

Distributed Cooperative Decision-Making in Multiarmed Bandits: Frequentist and Bayesian Algorithms

Peter Landgren, Vaibhav Srivastava, and Naomi Ehrlich Leonard

Abstract— We study distributed cooperative decision-making under the explore-exploit tradeoff in the multiarmed bandit (MAB) problem. We extend state-of-the-art frequentist and Bayesian algorithms for single-agent MAB problems to cooperative distributed algorithms for multi-agent MAB problems in which agents communicate according to a fixed network graph. We rely on a running consensus algorithm for each agent’s estimation of mean rewards from its own rewards and the estimated rewards of its neighbors. We prove the performance of these algorithms and show that they asymptotically recover the performance of a centralized agent. Further, we rigorously characterize the influence of the communication graph structure on the decision-making performance of the group.

I. INTRODUCTION

Cooperative decision-making under uncertainty is ubiquitous in natural systems as well as in engineering networks. A fundamental feature of decision-making under uncertainty is the *explore-exploit* tradeoff: the decision-making agent needs to learn the unknown system parameters (exploration), while maximizing its parameter-dependent decision-making objective (exploitation).

Multiarmed bandit (MAB) problems are canonical formulations of the explore-exploit tradeoff. In a stochastic MAB problem a set of options (arms) is given. A stochastic reward with an unknown mean is associated with each option. A player can pick only one option at a time, and the player’s objective is to maximize the cumulative expected reward over a sequence of choices. In an MAB problem, the player needs to balance the tradeoff between learning the mean reward at each arm (exploration), and picking the arm with maximum mean reward (exploitation).

MAB problems are pervasive across a variety of scientific communities and have found application in diverse areas including control and robotics [1], [2], ecology [3], [4], and communications [5]. Despite the prevalence of the MAB problem, the research on MAB problems has primarily focused on policies for a single agent. The increasing importance of networked systems warrants the development of distributed algorithms for multiple communicating agents

This revision provides a correction to the original paper, which appeared in the Proceedings of the 2016 IEEE Conference on Decision and Control (CDC). The second statement of Proposition 1 and Theorem 1 are new from [TODO: cite the arXiv update for ECC] and Lemma 1 is new. These are used to prove regret bounds in Theorems 2 and 3.

This research has been supported by ONR grant N00014-14-1-0635, ARO grant W911NF-14-1-0431, and the DoD through the NDSEG Program.

P. Landgren and N. E. Leonard are with the Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ, USA, {landgren, naomi}@princeton.edu.

V. Srivastava is with the Department of Electrical and Computer Engineering, Michigan State University, East Lansing, MI, USA, vaibhav@egr.msu.edu

faced with MAB problems. In this paper, we build upon previous work by extending two popular single-agent algorithms for the stochastic MAB problem to the distributed multiple agent setting and analyze decision-making performance as a function of the network structure.

The MAB problem has been extensively studied (see [6] for a survey). In their seminal work, Lai and Robbins [7] established a logarithmic lower bound on the expected number of times a sub-optimal arm needs to be selected by an optimal policy in a frequentist setting. In another seminal work, Auer *et al.* [8] developed the upper confidence bound (UCB) algorithm for the stochastic MAB problem, which achieves the lower bound in [7] uniformly in time.

The MAB problem has also been studied in the Bayesian setting. Kaufmann *et al.* [9] proposed the Bayes-UCB algorithm and showed that it achieves Lai-Robbins’ lower bound for Bernoulli rewards and uninformative priors. Reverdy *et al.* [10] developed and analyzed the upper credible limit (UCL) algorithm for correlated multiarmed bandits by applying the approach of [9] to the case of Gaussian rewards.

The classical single-agent MAB problem was extended by Anantharam *et al.* [11] to the setting of a single-agent with multiple plays. Recently, researchers [12], [13], [5] have studied the decentralized multi-player MAB problem with no communication among agents. Kar *et al.* [14] investigated the multi-agent MAB problem in a leader-follower setting. Here, we use a running consensus algorithm [15] for assimilation of information. Running consensus, also known as dynamic consensus, has been used to study related collective decision-making models in social networks [16].

