounds, denoted by $t_1, \ldots, t_{2K-2}$. Should a player be in state 'k', she selects arms in the following manner: (i) in a round $t_s \in \{t_1, \ldots, t_{2k}\} \cup \{t_{K+k}, \ldots, t_{2K-2}\}$, she selects arm k; (ii) otherwise, when $t_s \in \{t_{2k+1}, \ldots, t_{K+k-1}\}$ she selects arm $s - k$, which corresponds to selecting sequentially the arms $k+1, \ldots, K-1$.

It is easy to observe that with the above procedure, two players only collide once (see Appendix). Furthermore, to determine her rank, a player initially in state k just needs to count the number i_k of collisions experienced in the first $2k$ rounds. $i_k + 1$ becomes her rank. The rank-1 player is the leader.

It should be noted that the leader-follower structure is also adopted in Tibrewal et al. (2019). There, however, the orthogonalization phase could well stop before players have distinct ranks.

2.1.2 Exploration-exploitation phase

In DPE1, the leader is responsible for exploring and maintaining the set of the M best empirical arms. Exploration is conducted using the following KL-UCB indexes. The index of arm k in round t is

$$b_k(t) = \sup\{q \geq 0 : N_k(t)\mathrm{kl}(\hat{\mu}_k(t), q) \leq f(t)\},$$

where $f(t) = \log(t) + 4\log\log(t)$, $N_k(t)$ denotes the number of times the leader has played arm k up to round t, and $\hat{\mu}_k(t)$ is the empirical average of arm k based on the rewards obtained before round t. The leader is also responsible for communicating to the

followers when the set $\mathcal{M}(t)$ of the M best empirical arms changes. To this aim, she leverages collisions in the same manner as in SIC-MMAB. Each time $\mathcal{M}(t)$ changes, a communication phase is initiated by the leader, and this phase lasts a finite number of rounds. The algorithm is designed so that the expected number of times $\mathcal{M}(t)$ changes is finite. The followers just play different arms from $\mathcal{M}(t)$. Note that the followers do not need to communicate anything to the leader; in particular, the rewards they collect are not taken into account by the leader. Each communication phase has a fixed and finite duration and is known to all players – see Subsection 2.1.3 for detail. Hence without loss of generality, we ignore these periods of communication and we can assume that the leader communicates the new $\mathcal{M}(t)$ *instantaneously* whenever required. Communication phases will be, however, accounted for when deriving the regret of the algorithm.

The set of rounds is divided into blocks of MJ rounds where $J = \lceil K^{1/2} \rceil$. In rounds belonging to the same block, the empirical means of the arms, the KL-UCB indexes, and the set of best empirical arms are kept constant. More precisely, the decisions made in one block are based on: For each $k \in \{1, \ldots, K\}$,

$$\hat{\nu}_k(t) = \hat{\mu}_k\left(\left\lfloor\frac{t}{MJ}\right\rfloor MJ\right), \quad d_k(t) = b_k\left(\left\lfloor\frac{t}{MJ}\right\rfloor MJ\right),$$
$$\text{and } \mathcal{N}(t) = \mathcal{M}\left(\left\lfloor\frac{t}{MJ}\right\rfloor MJ\right).$$

At the beginning of a block, the leader updates the above variables. The block structure is designed so that (i) the leader gathers J samples of each of the $(M-1)$ best empirical arms, and in expectation $J/2$ samples from the M-th best empirical arm; (ii) each follower selects each arm in $\mathcal{N}(t)$ J times, and (iii) the leader explores only when the followers play the $(M-1)$ best empirical arms. In particular, a block can be split into J sub-block: each sub-block consists of M consecutive rounds and the follower plays each of the M empirical best arms once in a sub-block. Now $\mathcal{K}$ is decomposed into J subsets $\mathcal{K}_j = \{k \in \mathcal{K} : (j-1)J < k \leq jJ\}$. In the j-th sub-block, we impose that the leader may explore an arm from $\mathcal{K}_j$ only. We introduce in DPE1 the J sub-blocks so as to optimize the constant term (w.r.t. T) of our regret upper bound. We could well run DPE1 without sub-blocks ($J = 1$), it would be still asymptotically optimal, but the constant term in the regret upper bound would be multiplied by $\sqrt{K}$.

Next, we describe in detail the decisions taken by the leader and the followers under DPE1.

Leader. At the beginning of round t, if $t = 0(\mathrm{mod}\ MJ)$, the leader updates the vectors $\hat{\nu}(t)$, $d(t)$, and $\mathcal{N}(t)$. The set $\mathcal{N}(t)$ is ordered: $\mathcal{N}(t) = \{\ell_1(t), \ldots, \ell_M(t)\}$. This order is arbitrary, but indepen-

dent of the empirical means of the arms. In particular, the order is kept fixed even if the relative empirical means of the arms in $\mathcal{N}(t)$ evolve, so that the leader only needs to communicate to the followers when $\mathcal{N}(t)$ changes. Ordering $\mathcal{N}(t)$ is important to avoid collisions. In the following, we denote by $\hat{M}(t)$ the arm in $\mathcal{N}(t)$ with the smallest empirical mean.

If $\mathcal{N}(t) \neq \mathcal{N}(t-1)$, the leader communicates to the followers the identity of the arm leaving the set and that of the new arm that replaces it in $\mathcal{N}(t)$ (the rank of the new arm inherits that of the arm that left).

The sequential arm selections made by the leader are as follows. In round t, define $m = [(t-1)(\mathrm{mod}\ M)]+1$ and $j = \lfloor t/M \rfloor (\mathrm{mod}\ J) + 1$. If $\ell_m(t) \neq \hat{M}(t)$, then the leader selects $\rho(t) = \ell_m(t)$. If $\ell_m(t) = \hat{M}(t)$, then with probability $1/2$, the leader selects arm $\hat{M}(t)$, and with probability $1/2$, the leader plays an arm $k \notin \mathcal{N}(t)$ such that $d_k(t) > \hat{\nu}_{\hat{M}(t)}$, should such an arm exists in $\mathcal{K}_j \backslash \mathcal{N}(t)$, and plays $\hat{M}(t)$ otherwise.

Followers. The followers just exploit the knowledge of the leader: they play greedily different arms of $\mathcal{N}(t)$. More precisely, the follower with rank $i \in \{1, \ldots, M-1\}$ plays in round t the arm $\ell_{m_i}(t)$ where $m_i = [(t-1+i)(\mathrm{mod}\ M)]+1$.

The pseudo-code of the exploration-exploitation phase of the DPE1 algorithm is presented in Algorithm 1.

2.1.3 Communication phases

When $\mathcal{N}(t) \neq \mathcal{N}(t-1)$, the leader starts a communication of the new ordered set $\mathcal{M}(t)$ as follows. She uses 3 blocks: (i) a block of $M-1$ rounds to initiate a communication with the $M-1$ followers, (ii) a block of M rounds to inform the followers of the arm a^- to be removed from the list, and finally (iii) a block of K rounds to inform the followers of the arm a^+ to be added to the list. Hence a communication phase requires $K + 2M - 1$ rounds.

Initial block. In the first block of $(M-1)$ rounds, the leader sequentially selects the arms selected by the followers ranked $2, \ldots, M-1$. The rank-i follower experiences a collision indicating the communication phase, and she knows the round when the communication phase started (she is aware of her rank).

Second block: Removing a^- from $\mathcal{N}(t-1)$. In the next M rounds, the leader selects arm a^-. Followers continue the arm selection as in the exploration-exploitation phase: thus, at some round in this block, they select a^- and collide with the leader, which indicates the arm to be removed.

Third block: Adding a^+ to $\mathcal{N}(t-1)$. During the final

---

Algorithm 1: The DPE1 algorithm: Exploration-exploitation phase

Leader.

Initialization: Set $\hat{\nu}_k(1) = d_k(1) = 0$ for all k
Initialize the set of best empirical arms $\mathcal{N}(1)$ and $\hat{M}(1)$ arbitrarily
Communicate $\mathcal{N}(1) = \{\ell_1(1), \ell_2(1), \ldots, \ell_M(1)\}$ to the followers

For round $t \geq 1$:
1. $m \leftarrow [(t-1)(\mathrm{mod}\ M)]+1$,
   $j \leftarrow \lfloor t/M \rfloor (\mathrm{mod}\ J) + 1$
2. If $t > 1$ and $j = m = 1$, update $\hat{\nu}_k(t), d_k(t)$ for each arm
   k, $\hat{M}(t)$, and update the ordered set
   $\mathcal{N}(t) \leftarrow \{\ell_1(t), \ell_2(t), \ldots, \ell_M(t)\}$
   (the set of the M best empirical arms)
3. If $t > 1$ and $\mathcal{N}(t) \neq \mathcal{N}(t-1)$, communicate $\mathcal{N}(t)$ to the followers
4. For each $s = 1, 2, \ldots, J$,
   $\mathcal{D}_s(t) \leftarrow \left\{ k \notin \mathcal{N}(t) : d_k(t) \geq \hat{\nu}_{\hat{M}(t)}(t), \lceil \frac{k}{J} \rceil = s \right\}$,
   If $\mathcal{D}_j(t) = \emptyset$ or $\ell_m(t) \neq \hat{M}(t)$, $\rho(t) \leftarrow \ell_m(t)$
   Else
      w.p. $1/2$, $\rho(t) \leftarrow \hat{M}(t)$
      w.p. $1/2$, $\rho(t) \leftarrow k$ where $k \sim \mathcal{D}_j(t)$ uniformly
   Select arm $\rho(t)$
Follower with rank $i \in \{1, 2, \ldots, M-1\}$.
In round $t \geq 1$: $m_i \leftarrow [(t-1+i)(\mathrm{mod}\ M)]+1$,
   select arm $\ell_{m_i}(t)$

K rounds, the leader selects arm a^+. Followers change their arm selection, and select all arms during this block without colliding with each other (except with the leader). More precisely, in the m-th round of this block, the follower with rank i selects the arm $[(m + i)(\mathrm{mod}\ K)]+1$. Each follower will experience a collision when selecting a^+, and they can add it to $\mathcal{N}(t-1)$.

