

Towards Optimal Algorithms for Multi-Player Bandits without Collision Sensing Information

Wei Huang

WEI.HUANG@TELECOM-SUDPARIS.EU

Telecom SudParis, Institut Polytechnique de Paris, France

Richard Combes

RICHARD.COMBES@CENTRALESUPELEC.FR

Université Paris-Saclay, CNRS, CentraleSupélec, Laboratoire des signaux et systèmes, Gif-sur-Yvette, France.

Cindy Trinh

CINDY.TRINH.SRIDYKHAN@GMAIL.COM

Université Paris-Saclay, CNRS, CentraleSupélec, Laboratoire des signaux et systèmes, Gif-sur-Yvette, France.

Editors: Po-Ling Loh and Maxim Raginsky

Abstract

We propose a novel algorithm for multi-player multi-armed bandits without collision sensing information. Our algorithm circumvents two problems shared by all state-of-the-art algorithms: it does not need as an input a lower bound on the minimal expected reward of an arm, and its performance does not scale inversely proportionally to the minimal expected reward. We prove a theoretical regret upper bound to justify these claims. We complement our theoretical results with numerical experiments, showing that the proposed algorithm outperforms state-of-the-art.

Keywords: Multi-Armed Bandits, Distributed Bandits

1. Introduction and Contribution

1.1. Multi-player Bandits

We consider the multi-player multi-armed bandit problem with K arms and $M < K$ players. At each time $t = 1, \dots, T$, where T is the time horizon, each player $m = 1, \dots, M$ selects an arm denoted by $\pi_m(t) \in \{1, \dots, K\}$. If arm $k = 1, \dots, K$ has been selected by strictly more than one player we say that a collision occurs on arm k , and we denote by $\eta_k(t) = \mathbb{1} \left\{ \sum_{m=1}^M \mathbb{1} \{ \pi_m(t) = k \} \geq 2 \right\}$ the corresponding indicator function. The reward received by player $m = 1, \dots, M$ is denoted by $r_m(t) = X_{\pi_m(t)}(t)[1 - \eta_{\pi_m(t)}(t)]$ where $(X_k(t))_{t \geq 1}$ is i.i.d. Bernoulli with mean μ_k for $k = 1, \dots, K$. In other words, if a player m is involved in a collision then he receives a reward of 0, otherwise he receives a binary reward with mean $\mu_{\pi_m(t)}$.

1.2. No Collision Information

We consider the case where each player $m = 1, \dots, M$ only observes its own reward $r_m(t)$ and nothing else. He does not observe the rewards obtained by other players, he does not know which arms have been selected by other players, and he cannot observe the collision indicator $\eta_{\pi_m(t)}$. This case is called "without collision sensing", and is the most arduous because when player m gets a reward of $r_m(t) = 0$, he does not know whether this was caused by a collision so that $\eta_{\pi_m(t)}(t) = 1$ or whether this was caused because the reward of the chosen arm was null i.e. $X_{\pi_m(t)}(t) = 0$. Of course, when a reward of $r_m(t) = 1$ is obtained, he can however be sure that no collision has occurred. Therefore, the problem is fully decentralized, and no communication is allowed between the players.

1.3. Regret Minimization

Recall that μ_1, \dots, μ_K denotes the mean reward of each arm where no collision has occurred, and define $\mu_{(1)}, \dots, \mu_{(K)}$ its version sorted in descending order, so that $\mu_{(1)} \geq \mu_{(2)} \geq \dots \geq \mu_{(K)}$. The goal for the players is to maximize the expected sum of rewards, and to do so, each player must select a distinct arm amongst the M best arms to both avoid any collisions and maximize total reward. The regret (sometimes called pseudo-regret) is defined as the difference between the total expected reward obtained by an oracle that plays the optimal decision and the reward obtained by the players

$$R(T, \mu) = \sum_{t=1}^T \sum_{m=1}^M \mu_{(m)} - \mathbb{E}(r_m(t))$$

Our goal is to design distributed algorithms in order to minimize the regret.

1.4. Related Work

Single-player multi-armed bandits were originally proposed by [Thompson \(1933\)](#) and asymptotically optimal strategies proposed by [Lai and Robbins \(1985\)](#), [Cappé et al. \(2013\)](#) amongst others. Building on the work of [Lai and Robbins \(1985\)](#), [Anantharam et al. \(1987\)](#) considered a generalization to the case where a single player can choose several arms, and proposed asymptotically optimal strategies. In fact, the problem considered by [Anantharam et al. \(1987\)](#) can be seen as a centralized version of the distributed problem we consider here. Motivated by wireless communication and networking, more recent works have considered multi-player multi-armed bandits, [Liu and Zhao \(2010\)](#), [Rosenski et al. \(2016\)](#), for instance as a natural model for a scenario where several players must access a wireless channel in a decentralized manner [Besson and Kaufmann \(2018\)](#). When collision sensing is available, several authors [Boursier and Perchet \(2019\)](#), [Wang et al. \(2020\)](#) have shown that one can obtain roughly the same performance as in the centralized case by cleverly exploiting collisions. A harder problem is the case without collision sensing. The state-of-the-art known algorithms for this problem are SIC-MMAB2 [Boursier and Perchet \(2019\)](#), EC-SIC [Shi et al. \(2020\)](#) and the algorithms of [Lugosi and Mehrabian \(2018\)](#). Our work improves on these algorithms, both in terms of performance, and in terms of practical applicability. Some generalizations and alternative settings were also considered. [Bistriz and Leshem \(2018\)](#) and [Mehrabian et al. \(2020\)](#) consider the more general heterogenous case where the mean of arms vary among players. [Magesh and Veeravalli \(2021\)](#) and [Shi and Shen \(2021\)](#) consider a variant of the problem where collisions can yield a non-zero reward. [Bubeck et al. \(2021\)](#) consider the additional assumption that players have access to shared randomness, and study the minimax regret.

1.5. Our Contribution

The regret upper bounds for state-of-the-art algorithms and the prior information they require is presented in Table 1 ⁽¹⁾. We also present a regret lower bound which corresponds to the case with collision sensing, although we do not know if this bound is attainable without collision sensing. These algorithms suffer from two important issues: (i) They require a lower bound on the expected reward of the worst arm $\mu_{(K)}$, and cannot be used without this prior information. Since in practice this information is usually not available, this limits their practical applicability. In fact EC-SIC also

1. where E in the EC-SIC regret bound is the error exponent of the error correcting code used.

