
Decentralized Multi-player Multi-armed Bandits with No Collision Information

Chengshuai Shi
University of Virginia

Wei Xiong
University of Virginia

Cong Shen
University of Virginia

Jing Yang
Pennsylvania State University

Abstract

The decentralized stochastic multi-player multi-armed bandit (MP-MAB) problem, where the collision information is not available to the players, is studied in this paper. Building on the seminal work of Boursier and Perchet (2019), we propose error correction synchronization involving communication (EC-SIC), whose regret is shown to approach that of the centralized stochastic MP-MAB with collision information. By recognizing that the communication phase without collision information corresponds to the *Z-channel* model in information theory, the proposed EC-SIC algorithm applies optimal error correction coding for the communication of reward statistics. A fixed message length, as opposed to the logarithmically growing one in Boursier and Perchet (2019), also plays a crucial role in controlling the communication loss. Experiments with practical Z-channel codes, such as repetition code, flip code and modified Hamming code, demonstrate the superiority of EC-SIC in both synthetic and real-world datasets.

1 Introduction

Recent years have witnessed an increased interest in the multi-player multi-armed bandits (MP-MAB) problem, in which multiple players simultaneously play the bandit game and interact with each other through arm collisions. In particular, motivated by the practical application of cognitive radio (Anandkumar et al., 2011), *decentralized* stochastic MP-MAB problems have been widely studied. See Section 7 for a review of the related work.

Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS) 2020, Palermo, Italy. PMLR: Volume 108. Copyright 2020 by the author(s).

Since the stochastic MAB problem for a single player is well understood, a predominant approach in decentralized MP-MAB is to let each player play the single-player MAB game while avoiding collisions for as much as possible (Liu and Zhao, 2010; Avner and Mannor, 2014; Rosenski et al., 2016; Besson and Kaufmann, 2018). Intuitively, this would allow the algorithm to behave as the single-player MAB. Comparing to the centralized MP-MAB, there is a multiplicative factor M (number of players) increase in the regret coefficient of $\log(T)$. This has long been considered fundamental due to the lack of communication among players.

Recently, a surprising and inspiring approach, called SIC-MMAB, is proposed in Boursier and Perchet (2019) for the collision-sensing MP-MAB problem. Instead of viewing collision as detrimental, Boursier and Perchet (2019) purposely instigates collisions as a way to communicate between players. With a careful design, an internal rank can be assigned to each player and arm statistics can be completely shared among players at a communication cost that does not dominate the arm exploration regret, which leads to an overall regret approaching the centralized setting. The proposed communication phase transmits the *total reward* by using collision/no-collision to represent bit 1/0. The theoretical analysis shows, for the first time, that the regret of a decentralized MP-MAB algorithm can approach that of the centralized counterpart, which represents a significant progress in decentralized MP-MAB.

The no-sensing problem, on the other hand, represents arguably the most difficult setting in MP-MAB and there has been little progress in the literature. Boursier and Perchet (2019) makes two attempts to generalize the forced collision idea to this setting. Directly applying SIC-MMAB to this setting leads to a $O(\log(T) \log^2(\log(T)))$ regret, which has an additional $\log(T)$ multiplicative coefficient due to no sensing. In other words, a straightforward application of SIC-MMAB results in the communication loss dominating the total regret. Then, the authors propose a different approach: use communication only to exchange the

Table 1: Regret Upper Bounds of MP-MAB Algorithms

Model	Reference	Asymptotic Upper Bound (up to constant factor)
Centralized Multiplayer	Komiyama et al. (2015)	$\sum_{k>M} \frac{\log(T)}{\mu_{(M)} - \mu_{(k)}}$
Decentralized, Col. Sensing	SIC-MMAB (Boursier and Perchet, 2019)	$\sum_{k>M} \frac{\log(T)}{\mu_{(M)} - \mu_{(k)}} + MK \log(T)$
Decentralized, No Sensing	Lugosi and Mehrabian (2018)	$\frac{MK \log(T)}{M \log(T)}$
Decentralized, No Sensing	Lugosi and Mehrabian (2018)	$\frac{MK^2}{\Delta^2} \log^2(T) + MK \frac{\log(T)}{\Delta}$
Decentralized, No Sensing	ADAPTED SIC-MMAB (Boursier and Perchet, 2019)	$\sum_{k>M} \frac{\log(T)}{\mu_{(M)} - \mu_{(k)}} + \frac{M^3 K \log(T)}{\mu_{(K)}} \log^2(\log(T))$
Decentralized, No Sensing	SIC-MMAB2 (Boursier and Perchet, 2019)	$M \sum_{k>M} \frac{\log(T)}{\mu_{(M)} - \mu_{(k)}} + \frac{MK^2}{\mu_{(K)}} \log(T)$
Decentralized, No Sensing	EC-SIC (this paper)	$\sum_{k>M} \frac{\log(T)}{\mu_{(M)} - \mu_{(k)}} + \left(\frac{M^2 K}{E(\mu_{(K)})} \log\left(\frac{1}{\Delta}\right) + \frac{MK}{\mu_{(K)}} \right) \log(T)$

K : number of arms; $M \leq K$: number of players; $\mu_{(j)}$ is j -th order statistics of μ ; $\Delta := \mu_{(M)} - \mu_{(M+1)} > 0$;
 $\Delta' := \min\{\mu_{(M)} - \mu_{(i)} \mid \mu_{(M)} - \mu_{(i)} > 0\}$; $E(\mu_{(K)})$ is the random coding error exponent.

accepted and rejected arms, thus reducing the regret caused by communication. However, this approach, philosophically speaking, deviates from the core idea of SIC-MMAB and does not fully utilize the communication benefit of collision (arm statistics are not shared among players). This is also the reason why the multiplicative factor M reappears in the regret formula, which was eliminated in the collision-sensing case by SIC-MMAB. It remains an open problem *whether a decentralized MP-MAB algorithm without collision information can approach the performance of its centralized counterpart*.