2.2 Regret and communication complexity analysis

The next theorem provides a finite-time analysis of the regret of DPE1. It also gives an upper bound on the expected number of collisions involved in the algorithm, including those of the initialization and communication phases.

Theorem 1. *For any μ, $T \geq 3$, and $0 < \delta < \min_{1 \leq k \leq K-1}(\mu_k - \mu_{k+1})/2$, the regret of DPE1 satisfies:*

$$R^{\mathrm{DPE1}}(T) \leq K^2 M^2 \left[\frac{1}{(K-M)} + 284 K^{1/2} M (7 + \delta^{-2}) \right]$$

$$+ \sum_{k > M} \frac{(\mu_M - \mu_k) f(T)}{\mathrm{kl}\,(\mu_k + \delta, \mu_M - \delta)}.$$

Irrespective of the time horizon T, the expected number of collisions under DPE1 *is upper bounded by:*

$$K^2 M^2 \left[\frac{1}{(K-M)} + 242 K^{1/2}(7 + \delta^{-2}) \right].$$

By letting T tend to ∞, and then δ tend to 0 in the above regret upper bound, we simply deduce that DPE1 is asymptotically optimal:

$$\limsup_{T \to \infty} \frac{R^{\text{DPE1}}(T)}{\log T} \leq \sum_{k>M} \frac{\mu_M - \mu_k}{\text{kl}\,(\mu_k, \mu_M)}.$$

To establish Theorem 1, we prove that the expected number of rounds where $\mathcal{M}(t)$ does not correspond to the actual M best arms is finite. This is the key ingredient whose proof actually exploits the arguments used in Combes et al. (2015) to establish a regret upper bound of a centralized algorithm for some MAB problems with multiple plays. From this result, we know that the number of communication phases is finite in expectation, as well as the average regret experienced by the followers. The last term of the regret upper bound just corresponds to the regret paid by the leader.

3 Multiplayer MAB without collisions

This section is devoted to MMAB problems without collisions: players can select the same arm and collect the corresponding rewards. Players sit on the vertices of a graph G, and are assumed to be able to send messages to their neighbors in G in each round. We present DPE2, an algorithm similar to DPE1 and adapted to this new setting.

3.1 The DPE2 algorithm

DPE2 starts with a leader election phase. After the election, DPE2 alternates between exploration and exploitation. As in DPE1, under DPE2, the exploration is conducted by the leader only. The followers just play the best empirical arm seen by the leader. The latter needs to communicate to the followers only when her best empirical arm changes.

3.1.1 Leader election

The leader election problem is a well studied problem in the field of distributed computing Fokkink (2013); Tel (1994); Casteigts et al. (2019). To simplify, we assume here that each player has initially a unique id in $\{1, \ldots, P\}$, where P could be potentially very large[3]. For the analysis of the communication complexity of

the algorithm, we will take $P = M$ for simplicity. With this assumption, leader election can be performed in $O(D + \log M)$ rounds using $\mathcal{O}(1)$ bits per messages Casteigts et al. (2019), where D denotes the diameter of the graph. We propose a simpler alternative election process whose duration is enough for our purposes.

Every player initializes her state to her id. For $(D+1)$ consecutive rounds, every player sends her state to all her neighbours. When a player receives states from other players, she updates her state to the minimal value of the states received and her state. After $(D+1)$ rounds, the player whose state corresponds to her initial id is the leader. Note that the total number of bits sent during this procedure is at most $2|E|(D+1)\lceil \log_2 M \rceil$ (where $|E|$ is the number of edges in G).

3.1.2 Exploration-exploitation phase

In round t, the leader maintains the set $\mathcal{D}(t)$, which consists of those arms with a larger KL-UCB index than that of the best empirical arm $\ell_1(t)$. When this set is empty, she plays $\ell_1(t)$ and updates the empirical means $\hat{\nu}_k(t)$ of the arms and KL-UCB $d_k(t)$ indexes of the arms as well as $\mathcal{D}(t+1)$. She also communicates to her neighbours if $\ell_1(t)$ has changed.
If the set $\mathcal{D}(t)$ is not empty, a block of rounds starts. In the first round of this block, the leader play $\ell_1(t)$; in the subsequent rounds, she plays the arms in $\mathcal{D}(t)$ until this set is exhausted, which then ends the block.

Algorithm 2: The DPE2 algorithm: Exploration-exploitation phase

---

Initialization: Set $\hat{\nu}_k(1) = d_k(1) = 0$ for all k, $\mathcal{D}(1) = \emptyset$, and $s = 0$. Initialize best empirical arm $\ell_1(1)$ arbitrarily.
For round $t \geq 1$:
Leader.
If $\mathcal{D}(t) = \emptyset$,
    If $t > 1$, update $\hat{\nu}_k(t), d_k(t)$ for all k, and $\ell_1(t)$
        $\mathcal{D}(t) \leftarrow \{k \neq \ell_1(t) : d_k(t) \geq \hat{\nu}_{\ell_1(t)}(t)\}$,
        $s \leftarrow \mathbf{1}_{\{|\mathcal{D}(t)|>0\}}$
        If $\ell_1(t) \neq \ell_1(t-1)$, communicate $\ell_1(t)$ to
          the followers
    $\rho(t) \leftarrow \ell_1(t)$, select arm $\rho(t)$
Else
    If $s = 1$, $\rho(t) \leftarrow \ell_1(t)$, $s \leftarrow 0$, select arm $\rho(t)$
    Else, $\rho(t) \leftarrow \arg \max_{k \in \mathcal{D}(t)} d_k(t)$, select arm $\rho(t)$
        $\mathcal{D}(t+1) \leftarrow \mathcal{D}(t) \backslash \{\rho(t)\}$
Follower.
    Select arm $\ell_1(t)$

---

[3]This assumption is mild compared to those made in recent papers, e.g. Martínez-Rubio et al. (2019), addressing the MMAB problem without collision. There, players can

pass real numbers to their neighbours, in which case the leader election becomes trivial. Note also that anonymous leader election can be also performed in finite expected time Fokkink (2013).

3.1.3 Communication phases

When $\ell_1(t-1) \neq \ell_1(t)$, the leader sends the new best empirical arm to all her neighbours. In every round, each time a follower receives the id of an arm that does not correspond to the arm she is playing: (i) she starts playing the new arm from the next round, and (ii) forward the id of the new arm to her neighbours except to the player from whom she received the information.

Each communication phase takes at most D rounds, and the total number of bits transmitted during such a phase does not exceed $2|E|\lceil \log_2(K) \rceil$.

3.2 Regret and communication complexity analysis

The next theorem provides a finite-time analysis of the regret of DPE2, and an upper bound on the expected amount of communication involved in the algorithm.

Theorem 2. *For any μ, $T \geq 3$, and $0 < \delta < \min_{1 \leq k \leq K-1}(\mu_k - \mu_{k+1})/2$, the regret of DPE2 satisfies:*

$$R^{\mathrm{DPE2}}(T) \leq 9MDK(29 + K\delta^{-2})$$
$$+ \sum_{k>1} \frac{\mu_1 - \mu_k}{\mathrm{kl}(\mu_k + \delta, \mu_1 - \delta)} f(T).$$

The total expected number of bits sent under DPE2 is lower than:

$$4DM^2 \log_2(M) + 8KM^2 D \log_2(K)(29 + K\delta^{-2}).$$

The above theorem implies that DPE2 is asymptotically optimal. Its proof relies on similar arguments as those used to establish Theorem 1.

4 Numerical experiments

In this section, we provide initial experiments to illustrate the performance of DPE1 and DPE2. Further experiments are presented in the appendix.

4.1 DPE1

The experiments we run follow the same setting as in Boursier and Perchet (2019). We consider $K = 9$ arms with Bernoulli rewards, and a fixed number of players $M = 6$. We compare the regret of DPE1 with those of SIC-MMAB Boursier and Perchet (2019) and MCTopM Besson and Kaufmann (2018). All the regret and communication complexity values are averaged over 208 runs.

Figure 1 compares DPE1, MCTopM, and SIC-MMAB, over a time horizon of $T = 5 \cdot 10^5$ rounds. The means of the arms are linearly distributed between 0.9

and 0.89, so the gap is approximately $\Delta = 1.1 \cdot 10^{-3}$. DPE1 significantly outperforms both SIC-MMAB and MCTopM.


Figure 1: Regret evolution over a horizon of $5 \cdot 10^5$ rounds. The continuous curves represent the empirical average value, and the shadowed area 3 times the standard deviation.

In Figure 2, we plot the communication complexity vs. time, i.e., the average number of communication phases under DPE1 over time. Figure 2 confirms our theoretical insights: the expected number of times the set $\mathcal{M}(t)$ gets updated is finite. We show the communication complexity for several values of the gap Δ between two consecutive arms, keeping the average reward of the best arm equal to 0.9. Note that the communication complexity increases as Δ approaches 0. In the appendix we also show the average number of communication phases as a function of the gap Δ.