In the present paper we study the distributed cooperative MAB problem in which agents are faced with a stochastic MAB problem and communicate their information with their neighbors in an undirected and connected communication graph. We use a set of running consensus algorithms for cooperative estimation of the mean reward at each arm, and we design an arm selection heuristic that leads to an order-optimal performance for the group. The major contributions of this paper are as follows.

First, we propose and thoroughly analyze the coop-UCB2 and coop-UCL algorithms. We derive bounds on decision-making performance for the group and characterize the influence of the network structure on performance. To predict nodal performance, we propose a measure of nodal “explore-exploit centrality,” which depends on the location of the node in the graph.

Second, we demonstrate that the ordering of nodes by performance predicted by our explore-exploit centrality measure matches the order obtained using numerical simulations.

We also show that the incorporation of priors that are well-informative about the correlation structure markedly improve performance.

The remainder of the paper is organized as follows. In Section II we introduce the cooperative MAB problem. In Section III we recall a cooperative estimation algorithm. We review the coop-UCB algorithm in Section IV, and propose and analyze the improved coop-UCB2 and new coop-UCL algorithms. We illustrate our analytic results with numerical examples in Section VI. We conclude in Section VII.

II. COOPERATIVE MULTIARMED BANDIT PROBLEM

Consider an MAB problem with N arms and M decision-making agents. The reward associated with arm $i \in \{1, \dots, N\}$ is a random variable with an unknown mean m_i . Let the communication topology of agents be modeled by an undirected graph \mathcal{G} in which each node represents an agent and edges represent the communication between agents. Let $A \in \mathbb{R}^{M \times M}$ be the adjacency matrix associated with \mathcal{G} and let $L \in \mathbb{R}^{M \times M}$ be the corresponding Laplacian matrix. We assume that the graph \mathcal{G} is connected, i.e., there exists a path between each pair of nodes.

Let the k -th agent choose arm $i^k(t)$ at time $t \in \{1, \dots, T\}$ and receive a reward $r^k(t)$. The objective of each decision-maker k is to choose using its local information a sequence of arms $\{i^k(t)\}_{t \in \{1, \dots, T\}}$ such that the total expected cumulative reward $\sum_{k=1}^M \sum_{t=1}^T m_{i^k(t)}$ is maximized, where T is the horizon length of the sequential allocation process.

For an MAB problem, the expected *regret* of agent k at time t is defined by $R^k(t) = m_{i^*} - m_{i^k(t)}$, where $m_{i^*} = \max\{m_i \mid i \in \{1, \dots, N\}\}$. The collective objective of the M decision-makers can be equivalently defined as minimizing the expected cumulative regret defined by $\sum_{k=1}^M \sum_{t=1}^T R^k(t) = \sum_{k=1}^M \sum_{i=1}^N \Delta_i \mathbb{E}[n_i^k(T)]$, where $n_i^k(T)$ is the cumulative number of times arm i has been chosen by agent k until time T and $\Delta_i = m_{i^*} - m_i$ is the expected regret due to picking arm i instead of arm i^* . It is known that the regret of any algorithm for an MAB problem is asymptotically lower bounded by a logarithmic function of the horizon length T [7], [11], i.e., no algorithm can achieve an expected cumulative regret smaller than a logarithmic function of horizon length as $T \rightarrow \infty$.

In this paper, we focus on Gaussian rewards, i.e., the reward at arm i is sampled from a Gaussian distribution with mean m_i and variance σ_s^2 . We assume that the variance σ_s^2 is known and is the same at each arm. In the context of Gaussian rewards, the lower bound [11] on the expected number of times a suboptimal arm i is selected by a fusion center that has access to reward for each agent is

$$\sum_{k=1}^M \mathbb{E}[n_i^k(T)] \geq \left(\frac{2\sigma_s^2}{\Delta_i^2} + o(1) \right) \ln T. \quad (1)$$

In the following, we will design policies that sample a suboptimal arm i within a constant factor of the above bound.