In this work, we return to the original idea of utilizing collisions to communicate sampled arm rewards. By modelling no collision information as the *Z-channel* communication problem in information theory, we propose to incorporate optimal error correction coding in the communication phase to control the error rate of decoding the message. With this approach, we are able to transmit a *quantized* sample reward mean with a *fixed* length for each player without having the communication loss dominate the total regret. The resulting asymptotic regret improves the $\log(T)$ coefficients over SIC-MMAB2, and represents the best known regret in the no-sensing MP-MAB model to the best of the authors' knowledge. Table 1 compares the asymptotic regret upper bounds for different algorithms. We also propose two practical enhancements that significantly improve the algorithm's empirical performance. Numerical experiments on both synthetic and real-world datasets corroborate the analysis and offer interesting insights into EC-SIC.

2 The No-Sensing MP-MAB Problem

In the standard (single player) stochastic MAB setting, there are K arms, with rewards X_i of arm $i \in [K]$ sampled independently from a distribution ν_i on $[0, 1]$, where $\mathbb{E}[X_i] = \mu[i]$. At time t , a player chooses an arm $\pi(t)$ and the goal is to receive the highest mean cumulative reward in T rounds.

In this section, we introduce the no-sensing multi-

player MAB model with a known number of arms K but an unknown number of players $M \leq K$. The horizon T is known to the players. At each time step $t \in [T]$, all the players $j \in [M]$ simultaneously pull the arms $\pi^j(t)$ and receive the reward $r^j(t)$ such that

$$r^j(t) := X_{\pi^j(t)}(t)(1 - \eta_{\pi^j(t)}(t))$$

where $\eta_{\pi^j(t)}(t)$ is the collision indicator defined by $\eta_k(t) := \mathbb{1}_{|C_k(t)| > 1}$ with $C_k(t) := \{j \in [M] \mid \pi^j(t) = k\}$.

If players can observe both $r^j(t)$ and $\eta_{\pi^j(t)}(t)$, it is the collision-sensing problem. On the other hand, in the no-sensing case as in our paper, players can only access $r^j(t)$, i.e., a reward of 0 can indistinguishably come from a collision with another player or $X_{\pi^j(t)}(t) = 0$. Note that if $\mathbb{P}(\forall k \in [K], X_k = 0) = 0$, the no-sensing and collision-sensing models are equivalent¹.

The performance in the standard single-player MAB setting is usually measured by the regret:

$$R(T) := T\mu_* - \sum_{t=1}^T \mu[\pi(t)],$$

where $\mu_* = \max_i \mu[i]$ is the expected reward of the arm with the highest expected reward. As shown in the lower bound by Lai and Robbins (1985), the optimal order of the regret cannot be better than $O(\log(T))$.

In the multiplayer setting, the notion of regret can be generalized and defined with respect to the best allocation of players to arms, as follows:

$$R(T) := T \sum_{j \in [M]} \mu_{(j)} - \mathbb{E} \left[\sum_{t=1}^T \sum_{j \in [M]} r^j(t) \right],$$

where $\mu_{(j)}$ is j -th order statistics of μ , i.e. $\mu_{(1)} \geq \mu_{(2)} \geq \dots \geq \mu_{(K-1)} \geq \mu_{(K)}$.

Two technical assumptions are made in this paper, which are also widely used in the literature. The first is a strictly positive lower bound of $\mu_{(K)}$, which has

¹We further note that the no-sensing model can be generalized to an arbitrary but bounded reward support where collision results in the lowest value in the support.

been used by Lugosi and Mehrabian (2018) and Boursier and Perchet (2019) for the no-sensing model. The second assumption is a finite gap between the optimal and suboptimal (group of) arms; see Avner and Mannor (2014); Kalathil et al. (2014); Rosenski et al. (2016); Nayyar et al. (2016) for this assumption.

Assumption 1. 1) A positive lower bound μ_{\min} of $\mu_{(K)}$ is known to all players: $0 < \mu_{\min} \leq \min_{i \in [K]} \mu[i]$.
 2) There exists a positive gap $\Delta \doteq \mu_{(M)} - \mu_{(M+1)} > 0$, and it is known to all players.

Assumption 1.1 is equivalent to $\forall k \in [K], \mathbb{P}[X_k > 0] \geq \mu_{\min}$. This also bounds $\mathbb{P}[X_k = 0]$. Note that although μ_{\min} provides a lower bound for $\mu_{(K)}$, Assumption 1.1 does not require the exact value of $\mu_{(K)}$. The gap in Assumption 1.2 measures the difficulty of the bandit game and ensures the existence of only one optimal choice.

3 The EC-SIC Algorithm

The proposed error correction synchronization involving communication (EC-SIC) is compactly described in Algorithm 1. Similar to SIC-MMAB, the overall algorithm can be structurally divided into four phases: initialization phase, exploration phase, communication phase, and exploitation phase. It is important to note that all players are synchronized in running EC-SIC, i.e., they enter each phase at the same time (or at least with high probability in some cases) except the exploitation phase. Until a player fixates on a specific arm and enters the exploitation phase, the algorithm keeps iterating between the exploration and communication phases. Players that have (not) entered the exploitation phase are called inactive (active). We denote the set of active players during the p -th phase by $[M_p]$ and its cardinality by M_p . Similarly, arms that have not been decided to be optimal or sub-optimal are called active. The set of active arms during the p -th phase is denoted by $[K_p]$ with cardinality K_p .