Figure 2: Communication complexity, i.e, the expected number of communication phases. For smaller value of Δ, it is harder for the leader to identify the best arm, and the number of communication phases increases.

4.2 DPE2

To evaluate DPE2, we used the same communicating graph as in Landgren et al. (2016). The graph has 4 players – 3 players are connected to each other, and a fourth player is only connected to one player only. We consider 10 arms with Bernoulli rewards. The range of the rewards is $[0.1, 0.9]$, with a gap of $\Delta = 8/9$ between two consecutive arms. In our experiments, we noticed that the identity of the leader (for this scenario) does not impact regret significantly, so we choose one player arbitrarily as the leader. We use a time horizon of $T = 5 \cdot 10^4$ rounds, and average regrets over 1024 runs.


Figure 3: Regret evolution of a time horizon of $5 \cdot 10^4$ rounds. In *blue* is shown DPE2, whilst in *orange* is shown the regret of a single agent running KL-UCB.

In Figure 3, we compare the regret obtained under DPE2 to the regret that would be obtained by a single player implementing KL-UCB. The regret of DPE2 is not really worse than the regret experienced by a single player. The communication complexity of DPE2 is numerically illustrated in the appendix, and as predicted by our analysis, it is rather small.

5 Related Work

MMAB problems with collisions have attracted a lot of attention over the last decade due to their applicability to the design of decentralized channel selection schemes in cognitive radio systems. Early works include Liu and Zhao (2010) and Anandkumar et al. (2011), proposing various algorithms with regret guarantees far from the best possible regret achieved by DPE1. Even in more recent papers such as Rosenski et al. (2016) and Besson and Kaufmann (2018), the regret upper bounds were actually bigger than expressions of the form: $M \sum_{k>M} \frac{\log(T)}{\mu_M - \mu_k}$. The multiplicative factor M was considered unavoidable, until Boursier and Perchet Boursier and Perchet (2019) proposed to actually use

collisions as a way to communicate. In turn, SIC-MMAB exploits collisions to share the estimated mean rewards of arms. The idea of exploiting collisions had also been suggested in Lugosi and Mehrabian (2018). As already mentioned in the introduction, SIC-MMAB was, until now, the algorithm with the best regret guarantees. Our algorithm, DPE1, is simpler, has better regret guarantees, and as our numerical experiments suggest, outperforms other algorithms. Note that the leader-follower framework used by DPE1 had been also proposed in the algorithms presented in Hanawal and Darak (2018), but the latter exhibit worse regret guarantees than DPE1.

The literature on MMAB problems without collisions is not as abundant. For the case of subgaussian rewards with the same and known variance, Landgren et al. (2016) propose COOP-UCB, an algorithm where each player maintains an estimate of the mean reward of each arm using information obtained by their neighbours via a running consensus scheme Braca et al. (2008). Martínez-Rubio et al. (2019) recently managed to develop a similar algorithm, DD-UCB, with better regret guarantees than that of COOP-UCB. Under both COOP-UCB and DD-UCB, the players need to communicate real numbers to their neighbours (to run a consensus scheme). DPE2, our algorithm, has better regret guarantees, and keeps the total expected number of bits communicated finite.

6 Conclusion

In this paper, we have studied two Multiplayer MAB problems, where we were able to devise decentralized algorithms that achieve the same regret as the one we would obtain by using an optimal centralized algorithm. The design of these algorithms leveraged the critical observation that minimal regret can be achieved by letting only one player exploring arms, and by allowing other players to select the best empirical arms greedily. This observation also implies that the player exploring arms just needs to inform other players regarding the change of the best empirical arms. Moreover, our algorithms do require very little communication: the expected number of bits used for communication does not depend on the time horizon. Further research includes the case of heterogeneous rewards, i.e., the case where the average reward of an arm depends on the player selecting that arm. On the other hand, it would be interesting to further investigate scenarios where players sit on the vertices of a graph. For instance, we could consider that players who are neighbors in the graph collide (collision graphs are instrumental in modelling radio communication systems). In this case, several players would need to explore.

References

Anandkumar, A., Michael, N., Tang, A. K., and Swami, A. (2011). Distributed algorithms for learning and cognitive medium access with logarithmic regret. *IEEE Journal on Selected Areas in Communications*, 29(4):731–745.

Anantharam, V., Varaiya, P., and Walrand, J. (1987). Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays-part i: I.i.d. rewards. *IEEE Transactions on Automatic Control*, 32(11):968–976.

Besson, L. (2018). SMPyBandits: an Open-Source Research Framework for Single and Multi-Players Multi-Arms Bandits (MAB) Algorithms in Python. Online at: github.com/SMPyBandits/SMPyBandits. Code at https://github.com/SMPyBandits/SMPyBandits/, documentation at https://smpybandits.github.io/.

Besson, L. and Kaufmann, E. (2018). Multi-player bandits revisited. *Algorithmic Learning Theory*.

Boursier, E. and Perchet, V. (2019). SIC-MMAB: synchronisation involves communication in multiplayer multi-armed bandits. In *Advances in Neural Information Processing Systems 32*, pages 12048–12057.

Braca, P., Marano, S., and Matta, V. (2008). Enforcing consensus while monitoring the environment in wireless sensor networks. *IEEE Transactions on Signal Processing*, 56(7):3375–3380.

Casteigts, A., Métivier, Y., Robson, J. M., and Zemmari, A. (2019). Deterministic leader election takes $\Theta(D + \log(n))$ bit rounds. *Algorithmica*, 81(5):1901–1920.

Combes, R., Magureanu, S., Proutiere, A., and Laroche, C. (2015). Learning to rank: Regret lower bounds and efficient algorithms. *SIGMETRICS Perform. Eval. Rev.*, 43(1):231–244.

Combes, R. and Proutiere, A. (2014). Unimodal bandits: Regret lower bounds and optimal algorithms. In *International Conference on Machine Learning*, pages 521–529.

Fokkink, W. (2013). *Distributed algorithms: an intuitive approach*. MIT Press.

Garivier, A. and Cappé, O. (2011). The kl-ucb algorithm for bounded stochastic bandits and beyond. In *Proceedings of the 24th annual conference on learning theory*, pages 359–376.

Hanawal, M. K. and Darak, S. J. (2018). Multi-player bandits: A trekking approach. *arXiv preprint arXiv:1809.06040*.

Jouini, W., Ernst, D., Moy, C., and Palicot, J. (2009). Multi-armed bandit based policies for cognitive radio's decision making issues. In *2009 3rd International Conference on Signals, Circuits and Systems (SCS)*, pages 1–6. IEEE.

Lai, T. L. and Robbins, H. (1985). Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22.

Landgren, P., Srivastava, V., and Leonard, N. E. (2016). Distributed cooperative decision-making in multi-armed bandits: Frequentist and bayesian algorithms. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 167–172. IEEE.

Landgren, P., Srivastava, V., and Leonard, N. E. (2018). Social imitation in cooperative multiarmed bandits: Partition-based algorithms with strictly local information. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 5239–5244. IEEE.

Liu, K. and Zhao, Q. (2010). Distributed learning in multi-armed bandit with multiple players. *IEEE Transactions on Signal Processing*, 58(11):5667–5681.

Lugosi, G. and Mehrabian, A. (2018). Multiplayer bandits without observing collision information. *arXiv preprint arXiv:1808.08416*.

Martínez-Rubio, D., Kanade, V., and Rebeschini, P. (2019). Decentralized cooperative stochastic bandits. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R., editors, *Advances in Neural Information Processing Systems 32*, pages 4529–4540. Curran Associates, Inc.

Rosenski, J., Shamir, O., and Szlak, L. (2016). Multiplayer bandits–a musical chairs approach. In *International Conference on Machine Learning*, pages 155–163.

Tel, G. (1994). Network orientation. *International Journal of Foundations of Computer Science*, 5(01):23–57.

Tibrewal, H., Patchala, S., Hanawal, M. K., and Darak, S. J. (2019). Multi-player bandits for optimal assignment with heterogeneous rewards.

A   Initialization of DPE1

The initialization phase of DPE1 consists of two sub-phases, the orthogonalization sub-phase and the rank assignment sub-phase. In this section, we provide Lemma 1 stating that the expected duration of the orthogonalization phase is smaller than $M(K-1)(K+1)/(K-M)$, and Lemma 2 proving the correctness of the rank assignment sub-phase. We also present numerical experiments illustrating the durations of the two sub-phases.