Algorithm	Regret Upper Bound
SIC-MMAB2	$O\left(M \sum_{k>M} \frac{\ln T}{\mu_{(M)} - \mu_{(k)}} + \frac{K^2 M}{\mu_{(K)}} \ln T\right)$
EC-SIC	$O\left(\sum_{k>M} \frac{\ln T}{\mu_{(M)} - \mu_{(k)}} + \frac{KM}{\mu_{(K)}} \ln T + \frac{KM^2}{E(\mu_{(K)})} \ln\left(\frac{1}{\mu_{(M)} - \mu_{(M+1)}}\right) \ln T\right)$
Lugosi et al.	$O\left(\frac{KM}{(\mu_{(M)} - \mu_{(M+1)})^2} \ln T\right)$ (first algorithm)
Lugosi et al.	$O\left(\frac{KM}{\mu_{(M)} - \mu_{(M+1)}} \ln T + \frac{K^2 M}{\mu_{(M)}} (\ln T)^2\right)$ (second algorithm)
Our work	$O\left(\sum_{k>M} \frac{\ln T}{\mu_{(M)} - \mu_{(k)}} + K^2 M \ln T + KM^2 \ln\left(\frac{1}{\mu_{(M)} - \mu_{(M+1)}}\right)^2 \ln T\right)$
Lower bound	$\Omega\left(\sum_{k>M} \frac{\ln T}{\mu_{(M)} - \mu_{(k)}}\right)$

Algorithm	Prior Information
SIC-MMAB2	$\mu_{(K)}$
EC-SIC	$\mu_{(K)}$ and $\mu_{(M)} - \mu_{(M+1)}$
Lugosi et al.	$\mu_{(M)}$ and M
Our work	\emptyset

Table 1: State-of-the-art without collision sensing: regret upper bounds and prior information

requires as prior information the gap $\mu_{(M)} - \mu_{(M+1)}$ which is also typically unknown (ii) Their regret is proportional to $O(1/\mu_{(K)})$ which can be arbitrarily large. Namely $\mu_{(K)}$ may be very small, for instance, in wireless applications where arms represent channels, where μ_k is an increasing function of the signal-to-noise ratio on channel k , and where μ_1, \dots, μ_K are weakly correlated due to frequency selective fading, this issue occurs. Indeed, if we assume that μ_1, \dots, μ_K are drawn independently from some distribution over $[0, 1]$, then $\mu_{(K)} = \min_{k=1, \dots, K} \mu_k$ is close to 0 with overwhelming probability when K is large.

Our main contribution is to propose an algorithm correcting issues (i) and (ii). The proposed algorithm requires no input information besides the number of arms K and the time horizon T and its regret is upper bounded as:

$$R(T, \mu) = O\left(\sum_{k>M} \frac{\ln T}{\mu_{(M)} - \mu_{(k)}} + K^2 M \ln T + KM^2 \ln\left(\frac{1}{\mu_{(M)} - \mu_{(M+1)}}\right)^2 \ln T\right)$$

which does not depend on $\mu_{(K)}$. We also present numerical experiments which confirm our theoretical predictions. We believe that our work is a step towards closing the gap between the best known upper bound and lower bound, by showing that, just like in the case with collision sensing, the regret need not depend on $\mu_{(K)}$, and that is it not necessary to have prior information about $\mu_{(K)}$ either.

2. Algorithm and Analysis

The proposed algorithm can be broken down into three steps and corresponding subroutines. Of course, since the proposed algorithm is distributed, all players $m = 1, \dots, M$ apply the same algorithm based on their own observations. When providing pseudo-code, it is understood that each player

follows the same pseudo code, in a distributed fashion. When stating our algorithms we use the notation $X += x$ to denote $X \leftarrow X + x$, i.e. add x to X . All additional subroutines as well as proofs are presented in the appendix.

2.1. Step 1: Agreement on a Good Arm

The first step of the proposed algorithm is the FindGoodArm subroutine. The goal of this subroutine is for all players $m = 1, \dots, M$ to agree on a "good" arm \tilde{k} , in the sense that $\mu_{\tilde{k}} \geq (1/8)\mu_{(1)}$, as well as to output a lower bound on $\mu_{\tilde{k}}$.

This subroutine is the cornerstone of our algorithm, and the main reason why our algorithm does not suffer from (i) a regret scaling inversely proportionally to the reward of the first arm $O(1/\mu_{(K)})$ (ii) requiring a non trivial lower bound of $\mu_{(K)}$ as an input parameter, while all state-of-the-art algorithms suffer from problems (i) and (ii). Indeed, once a good arm is identified, one can use this arm to exchange information between players, and arms with very low rewards have no bearing on performance and can be mostly ignored.

The FindGoodArm subroutine is designed such that, with high probability, not only will all players agree on the same good arm, but also that they will all terminate the procedure at the same time. This is important to maintain synchronization, and achieving this synchronization is not trivial as shown next.

The FindGoodArm subroutine is broken in successive phases, and each phase $p \geq 1$ has two sub-phases. In the first sub-phase of phase p , each player selects an arm uniformly at random $6K2^p \ln \frac{2}{\delta}$ times and observes the resulting rewards. For each arm $k = 1, \dots, K$, each player computes $\hat{\mu}_k^p$ which is the ratio between the sum of rewards obtained from arm k and the number of times that arm k has been selected. If $\hat{\mu}_k^p \geq 2^{1-p}$, we say that the player accepts arm k at phase p (i.e. he believes that μ_k is greater than 2^{-p}) and otherwise he rejects arm k at phase p . In the second sub-phase, the players attempt to reach agreement on one good arm. For all arms $\ell = 1, \dots, K$, if a player accepts arm ℓ at phase p , he chooses an arm uniformly at random $K2^p \ln \frac{2}{\delta}$ times and if he does not accept ℓ then he chooses arm ℓ for $K2^p \ln \frac{2}{\delta}$ times. If arm ℓ yields at least one non-zero reward then we say that the player confirms arm ℓ in phase p . If there exists \tilde{k} such that the player confirms \tilde{k} , then the procedure stops and \tilde{k} is output. This mechanism is designed so that, with high probability, the procedure does not exit until all players have agreed on a good arm. Indeed, if at least one of the players rejects arm ℓ , then he will select arm ℓ all the time, and no non-zero reward can be obtained from ℓ , preventing other players from confirming arm ℓ .

To summarize the FindGoodArm subroutine involves sampling arms uniformly at random more and more, until at least one of them has a large enough reward, and periodically exchanging information between players in order to reach agreement over one good arm, and once agreement is reached, the procedure stops. Lemma 1 provides an analysis of FindGoodArm and Algorithm 1 provides its pseudo-code.

Lemma 1 *Consider the FindGoodArm subroutine with input parameters K and δ . Then, with probability greater than $1 - C_1KM(\ln \frac{1}{\mu_{(1)}})\delta$: (i) The subroutine lasts at most $C_2 \frac{K^2}{\mu_{(1)}} \ln \frac{1}{\delta}$ time slots causing at most $C_2K^2M \ln \frac{1}{\delta}$ regret. (ii) The output \tilde{k} is a "good" arm i.e. $\mu_{\tilde{k}} \geq \frac{1}{8}\mu_{(1)}$ (iii) The output $\tilde{\mu}$ is a lower bound on the expected reward of \tilde{k} i.e. $\mu_{\tilde{k}} \geq \tilde{\mu}$ (vi) All players exit the subroutine at exactly the same time and output the same arm \tilde{k} and estimate $\tilde{\mu}$, where C_1, C_2 are universal constants.*

Algorithm 1 FindGoodArm (for player $m = 1, \dots, M$)

Input: K : number of arms, δ : confidence parameter**Output:** \tilde{k} : good arm, $\tilde{\mu}$: lower bound on reward of \tilde{k} $p \leftarrow 0, \tilde{k} \leftarrow -1$ # initialization**while** $\tilde{k} = -1$ **do** $p += 1, R[k], N[k] \leftarrow 0$ for $k = 1, \dots, K$ # current phase, rewards and number of samples

sub-phase 1: explore arms uniformly at random

for $t = 1, \dots, 6K2^p \ln \frac{2}{\delta}$ **do**Select arm $k \in \{1, \dots, K\}$ uniformly at random, observe reward $r, R[k] += r, N[k] += 1$ **end for**

sub-phase 2: confirm accepted arms

for $\ell \leftarrow 1, \dots, K$ **do** $R'[\ell] \leftarrow 0$ for $k = 1, \dots, K$ # rewards of samples# if arm ℓ was accepted sample arms uniformly**if** $\frac{R[\ell]}{N[\ell]} \geq 2^{1-p}$ **then****for** $t = 1, \dots, 2^p K \ln \frac{2}{\delta}$ **do**Select arm $k \in \{1, \dots, K\}$ uniformly at random and observe reward $r, R'[k] += r$ **end for**# if a non-zero reward is obtained, confirm arm ℓ **if** $R'[\ell] \geq 1$ **then** $\tilde{k} \leftarrow \ell, \tilde{\mu} \leftarrow 2^{-p}$ **break end if**# if arm ℓ was rejected only sample ℓ **else****for** $t = 1, \dots, 2^p K \ln \frac{2}{\delta}$ **do** select arm $k = \ell$, observe reward $r, R'[k] += r$ **end for****end if****end for****end while**

2.2. Step 2: Rank Assignment

The second part of our algorithm is the subroutine `VirtualMusicalChairs` which performs rank assignment, that is, assigning a different number to each of the players. This allows to perform orthogonalization and explore arms without collisions by using sequential hopping as done in [Boursier and Perchet \(2019\)](#).