3.1 Initialization phase

The same structure of the initialization phase as Boursier and Perchet (2019) is used in EC-SIC, which outputs an internal rank $j \in \{1, \dots, M\}$ for each player as well as the estimated value of M . It starts with a ‘‘Musical Chair’’ phase and is followed by a so-called Sequential Hopping protocol. The full procedure is described in Appendix B.1 for completeness.

3.2 Exploration phase

During the p -th exploration phase, active players sequentially hop among the active arms for

Algorithm 1 The EC-SIC Algorithm

Input: $T, K, \Delta, \epsilon, \mu_{\min}$
 1: Initialize $p \leftarrow 1$; $F \leftarrow -1$; $T_0, T_0^j \leftarrow 0$; $[K_p] \leftarrow [K]$;
 $Q \leftarrow \max\{\lceil \log_2 \frac{1}{\Delta - \epsilon} \rceil, \lceil \log_2(K + 1) \rceil\}$; $T_c \leftarrow \lceil \frac{\log(T)}{\mu_{\min}} \rceil$
 2: Select an error-correction code (N', Q) with code length N' defined in Theorem 2
 \triangleright **Initialization Phase:**
 3: $k \leftarrow \text{Musical_Chair}([K], KT_c)$
 4: $(M, j) \leftarrow \text{Estimate_M_NoSensing}(k, T_c)$
 5: **while** $F = -1$ **do**
 \triangleright **Exploration Phase:**
 6: $\pi \leftarrow j$ -th active arm
 7: **for** $K_p 2^p \lceil \log(T) \rceil$ time steps **do**
 8: $\pi \leftarrow \pi + 1 \pmod{K_p}$ and play arm π
 9: $s[\pi] \leftarrow s[\pi] + r^j(t)$
 10: **end for**
 11: $T_p^j = T_{p-1}^j + 2^p \lceil \log(T) \rceil$
 12: $\hat{\mu}_j = s/T_p^j$
 \triangleright **Communication Phase:**
 13: **if** $j = 1$ **then** $(F, M_{p+1}, [K_{p+1}]) \leftarrow \text{Communication Leader}(\hat{\mu}_1, p, [K_p], M_p, Q, N')$
 14: **else** $(F, M_{p+1}, [K_{p+1}]) \leftarrow \text{Communication Follower}(\hat{\mu}_j, j, p, [K_p], M_p, Q, N')$
 15: **end if**
 16: $p \leftarrow p + 1$
 17: **end while**
 18: \triangleright **Exploitation phase:** Pull F until T

$K_p 2^p \lceil \log(T) \rceil$ steps, and any active arm is pulled $2^p \lceil \log(T) \rceil$ times by each active player. Since the hopping is based on each player’s internal rank, the exploration phase is collision-free.

We note that the length of an exploration phase is different from Boursier and Perchet (2019), which is a key component of the performance improvement. The difference of a $\lceil \log(T) \rceil$ factor, in fact, results in an overall $O(\log(\log(T)))$ rounds of exploration and communication phases in the ADAPTED SIC-MMAB algorithm of Boursier and Perchet (2019). This directly leads to a dominating communication loss that breaks the order-optimality. With an expansion of length by $\lceil \log(T) \rceil$ in EC-SIC, the overall rounds become a constant, and the communication regret can be better controlled as shown in Section 4.

3.3 Communication phase

In the communication phase, all players attempt to exchange their sampled reward information in a synchronized and distributed manner. The communication takes place via a careful collision design. All players enter this phase synchronously and, by default, keep pulling different arms based on their internal ranks. Then, when it is player i ’s turn to communicate with player j , she would purposely pull (not pull) player j ’s arm as a way to communicate bit 1 (0). If player j can fully access the collision information, i.e., knowing

Algorithm 2 Communication Leader

Input: $\hat{\mu}_1, p, [K_p], M_p, Q, N'$
Output: $F, M_{p+1}, [K_{p+1}]$

- 1: Initialize $T_p = T_{p-1} + M_p 2^p \lceil \log(T) \rceil$; $T_p^i = T_p^1, i = 1, \dots, M_p$; $\bar{\mu}_i^p = \bar{\mu}_i^{p-1}, T_p^i = T_{p-1}^i, i = M_p + 1, \dots, M$
 \triangleright Gather information from followers:
- 2: **for** $i = 2, \dots, M_p$ **do** \triangleright Receive arm statistics
- 3: **for** $k \in [K_p]$ **do**
- 4: $\bar{\mu}_i^p[k] \leftarrow \text{Decoder}(\text{Receive}(1, i, N'))$
- 5: **end for**
- 6: **end for**
- 7: $\bar{\mu}^p = \sum_{i=1}^M \bar{\mu}_i^p \cdot T_p^i / T_p$; $B_{T_p} = \sqrt{\frac{2 \log(T)}{T_p}} + (\frac{\Delta}{4} - \epsilon)$
 \triangleright Update statistics:
- 8: $\text{Rej} \leftarrow$ set of active arms k satisfying $|\{i \in [K_p] | \bar{\mu}^p[i] - B_{T_p} \geq \bar{\mu}^p[k] + B_{T_p}\}| \geq M_p$
- 9: $\text{Acc} \leftarrow$ set of active arms k satisfying $|\{i \in [K_p] | \bar{\mu}^p[i] - B_{T_p} \geq \bar{\mu}^p[k] + B_{T_p}\}| \geq K_p - M_p$, ordered according to their indices
 \triangleright Transmit acc\rej arms to followers:
- 10: **for** $i = 2, \dots, M_p$ **do** \triangleright Send acc\rej set size
- 11: $\text{Send}(1, i, N', \text{Encoder}(|\text{Rej}|, Q))$
- 12: $\text{Send}(1, i, N', \text{Encoder}(|\text{Acc}|, Q))$
- 13: **end for**
- 14: **for** $i = 2, \dots, M_p$ **do** \triangleright Send acc\rej set content
- 15: $\text{Send}(1, i, N', \text{Encoder}(k, Q))$ for $k \in \text{Rej}$
- 16: $\text{Send}(1, i, N', \text{Encoder}(k, Q))$ for $k \in \text{Acc}$
- 17: **end for**
- 18: **if** $M_p \leq |\text{Acc}|$ **then**
- 19: $F \leftarrow \text{Acc}[M_p]$
- 20: **else** $M_p \leftarrow M_p - |\text{Acc}|$
- 21: $[K_{p+1}] \leftarrow [K_p] \setminus (\text{Acc} \cup \text{Rej})$
- 22: **end if**