A.1   Orthogonalization

Lemma 1 (Complexity of the orthogonalization sub-phase)**.** *Let τ denotes the number of rounds required to end the orthogonalization sub-phase. Then*

$$\mathbb{E}[\tau] \leq \frac{M(K-1)(K+1)}{K-M}.$$

Proof. Let $\mathcal{E}_t$ denote the event that in the end of the t-th block of rounds, some players are still unsatisfied (i.e. still in state 0). Let $\mathcal{A}_t^k$ denote the event that player k experiences a collision in the first round of the t-th block. Note that $\mathcal{E}_t = \cup_{k=1}^M \cap_{s \leq t} \mathcal{A}_s^k$. The probability that at some block $s \leq t$, player k is satisfied satisfies:

$$\mathbb{P}(\bar{\mathcal{A}}_s^k) \geq \frac{K-M}{K-1}.$$

Indeed, there are $K-1$ arms to choose from; there are $M-1$ other players and each of them is either satisfied in block s or unsatisfied in which case, she samples an arm uniformly at random in $[K-1]$. In all the cases, at least $K-M$ arms are not selected by the other players. We deduce that:

$$\mathbb{P}(\mathcal{A}_s^k) \leq 1 - \frac{K-M}{K-1} = \frac{M-1}{K-1}.$$

Hence, by the union bound, we get:

$$\mathbb{P}(\mathcal{E}_t) \leq \sum_{k=1}^M \mathbb{P}(\cap_{s \leq t} \mathcal{A}_s^k) \leq M\Big(\frac{M-1}{K-1}\Big)^t.$$

Denote by τ_b the number of blocks required to end the orthogonalization sub-phase. We have:

$$\begin{aligned}
\mathbb{E}[\tau_b] &= \sum_{t=1}^\infty \mathbb{P}(\tau_b \geq t), \\
&\leq \sum_{t=1}^\infty \mathbb{P}(\mathcal{E}_{t-1}), \\
&\leq M \sum_{t=1}^\infty \Big(\frac{M-1}{K-1}\Big)^{t-1}, \\
&= M\frac{K-1}{K-M}.
\end{aligned}$$

Note that $\tau = (K+1)\tau_b$, and thus

$$\mathbb{E}[\tau] \leq \frac{M(K-1)(K+1)}{K-M}.$$

$\square$

A.2   Rank assignment

The following lemma guarantees the correctness of the rank assignment sub-phase. Indeed, a direct consequence of the lemma is that a player in state k will collide once during the $2k$ first rounds with each player with state smaller than k. The player in state k will not collide with players with state higher than k. Thus, if i_k denotes the number of collisions seen by this player in the $2k$ first rounds, then her rank is simply $i_k + 1$.

Lemma 2 (Pairwise collisions). *The arm selection scheme of the rank assignment sub-phase guarantees that any two players with distinct states k and j such that $k < j$, only collide once during the rank assignment sub-phase. More precisely, they collide at the round t_{k+j} over the selection of arm j.*

Proof of lemma 2. In the following we say player 'k' to refer to the player with state 'k'. Let I_1^k, I_2^k, I_3^k be the partition of $\{t_1, \ldots, t_{2K-2}\}$ parametrized by the state 'k', where

$$I_1^k \triangleq \{t_1, \ldots, t_{2k}\},$$
$$I_2^k \triangleq \{t_{2k+1}, \ldots, t_{K+k-1}\},$$
$$I_3^k \triangleq \{t_{K+k}, \ldots, t_{2K-2}\}.$$

We list all the possible cases:

(i) If $t_s \in \left(I_1^k \cup I_3^k\right) \cap \left(I_1^j \cup I_3^j\right)$, then player '$k$' selects arm k, and player 'j' selects j, thus the two players do not collide.

(ii) If $t_s \in \left(I_1^k \cup I_3^k\right) \cap I_2^j$, then player '$k$' selects arm k, and player 'j' selects $a \in \{j+1, \ldots, K-1\}$, thus the two players do not collide.

(iii) If $t_s \in I_2^k \cap I_1^j$, then player k selects $s - k$, and player j selects j. Note that the two players collide iff $s = k + j$ and $k + j \in I_2^k \cap I_1^j$.

(iv) If $t_s \in I_2^k \cap I_2^j$, then player k selects $s - k$, and player j selects arm $s - j$ thus they do not collide.

(v) Finally, note that $I_2^k \cap I_3^j = \emptyset$ because $K + k - 1 < K + j$.

As a conclusion the players 'k' and 'j' only collide once throughout the rank assignment sub-phase, and this occurs at round t_{k+j} over the selection of arm j. This concludes the proof. $\square$

A.3 Numerical experiments

We compare in Figure 4 the upper bound derived in Lemma 1 with empirical data. In Figure 4, the black dashed line shows the average duration of the initialization phase, whilst blue dashed line shows the upper-bound derived in Lemma 1. Notice that all the values are normalized by $K + 1$, and the curves were computed using data from 2000 runs.
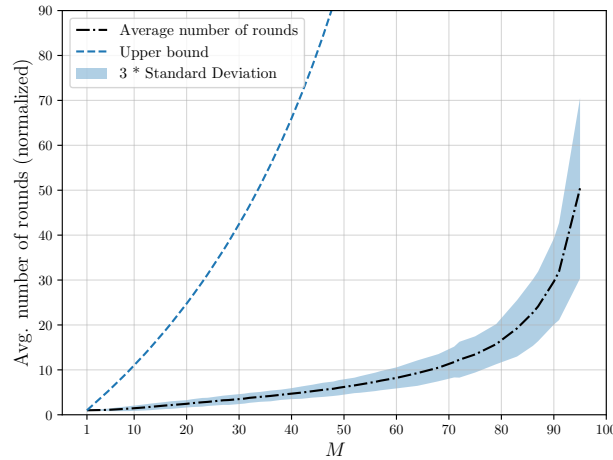

Figure 4: Average duration of the initialization phase with $K = 100$. All results are normalized by $K + 1$.

It is apparent that the upper bound derived in Lemma 1 is loose. Furthermore, the average duration of the initialization phase seems to be linear in M until M becomes too large (here 60).

Figure 5 and 6 present the histograms of the number of rounds of the initialization phase for $M = 20$ and $M = 90$, respectively, with $K = 100$. As M increases, the bell shape is less distinguishable, but still apparent. The smoothed curve is a Gaussian kernel density estimate, computed using the default parameters of the DISTPLOT function in the SEABORN library.
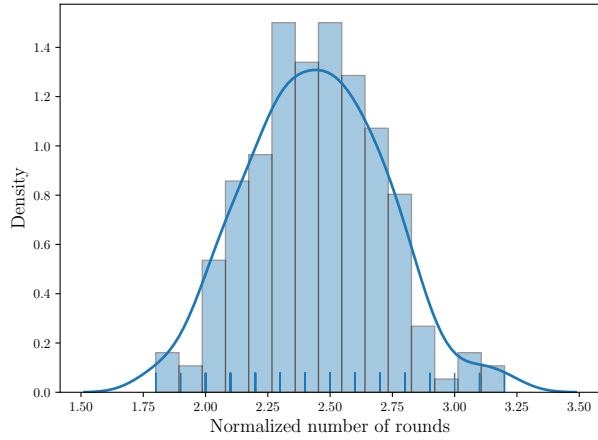

Figure 5: Histogram of the number of rounds of the initialization phase with $M = 20$ and $K = 100$.
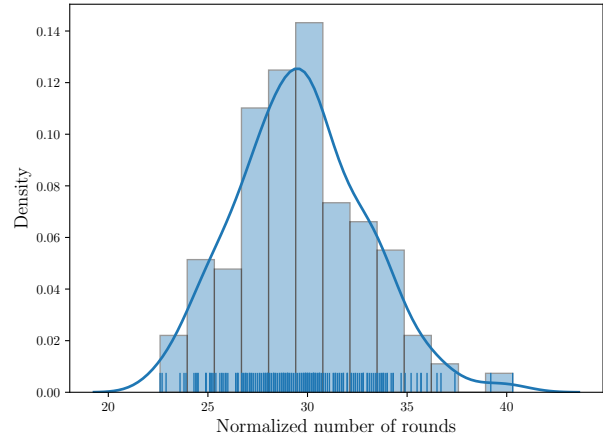
Figure 6: Histogram of the number of rounds of the initialization phase with $M = 90$ and $K = 100$.

B   Performance analysis of DPE1

B.1   Preliminaries

In the proof of Theorem 1, we repeatedly use the following lemma. The latter is a simplified version of Lemma 5 in Combes et al. (2015)). In what follows, $\mathcal{F}_n$ denotes the σ-algebra generated by $(X_k(t), k \in [K], t \leq n)$.

Lemma 3. *Let $k \in [K]$, and $c > 0$. Let H be a random set of rounds such that for all n, $\{n \in H\} \in \mathcal{F}_{n-1}$. Assume that there exists $(C_n)_{n \geq 0}$, a sequence of independent binary random variables such that for any $n \geq 1$, C_n is $\mathcal{F}_n$-measurable and $\mathbb{P}[C_n = 1] \geq c$. Further assume for any $n \in H$, k is selected ($\rho(n) = k$) if $C_n = 1$. Then:*

$$\sum_{n \geq 1} \mathbb{P}[n \in H, |\hat{\mu}_k(n) - \mu_k| \geq \delta\}] \leq 2c^{-1} \left(2c^{-1} + \delta^{-2}\right).$$

In addition, we need some known results about the KL-UCB indexes. The following lemma is a direct consequence of Theorem 10 in Garivier and Cappé (2011).