III. COOPERATIVE ESTIMATION OF MEAN REWARDS

In this section we recall the algorithm for cooperative estimation of mean rewards proposed in our earlier work [17].

A. Cooperative Estimation Algorithm

For distributed cooperative estimation of the mean reward at each arm i , we employ two running consensus algorithms to estimate (i) total reward provided at the arm, and (ii) the total number of times the arm has been sampled.

Let $\hat{s}_i^k(t)$ and $\hat{n}_i^k(t)$ be agent k 's estimate of the total reward provided at arm i per unit agent and the total number of times arm i has been selected until time t per unit agent, respectively. Using $\hat{s}_i^k(t)$ and $\hat{n}_i^k(t)$ agent k can calculate $\hat{\mu}_i^k(t)$, the estimated empirical mean of arm i at time t as

$$\hat{\mu}_i^k(t) = \frac{\hat{s}_i^k(t)}{\hat{n}_i^k(t)}. \quad (2)$$

Let $i^k(t)$ be the arm sampled by agent k at time t and let $\xi_i^k(t) = \mathbb{1}(i^k(t) = i)$. $\mathbb{1}(\cdot)$ is the indicator function, here equal to 1 if $i^k(t) = i$ and 0 otherwise. For simplicity of notation we define $r_i^k(t)$ as the realized reward at arm i for agent k , which is a random variable sampled from $\mathcal{N}(m_i, \sigma_s^2)$, and the corresponding accumulated reward is $r^k(t) = r_i^k(t) \cdot \mathbb{1}(i^k(t) = i)$. Let P be a row stochastic matrix given by

$$P = \mathcal{I}_M - \frac{\kappa}{d_{\max}} L, \quad (3)$$

where \mathcal{I}_M is the identity matrix of order M , $\kappa \in (0, 1]$ is a step size parameter [18], $d_{\max} = \max\{\deg(k) \mid k \in \{1, \dots, M\}\}$, and $\deg(k)$ is the degree of node k .

The estimates $\hat{n}_i^k(t)$ and $\hat{s}_i^k(t)$ are updated locally using running consensus [15] as follows:

$$\hat{\mathbf{n}}_i(t) = P\hat{\mathbf{n}}_i(t-1) + P\xi_i(t), \quad (4)$$

$$\text{and } \hat{\mathbf{s}}_i(t) = P\hat{\mathbf{s}}_i(t-1) + P(\mathbf{r}_i(t) \circ \xi_i(t)), \quad (5)$$

where $\hat{\mathbf{n}}_i(t)$, $\hat{\mathbf{s}}_i(t)$, $\xi_i(t)$, and $\mathbf{r}_i(t)$ are vectors of $\hat{n}_i^k(t)$, $\hat{s}_i^k(t)$, $\xi_i^k(t)$, and $r_i^k(t)$, $k \in \{1, \dots, M\}$, respectively, and \circ denotes element-wise multiplication (Hadamard product).

B. Analysis of the Cooperative Estimation Algorithm

We now recall the performance of the estimation algorithm defined by (2–5). Let $n_i^{\text{cent}}(t) \equiv \frac{1}{M} \sum_{\tau=1}^t \mathbf{1}_M^\top \xi_i(\tau)$ be the total number of times arm i has been selected per unit agent until time t , and let $s_i^{\text{cent}}(t) \equiv \frac{1}{M} \sum_{\tau=1}^t \xi_i^\top(\tau) \mathbf{r}_i(\tau)$ be the total reward provided at arm i per unit agent until time t . Also, let λ_i denote the i -th largest eigenvalue of P , \mathbf{u}_i the eigenvector corresponding to λ_i , u_i^d the d -th entry of \mathbf{u}_i , and

$$\epsilon_n = \sqrt{M} \sum_{p=2}^M \frac{|\lambda_p|}{1 - |\lambda_p|}. \quad (6)$$