whether collision happens or not at each time step, she will be able to receive the bit sequence successfully, which conveys player i 's sample reward statistics. However, for the no-sensing model, such error-free communication becomes impossible.

Three new ideas are used in the communication phase of EC-SIC. The first is the introduction of *Z-channel coding*. In the no-sensing scenario, players cannot directly identify collision. If the same communication protocol in Boursier and Perchet (2019) (representing 1 or 0 by collision or no collision) is used, the confusion may mislead the player to believe that collision has occurred (bit 1) while it is actually a null statistic of reward sampling (bit 0). This error has a catastrophic consequence in that it breaks the essential synchronization between players. We are thus facing the challenge of communicating the reward statistics to other users while controlling the error rate for the overall communication loss to *not* dominate the regret.

Luckily, this is the well-known *reliable communication over a noisy channel* problem, one of the foundations in information theory. In particular, our communication channel is asymmetric: 1 (collision) is always received correctly and 0 (no collision) may be received incorrectly with a certain probability $P(X_k = 0)$. This

Algorithm 3 Communication Follower

Input: $\hat{\mu}_j, j, p, [K_p], M_p, Q, N'$
Output: $F, M_{p+1}, [K_{p+1}]$

\triangleright Transmit information to the leader:

- 1: **for** $i = 2, \dots, M_p$ **do** \triangleright Send arm statistics
- 2: **if** $j = i$ **then**
- 3: $\text{Send}(j, 1, N', \text{Encoder}(\hat{\mu}_j[k], Q))$ for $k \in [K_p]$
- 4: **else** pull the j -th active arm for $K_p N'$ steps
- 5: **end if**
- 6: **end for**

\triangleright Receive acc\rej arms from the leader:

- 7: **for** $i = 2, \dots, M_p$ **do** \triangleright Receive acc\rej set size
- 8: **if** $j = i$ **then**
- 9: $N_{\text{rej}} \leftarrow \text{Decoder}(\text{Receive}(j, 1, N'))$
- 10: $N_{\text{acc}} \leftarrow \text{Decoder}(\text{Receive}(j, 1, N'))$
- 11: **else** pull j -th active arm for $2N'$ steps
- 12: **end if**
- 13: **end for**
- 14: **for** $i = 2, \dots, M_p$ **do** \triangleright Receive acc\rej set content
- 15: **if** $j = i$ **then**
- 16: $\mathbf{w}[k] \leftarrow \text{Receive}(j, 1, N')$ and
- 17: $\text{Rej}[k] \leftarrow \text{Decoder}(\mathbf{w}[k])$ for $k = 1, \dots, N_{\text{rej}}$
- 18: $\mathbf{w}[k] \leftarrow \text{Receive}(j, 1, N')$ and
- 19: $\text{Acc}[k] \leftarrow \text{Decoder}(\mathbf{w}[k])$ for $k = 1, \dots, N_{\text{acc}}$
- 20: **else** pull j -th active arm for $(N_{\text{rej}} + N_{\text{acc}})N'$ steps
- 21: **end if**
- 22: **end for**
- 23: **if** $M_p - j + 1 \leq |\text{Acc}|$ **then**
- 24: $F \leftarrow \text{Acc}[M_p - j + 1]$
- 25: **else** $M_p \leftarrow M_p - |\text{Acc}|$
- 26: $[K_{p+1}] \leftarrow [K_p] \setminus (\text{Acc} \cup \text{Rej})$
- 27: **end if**




Figure 1: The Z-channel model

corresponds to the **Z-channel** model (see Fig.1) in information theory (Tallini et al., 2002), which represents a broad class of asymmetric channels. The Z-channel has a crossover probability q of $0 \rightarrow 1$ that corresponds to $P(X_k = 0)^2$.

The Z-channel capacity is derived in Tallini et al. (2002) as follows.