Lemma 4.   *Under the DPE1 algorithm, for any $k \in \mathcal{K}$, we have:*

$$\sum_{n \geq 1} \mathbb{P}\left[d_k(n) < \mu_k\right] \leq 114 K^{1/2} M$$

Proof. By the design of DPE1, we update $d_k(n)$ only if n is the first round of the block. Furthermore, since the size of each block is always MJ, we have

$$\sum_{n \geq 1} \mathbb{P}\left[d_k(n) < \mu_k\right] \leq \sum_{s \geq 0} (MJ)\mathbb{P}\left[d_k(sMJ + 1) < \mu_k\right].$$

Additionally, as a consequence of Theorem 10 in Garivier and Cappé (2011), we have for each $s \geq 2$, $\mathbb{P}\left[d_k(sMJ + 1) < \mu_k\right] \leq e\lceil f(sMJ + 1)\log(sMJ + 1)\rceil e^{-f(sMJ+1)}$. Hence, we can conclude

$$\sum_{n \geq 1} \mathbb{P}\left[d_k(n) < \mu_k\right] \leq MJ(2 + \sum_{s \geq 2} e\lceil f(sMJ + 1)\log(sMJ + 1)\rceil e^{-f(sMJ+1)})$$

$$\leq MJ(2 + \sum_{s \geq 3} e\lceil f(s)\log(s)\rceil e^{-f(s)})$$

$$= MJ(2 + \sum_{s \geq 3} \frac{e\lceil \log(s)^2 + 4\log(s)\log(\log(s))\rceil}{s\log(s)^4})$$

$$\leq MJ\left[2 + e\left(\int_e^\infty \frac{1}{s(\log s)^2} + \int_e^\infty \frac{4\log(\log s)}{s(\log s)^3} + \int_e^\infty \frac{1}{s(\log s)^4}\right)\right]$$

$$\leq 114 K^{1/2} M.$$

The second inequality follows from the fact that $\{sMJ + 1\}_{s \geq 1} \subset \{s \in \mathbb{N} : s \geq 3\}$ and $J = \lceil K^{1/2} \rceil \leq 2K^{1/2}$.   □

B.2   Proof of Theorem 1

Let $\mathcal{M}^* = \{1, \ldots, M\}$ be the set of the M best arms. Define $\delta_0 = \min_{1 \leq k \leq K-1} \frac{\mu_k - \mu_{k+1}}{2}$ as half of the minimum gap between the expected rewards of the arms. In what follows we choose $0 < \delta < \delta_0$. We define for any $t \geq 1$, $m(t) = [(t-1)(\text{mod } M)] + 1$ and $j(t) = \lfloor t/M \rfloor (\text{mod } J) + 1$.

We define the following sets of rounds:

$$\mathcal{A} = \{n \geq 1 : \mathcal{N}(n) \neq \mathcal{M}^*\},$$
$$\mathcal{B} = \{n \geq 1 : \exists k \in \mathcal{N}(n) \text{ s.t. } |\hat{\nu}_k(n) - \mu_k| \geq \delta\},$$
$$\mathcal{C} = \{n \geq 1 : \exists k \in \mathcal{M}^*, d_k(n) < \mu_k\},$$
$$\mathcal{E} = \{n \geq 1 : n \in \mathcal{A} \backslash (\mathcal{B} \cup \mathcal{C}), \exists k \in \mathcal{M}^* \backslash \mathcal{N}(n) \text{ s.t. } |\hat{\nu}_k(n) - \mu_k| \geq \delta\}.$$

Lemma 5. $\mathcal{A} \cup \mathcal{B} \subseteq \mathcal{B} \cup \mathcal{C} \cup \mathcal{E}$. *As a consequence, we have*

$$\mathbb{E}\left[|\mathcal{A} \cup \mathcal{B}|\right] \leq \mathbb{E}\left[|\mathcal{B}|\right] + \mathbb{E}\left[|\mathcal{C}|\right] + \mathbb{E}\left[|\mathcal{E}|\right].$$

Proof. Let $n \in \mathcal{A} \backslash (\mathcal{B} \cup \mathcal{C})$. We show that $n \in \mathcal{E}$. Since $n \notin \mathcal{B}$, $\forall k \in \mathcal{N}(n)$, we have

$$|\hat{\nu}_k(n) - \mu_k| < \delta. \tag{3}$$

Moreover, $n \in \mathcal{A}$. Hence there exists $i \in \mathcal{M}^* \backslash \mathcal{N}(n)$ such that

$$\hat{\nu}_i(n) < \hat{\nu}_k(n) \text{ for some } k \in \mathcal{N}(n) \backslash \mathcal{M}^*. \tag{4}$$

Combining (3) and (4) leads to $\hat{\nu}_i(n) < \hat{\nu}_k(n) \leq \mu_k + \delta \leq \mu_M - \delta \leq \mu_i - \delta$. The last two inequalities are due to our assumption that $i \leq M < k$ and $\delta < \delta_0$. It implies $|\hat{\nu}_i(n) - \mu_i| \geq \delta$ and thus, $n \in \mathcal{E}$. Therefore, $\mathcal{A} \cup \mathcal{B} \subseteq \mathcal{B} \cup \mathcal{C} \cup \mathcal{E}$. □

Lemma 6. *We have:* $\mathbb{E}\left[|\mathcal{B}|\right] + \mathbb{E}\left[|\mathcal{C}|\right] + \mathbb{E}\left[|\mathcal{E}|\right] \leq 40K^{3/2}M^2(7 + \delta^{-2})$.

Proof. We upper bound each term.
(a) We show that $\mathbb{E}\left[|\mathcal{B}|\right] < 8MK^{3/2}(4 + \delta^{-2})$.
For each $k \in \{1, 2, \ldots, K\}$, define $\mathcal{B}_k = \{n \geq 1 : k \in \mathcal{N}(n), |\hat{\nu}_k(n) - \mu_k| \geq \delta\}$. We have $\mathcal{B} = \cup_{1 \leq k \leq K} \mathcal{B}_k$. Let us fix $k \in \{1, 2, \ldots, K\}$ and split $\mathcal{B}_k$ into two sets,

$$\mathcal{B}_{k,1} = \left\{n \in \mathcal{B}_k : \ell_{m(n)}(n) = k, j(n) = 1\right\},$$
$$\mathcal{B}_{k,2} = \mathcal{B}_k \backslash \mathcal{B}_{k,1}.$$

Let us upper bound the expected cardinality of $\mathcal{B}_{k,1}$. Let $n \in \mathcal{B}_{k,1}$. Observe that arm k is not selected in the sub-block where round n lies until round n. Hence we have $\hat{\nu}_k(n) = \hat{\mu}_k(n)$ (because n is in the first sub-block of a block of rounds, i.e., $j(n) = 1$). We may now apply Lemma 3 with $H = \{n \geq 1 : \ell_{m(n)}(n) = k, j(n) = 1\}$ and $C_n = \mathbb{1}\{\rho(n) = k\}$. Note that $\{n \in H\} \in \mathcal{F}_{n-1}$. By design of the algorithm, for $n \in H$, k is selected with probability at least $c = 1/2$. Thus: $\mathbb{E}\left[|\mathcal{B}_{k,1}|\right] \leq 4(4 + \delta^{-2})$.

To upper bound the expected cardinality of $\mathcal{B}_{k,2}$, it is observed that since the DPE1 algorithm operates by blocks of MJ rounds (i.e., $\hat{\nu}(t)$, $d(t)$, and $\mathcal{N}(t)$ do not change over MJ consecutive rounds), when $n \in \mathcal{B}_{k,2}$, there exists a round p such that $|n - p| < MJ$ (p belongs to the same block as n) and such that $p \in \mathcal{B}_{k,1}$. Hence $|\mathcal{B}_{k,2}| \leq (MJ - 1)|\mathcal{B}_{k,1}|$.

We have established that $\mathbb{E}\left[|\mathcal{B}_k|\right] \leq 4MJ(4 + \delta^{-2})$, and thus $\mathbb{E}\left[|\mathcal{B}|\right] < 4MJK(4 + \delta^{-2}) = 8MK^{3/2}(4 + \delta^{-2})$ (because $J = \lceil K^{1/2} \rceil \leq 2K^{1/2}$).

(b) We show that $\mathbb{E}\left[|\mathcal{C}|\right] < 114K^{1/2}M$.
We apply Lemma 4 for each arm $k \in \{1, 2, \ldots, M\}$, so $\sum_{n \geq 1} \mathbb{P}\left[d_k(n) < \mu_k\right] \leq 114K^{1/2}M$. We simply deduce that:

$$\mathbb{E}\left[|\mathcal{C}|\right] \leq 114K^{1/2}M.$$

(c) We show that $\mathbb{E}\left[|\mathcal{E}|\right] < 32K^{3/2}M^2(4 + \delta^{-2})$.
Define

$$\mathcal{E}_k = \{n \geq 1 : n \in \mathcal{A} \backslash (\mathcal{B} \cup \mathcal{C}), k \notin \mathcal{N}(n) \text{ and } |\hat{\nu}_k(n) - \mu_k| \geq \delta\},$$

for all $k \in \mathcal{M}^*$. Then $\mathcal{E} \subseteq \bigcup_{k \leq M} \mathcal{E}_k$.
Fix $k \in \mathcal{M}^*$, and let $n \in \mathcal{E}_k$. Since $n \notin \mathcal{B}$, we have for all $i \in \mathcal{N}(n)$, $|\hat{\nu}_i(n) - \mu_i| < \delta$. Now let $i^* = \max\{i : i \in \mathcal{N}(n)\}$. We have $i^* > M$ (since $n \in \mathcal{A}$), which implies that:

$$\hat{\nu}_{i^*}(n) < \mu_{i^*} + \delta \leq \mu_{M+1} + \delta < \frac{\mu_{M+1} + \mu_M}{2}. \tag{5}$$

The last inequality follows from the definition of $\delta < \frac{\mu_{M+1}-\mu_M}{2}$. Furthermore, since $n \notin \mathcal{C}$,

$$d_k(n) \geq \mu_k. \tag{6}$$

Combining (5) and (6), we get

$$d_k(n) \geq \mu_k \geq (\mu_{M+1} + \mu_M)/2 > \hat{\nu}_{i^*}(n) \geq \hat{\nu}_{\hat{M}(n)}(n),$$

where the last inequality stems from the fact that $i^* \in \mathcal{N}(n)$. Observe then that such a round n, by the design of algorithm, arm k will be selected with probability at least $1/(2J)$ when $\ell_{m(n)}(n) = \hat{M}(n)$ (exploration rounds). Next we split $\mathcal{E}_k$ into the following two sets:

$$\mathcal{E}_{k,1} = \left\{ n \in \mathcal{E}_k : \ell_{m(n)}(n) = \hat{M}(n), k \in \mathcal{K}_{j(n)} \right\},$$
$$\mathcal{E}_{k,2} = \mathcal{E}_k \backslash \mathcal{E}_{k,1}.$$

For the set $\mathcal{E}_{k,1}$, since $k \notin \mathcal{M}(n)$ and $k \in \mathcal{K}_{j(n)}$, the real empirical mean $\hat{\mu}_k(t)$ remains as $\hat{\nu}_k(t)$ for all $t \leq n$ in the block n belongs to. We apply Lemma 3 with $H = \mathcal{E}_{k,1}$, $C_n = \mathbb{1}\{\rho(n) = k\}$ and $c = 1/(2J)$, and conclude that $\mathbb{E}\left[|\mathcal{E}_{k,1}|\right] \leq 4J(4J + \delta^{-2}) \leq 16K(4 + \delta^{-2})$.