Note that $\lambda_1 = 1$ and $\mathbf{u}_1 = \mathbf{1}_M / \sqrt{M}$. Let us define

$$\nu_{pj}^{+\text{sum}} = \sum_{d=1}^M u_p^d u_j^d \mathbb{1}(u_p^k u_j^k \geq 0)$$

$$\text{and } \nu_{pj}^{-\text{sum}} = \sum_{d=1}^M u_p^d u_j^d \mathbb{1}(u_p^k u_j^k \leq 0).$$

We also define

$$a_{pj}(k) = \begin{cases} \nu_{pj}^{+\text{sum}} u_p^k u_j^k, & \text{if } \lambda_p \lambda_j \geq 0 \ \& \ u_p^k u_j^k \geq 0, \\ \nu_{pj}^{-\text{sum}} u_p^k u_j^k, & \text{if } \lambda_p \lambda_j \geq 0 \ \& \ u_p^k u_j^k \leq 0, \\ \nu_{pj}^{\text{max}} |u_p^k u_j^k|, & \text{if } \lambda_p \lambda_j < 0, \end{cases} \quad (7)$$

where $\nu_{pj}^{\text{max}} = \max\{|\nu_{pj}^{-\text{sum}}|, \nu_{pj}^{+\text{sum}}\}$. Furthermore, let

$$\epsilon_c^k = M \sum_{p=1}^M \sum_{j=2}^M \frac{|\lambda_p \lambda_j|}{1 - |\lambda_p \lambda_j|} a_{pj}(k). \quad (8)$$

We note that both ϵ_n and ϵ_c^k depend only on the topology of the communication graph and are measures of distributed cooperative estimation performance. We now recall the following results from [17].

Proposition 1 (Performance of cooperative estimation):

For the distributed estimation algorithm defined in (2–5), the following statements hold

(i) the estimate $\hat{n}_i^k(t)$ satisfies

$$n_i^{\text{cent}}(t) - \epsilon_n \leq \hat{n}_i^k(t) \leq n_i^{\text{cent}}(t) + \epsilon_n;$$

(ii) the following inequality holds for the estimate $\hat{n}_i^k(t)$ and the sequence $\{\xi_i^j(\tau)\}_{\tau \in \{1, \dots, t\}}$, $j \in \{1, \dots, M\}$

$$\sum_{\tau=1}^t \sum_{j=1}^M \left(\sum_{p=1}^M \lambda_p^{t-\tau+1} u_p^k u_p^j \right)^2 \xi_i^j(\tau) \leq \frac{\hat{n}_i^k(t) + \epsilon_c^k}{M}.$$

Theorem 1 (Estimator Deviation Bounds): For the estimates $\hat{s}_i^k(t)$ and $\hat{n}_i^k(t)$ obtained using equations (4) and (5), the following concentration inequality holds

$$\mathbb{P} \left(\frac{\hat{s}_i^k(t) - m_i \hat{n}_i^k(t)}{\left(\frac{1}{M} (\hat{n}_i^k(t) + \epsilon_c^k) \right)^{1/2}} > \delta \right) < \left[\frac{\ln(t + \epsilon_n)}{\ln(1 + \eta)} \right] \exp \left(\frac{-\delta^2}{2\sigma_s^2} G(\eta) \right), \quad (9)$$

where $\delta > 0$, $\eta > 0$, $G(\eta) = (1 - \eta^2/16)$, and ϵ_n and ϵ_c^k are defined in (6) and (8), respectively.

Proof: See [17]. ■

IV. FREQUENTIST COOPERATIVE DECISION-MAKING

In this section, we first review the coop-UCB algorithm proposed in our earlier work [17]. We then improve on this algorithm with a new algorithm: coop-UCB2. Unlike coop-UCB the improved algorithm does not require each agent to know the global graph structure. Finally, we compute bounds on the performance of the group for this algorithm as a function of the graph structure.