Theorem 1. *The capacity $C_Z(q)$ of a Z-channel with crossover $0 \rightarrow 1$ probability q is:*

$$C_Z(q) = \log_2(1 + (1 - q)q^{q/(1-q)}). \quad (1)$$

Shannon theory guarantees that as long as the coding rate R is below $C_Z(q)$ in Eqn. (1), there exists at least one code that allows for an arbitrarily low error

²Since the crossover probability $P(X_k = 0)$ is unknown and varies for different arm k , $1 - \mu_{\min}$ is used to capture the worst case.

rate asymptotically. This means that theoretically, for this Z-channel, it is possible to transmit information nearly error-free when the rate is close to $C_Z(q)$ bits per channel use. In reality, however, different finite block-length channel codes may have different performances; we thus evaluate several practical codes both theoretically (in Section 4) and experimentally (in Section 6). For simplicity, Functions `Send()`, `Receive()`, `Encoder()` and `Decoder()` are used in the algorithm as the sending and receiving protocol and the encoder and corresponding decoder, respectively.

The second enhancement is to *transmit each arm's quantized sample reward mean with a fixed length*. The reason not to use the *total* reward as Boursier and Perchet (2019) is that the gradually increasing total reward leads to a message length $O(\log(\log(T)))$, which cannot be transmitted efficiently in the no-sensing case with $O(\log(T))$ bits. However, with a finite gap, a less precise statistics sharing is tolerable as long as it does not affect the choice of optimal arms. For a quantized sample mean of length Q , the error is at most 2^{-Q} . We thus control the length Q such that $2^{-Q} < \frac{\Delta}{4} - \epsilon$, where $\epsilon \in (0, \frac{\Delta}{4})$ is a pre-defined constant. By constructing $B_s = \sqrt{\frac{2\log(T)}{s}} + (\frac{\Delta}{4} - \epsilon)$ as the confidence bound, analysis in Section 4 shows that acceptance and rejection maintain a high probability of success.

Lastly, compared to the mesh-structured communication in Boursier and Perchet (2019), it is more efficient to form a tree structure that one player (“leader”) gathers all statistics and makes decisions for others (“followers”). The player with internal rank 1 becomes the leader and the rest become the followers (Kaufmann and Mehrabian, 2019). Statistics of arms are transmitted from followers to the leader. The leader decides the set of arms to be accepted or rejected by comparing their upper confidence bounds and lower confidence bounds with each other, and sends back to the followers. Upon reception, active players either enter another iteration of exploration and communication, or begin exploitation. This process utilizes reward statistics from all players and has better communication efficiency. Procedures for the leader and followers are given in Algorithm 2 and 3, respectively.

4 Theoretical Analysis

The overall regret of EC-SIC can be decomposed as $R(T) = R^{init} + R^{expl} + R^{comm}$. The first, second and third term refers to the regret caused by the initialization, exploration, and communication phase, respectively. The main result is presented in Theorem 2, and each component regret is subsequently analyzed. Detailed proofs can be found in Appendix B.

Theorem 2. *With an optimal coding technique that achieves Gallager's error exponent $E(\mu_{min})$ for the corresponding Z-channel with crossover probability $1 - \mu_{min}$, for any $\epsilon \in (0, \frac{\Delta}{4})$, we have*

$$\begin{aligned} R(T) &\leq c_1 MK \frac{\log(T)}{\mu_{min}} \\ &+ c_2 \frac{\Delta}{4\epsilon} \left(\sum_{k>M} \min \left\{ \frac{\log(T)}{\mu_{(M+1)} - \mu_{(k)} + 4\epsilon}, \sqrt{T \log(T)} \right\} \right) \\ &+ c_3 N' \left(M^2(K+2) \log \left(\min \left\{ \frac{1}{4\epsilon}, T \right\} \right) + M^2 K \right) \end{aligned} \quad (2)$$

where c_1 , c_2 and c_3 are constants and $N' = \max\{\frac{Q}{C_Z(1-\mu_{min})}, \frac{1}{E(\mu_{min})} \log(T)\}$.

Theorem 2 involves an information-theoretic concept called *error exponent*, which is explained in Theorem 3 in Section 4.2 but more details can be found in (Gallager, 1968).

An asymptotic upper bound can be obtained from (2) with $\epsilon = \frac{\Delta}{8}$:

$$R(T) = O \left(\sum_{k>M} \frac{\log(T)}{\mu_{(M)} - \mu_{(k)}} + \left(\frac{M^2 K \log(\frac{1}{\Delta})}{E(\mu_{min})} + \frac{MK}{\mu_{min}} \right) \log(T) \right). \quad (3)$$

Compared to SIC-MMAB2, we have successfully removed the multiplicative factor of M in the first $\log(T)$ term. This is due to the efficient communication phase that transmits the reward statistics. In addition, we have a $M^2 K$ factor in the second $\log(T)$ term, as opposed to MK^2 in SIC-MMAB2. This is also an improvement since $M < K$. We also note that Eqn. (2) and Eqn. (3) hold when $\mu_{(min)}$ is replaced by $\mu_{(K)}$.

To prove Theorem 2, we first define the “typical event” as the success of initialization, communication and exploration throughout the entire horizon T . More specifically, we define three events: $A_1 = \{\text{each player has a correct estimation of } M \text{ and an orthogonal internal rank after initialization}\}$; $A_2 = \{\text{messages are decoded correctly in all communication phases}\}$; $A_3 = \{|\bar{\mu}^p[k] - \mu[k]| \leq B_{T_p} \text{ holds for phase } p, \forall k \in [K_p], \forall p\}$. We use P_s to denote the probability that the typical event happens, which is $A_1 \cap A_2 \cap A_3$. The regret caused by the “atypical event” can be simply bounded by a linear regret $O(MT)$. Then the result of (2) can be proved by controlling P_s to balance both events.

4.1 Initialization phase

Similar to Lemma 11 in Boursier and Perchet (2019), we can bound the regret of initialization as follows.