For the set $\mathcal{E}_{k,2}$, again since the algorithm works in blocks, when $n \in \mathcal{E}_{k,2}$, there exists a round p such that $|n - p| < MJ$ (p belongs to same block as n) and such that $p \in \mathcal{E}_{k,1}$. Hence, $\mathbb{E}\left[|\mathcal{E}_{k,2}|\right] \leq (MJ - 1)\mathbb{E}\left[|\mathcal{E}_{k,1}|\right]$, and

$$\mathbb{E}\left[|\mathcal{E}|\right] \leq \sum_{k=1}^{M}\left(\mathbb{E}\left[|\mathcal{E}_{k,1}|\right] + \mathbb{E}\left[|\mathcal{E}_{k,2}|\right]\right) \leq 16KM^2J(4 + \delta^{-2}),$$
$$\leq 32K^{3/2}M^2(4 + \delta^{-2}).$$

$\square$

Lemma 7. *Given $T \geq 3$, we define $\mathcal{G}_k = \{n \leq T, n \notin \mathcal{A} \cup \mathcal{B}, \rho(n) = k\}$ for some $k \notin \mathcal{M}^*$. We show that*

$$\mathbb{E}\left[|\mathcal{G}_k|\right] \leq \frac{\log T + 4\log(\log T)}{\mathrm{kl}\left(\mu_k + \delta, \mu_M - \delta\right)} + 4 + 2\delta^{-2}.$$

Proof. Define the counter $c(n) = \sum_{t=1}^{n} \mathbb{1}_{\{t \in \mathcal{G}_k\}}$, which is the number of rounds in $\mathcal{G}_k$ up to round n and define $t_0 = (\log T + 4\log(\log T))/\mathrm{kl}\left(\mu_k + \delta, \mu_M - \delta\right)$.
Define two subsets of $\mathcal{G}_k$ as

$$\mathcal{G}_{k,1} = \{n \in \mathcal{G}_k : |\hat{\nu}_k(n) - \mu_k| \geq \delta\},$$
$$\mathcal{G}_{k,2} = \{n \in \mathcal{G}_k : c(n) < t_0\}.$$

We first show that $\mathcal{G}_k \subseteq \mathcal{G}_{k,1} \cup \mathcal{G}_{k,2}$. Let $n \in \mathcal{G}_k \backslash (\mathcal{G}_{k,1} \cup \mathcal{G}_{k,2})$. By definition of $\mathcal{A}$, since $n \notin \mathcal{A}$, $\mathcal{N}(n) = \mathcal{M}^*$. We deduce that $k \notin \mathcal{N}(n)$.

The arm selection scheme of DPE1 is such that, during the block to which round n belongs, arm k is only selected at round n (indeed note that the algorithm explores at most once per sub-block). Hence, $N_k(n) = N_k(\lfloor n/(MJ)\rfloor MJ)$. Since $n \notin \mathcal{G}_{k,2}$,

$$N_k(n) \geq c(n) \geq t_0. \tag{7}$$

Observe that $\rho(n) = k$ can only happen when

$$d_k(n) \geq \hat{\nu}_{\hat{M}(n)}(n) = \hat{\nu}_M(n), \tag{8}$$

where the last equality stems from the fact that $n \notin \mathcal{B}$ (indeed, all arms in $\mathcal{N}(n) = \mathcal{M}^*$ are well estimated, so that $M = \hat{M}(n)$). Moreover, $n \notin \mathcal{B}$ implies that

$$\hat{\nu}_M(n) > \mu_M - \delta. \tag{9}$$

Finally $n \notin \mathcal{G}_{k,1}$ and $\delta < \min_{1 \le k \le K-1} \frac{\mu_k - \mu_{k+1}}{2}$ imply that

$$\hat{\nu}_k(n) < \mu_k + \delta < \mu_M - \delta. \tag{10}$$

Combining the above arguments, we get:

$$
\begin{aligned}
t_0 \mathrm{kl}\left(\hat{\nu}_k(n), \mu_M - \delta\right) &\le N_k(n)\mathrm{kl}\left(\hat{\nu}_k(n), \mu_M - \delta\right), \\
&< N_k(n)\mathrm{kl}\left(\hat{\nu}_k(n), d_k(n)\right), \\
&\le \log T + 4\log(\log T).
\end{aligned} \tag{11}
$$

The first inequality follows from (16); the second inequality stems from (17)-(18)-(19) and the fact that $y \mapsto \mathrm{kl}\,(x, y)$ is an increasing function when $0 < x < y < 1$; the last inequality is obtained by definition of $d_k(n)$. Replacing t_0 by its value in the above inequality, we finally obtain:

$$\mathrm{kl}\left(\hat{\nu}_k(n), \mu_M - \delta\right) \le \mathrm{kl}\left(\mu_k + \delta, \mu_M - \delta\right).$$

Now observe that $x \mapsto \mathrm{kl}\,(x, y)$ is a decreasing function when $0 < x < y < 1$. We conclude that $\hat{\nu}_k(n) \ge \mu_k + \delta$ which contradicts the assumption that $n \notin \mathcal{G}_{k,1}$. Hence, $\mathcal{G}_k = \mathcal{G}_{k,1} \cup \mathcal{G}_{k,2}$.

To complete the proof of the lemma, we upper bound $\mathbb{E}\left[|\mathcal{G}_{k,1}|\right]$ and $\mathbb{E}\left[|\mathcal{G}_{k,2}|\right]$.
For $\mathbb{E}\left[|\mathcal{G}_{k,1}|\right]$. Let $n \in \mathcal{G}_{k,1}$. Since $k \notin \mathcal{N}(n)$, k is selected for the first time in the block to which n belongs precisely in round n. Hence, we have $\hat{\mu}_k(n) = \hat{\nu}_k(n)$. We apply Lemma 3 with $C_n = 1$, $c = 1$ and $H = \mathcal{G}_{k,1}$, and get $\mathbb{E}\left[|\mathcal{G}_{k,1}|\right] \le 4 + 2\delta^{-2}$.

For $\mathbb{E}\left[|\mathcal{G}_{k,2}|\right]$. If $n \in \mathcal{G}_{k,2}$, $c(n) \le t_0$ and $c(n)$ is incremented by $+1$. Therefore,

$$\mathbb{E}\left[|\mathcal{G}_{k,2}|\right] < t_0 = \frac{\log T + 4\log(\log T)}{\mathrm{kl}\left(\mu_k + \delta, \mu_M - \delta\right)}.$$

We have proved that:

$$\mathbb{E}\left[|\mathcal{G}_k|\right] < \frac{\log T + 4\log(\log T)}{\mathrm{kl}\left(\mu_k + \delta, \mu_M - \delta\right)} + 4 + 2\delta^{-2}.$$

$\square$

Proof of Theorem 1. The regret can be bounded as follows:

$$R^\pi(T) \le M\left(\frac{K^2 M}{K - M} + 2K\right) + 6KM\mathbb{E}\left[|\mathcal{A}|\right] + M\mathbb{E}\left[|\mathcal{A} \cup \mathcal{B}|\right] + \sum_{k > M}(\mu_M - \mu_k)\mathbb{E}\left[|\mathcal{G}_k|\right].$$

The first term and second term correspond to an upper bound of the regret induced by initialization and communication rounds respectively. Indeed, the number of rounds per communication phase is $K + 2M - 1 \le 3K$. In addition, note that the number of communication phases can be bounded as

$$|\{t \ge 2 : \mathcal{N}(t) \ne \mathcal{N}(t-1)\}| \le 2|\mathcal{A}|.$$

Applying Lemmas 5 and 6, we get:

$$6KM\mathbb{E}\left[|\mathcal{A}|\right] + M\mathbb{E}\left[|\mathcal{A} \cup \mathcal{B}|\right] \le 40K^{3/2}M^3(6K + 1)(7 + \delta^{-2}). \tag{12}$$

Hence, Lemma 7 yields:

$$R^\pi(T) \le K^2 M^2 \left[\frac{1}{K - M} + 284K^{1/2}M(7 + \delta^{-2})\right] + \sum_{k > M}\frac{(\mu_M - \mu_k)(\log T + 4\log\log T)}{\mathrm{kl}\left(\mu_k + \delta, \mu_M - \delta\right)}.$$

The second statement stems from the fact that each communication phase lasts less than $3K$ rounds. Therefore, the total number of collisions can be bounded by

$$M\left(\frac{K^2 M}{K - M} + 2K\right) + 6K\mathbb{E}\left[|\mathcal{A}|\right] \le K^2 M^2 \left(\frac{1}{K - M} + 242K^{1/2}(7 + \delta^{-2})\right).$$

$\square$

C   Performance analysis of DPE2

C.1   Preliminary

In the proof of Theorem 2, we repeatedly use the following lemma from Combes and Proutiere (2014).