Our subroutine `VirtualMusicalChairs` is based on the so-called Musical Chairs algorithm where players repeatedly choose an arm at random until they obtain a non-zero reward and then keep sampling this very same arm, which will also be their external rank. However, the Musical Chairs used in [Boursier and Perchet \(2019\)](#) is not satisfactory since it samples from all arms $k = 1, \dots, K$. Its performance is therefore limited by the worst arm leading to a regret which scales proportionally to $O(1/\mu_{(K)})$, in addition to requiring a lower bound on $\mu_{(K)}$ as an input parameter.

Instead we propose to adapt this algorithm to only sample from the unique good arm \tilde{k} derived by the `FindGoodArm` subroutine. To do so we map this unique good arm into K "virtual" arms by using a time division technique depicted in [Figure 1](#). We divide time in τ blocks of length K . The ℓ -th time slot of each block is mapped to the ℓ -th virtual arm. At the beginning of a block, player j chooses a virtual arm ℓ at random (without pulling it). At the ℓ -th time slot of the block, he pulls




Figure 1: Mapping between a single real arm (below) and K identical virtual arms (above).

the good arm \tilde{k} , and if the reward is positive, he settles on virtual arm ℓ until the end. Note that a player actually selects one arm every K time steps on average. Note also that, due to this idea of virtual arms, `VirtualMusicalChairs` can be analyzed simply by considering the classical musical chairs algorithm applied to K identical arms with mean reward $\mu_{\tilde{k}}$. Lemma 2 provides an analysis of `VirtualMusicalChairs` and Algorithm 2 provides its pseudo-code.

Lemma 2 Consider $\delta > 0$. Consider the subroutine `VirtualMusicalChairs` with input parameters K , a good arm \tilde{k} and $\tau = \frac{K}{\mu_{\tilde{k}}} \ln \frac{1}{\delta}$. Then the subroutine exits after $K\tau$ time slots, incurring regret at most $8K^2M \ln \frac{1}{\delta}$ and with probability greater than $1 - C_3KM\delta$ it assigns a distinct external rank to all of the M players where C_3 is a universal constant.

2.3. Step 3: Computation of the Number of Players

The third step of the proposed algorithm is the `VirtualNumberPlayers` subroutine which allows players to estimate the number of players M . `VirtualNumberPlayers` takes as input K the number of arms, \tilde{k} a good arm, s the external rank of a player, τ a sampling time, all of which are available from the previous subroutines.

Once again we combine the sequential hopping technique used by Boursier and Perchet (2019) to estimate M along with the idea of virtual arms, so that the procedure only samples from a good arm \tilde{k} , and hence is not impacted by the presence of some arms with low reward.

To estimate M , we initialize our estimate as $\hat{M} = 1$, and perform sequential hopping for $2K$ rounds to make sure that every two different players only collide one time. Since players can be orthogonalized thanks to their external rank, a player increments \hat{M} only if he gets all zero rewards during a specific period (this suggests that some collisions have occurred). As a by-product, we also output the internal rank $j \in \{1, \dots, M\}$, not to be confused with the external rank $s \in \{1, \dots, K\}$. Each player is assigned a different internal rank, and it will serve for the last step of our algorithm, in order to assign roles when exploring the various arms.

Lemma 3 Consider $\delta > 0$. Consider the subroutine `VirtualNumberPlayers` with input parameters K , a good arm \tilde{k} and $\tau = \frac{1}{\mu_{\tilde{k}}} \ln \frac{1}{\delta}$. Then the subroutine exits after $K^2\tau$ time slots causing at most $C_4K^2M \ln \frac{1}{\delta}$ regret, and with probability greater than $1 - C_5\delta$, it both outputs a correct

Algorithm 2 VirtualMusicalChairs (for player $m = 1, \dots, M$)

Input: K : number of arms, \tilde{k} : a good arm, τ : sampling times
Output: s : external rank of the player
 $s \leftarrow -1$; # rank of the player is initially unset
 # musical chairs on the arm \tilde{k}
for $t \leftarrow 1, \dots, K\tau$ **do**
 # time is split in blocks of size K and we select when to sample at the start of a block.
if $\text{mod}(t, K) = 1$ **then**
if $s = -1$ **then**
 Draw $\ell \in \{1, \dots, K\}$ uniformly at random # Choose a random slot if rank is unset
else
 $\ell \leftarrow s$ # Choose the rank as a slot if it is set
end if
end if
 # sample the corresponding time slot
if $\text{mod}(t, K) = \ell$ **then**
 Select arm \tilde{k} , and observe reward r
 # set rank if it was not set yet and a non zero reward was obtained
if $r > 0$ and $s = -1$ **then** $s \leftarrow \ell$ **end if**
else
 Select an arbitrary arm in $\{1, \dots, K\} \setminus \{\tilde{k}\}$
end if
end for

estimate of the number of players $\hat{M} = M$, and it assigns a distinct internal rank $j \in \{1, \dots, M\}$ to all of the M players. Both C_4 and C_5 are universal constants.

Algorithm 3 VirtualNumberPlayers (for player $m = 1, \dots, M$)

Input: K : number of arms, \tilde{k} a good arm, s : external rank of a player, τ : sampling times

Output: \hat{M} : estimated number of players, j : internal rank of the player

$\hat{M} \leftarrow 1, \ell \leftarrow s, j \leftarrow 1$ # initialization

for $n = 1, \dots, 2K$ **do**

if $n > 2s$ **then** $\ell \leftarrow \text{mod}(\ell + 1, K)$ # sequential hopping **end if**

$R \leftarrow 0$ # sum of rewards from the good arm

 # sample from the good arm

for $k = 1, \dots, K$ **do**

if $\ell \neq k$ **then**

for $t = 1, \dots, \tau$ **do** Select an arbitrary arm in $\{1, \dots, K\} \setminus \{\tilde{k}\}$ **end for**

else

 # sample from virtual arm τ times

for $t = 1, \dots, \tau$ **do** Select arm \tilde{k} , observe reward r , $R += r$ **end for**

 # if no non-zero reward was obtained increase the estimated number of players

if $R = 0$ **then**

$\hat{M} += 1$,

if $n \leq 2s$ **then** $j += 1$ **end if**

end if

end if

end for

end for

2.4. Step 4: Distributed Exploration and Finding the Best Arms

The last step is the subroutine DistributedExploration, which takes as input K the number of arms, j the internal rank of the player, τ the sampling times, \hat{M} the estimated number of players, and outputs one arm amongst the M best arms, which is assigned to the player.