Lemma 1. *With probability $P_i = 1 - O(\frac{MK}{T})$, event A_1 happens. Furthermore, the regret of the initialization phase satisfies:*

$$R^{init} < 3MK \left\lceil \frac{\log(T)}{\mu_{min}} \right\rceil.$$

4.2 Exploration phase

The regret due to exploration is bounded in Lemma 2.

Lemma 2. *With probability $P_s = 1 - O(\frac{MK \log(T)}{T})$, the typical event happens and the exploration regret conditioned on the typical event satisfies:*

$$R^{expl} = O\left(\frac{\Delta}{4\epsilon} \sum_{k>M} \min\left\{\frac{\log(T)}{\mu_{(M+1)} - \mu_{(k)} + 4\epsilon}, \sqrt{T \log(T)}\right\}\right).$$

We first present a fundamental result of channel coding for communication in a noisy channel, known as the *error exponent* (Gallager, 1968).

Theorem 3. *For a discrete memoryless channel, if $R < C$, there exists a code of block length N without feedback such that the error probability is bounded by*

$$P_e \leq \exp[-NE_r(R)],$$

where $E_r(R)$ is the random coding error exponent with rate R .

We note that the error exponent used in Theorem 2 corresponds to $E(\mu_{\min}) = E_r(C_Z(1 - \mu_{\min}))$.

Theorem 3 suggests that, to transmit a Q -bit message over a Z-channel, there exists an optimal coding scheme with length $N' = \max\{\frac{Q}{C_Z(1-\mu_{\min})}, N\}$ to achieve an error rate less than $\frac{1}{T}$, where $N = \frac{1}{E(\mu_{\min})} \log(T)$. Several of the existing coding techniques, although not optimal, can achieve this error rate with $N = \Theta(\log(T))$, which only leads to a multiplicative factor larger than $\frac{1}{E(\mu_{\min})}$ but does not change the regret order. For example, with repetition code, flip code and modified Hamming code, we have $N_{rep} = Q \lceil \frac{\log(QT)}{\mu_{\min}} \rceil$, $N_{flip} = Q \lceil \frac{\log(QT/2)}{\mu_{\min}} \rceil$, $N_{ham} = \frac{7Q}{8} \lceil \frac{\log(7QT/8)}{\mu_{\min}} \rceil$ respectively (see Appendix A for detailed analysis of practical codes). The remaining analysis will be based on the optimal channel coding with the caveat that a ‘‘good’’ Z-channel code should be applied in practice.

With at most $\log(T)$ exploration and communication phases and K arms to be accepted or rejected, there are at most $MK \log(T)$ communication instances on arm statistics, $2M \log(T)$ communication instances on the number of acc/rej arms, and KM communication instances on the index of acc/rej arms. A simple union bound analysis leads to the following result.

Lemma 3. *Denoting the probability that event A_2 holds by P_r , with an optimal Z-channel code of $N' = \max\{\frac{Q}{C_Z(1-\mu_{\min})}, \frac{\log(T)}{E(\mu_{\min})}\}$, we have*

$$P_r = 1 - O\left(\frac{MK \log(T)}{T}\right).$$

Lemma 3 guarantees all communications are correct. To bound the probability that all arms are correctly estimated, we have the following result.

Lemma 4. *In phase p , for any active arm $k \in [K_p]$,*

$$P\{|\bar{\mu}^p[k] - \mu[k]| \geq B_{T_p}\} \leq \frac{2}{T}.$$

With at most $\log(T)$ exploration-communication phases, event A_3 happens with probability:

$$P_c = 1 - O\left(\frac{K \log(T)}{T}\right). \quad (4)$$

A union bound argument leveraging P_i , P_r and P_c leads to probability P_s for the typical event to happen, as defined in Lemma 2. Finally, for the exploration phases, the number of times that an arm is pulled before being accepted or rejected are well controlled.

Lemma 5. *In the typical event, every optimal arm is accepted after at most $O\left(\frac{\log(T)}{(\mu_{(k)} - \mu_{(M)} + 4\epsilon)^2}\right)$ pulls, and every sub-optimal arm is rejected after at most $O\left(\frac{\log(T)}{(\mu_{(M+1)} - \mu_{(k)} + 4\epsilon)^2}\right)$ pulls.*

Denote T^{expl} as the overall time of exploration and exploitation phase and $T_{(k)}^{expl}(T)$ as the number of time steps where the k -th best arm is pulled during these two phases. With no collision in exploration and exploitation, the exploration regret can be decomposed as (Anantharam et al., 1987)

$$R^{expl} = \sum_{k>M} (\mu_{(M)} - \mu_{(k)}) T_{(k)}^{expl}(T) + \sum_{k \leq M} (\mu_{(k)} - \mu_{(M)}) (T^{expl} - T_{(k)}^{expl}(T)), \quad (5)$$

Both components in (5) can be upper bounded by Lemma 7 in Appendix B.2.4, which proves Lemma 2.

4.3 Communication phase

Thanks to the expanded length of each exploration phase and the fixed-length quantization of arm statistics, the regret R^{comm} does not dominate the overall regret, as stated in Lemma 6.

Lemma 6. *In the typical event,*

$$R^{comm} = O\left(N' (M^2 (K+2) \log(\min\{\frac{1}{4\epsilon}, T\}) + M^2 K)\right).$$

We note that $\log(\min\{\frac{1}{4\epsilon}, T\})$ becomes a constant when T is sufficiently large. Noting that $N' = \max\{\frac{Q}{C_Z(1-\mu_{\min})}, \frac{\log(T)}{E(\mu_{\min})}\}$, the communication loss has the same order as other phases.