Lemma 8. *Let $k \in \{1, \ldots, K\}$, and $\zeta > 0$. Define $\mathcal{F}_n$ the σ-algebra generated by $(X_k(t))_{1 \le t \le n, 1 \le k \le K}$. Let $H \subset \mathbb{N}$ be a (random) set of instants such that $\{n \in H\}$ is $\mathcal{F}_{n-1}$-measurable for all $n \ge 1$. Furthermore, we assume for each $n \in H$, we have $N_k(n) \ge \zeta \sum_{t=1}^{n} \mathbb{1}_{\{t \in H\}}$. Then for all $\delta > 0$:*

$$\mathbb{E}[\sum_{n \ge 1} \mathbb{1}\{n \in H, |\hat{\mu}_k(n) - \mu_k| > \delta\}] \le \frac{1}{\zeta \delta^2}. \tag{13}$$

Moreover, as Lemma 4, the following lemma is a consequence of Theorem 10 in Garivier and Cappé (2011).

Lemma 9. *Under the DPE2 algorithm, we have:*

$$\sum_{n \ge 1} \mathbb{P}[d_1(n) < \mu_1] \le 57K.$$

Proof. Suppose n_s is the first round of s-th block for any $s \ge 1$. By the design of DPE2, we only update $d_1(n)$ when $n = n_s$ for some $s \ge 1$. Therefore, for any $s \ge 1$, we have $\mathbb{P}[d_1(n) < \mu_1] = \mathbb{P}[d_1(n_{s+1}) < \mu_1]$ for all $n_s < n \le n_{s+1}$. Furthermore, the size of each block is always smaller than K, so $n_{s+1} - n_s \le K$. We deduce that

$$\sum_{n \ge 1} \mathbb{P}[d_1(n) < \mu_1] \le \mathbb{E}\left[\sum_{s \ge 1}(n_{s+1} - n_s)\mathbb{1}\{d_1(n_s) < \mu_1\}\right] \le K\sum_{s \ge 1}\mathbb{P}[d_1(n_s) < \mu_1].$$

Additionally, as a consequence of Theorem 10 in Garivier and Cappé (2011), we have for any $n \ge 3$, $\mathbb{P}[b_1(n) < \mu_1] \le e\lceil f(n)\log(n)\rceil e^{-f(n)}$. Hence, from the fact that $d_1(n_s) = b_1(n_s)$ for any $s \ge 1$, we can conclude

$$
\begin{aligned}
\sum_{n \ge 1} \mathbb{P}[d_1(n) < \mu_1] &\le K\left(2 + \sum_{s \ge 3}\mathbb{P}[b_1(n_s) < \mu_1]\right) \\
&\le K\left(2 + \sum_{s \ge 3}\mathbb{P}[b_1(s) < \mu_1]\right) \\
&\le K(2 + \sum_{s \ge 3}e\lceil f(s)\log(s)\rceil e^{-f(s)}) \\
&= K(2 + \sum_{s \ge 3}\frac{e\lceil \log(s)^2 + 4\log(s)\log(\log(s))\rceil}{s\log(s)^4}) \\
&\le K\left[2 + e\left(\int_e^{\infty}\frac{1}{s(\log s)^2} + \int_e^{\infty}\frac{4\log(\log s)}{s(\log s)^3} + \int_e^{\infty}\frac{1}{s(\log s)^4}\right)\right] \\
&= 57K.
\end{aligned}
$$

The second inequality is due to $\{n_s\}_{s \ge 3} \subset \{s \in \mathbb{N} : s \ge 3\}$. $\square$

C.2   Regret Analysis

Recall that we denote $\delta_0 = \min_{1 \le k \le K-1}\frac{\mu_k - \mu_{k+1}}{2}$ as half of the minimum gap between the expected rewards of the arms. In what follows we choose $0 < \delta < \delta_0$. Let us introduce the following sets of events:

$$
\begin{aligned}
\mathcal{A} &= \{n \ge 1 : \ell_1(n) \ne 1\}, \\
\mathcal{B} &= \{n \ge 1 : |\hat{\nu}_{\ell_1(n)}(n) - \mu_{\ell_1(n)}| \ge \delta\}, \\
\mathcal{C} &= \{n \ge 1 : d_1(n) < \mu_1\}, \\
\mathcal{E} &= \{n \ge 1 : n \in \mathcal{A}\backslash(\mathcal{B} \cup \mathcal{C}), |\hat{\nu}_1(n) - \mu_1| \ge \delta\}.
\end{aligned}
$$

It follows from exactly the same arguments as those used to prove Lemma 5 that:

$$\mathbb{E}\left[|\mathcal{A}\cup\mathcal{B}|\right] \leq \mathbb{E}\left[|\mathcal{B}|\right] + \mathbb{E}\left[|\mathcal{C}|\right] + \mathbb{E}\left[|\mathcal{E}|\right].$$

Lemma 10. *We have:* $\mathbb{E}\left[|\mathcal{B}|\right] + \mathbb{E}\left[|\mathcal{C}|\right] + \mathbb{E}\left[|\mathcal{E}|\right] \leq 2K^2\delta^{-2} + 57K.$

Proof. We upper bound each term:
(a) We show that $\mathbb{E}\left[|\mathcal{B}|\right] < K^2\delta^{-2}$. We decompose $\mathcal{B} = \cup_{k=1}^{K}\mathcal{B}_k$ where $\mathcal{B}_k = \mathcal{B}\cap\{n \geq 1 : \ell_1(n) = k\}$ and bound each $\mathbb{E}\left[|\mathcal{B}_k|\right]$. Note that $|\mathcal{D}(t)| \leq K-1$, thus the length of each block is bounded by K. Now if $n \in \mathcal{B}_k$, $\ell_1(n) = k$ which implies that k is selected once in the block to which n belongs. We conclude that for $n \in \mathcal{B}_k$, $N_k(n) \geq \frac{1}{K}\sum_{t=1}^{n}\mathbb{1}_{\{t\in\mathcal{B}_k\}}$. We can apply Lemma 8 with $H = \mathcal{B}_k$, $\zeta = 1/K$, and conclude that:

$$\mathbb{E}\left[|\mathcal{B}_k|\right] \leq K\delta^{-2}.$$

Summing over k yields the result.

(b) We show that $\mathbb{E}\left[|\mathcal{C}|\right] \leq 57K$. In fact, this follows directly from Lemma 9.

(c) We show that $\mathbb{E}\left[|\mathcal{E}|\right] < K\delta^{-2}$. Let $n \in \mathcal{E}$. Note that $n \notin \mathcal{B}$, thus $|\hat{\nu}_{\ell_1(n)}(n) - \mu_{\ell_1(n)}| < \delta$. Additionally, $n \in \mathcal{A}$, thus $\ell_1(n) > 1$. Therefore,

$$\hat{\nu}_{\ell_1(n)}(n) < \mu_{\ell_1(n)} + \delta \leq \mu_2 + \delta < \frac{\mu_1 + \mu_2}{2}. \tag{14}$$

The last inequality follows from the definition of $\delta < \frac{\mu_1-\mu_2}{2}$. Furthermore, since $n \notin \mathcal{C}$,

$$d_1(n) \geq \mu_1. \tag{15}$$

Combining (14) and (15), we get

$$d_1(n) \geq \mu_1 \geq (\mu_1 + \mu_2)/2 > \hat{\nu}_{\ell_1(n)}(n).$$

From the above inequality, we know that arm 1 belongs to $\mathcal{D}(n)$ for $n \in \mathcal{E}$. By design of DPE2, arm 1 will hence be selected in the block containing round $n \in \mathcal{E}$. Again since each block has a length less than K, we deduce that: for $n \in \mathcal{E}$, $N_1(n) \geq \frac{1}{K}\sum_{t=1}^{n}\mathbb{1}_{\{t\in\mathcal{E}\}}$. Now, applying Lemma 8 with $H = \mathcal{E}$, $\zeta = 1/K$ gives

$$\mathbb{E}\left[|\mathcal{E}|\right] < K\delta^{-2}.$$

$\square$

Lemma 11. *Given* $T \geq 3$, *we define* $\mathcal{G}_k = \{n \leq T, n \notin \mathcal{A}\cup\mathcal{B}, \rho(n) = k\}$ *for some* $k \neq 1$. *We show that*

$$\mathbb{E}\left[|\mathcal{G}_k|\right] \leq \frac{\log T + 4\log(\log T)}{\mathrm{kl}\left(\mu_k + \delta, \mu_1 - \delta\right)} + \delta^{-2}.$$

Proof. Define the counter $c(n) = \sum_{t=1}^{n}\mathbb{1}_{\{t\in\mathcal{G}_k\}}$, which is the number of rounds in $\mathcal{G}_k$ up to round n and define $t_0 = (\log T + 4\log(\log T))/\mathrm{kl}\left(\mu_k + \delta, \mu_1 - \delta\right)$.
Define two subsets of $\mathcal{G}_k$ as

$$\mathcal{G}_{k,1} = \{n \in \mathcal{G}_k : |\hat{\nu}_k(n) - \mu_k| \geq \delta\},$$
$$\mathcal{G}_{k,2} = \{n \in \mathcal{G}_k : c(n) < t_0\}.$$

We first show that $\mathcal{G}_k \subseteq \mathcal{G}_{k,1}\cup\mathcal{G}_{k,2}$. Let $n \in \mathcal{G}_k\backslash(\mathcal{G}_{k,1}\cup\mathcal{G}_{k,2})$. By definition of $\mathcal{A}$, since $n \notin \mathcal{A}$, $\ell_1(n) = 1$.