DistributedExploration is similar to the exploration strategy used in Shi et al. (2020) with two key improvements. First, to avoid choosing arms with low expected reward, a player sends data only through arm \tilde{k} rather than through the j -th arm as done in Shi et al. (2020). Second, we set the quantization error to be dependent on phase p instead of the fixed quantity $\mu_M - \mu_{M+1}$. By doing so, the number of pulls is reduced, and the subroutine does not require this prior knowledge anymore. Its analysis, provided in lemma 4 follows along similar lines as Shi et al. (2020), so we simply highlight the main steps of this procedure. The player with internal rank $j = 1$ is called the leader, and other players are called the followers. The procedure operates in phases, and in phase p , the active players (the players that have not been assigned an arm yet) sample each active arm (the arms for which players have not figured out whether or not they are amongst the M best arms) $2^p \lceil \ln \frac{1}{\delta} \rceil$ times using sequential hopping (so that no collisions occur). Then all followers send their estimates to the leader, the leader aggregates their estimates, determines which arms can be accepted (the arms for which one is sure that they are amongst the M best arms) and rejected (the arms for which one is sure that

they are not amongst the M best arms) and then sends the sets of accepted and rejected arms back to the followers. If an arm is accepted then it is assigned to an active player, and both the arm and the player become inactive, and the player simply selects the assigned arm until the end of the procedure. If an arm is rejected then it becomes inactive. The strategy on how leader and followers act in this information exchange step is given in subroutines `ComFollow` and `ComLead` in appendix. In essence, the leader coordinates and takes all of the decisions, while followers simply collect samples, transmit their estimates and receive orders on what to do next. The strategy to send and receive data between leader and followers is given in subroutines `EncoderSendFloat` and `DecoderReceiveFloat` in appendix.

Lemma 4 *Consider $\delta > 0$ and $T \geq 0$. Consider the subroutine `DistributedExploration` with input parameters K , \bar{k} and $\tau = \frac{1}{\mu_{\bar{k}}} \ln \frac{1}{\delta}$ applied with time horizon T ⁽²⁾. Then with probability greater than $1 - C_6 \delta K M (\ln T)^2$ the subroutine incurs a regret upper bounded by*

$$C_7 \sum_{k>M} \frac{\ln \frac{1}{\delta}}{\mu_{(M)} - \mu_{(k)}} + C_8 K M^2 \ln \left(\frac{1}{\mu_{(M)} - \mu_{(M+1)}} \right)^2 \ln \frac{1}{\delta}$$

and assigns one distinct arm to each player from the set of the M best arms, where C_6, C_7, C_8 are universal constants.

2.5. Putting it all Together

The complete proposed algorithm is presented in algorithm 5, and combines the four steps above. It is noted that our proposed algorithm is in fact a procedure to identify the set of best arms. To turn it into an algorithm for minimizing regret over a time horizon of T , one can simply run the proposed algorithm with confidence parameter $\delta = \frac{1}{T}$, and then let players select the arm that has been assigned to them by the algorithm until the time horizon runs out. The regret upper bound of Theorem 5 is our main result, and is simply proven by combining the four previous lemmas. As promised, unlike state-of-the-art algorithms, our algorithm does not need any prior information such as the number of arms K or the reward of the worst arm $\mu_{(K)}$, and its regret does not depend on $\mu_{(K)}$. The regret bound is better than that of SIC-MMAB2 and EC-SIC, in the sense that the term proportional to $1/\mu_{(K)}$ (which can be arbitrarily large), has been eliminated. In other words, performance is not limited by the worst arm anymore. This causes a dramatic performance gain which is seen in numerical experiments shown below.

Theorem 5 *Consider $T \geq 0$, $\mu_M - \mu_{M+1} > 0$. First apply the proposed algorithm with input parameters K and $\delta = \frac{1}{T(\ln T)}$ with time horizon T , let \bar{k} denote its output, and select arm \bar{k} for the remaining time steps. Then the expected regret of this procedure is*

$$R(T) \leq C_9 \sum_{k>M} \frac{\ln T}{\mu_{(M)} - \mu_{(k)}} + C_{10} K^2 M \ln T + C_{11} K M^2 \ln \left(\frac{1}{\mu_{(M)} - \mu_{(M+1)}} \right)^2 \ln T$$

with C_9, C_{10}, C_{11} three universal constants.

2. Applying a procedure with time horizon T simply means that, if it does not terminate before T on its own, then its execution is immediately stopped when time T is reached.

Algorithm 4 DistributedExploration (for players $m = 1, \dots, M$)

Input: K : number of arms, j : internal rank of a player, \hat{M} : number of players, \tilde{k} : a good arm, τ : sampling times

Output: f an arm amongst the M best arms assigned to the player

Initialize $p \leftarrow 0$; $f \leftarrow -1$; # initialization

$R[k], v[k] \leftarrow 0$ for $k = 1, \dots, K$ # rewards and number of samples for each arm

rewards and number of samples for each arm held by each players, only stored by the leader

if $j = 1$ **then**

for $m = 1, \dots, \hat{M}$ and $k = 1, \dots, K$ **do** $\hat{\mu}[k, m], N[k, m] \leftarrow 0$ **end for**

end if

$M' \leftarrow \hat{M}$, $\mathcal{K} \leftarrow \{1, \dots, K\}$ # number of active players and set of active arms

while $f = -1$ **do**

$p += 1$ # start phase p

$k \leftarrow j$ # first sub-phase explore arms by sequential hopping

for $t \leftarrow 1, \dots, |\mathcal{K}|2^p \lceil \ln \frac{1}{\delta} \rceil$ **do**

$k \leftarrow (k + 1) \bmod |\mathcal{K}|$

 Select arm k , observe reward r , $R[k] += r$, $v[k] += 1$, $E[k] \leftarrow \frac{R[k]}{v[k]}$

end for

$Q \leftarrow \lceil \frac{p}{2} + 3 \rceil$ # second sub-phase: share estimates between players

if $j = 1$ **then**

$(f, \mathcal{K}, M', \hat{\mu}, N) \leftarrow \text{ComLeader}(\hat{\mu}, N, \mathcal{K}, M', Q, \tau, \tilde{k}, p, \delta)$ # player is a leader

else

$(f, \mathcal{K}, M') \leftarrow \text{ComFollow}(E, j, \mathcal{K}, M', Q, \tau, \tilde{k})$ # player is a follower

end if

end while

Algorithm 5 Proposed algorithm (for player $m = 1, \dots, M$)

Input: K : number of arms, δ : confidence level

Output: \bar{k} an arm amongst the M best arms assigned to the player

$(\tilde{k}, \tilde{\mu}) \leftarrow \text{FindGoodArm}(K, \delta)$ # find a good arm and a lower bound on its reward

$s \leftarrow \text{VirtualMusicalChairs}(K, \tilde{k}, K \ln(\frac{1}{\delta}) / \tilde{\mu})$ # assign external rank to each player

$(\hat{M}, j) \leftarrow \text{VirtualNumberPlayers}(K, \tilde{k}, s, \ln(\frac{1}{\delta}) / \tilde{\mu})$ # estimate the number of players and assign internal rank




$\bar{k} \leftarrow \text{DistributedExploration}(K, j, \hat{M}, \tilde{k}, \ln(\frac{1}{\delta}) / \tilde{\mu})$ # find one arms out of the M best arms

3. Numerical Experiments

We now compare the empirical performance of the proposed algorithm against the state-of-the-art algorithms SIC-MMAB2 and EC-SIC. For simplicity we assume that rewards decrease from best to worst in a linear fashion $\mu_{(k)} = \mu_{(1)} + \frac{k-1}{K-1}(\mu_{(K)} - \mu_{(1)})$. The regret of algorithms is averaged over 20 (or more) independent runs, and 95% confidence intervals are presented.