4.4 Overall regret

When the typical event happens, the overall regret is bounded by the sum of R^{init} , R^{comm} and R^{expl} ; otherwise, for the atypical event, the regret can be upper bounded as MT . Thus, the overall regret satisfies

$$R(T) \leq R^{init} + R^{expl} + R^{comm} + O(M^2 K \log(T)).$$

With Lemmas 1, 2 and 6, Theorem 2 can be proven.




Figure 2: Easy game




Figure 3: Hard game




Figure 4: Different game difficulties




Figure 5: Different coding techniques




Figure 6: Different codeword lengths




Figure 7: The MovieLens dataset

5 Algorithm Enhancement

EC-SIC has nice theoretical performance guarantees, but we have noticed that in practice some minor enhancements improve its performance significantly, which are shown in the next section. First, after each exploration and communication phase, player j can use the active arm with the j -th best empirical mean (sent by the leader) as her communication arm for the next around, instead of the j -th active arm. Since players keep receiving rewards from their communication arms while waiting for communication or receiving bit 0, using an arm with higher empirical mean can lead to a lower loss in these time steps.

Second, we have observed in practice that the first one or two explorations do not lead to effective acceptance or rejection even when the game is easy, which means all the communication losses during these phases are incurred with no benefit (this is much larger than the exploration loss). Thus, p can be initialized to a larger integer (e.g. 5), which leads to a longer exploration to start with and less ineffective communication.

Lastly, if μ_{\min} and Δ in Assumption 1 are not available, *adaptive estimation with confidence intervals* can be used to replace the true μ_{\min} and Δ in EC-SIC. The influence of mismatched μ_{\min} and Δ are evaluated in the experiments and reported in Appendix C.2.

6 Experiments

Numerical experiments have been carried out to verify the analysis of EC-SIC and compare its empirical performance to other methods. All rewards follow the

Bernoulli distributions with $\mu_{\min} = 0.3$, and we set $\epsilon = \Delta/8$. Results are obtained by averaging over 500 experiments. More detailed discussions and additional results can be found in Appendix C.

We compare state of the art algorithms under both easy and difficult bandit game settings. EC-SIC (with repetition code), ADAPTED SIC-MMAB, SIC-MMAB2, and the algorithm proposed by Lugosi and Mehrabian (2018) (labeled as “no-sensing-MC”) are first compared in a relatively easy game ($\Delta = 0.06$). Fig. 2 shows that even in an easy game, no-sensing-MC could not finish exploration within 10^6 time steps, and ADAPTED SIC-MMAB has poor performance compared to the other two. Both EC-SIC and SIC-MMAB2 converge to the optimal arm set quickly, but the overall regret of EC-SIC is smaller. For a hard game with $\Delta = 0.01$, Fig. 3 shows that EC-SIC is superior to SIC-MMAB2.

A detailed comparison of EC-SIC with SIC-MMAB2 is done by comparing their regrets as a function of the gap Δ in Fig.4. We see that when the game is not extremely difficult ($\Delta > 10^{-4}$), EC-SIC has better performance since players benefit from sharing statistics. When Δ becomes extremely small, the required communication length increases significantly, leading to a dominating communication regret in EC-SIC that cannot be offset by the benefits of sharing statistics.

Fig. 5 reports the performance while using different Z-channel codes in communication. We observe that modified Hamming Code has the best performance, which is due to its superior error correction capability. This observation also implies that with a near-optimal code that is specifically designed for Z-channel, perfor-

mance of EC-SIC can be further improved.

We also evaluate the impact of codeword length on the regret. For our simulation setting, the theoretical analysis requires a repetition code length $N = 53$ to transmit one bit, in order to achieve an error rate of $\frac{1}{T}$. We are interested in evaluating whether the theoretically required code length can be shortened in practice. Under the easy game setting of Fig. 2 with 2000 rounds averaging, Fig. 6 shows that with N decreasing from 53 to 35, the regret decreases 20%. More importantly, it shows that the convergence of EC-SIC does not change. When further reducing N to 25, we see the regret curve trends upward at large t , which represents a non-negligible loss due to unsuccessful communications. With $N = 15$, the regret increases rapidly, indicating that players suffer from an increased error rate. It is thus essential to strike a balance between error rate and communication loss.

Lastly, we evaluate EC-SIC on a real world dataset: the movie watching dataset (ml-20m) from MovieLens (Harper and Konstan, 2015). It consists of watching data of more than 2×10^4 movies from over 10^5 users between January 09, 1995 and March 31, 2015. In the pre-processing, we group these movies into $K = 40$ categories by their total number of views from high to low. The binary reward at time t (hour) is defined as whether there are users watching films in this group, and we replicate it to a final reward sequence of length $T = 2 \times 10^7$. $M = 20$ players are assumed to engage in the game. This final sequence has $\Delta \approx 0.007$ and $\mu_{\min} \approx 0.6$. Compared to synthetic datasets, this setting poses a larger and more difficult game. For each experiment, the reward sequence is randomly shuffled. We report the cumulative regret of EC-SIC and SIC-MMAB2, averaged over 100 experiments, in Figure 7. One can see that the advantage of EC-SIC over SIC-MMAB2 is significant for this real-world dataset. Intuitively, this is because the game is hard ($\Delta \approx 0.007$)³, and M and K are also large.