The arm selection scheme of DPE2 is such that, during the block to which round n belongs, arm k is only selected at round n (indeed note that the algorithm explores at most once per block). Since $n \notin \mathcal{G}_{k,2}$,

$$N_k(n) \geq c(n) \geq t_0. \tag{16}$$

Observe that $\rho(n) = k$ can only happen when

$$d_k(n) \geq \hat{\nu}_{\ell_1(n)}(n) = \hat{\nu}_1(n). \tag{17}$$

Moreover, $n \notin \mathcal{B}$ implies that

$$\hat{\nu}_1(n) > \mu_1 - \delta. \tag{18}$$

Finally $n \notin \mathcal{G}_{k,1}$ and $\delta < \min_{1 \leq k \leq K-1} \frac{\mu_k - \mu_{k+1}}{2}$ imply that

$$\hat{\nu}_k(n) < \mu_k + \delta < \mu_1 - \delta. \tag{19}$$

Combining the above arguments, we get:

$$
\begin{aligned}
t_0 \mathrm{kl}\left(\hat{\nu}_k(n), \mu_1 - \delta\right) &\leq N_k(n)\mathrm{kl}\left(\hat{\nu}_k(n), \mu_1 - \delta\right), \\
&< N_k(n)\mathrm{kl}\left(\hat{\nu}_k(n), d_k(n)\right), \\
&\leq \log T + 4\log(\log T).
\end{aligned}
\tag{20}
$$

The first inequality follows from (16); the second inequality stems from (17)-(18)-(19) and the fact that $y \mapsto \mathrm{kl}\left(x, y\right)$ is an increasing function when $0 < x < y < 1$; the last inequality is obtained by definition of $d_k(n)$. Replacing t_0 by its value in the above inequality, we finally obtain:

$$\mathrm{kl}\left(\hat{\nu}_k(n), \mu_1 - \delta\right) \leq \mathrm{kl}\left(\mu_k + \delta, \mu_1 - \delta\right).$$

Now observe that $x \mapsto \mathrm{kl}\left(x, y\right)$ is a decreasing function when $0 < x < y < 1$. We conclude that $\hat{\nu}_k(n) \geq \mu_k + \delta$ which contradicts the assumption that $n \notin \mathcal{G}_{k,1}$. Hence, $\mathcal{G}_k = \mathcal{G}_{k,1} \cup \mathcal{G}_{k,2}$.

To complete the proof of the lemma, we upper bound $\mathbb{E}\left[|\mathcal{G}_{k,1}|\right]$ and $\mathbb{E}\left[|\mathcal{G}_{k,2}|\right]$.
For $\mathbb{E}\left[|\mathcal{G}_{k,1}|\right]$. Let $n \in \mathcal{G}_{k,1}$. Since we play each arm at most once per block, k is selected for the first time in the block to which n belongs precisely in round n. Hence, we have $\hat{\mu}_k(n) = \hat{\nu}_k(n)$. We apply Lemma 8 with $\zeta = 1$ and $H = \mathcal{G}_{k,1}$, and get $\mathbb{E}\left[|\mathcal{G}_{k,1}|\right] \leq \delta^{-2}$.

For $\mathbb{E}\left[|\mathcal{G}_{k,2}|\right]$. If $n \in \mathcal{G}_{k,2}$, $c(n) \leq t_0$ and $c(n)$ is incremented by $+1$. Therefore,

$$\mathbb{E}\left[|\mathcal{G}_{k,2}|\right] < t_0 = \frac{\log T + 4\log(\log T)}{\mathrm{kl}\left(\mu_k + \delta, \mu_1 - \delta\right)}.$$

We have proved that:

$$\mathbb{E}\left[|\mathcal{G}_k|\right] < \frac{\log T + 4\log(\log T)}{\mathrm{kl}\left(\mu_k + \delta, \mu_1 - \delta\right)} + \delta^{-2}.$$

$\square$

Proof of Theorem 2. First, we prove the regret upper bound. Recall that the initialization phase takes $D + 1$ rounds. In addition, observe that the number of communication phases can be bounded as follows

$$\left|\{t \geq 2 : \mathcal{N}(t) \neq \mathcal{N}(t-1)\}\right| \leq 2\left|\mathcal{A}\right|.$$

Each communication takes at most D rounds. Thus, the regret bound can be bounded as follows:

$$R^\pi(T) \leq M(D+1) + 2DM\mathbb{E}\left[|\mathcal{A}|\right] + M\mathbb{E}\left[|\mathcal{A} \cup \mathcal{B}|\right] + \sum_{k>1}(\mu_1 - \mu_k)\mathbb{E}\left[|\mathcal{G}_k|\right],$$

where $\mathcal{G}_k = \{n \leq T, n \notin \mathcal{A} \cup \mathcal{B}, \rho(n) = k\}$ for $k \neq 1$. Additionally, Lemma 10 gives $\mathbb{E}\left[|\mathcal{B}|\right] + \mathbb{E}\left[|\mathcal{C}|\right] + \mathbb{E}\left[|\mathcal{E}|\right] \leq K(2K\delta^{-2} + 57)$. Thus, in the view of Lemma 11, we have:

$$
\begin{aligned}
R^\pi(T) &\leq M(D+1) + 2MD(2K^2\delta^{-2} + 57K) + M(2K^2\delta^{-2} + 57K) + \sum_{k>1}[\frac{(\mu_1 - \mu_k)(\log T + 4\log\log T)}{\mathrm{kl}(\mu_k + \delta, \mu_1 - \delta)} + (\mu_1 - \mu_k)\delta^{-2}] \\
&\leq 2MD + 6MD(K^2\delta^{-2} + 29K) + K\delta^{-2} + \sum_{k>1}\frac{\mu_1 - \mu_k}{\mathrm{kl}(\mu_k + \delta, \mu_1 - \delta)}(\log T + 4\log\log T) \\
&\leq 9MDK(29 + K\delta^{-2}) + \sum_{k>1}\frac{\mu_1 - \mu_k}{\mathrm{kl}(\mu_k + \delta, \mu_1 - \delta)}(\log T + 4\log\log T).
\end{aligned}
$$

The second inequality follows from $\sum_{k\geq 1} \mu_1 - \mu_k \leq K$ and $M(D+1) \leq 2MD$. This concludes the proof of the regret upper bound.

Next, we prove the upper bound of the total expected number of bits sent under DPE2. The leader election phase requires the transmission of $2|E|(D+1)\log_2(M)$ bits. Note that:

$$2|E|(D+1)\log_2(M) \leq 4M^2 D \log_2(M).$$

Now, each communication phase requires at most $2|E|D\log_2(K)$ transmitted bits. Moreover, as the proof of Theorem 2, the expected number of times $\ell_1(t)$ changes is bounded by $2\mathbb{E}[|\mathcal{A}|]$. Thus, the expected number of bits sent in communication phases is at most

$$4M^2 D \log_2(K)\mathbb{E}[|\mathcal{A}|] \leq 8KM^2 D \log_2(K)(29 + K\delta^{-2}).$$

$\square$

D   Additional numerical experiments

D.1   DPE1

Figures 7 and 8 illustrate the regret in round $T = 5 \cdot 10^5$ and the communication complexity of the various algorithms as a function of Δ, the gap between two consecutive arms. We chose the usual parameters $K = 9, M = 6$. We consider arms with Bernoulli rewards, with range $[0.9 - K \cdot \Delta, 0.9]$ and $\Delta \in [10^{-4}, 10^{-1}]$. All the values were computed over 960 runs.

In Figure 7 we can confirm what already has been observed in Boursier and Perchet (2019), that MCTopM seems to provide better results for larger values of Δ. On the other hand, both SIC-MMAB and DPE1 have smaller dependency in $1/\Delta$, with DPE1 having an even smaller dependency than SIC-MMAB.

Figure 8 shows the average number of communication phases of DPE1 as a function of the gap Δ, in logarithmic scale. We can distinguish a linear behavior in $1/\Delta$, that plateaus for very small values of Δ.


Figure 7: Regret at time $T = 5 \cdot 10^5$ as a function of the gap Δ.

Figure 8: Average number of communication phases at time $T = 5 \cdot 10^5$ as a function of the gap Δ.

D.2   DPE2

In Figure 9, we show the number of times the leader needs to communicate to the other players for DPE2. As suggested by our analysis, this number is small. Values are averaged over 1024 runs, and we used the same settings that are described in Subsection 4.2, i.e., 10 arms with Bernoulli rewards, gap $\Delta = 8/9$ and range $[0.1, 0.9]$.


Figure 9: Number of times the leader updates the best empirical arm.

E   Experimental setup

Hardware and software setup. All experiments were executed on a stationary desktop computer, featuring an Intel Xeon Silver 4110 CPU, 48GB of RAM and a GeForce GTX 1080 graphical card. Ubuntu 18.04 was installed on the computer.

Code and libraries. We set up our experiments using Python 3.7.1. The code from Boursier and Perchet (2019) was used to run SIC-MMAB and MCTopM. We also made use of some Kullback-Leibler diverge functions from the library SMPyBandits Besson (2018). The code is available at https://github.com/rssalessio/dpe.