Influence of the number of players In our first set of experiments, we consider $\mu_{(1)} = 1$ and $\mu_{(K)} = 0.01$, $M = \lfloor K/2 \rfloor$. We plot the expected regret of the various algorithms for $K = 5, 10, 20$ in figures 2,3 and 4 respectively. Overall, the proposed algorithm clearly outperforms EC-SIC and SIC-MMAB2, sometimes by several orders of magnitude, and the difference seems more and more severe when K increases.

Influence of the gap Figure 2 and 3 consider $\mu_M - \mu_{M+1} = 0.2$ and $\mu_M - \mu_{M+1} = 0.1$. We see that sometimes the proposed algorithm performs a bit worse than other algorithms for very small time horizons, however for larger time horizons it does outperform SIC-MMAB2 and EC-SIC which seem not to converge quickly enough.

Figure 2: $K=5, M=2, T=10^5$ Figure 3: $K=10, M=5, T=10^6$ Figure 4: $K=20, M=10, T=10^7$

Impact of the number of players In figures 5 and 6 we compare between $M = 2$ players and $M = 8$ players. For $M = 2$, both the proposed algorithm and EC-SIC do converge. However, when M increases, the collision probability becomes larger, causing EC-SIC to spend significant time performing musical chairs to estimate M .




Figure 5: $K=10, M=2, T=10^6$




Figure 6: $K=10, M=8, T=10^6$




Figure 7: $K=5, M=2, T=10^5$
with different $\mu(K)$

Impact of the worst arm We present the regret for different values of $\mu(K)$ in figure 7. When $\mu(K)$ decreases, we can see that the performance of SIC-MMAB2 and EC-SIC is greatly affected, and for $\mu(K) = 0.001$ they do not even start the exploration phase. On the other hand, no matter how small $\mu(K)$ is, the cumulative regret of the proposed algorithm stays similar, which confirms our theoretical predictions.

4. Conclusion

In this work we have proposed a new algorithm for multi-player multi-armed bandits without collision sensing information. Through both analysis and numerical experiments, we have proven that it performs significantly better than the state-of-the-art algorithms, while at the same time requiring no input parameter other than the number of arms. The cornerstone of our approach is a novel, non trivial distributed procedure to enable players to discover a good arm without any prior information other than the number of arms. We believe that this result is a significant contribution to the problem as it constitutes a step towards being able to solve multi-player multi-armed bandits in a practical setting (for instance cognitive radios) where no prior information is known about the expected rewards of the arms.

Acknowledgments

The work of Wei Huang was supported by Beyond 5G, a project of the French Government’s recovery plan “France Relance”.

References

- Venkatachalam Anantharam, Pravin Varaiya, and Jean Walrand. Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays-part i: iid rewards. *IEEE Transactions on Automatic Control*, 32(11):968–976, 1987.
- Lilian Besson and Emilie Kaufmann. Multi-player bandits revisited. In *Proc. of AISTATS*, 2018.
- Itai Bistriz and Amir Leshem. Distributed multi-player bandits - a game of thrones approach. In *Proc. of NeurIPS*, 2018.
- Etienne Boursier and Vianney Perchet. SIC-MMAB: Synchronisation involves communication in multiplayer multi-armed bandits. In *Proc. of NeurIPS*, 2019.
- Sébastien Bubeck, Thomas Budzinski, and Mark Sellke. Cooperative and stochastic multi-player multi-armed bandit: Optimal regret with neither communication nor collisions. In *Proc. of COLT*, 2021.
- Olivier Cappé, Aurélien Garivier, Odalric-Ambrym Maillard, Rémi Munos, and Gilles Stoltz. Kullback–leibler upper confidence bounds for optimal sequential allocation. *The Annals of Statistics*, 41(3):1516–1541, Jun 2013.
- Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- Keqin Liu and Qing Zhao. Distributed learning in multi-armed bandit with multiple players. *IEEE Transactions on Signal Processing*, 58(11):5667–5681, 2010.
- Gabor Lugosi and Abbas Mehrabian. Multiplayer bandits without observing collision information. arXiv 1808.08416, 2018.
- Akshayaa Magesh and Venugopal V. Veeravalli. Decentralized heterogeneous multi-player multi-armed bandits with non-zero rewards on collisions. *IEEE Transactions on Information Theory*, 2021.
- Abbas Mehrabian, Etienne Boursier, Emilie Kaufmann, and Vianney Perchet. A practical algorithm for multiplayer bandits when arm means vary among players. In *Proc. of AISTATS*, 2020.
- Jonathan Rosenski, Ohad Shamir, and Liran Szlak. Multi-player bandits—a musical chairs approach. In *Proc. of ICML*, 2016.
- Chengshuai Shi and Cong Shen. Multi-player multi-armed bandits with collision-dependent reward distributions. *IEEE Transactions on Signal Processing*, 69:4385–4402, 2021.
- Chengshuai Shi, Wei Xiong, Cong Shen, and Jing Yang. Decentralized multi-player multi-armed bandits with no collision information. In *Proc. of AISTATS*, 2020.
- William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3-4):285–294, 12 1933.
- Po-An Wang, Alexandre Proutiere, Kaito Ariu, Yassir Jedra, and Alessio Russo. Optimal algorithms for multiplayer multi-armed bandits. In *Proc. of AISTATS*, 2020.

Appendix A. Proofs

A.1. Technical results

We recall some well known concentration inequalities used for our analysis.

Proposition 6 (Multiplicative Chernoff Bound) Consider X_1, \dots, X_t i.i.d. random variables in $[0, 1]$ with expectation μ , then for all $\epsilon \in [0, 1]$

$$\mathbb{P}\left(\frac{1}{t} \sum_{s=1}^t X_s \leq (1 - \epsilon)\mu\right) \leq \exp\left(-\frac{\epsilon^2 t \mu}{2}\right) \text{ and } \mathbb{P}\left(\frac{1}{t} \sum_{s=1}^t X_s \geq (1 + \epsilon)\mu\right) \leq \exp\left(-\frac{\epsilon^2 t \mu}{2 + \epsilon}\right)$$

Proposition 7 (Hoeffding's Inequality) Consider X_1, \dots, X_t i.i.d. random variables in $[0, 1]$ with expectation μ , then for all $\epsilon \in [0, 1]$

$$\mathbb{P}\left(\frac{1}{t} \sum_{s=1}^t X_s \leq \mu - \epsilon\right) \leq \exp(-2t\epsilon^2) \text{ and } \mathbb{P}\left(\frac{1}{t} \sum_{s=1}^t X_s \geq \mu + \epsilon\right) \leq \exp(-2t\epsilon^2)$$

A.2. Proof of Lemma 1

When all players sample an arm uniformly at random, the probability that a player m does not experience a collision is $(1 - \frac{1}{K})^{M-1}$, and the expected reward he obtains when selecting arm k is therefore $\rho_k = (1 - \frac{1}{K})^{M-1} \mu_k$.

For phase $p \geq 1$, player $m = 1, \dots, M$ and arm k we denote by $\hat{\mu}_{k,m}^p$ the empirical mean reward received by player m from arm k during phase p , denote by $t_{k,m}^p$ the corresponding number of samples.