A. The coop-UCB Algorithm

The coop-UCB algorithm is analogous to the UCB algorithm [8], and uses a modified decision-making heuristic that captures the effect of the additional information an agent receives through communication with other agents as well as the rate of information propagation through the network.

The coop-UCB algorithm is initialized by each agent sampling each arm once and proceeds as follows. At time t each agent k selects the arm with maximum $Q_i^k(t-1) = \hat{\mu}_i^k(t-1) + C_i^k(t-1)$, where

$$C_i^k(t-1) = \sigma_s \sqrt{\frac{2\gamma}{G(\eta)} \cdot \frac{\hat{n}_i^k(t-1) + \epsilon_c^k}{M \hat{n}_i^k(t-1)} \cdot \frac{\ln(t-1)}{\hat{n}_i^k(t-1)}}, \quad (10)$$

$\gamma > 1$, $\eta \in (0, 4)$, and $G(\eta) = 1 - \eta^2/16$. Then, at each time t , each agent k updates its cooperative estimate of the mean reward at each arm using the distributed cooperative estimation algorithm described in (2–5).

The coop-UCB provides a distributed, cooperative solution to the MAB problem such that every agent in the network achieves logarithmic regret. However, the heuristic Q_i^k may be overly restrictive in the sense that it requires the agent k to know ϵ_c^k , which depends on the global graph structure. Further, agents with a relatively high ϵ_c^k are essentially forced to explore more while better positioned agents exploit, leading to wide disparities in performance across some networks. We will develop the coop-UCB2 algorithm that addresses these issues in the next section.

B. The coop-UCB2 Algorithm

The coop-UCB2 algorithm is initialized by each agent sampling each arm once and proceeds as follows. At time t each agent k selects the arm with maximum $Q_i^k(t-1) = \hat{\mu}_i^k(t-1) + C_i^k(t-1)$, where

$$C_i^k(t-1) = \sigma_s \sqrt{\frac{2\gamma}{G(\eta)} \cdot \frac{\hat{n}_i^k(t-1) + f(t-1)}{M \hat{n}_i^k(t-1)} \cdot \frac{\ln(t-1)}{\hat{n}_i^k(t-1)}}, \quad (11)$$

$f(t)$ is an increasing sublogarithmic function, $\gamma > 1$, $\eta \in (0, 4)$, and $G(\eta) = 1 - \eta^2/16$. Then, at each time t , each agent k updates its cooperative estimate of the mean reward at each arm using the distributed cooperative estimation algorithm described in (2–5). Note that the heuristic Q_i^k requires the agent k to know the total number of agents M , but not the global graph structure.

Theorem 2 (Regret of the coop-UCB2 Algorithm): For the coop-UCB2 algorithm and the cooperative Gaussian MAB problem, the number of times a suboptimal arm i is selected by all agents until time T satisfies

$$\begin{aligned} \sum_{k=1}^M \mathbb{E}[n_i^k(T)] &\leq 2 \sum_{k=1}^M (t_k^\dagger - 1) + \max \left\{ M, \right. \\ &\left. \left[M \epsilon_n + \frac{4\sigma_s^2 \gamma \ln T}{\Delta_i^2 G(\eta)} \left(1 + \sqrt{1 + \frac{\Delta_i^2 M G(\eta) f(T)}{2\gamma \sigma_s^2 \ln T}} \right) \right] \right\} \\ &\quad + \frac{2M}{\ln(1+\eta)} \left(\frac{1}{(\gamma-1)^2} + \frac{\gamma \ln(1+\epsilon_n)(1+\eta)}{\gamma-1} + 1 \right), \end{aligned}$$

where $t_k^\dagger = f^{-1}(\epsilon_c^k)$.