7 Related Work

Depending on how information is shared and actions are determined, existing literature can be categorized into centralized or decentralized (distributed) MP-MAB problems. The centralized scenario can be viewed as an application of the multiple-play bandit (Anantharam et al., 1987; Komiyama et al., 2015). A more interesting and challenging problem, introduced by Liu and Zhao (2010) and Anandkumar et al. (2011), lies in the decentralized scenario where explicit communications between players are not allowed and thus

³However, the game is also not too hard for communication to be ineffective as the case of Fig. 4.

collision may happen. For the collision-sensing MP-MAB problem, earlier works attempt to let each player play the single-player MAB game while avoiding collisions for as much as possible; see (Liu and Zhao, 2010; Avner and Mannor, 2014; Rosenski et al., 2016) for some representative approaches.

The SIC-MMAB algorithm in Boursier and Perchet (2019) is closely relevant to our work, which proposes to *exploit* collisions as opposed to *avoiding* them. Proutiere and Wang (2019) further refines this idea and decreases the communication regret so that the lower bound of the centralized setting can be approached asymptotically for Bernoulli distributed rewards. The SIC-MMAB principle has subsequently been applied to other multi-player settings. For example, Kalathil et al. (2014) considers an extended multiplayer model where reward distribution varies for each player. Bistriz and Leshem (2018) proposes a Game of Thrones algorithm that achieves a regret of $O(\log^{2+\kappa}(T))$. This is further improved by Kaufmann and Mehrabian (2019), leading to an improved regret of $O(\log^{1+\kappa}(T))$.

The no-sensing model, on the other hand, is very challenging and limited progress has been made so far. In Lugosi and Mehrabian (2018), sample means are rectified by the probability of collision and then the same Musical Chair approach is adopted. As discussed in Section 1, Boursier and Perchet (2019) touches upon the no-sensing model with ADAPTED SIC-MMAB and SIC-MMAB2. However, the former has a communication loss of $O(\log(T) \log^2(\log(T)))$ that dominates the total regret, while the latter drifts away from communication of full statistics, thus is fundamentally incapable of approaching the centralized performance.

8 Conclusion

In this work, we have proposed the EC-SIC algorithm for the no-sensing MP-MAB problem with forced collision. We proved that it is possible for a decentralized MP-MAB algorithm without collision information to approach the performance of its centralized counterpart. Recognizing that communication under the no-sensing setting corresponds to the Z-channel model in information theory, optimal error correction codes are applied for reliable communication via collision. With this tool, we return to the original idea of utilizing forced collisions to share complete arm statistics among players. By expanding exploration phases and fixing the message length, an order-optimal communication loss is achieved. Practical simulation results with several Z-channel codes have proved the superiority of EC-SIC algorithm under different bandit game settings, using both synthetic and real-world datasets.

Acknowledgements

JY acknowledges the support from U.S. National Science Foundation under Grant ECCS-1650299.

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 multi-armed bandit problem with multiple plays - part I: IID rewards. *IEEE Trans. Autom. Control*, 32(11):968–976.
- Avner, O. and Mannor, S. (2014). Concurrent bandits and cognitive radio networks. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 66–81. Springer.
- Barbero, A., Ellingsen, P., Spinsante, S., and Ytrehus, O. (2006). Maximum likelihood decoding of codes on the Z-channel. In *IEEE International Conference on Communications*, volume 3, pages 1200–1205. IEEE.
- Besson, L. and Kaufmann, E. (2018). Multi-player bandits revisited. In *Proceedings of Algorithmic Learning Theory*, pages 56–92.
- Bistritz, I. and Leshem, A. (2018). Distributed multi-player bandits - a game of thrones approach. In *Advances in Neural Information Processing Systems*, pages 7222–7232.
- Boursier, E. and Perchet, V. (2019). SIC-MMAB: Synchronisation involves communication in multiplayer multi-armed bandits. In *Advances in Neural Information Processing Systems*, pages 12071–12080.
- Chen, P.-N., Lin, H.-Y., and Moser, S. M. (2013). Optimal ultrasmall block-codes for binary discrete memoryless channels. *IEEE Transactions on Information Theory*, 59(11):7346–7378.
- Gallager, R. G. (1968). *Information theory and reliable communication*. John Wiley & Sons, USA.
- Harper, F. M. and Konstan, J. A. (2015). The MovieLens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems*, 5(4).
- Kalathil, D., Nayyar, N., and Jain, R. (2014). Decentralized learning for multiplayer multiarmed bandits. *IEEE Transactions on Information Theory*, 60(4):2331–2345.
- Kaufmann, E. and Mehrabian, A. (2019). New algorithms for multiplayer bandits when arm means vary among players. *arXiv preprint arXiv:1902.01239*.
- Komiyama, J., Honda, J., and Nakagawa, H. (2015). Optimal regret analysis of Thompson sampling in stochastic multi-armed bandit problem with multiple plays. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 1152–1161, Lille, France.
- Lai, T. L. and Robbins, H. (1985). Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22.
- 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*.
- Nayyar, N., Kalathil, D., and Jain, R. (2016). On regret-optimal learning in decentralized multiplayer multiarmed bandits. *IEEE Transactions on Control of Network Systems*, 5(1):597–606.
- Proutiere, A. and Wang, P.-A. (2019). An optimal algorithm in multiplayer multi-armed bandits. *arXiv preprint arXiv:1909.13079*.
- Rosenski, J., Shamir, O., and Szlak, L. (2016). Multi-player bandits – a musical chairs approach. In *Proceedings of The 33rd International Conference on Machine Learning*, pages 155–163.
- Tallini, L. G., Al-Bassam, S., and Bose, B. (2002). On the capacity and codes for the Z-channel. In *Proceedings of the IEEE International Symposium on Information Theory*, page 422.