Acceptance in the first sub-phase Assume that $\rho_k \geq 2^{2-p}$ then:

$$\mathbb{P}(\hat{\mu}_{k,m}^p \leq 2^{1-p}) \leq \mathbb{P}(\hat{\mu}_{k,m}^p \leq 2^{1-p}, t_{k,m}^p \geq 3 \times 2^p \ln \frac{2}{\delta}) + \mathbb{P}(t_{k,m}^p \leq 3 \times 2^p \ln \frac{2}{\delta})$$

Since $\hat{\mu}_{k,m}^p$ is an empirical mean of $t_{k,m}^p$ i.i.d. random variables in $[0, 1]$ with expectation $\rho_k \geq 2^{2-p}$ a Chernoff bound yields:

$$\mathbb{P}(\hat{\mu}_{k,m}^p \leq 2^{1-p}, t_{k,m}^p \geq 3 \times 2^p \ln \frac{2}{\delta}) \leq \exp(-\frac{4}{3} \ln \frac{2}{\delta}) \leq \frac{\delta}{2}$$

Since $t_{k,m}^p$ is an empirical mean of $6K2^p \ln \frac{2}{\delta}$ i.i.d. random variables in $[0, 1]$ with expectation $\frac{1}{K}$ a Chernoff bound yields:

$$\mathbb{P}(t_{k,m}^p \leq 3 \times 2^p \ln \frac{2}{\delta}) \leq \exp(-3 \times 2^{p-2} \ln \frac{2}{\delta}) \leq \frac{\delta}{2}$$

We have proven that if $\rho_k \geq 2^{2-p}$ then $\mathbb{P}(\hat{\mu}_{k,m}^p \leq 2^{1-p}) \leq \delta$, so that arm k will be accepted at phase p by player m with probability greater than $1 - \delta$.

Rejection in the first sub-phase Assume that $\rho_k \leq 2^{-p}$ then once again:

$$\mathbb{P}(\hat{\mu}_{k,m}^p \geq 2^{1-p}) \leq \mathbb{P}(\hat{\mu}_{k,m}^p \geq 2^{1-p}, t_{k,m}^p \geq 3 \times 2^p \ln \frac{2}{\delta}) + \mathbb{P}(t_{k,m}^p \leq 3 \times 2^p \ln \frac{2}{\delta})$$

Since $\hat{\mu}_{k,m}^p$ is an empirical mean of $t_{k,m}^p$ i.i.d. random variables in $[0, 1]$ with expectation $\rho_k \leq 2^{-p}$ a Chernoff bound yields:

$$\mathbb{P}(\hat{\mu}_{k,m}^p \geq 2^{1-p}, t_{k,m}^p \geq 3 \times 2^p \ln \frac{2}{\delta}) \leq \exp(-\ln \frac{2}{\delta}) = \frac{\delta}{2}$$

Since $t_{k,m}^p$ is an empirical mean of $6K2^p \ln \frac{2}{\delta}$ i.i.d. random variables in $[0, 1]$ with expectation $\frac{1}{K}$ a Chernoff bound yields:

$$\mathbb{P}(t_{k,m}^p \leq 3 \times 2^p \ln \frac{2}{\delta}) \leq \exp(-3 \times 2^{p-2} \ln \frac{2}{\delta}) \leq \frac{\delta}{2}$$

We have proven that if $\rho_k \leq 2^{-p}$ then $\mathbb{P}(\hat{\mu}_{k,m}^p \geq 2^{1-p}) \leq \delta$, so that arm k will be rejected by player m at phase p with probability greater than $1 - \delta$.

Confirmation in the second sub-phase Consider arm k . If there exists at least one player whom did not accept this arm in the first phase then, by design, this arm cannot be confirmed. On the other hand, assume that $\rho_k \geq 2^{-p}$ and assume that all players accept arm k . Denote by $R_{k,m}^p$ the total reward received by player m when attempting to confirm arm k , which is a sum of $2^p K \ln \frac{2}{\delta}$ Bernoulli random variables with mean $\frac{\rho_k}{K}$ so that:

$$\mathbb{P}(R_{k,m}^p = 0) = (1 - \frac{\rho_k}{K})^{K2^p \ln \frac{2}{\delta}} \leq \exp(-\rho_k 2^p \ln \frac{2}{\delta}) \leq \exp(-\ln \frac{2}{\delta}) = \frac{\delta}{2}$$

where we used the inequality $1 - x \leq e^{-x}$. We have proven that, if $\rho_k \geq 2^{-p}$ and that all players accept arm k , then arm k will be confirmed by player m in the second subphase with probability greater than $1 - \frac{\delta}{2}$ and the procedure will terminate with probability greater than $1 - \frac{\delta}{2}$.

Putting things together We can now check that all statements of the theorem are true. We will use the fact that since $\rho_k = \mu_k(1 - 1/K)^{M-1}$ and we have

$$(1 - 1/e)\mu_k \leq (1 - 1/K)^{K-1}\mu_k \leq \rho_k \leq \mu_k.$$

Statement (i) Define $k^* \in \arg \max_{k=1,\dots,K} \rho_k$. Define p^* the smallest p such that $\rho_{k^*} \geq 2^{2-p}$. With probability greater than $1 - O(M\delta)$ all players will accept arm k^* in phase p^* and with probability greater than $1 - O(M\delta)$ all players will confirm arm k^* so that the procedure will terminate. So, with probability greater than $1 - O(M\delta)$, the procedure terminates at or before phase p^* . The duration of phase p is $(6 + K)K2^p \ln \frac{1}{\delta}$, so that the duration of the procedure up to phase p^* is:

$$(6 + K)K \ln \frac{1}{\delta} \sum_{p=1}^{p^*} 2^p = O(K^2 2^{p^*} \ln \frac{1}{\delta}) = O(\frac{K^2}{\rho_{k^*}} \ln \frac{1}{\delta}) = O(\frac{K^2}{\mu_{k^*}} \ln \frac{1}{\delta})$$

using the fact that $\rho_k \geq (1 - 1/e)\mu_k$ which is the announced duration. Furthermore, the amount of regret caused by any procedure is at most $M\mu_{k^*}$ times its duration, therefore the procedure incurs regret at most

$$O(K^2 M \ln \frac{1}{\delta})$$

Statement (ii) Denote by \tilde{k} the output of the procedure. Assume that $\rho_{\tilde{k}} \leq 2^{-p^*}$ then There must exist some player m that accepts this arm at some phase $1 \leq p \leq p^*$, and since $\rho_{\tilde{k}} \leq 2^{-p}$, this happens with probability at most $O(\delta p^* K M)$. On the other hand if $\rho_{\tilde{k}} \geq 2^{-p^*}$ we have $\rho_{\tilde{k}} \geq \frac{1}{8}\rho_{k^*}$

so that $\mu_{\tilde{k}} \geq \frac{1}{8}\mu_{k^*}$ and \tilde{k} is a good arm. Hence with probability at least $1 - O(MKp^*\delta)$ the procedure outputs a good arm.

Statement (iii) Assume that the procedure terminates at phase $1 \leq p \leq p^*$. With probability at least $1 - O(MKp^*\delta)$ we have that $\rho_{\tilde{k}} \geq 2^{-p}$ since \tilde{k} must have been accepted by at least one player. Therefore $\mu_{\tilde{k}} \geq \rho_{\tilde{k}} \geq 2^{-p} = \bar{\mu}$, and the output $\bar{\mu}$ is indeed a lower bound on the expected reward of \tilde{k} with probability at least $1 - O(MKp^*\delta)$.

Statement (vi) By design of the procedure, all players must output the same arm unless they do not exit simultaneously. Denote by \tilde{k} the output of the procedure. If not all players exit the procedure simultaneously, then there must exist a phase $p = 1, \dots, p^*$ and an arm k such that all players accept arm k but at least one of them does not confirm it. This happens with probability at most $O(Mp^*\delta)$, therefore, with probability at least $1 - O(Mp^*\delta)$ all players exit the procedure simultaneously.

We also note that $p^* = O(\ln \frac{1}{\mu_{k^*}})$. This completes the proof.

A.3. Proof of Lemmas 2 and 3

By design `VirtualMusicalChairs` and `VirtualNumberPlayers` terminate after exactly $K\tau$ and $2K\tau$ time steps respectively. Furthermore applying those two procedures is equivalent to running the initialization phase of SIC-MMAB used by [Boursier and Perchet \(2019\)](#), in a system with K identical arms with mean reward $\mu_{\tilde{k}}$. Therefore, from [Boursier and Perchet \(2019\)](#)[Lemma 11], we have that the both of those procedures succeed with probability at least $1 - O(MK\delta)$. Also, as stated above, the amount of regret caused by any procedure is at most $M\mu_{k^*}$ times its duration, therefore both procedures incur regret at most

$$O(K\tau M\mu_{k^*}) = O(KM \ln \frac{1}{\delta})$$

A.4. Proof of Lemma 4

We now analyze the `DistributedExploration` subroutine.

Maximal number of phases It is noted that phase p lasts $|\mathcal{K}|2^p \lceil \ln \frac{1}{\delta} \rceil \geq 2^p \lceil \ln \frac{1}{\delta} \rceil$. Since the time horizon is T , the maximal number of phases p^* must satisfy

$$2^{p^*} \leq 2^{p^*+1} \ln \frac{1}{\delta} = \ln \frac{1}{\delta} \sum_{p=1}^{p^*} 2^p \leq T$$

therefore the number of phases is upper bounded by $p^* \leq \ln_2 T$.

Communication errors Consider the subroutines `EncoderSendFloat` and `DecoderReceiveFloat` float, which allow to send/receive Q bits each time they are called. By design, in phase p the total number of bits sent is at most

$$3KMQ = 3KM \lceil \frac{p}{2} + 3 \rceil$$

Hence the total number of bits sent across phases is

$$O(KM(p^*)^2)$$

Since, by design of the communication protocol and the fact that $\tau = \frac{1}{\mu_{\tilde{k}}} \ln \frac{1}{\delta}$, each time a bit is sent, it is decoded correctly with probability greater than $1 - \delta$. Therefore, with probability greater than $1 - O(\delta KM(p^*)^2)$, there are no communication errors.

Estimation errors Now assume that no communication errors have taken place, and denote by $\hat{\mu}_k^p$ the estimate of the mean reward of arm k computed by the leader in phase k , t_k^p the corresponding number of samples and

$$B_k^p = \sqrt{(2 \ln \frac{1}{\delta})/t_k^p} + 2^{-\frac{p}{2}-3}$$

We say that there is an estimation error for k in phase p if $|\hat{\mu}_k^p - \mu_k| \geq B_k^p$. We have:

$$|\hat{\mu}_k^p - \mu_k| \leq |\hat{\mu}_k^p - \mathbb{E}(\hat{\mu}_k^p)| + |\mu_k - \mathbb{E}(\hat{\mu}_k^p)|$$

The first term is a random fluctuation and the second term is a quantification error due to the fact that players send quantified estimates. Since $Q = \frac{p}{2} + 3$ the error due to quantification is smaller than

$$|\mu_k - \mathbb{E}(\hat{\mu}_k^p)| \leq 2^{-\frac{p}{2}-3}$$

And using Hoeffding's inequality yields an upper bound for the estimation error probability:

$$\mathbb{P}(|\hat{\mu}_k^p - \mu_k| \geq B_k^p) \leq \mathbb{P}(|\hat{\mu}_k^p - \mathbb{E}(\hat{\mu}_k^p)| \geq \sqrt{(2 \ln \frac{1}{\delta})/t_k^p}) \leq \delta$$

Considering all arms and all phases, we have proven that with probability at least $1 - O(\delta K p^*)$, no estimation error occurs during a run of the algorithm.

Behaviour on a clean run We say that we have a clean run if no communication errors and no estimation errors occur throughout the run of the algorithm. We have already established that a clean run occurs with probability at least $1 - O(\delta K M (p^*)^2)$. On a clean run, we always have that

$$\hat{\mu}_k^p + B_k^p \leq \mu_k$$

and

$$\hat{\mu}_k^p - B_k^p \geq \mu_k$$

so that one can readily check that if k is amongst the M best arm, it can never be rejected, and that if k is not amongst the M best arms, if it is still active at phase p it will be rejected if $4B_k^p \leq \mu_{(M)} - \mu_k$. Following the same logic as [Boursier and Perchet \(2019\)](#), one can establish that k is selected at most $O(\frac{\ln \frac{1}{\delta}}{(\mu_{(M)} - \mu_k)^2})$ times, and that the maximal number of phases is upper bounded by:

$$p^* \leq O\left(\ln\left(\frac{1}{\mu_{(M)} - \mu_{(M+1)}}\right)\right)$$

Regret on a clean run We can complete the proof by upper bounding the regret on a clean run.

We first turn to the regret caused by communication. Sending a bit using the described protocol incurs at most $M\mu_{(1)}\tau = M \ln \frac{1}{\delta}$ regret and $O(KM(p^*)^2)$ bits must be sent, therefore the regret caused by communication is upper bounded by

$$O\left(KM^2 \ln\left(\frac{1}{\mu_{(M)} - \mu_{(M+1)}}\right)^2 \ln \frac{1}{\delta}\right)$$

Since no collisions occur, optimal arms are never eliminated, and any suboptimal arm k is selected at most $O(\frac{\ln(1/\delta)}{(\mu_{(M)} - \mu_k)^2})$ times, the total regret caused by exploration is $O(\sum_{k>M} \frac{\ln \frac{1}{\delta}}{\mu_{(M)} - \mu_k})$. Therefore,

as announced, with probability at least $1 - O(\delta KM(\ln T)^2)$ a clean run occurs, and the resulting regret is

$$O\left(KM^2 \ln\left(\frac{1}{\mu_{(M)} - \mu_{(M+1)}}\right)^2 \ln \frac{1}{\delta} + \sum_{k>M} \frac{\ln \frac{1}{\delta}}{\mu_{(M)} - \mu_k}\right)$$

which completes the proof.

Appendix B. Additional Subroutines

We provide here the pseudo-codes for the remaining subroutines. These subroutines are, we recall, described from the point of view of a single player, which is the natural way to describe a decentralized algorithm. It is noted that some subroutines are broken over several pages. `EncoderSendFloat` and `DecoderReceiveFloat` allow to exchange floating point numbers, and can be modified in a straightforward manner to deal with integers, yielding `EncoderSendInt` and `DecoderReceiveInt`, whose pseudo-codes are omitted.

Algorithm 6 ComLeader

Input: $\hat{\mu}$: estimates held by each player, N : number of samples held by players, \mathcal{K} : set of active arms, M' : number of active players, Q message length, τ : sampling time, \tilde{k} : good arm, p : phase, δ : confidence parameter

Output: f one of the M best arms to be assigned to player, $\tilde{\mathcal{K}}$ updated version set of active arms, \bar{M}' updated number of active players, $\bar{\mu}$ updated estimates held by players, \bar{N} updates number of samples held by players

$\bar{\mu} \leftarrow \hat{\mu}, \bar{N} \leftarrow N$

receive the updated values of the estimates held by active players

for $i \leftarrow 2, \dots, M$ **do**

for $k \in \mathcal{K}$ **do**

$\bar{\mu}[k, i] \leftarrow \text{DecoderReceiveFloat}(\tilde{k}, \tau, Q)$

$\bar{N}[k, i] \leftarrow \bar{N}[k, i] + 2^p \lceil \ln \frac{1}{\delta} \rceil$

end for

end for

compute the estimates aggregated across all players decide which arms to accept / reject

for $k \in \mathcal{K}$ **do**

compute the estimate aggregated across all players

$\rho[k] \leftarrow (\sum_{i=1}^M \bar{\mu}[k, i] \bar{N}) / (\sum_{i=1}^M \bar{N}[k, i])$

compute confidence radius

$B[k] \leftarrow \sqrt{(2 \ln \frac{1}{\delta}) / (\sum_{i=1}^M \bar{N}[k, i])} + 2^{-\frac{p}{2}-3}$

accept arm

if $|\{i \in \mathcal{K} : \rho[k] - B[k] \geq \rho[i] + B[i]\}| \geq |\mathcal{K}| - M'$ **then**

 Add k to $C[., 1]$

end if

reject arm

if $|\{i \in \mathcal{K} : \rho[i] - B[i] \geq \rho[k] + B[k]\}| \geq M'$ **then**

 Add k to $C[., 2]$

end if

end for

message size to send accepted / rejected arms

$Q' \leftarrow \lceil \log_2 |\mathcal{K}| \rceil$

send the size of the sets of accepted / rejected arms

for $i \leftarrow 2, \dots, M'$ and $s = 1, 2$ **do**

 EncoderSendInt($\mathcal{K}, \tilde{k}, \tau, Q', \text{length}(C[., s])$)

end for

send the contents of the sets of accepted / rejected arms

for $i \leftarrow 2, \dots, M'$ and $s = 1, 2$ and $k \in C[., s]$ **do**

 EncoderSendInt($\mathcal{K}, \tilde{k}, \tau, Q', k$)

end for

$C' \leftarrow C$ and remove \tilde{k} from $C'[., 1]$

```

if  $M' = \text{length}(C'[:, 1])$  and  $C = C'$  then
     $f \leftarrow C[M', 1]$ 
else if  $M' - 1 = \text{length}(C'[:, 1])$  and  $C \neq C'$  then
    # assign the good arm to the leader
     $f \leftarrow \hat{k}$ 
else
    # otherwise make accepted and rejected arms inactive and update number of active players
     $\bar{M} \leftarrow M' - \text{length}(C[:, 1])$ 
     $\mathcal{K}' \leftarrow \mathcal{K}$ 
    for  $k \in C'$  do
        Remove  $k$  from  $\mathcal{K}'$ 
    end for
end if

```

Algorithm 7 ComFollow

Input: E : estimated reward of each arm by the player, j : rank of the player, \mathcal{K} : set of active arms, M' : number of active players, Q : message size, τ : sampling time, \tilde{k} : a good arm

Output: f one of the M best arms to be assigned to player, $\tilde{\mathcal{K}}$ updated version set of active arms, \bar{M}' updated number of active players

```

 $\mathcal{K}' \leftarrow \mathcal{K} \setminus \{\tilde{k}\}$ 
# send reward estimates of active arms to the leader
for  $i = 2, \dots, M'$  do
  if  $j = i$  then
    for  $k \in \mathcal{K}$  do
      EncoderSendFloat( $\mathcal{K}, \tilde{k}, \tau, Q, E[k]$ )
    end for
  else
    for  $t = 1, \dots, |\mathcal{K}| \tau Q$  do
      Select the  $(j \bmod |\mathcal{K}'|)$ -th arm in set  $\mathcal{K}'$ 
    end for
  end if
end for
# message size to send accepted / rejected arms
 $Q' \leftarrow \lceil \log_2 |\mathcal{K}| \rceil$ 
# receive the sizes of the sets of accepted and rejected arms
for  $i \leftarrow 2, \dots, M'$  do
  if  $j = i$  then
    for  $s = 1, 2$  do
       $N[s] \leftarrow \text{DecoderReceiveInt}(\tilde{k}, \tau, Q')$ 
    end for
  else
    for  $t = 1, \dots, 2\tau Q'$  do
      Select the  $(j \bmod |\mathcal{K}'|)$ -th arm in set  $\mathcal{K}'$ 
    end for
  end if
end for
# receive the contents of the sets of accepted and rejected arms
for  $i \leftarrow 2, \dots, M'$  do
  if  $j = i$  then
    for  $s = 1, 2$  and  $q = 1, \dots, N[s]$  do
       $C[q, s] \leftarrow \text{DecoderReceiveInt}(\tilde{k}, \tau, Q')$ 
    end for
  else
    for  $t = 1, \dots, \tau Q'(N[1] + N[2])$  do
      Select the  $(j \bmod |\mathcal{K}'|)$ -th arm in set  $\mathcal{K}'$ 
    end for
  end if
end for

```

```

# update the set of active arms, and the active players
# avoid assigning the good arm to followers
 $C \leftarrow C'$ 
Remove  $\tilde{k}$  from  $C'[:, 1]$ 
# if an accepted arm can be assigned to player then do so
if  $M' - j + 1 \leq \text{length}(C'[:, 1])$  then
     $f \leftarrow C'[M' - j + 1, 1]$ 
    # otherwise make accepted and rejected arms inactive and update number of active players
else
     $\bar{M} \leftarrow M' - \text{length}(C'[:, 1])$ 
     $\bar{\mathcal{K}} \leftarrow \mathcal{K}$ 
    for  $k \in C'$  do
        Remove  $k$  from  $\bar{\mathcal{K}}$ 
    end for
end if

```

Algorithm 8 EncoderSendFloat

Input: \mathcal{K} : a subset of arms, \tilde{k} : a good arm, τ : a sampling time, Q : the message size, $\mu \in [0, 1]$ a real number to send

```

# convert the number to send to a binary message of size  $Q$ 
 $S \leftarrow \text{FloatToBinary}(\mu, Q)$ 
# send each bit of the binary message
for  $q = 1, \dots, Q$  do
    if  $S[q] = 1$  then
        # send a 1 bit
         $\ell \leftarrow q \bmod |\mathcal{K} \setminus \tilde{k}|$ 
         $k \leftarrow$  the  $\ell$ -th arm in set  $\mathcal{K} \setminus \tilde{k}$ 
    else
        # send a 0 bit
         $k \leftarrow \tilde{k}$ 
    end if
    for  $t = 1, \dots, \tau$  do
        Select arm  $k$ 
    end for
end for

```

Algorithm 9 DecoderReceiveFloat

Input: \tilde{k} : a good arm, τ : a sampling time, Q : the message size**Output:** $\mu \in [0, 1]$ a received real number to send*# decode each bit of the binary message***for** $q = 1, \dots, Q$ **do** $B[q] \leftarrow 0$ **for** $t = 1, \dots, \tau$ **do**Select arm \tilde{k} and observe reward r *# decode a 1 bit if at least one non zero reward is obtained***if** $r > 0$ **then** $B[q] \leftarrow 1$ **end if****end for****end for***# convert the received binary message to a real number* $\mu \leftarrow \text{BinaryToFloat}(B, Q)$