Proof: We proceed similarly to [8]. The number of selections of a suboptimal arm i by all agents until time T is

$$\begin{aligned} \sum_{k=1}^M n_i^k(T) &\leq \sum_{k=1}^M \sum_{t=N+1}^T \mathbb{1}(Q_i^k(t-1) \geq Q_{i^*}^k(t-1)) + M \\ &\leq \max\{M, A\} + \sum_{k=1}^M \sum_{t=N}^{T-1} \mathbb{1}(Q_i^k(t) \geq Q_{i^*}^k(t), M n_i^{\text{cent}}(t) \geq A), \end{aligned} \quad (12)$$

where $A > 0$ is a constant that will be chosen later.

At a given time $t+1$ an individual agent k will choose a suboptimal arm only if $Q_i^k(t) \geq Q_{i^*}^k(t)$. For this condition

to be true at least one of the following three conditions must hold:

$$\hat{\mu}_{i^*}^k(t) \leq m_{i^*} - C_{i^*}^k(t) \quad (13)$$

$$\hat{\mu}_i^k(t) \geq m_i + C_i^k(t) \quad (14)$$

$$m_{i^*} < m_i + 2C_i^k(t). \quad (15)$$

We now bound the probability that (14) holds using Theorem 1:

$$\begin{aligned} & \mathbb{P}\left((14) \text{ holds} \mid t \geq t_k^\dagger\right) \\ &= \mathbb{P}\left(\left.\frac{\hat{s}_i^k(t) - m_i \hat{n}_i^k(t)}{\sqrt{\frac{1}{M}(\hat{n}_i^k(t) + f(t))}} \geq \sigma_s \sqrt{\frac{2\gamma \ln(t)}{G(\eta)}} \right| t \geq t_k^\dagger\right) \\ &\leq \mathbb{P}\left(\left.\frac{\hat{s}_i^k(t) - m_i \hat{n}_i^k(t)}{\sqrt{\frac{1}{M}(\hat{n}_i^k(t) + \epsilon_c^k)}} \geq \sigma_s \sqrt{\frac{2\gamma \ln(t)}{G(\eta)}} \right| t \geq t_k^\dagger\right) \\ &\leq \left(\frac{\ln(t)}{\ln(1+\eta)} + \frac{\ln(1+\epsilon_n)}{\ln(1+\eta)} + 1\right) \frac{1}{t^\gamma}. \end{aligned}$$

It also follows analogously that

$$\mathbb{P}\left((13) \text{ holds} \mid t \geq t_k^\dagger\right) \leq \left(\frac{\ln(t)}{\ln(1+\eta)} + \frac{\ln(1+\epsilon_n)}{\ln(1+\eta)} + 1\right) \frac{1}{t^\gamma}.$$

We now examine the event (15).

$$\begin{aligned} & m_{i^*} < m_i + 2C_i^k(t) \\ \implies & \hat{n}_i^k(t)^2 \frac{\Delta_i^2 MG(\eta)}{8\sigma_s^2} - \gamma \hat{n}_i^k(t) \ln(t) - \gamma f(t) \ln(t) < 0. \end{aligned} \quad (16)$$

The quadratic equation (16) can be solved to find the roots, and if $\hat{n}_i^k(t)$ is greater than the larger root the inequality will never hold. Solving the quadratic equation (16), we obtain that event (15) does not hold if

$$\begin{aligned} \hat{n}_i^k(t) &\geq \frac{4\sigma_s^2 \gamma \ln(t)}{\Delta_i^2 MG(\eta)} + \sqrt{\left(\frac{4\sigma_s^2 \gamma \ln(t)}{\Delta_i^2 MG(\eta)}\right)^2 + \frac{8\sigma_s^2 f(t) \gamma \ln(t)}{\Delta_i^2 MG(\eta)}} \\ &= \frac{4\sigma_s^2 \gamma \ln t}{\Delta_i^2 MG(\eta)} \left(1 + \sqrt{1 + \frac{\Delta_i^2 MG(\eta) f(t)}{2\sigma_s^2 \gamma \ln t}}\right). \end{aligned}$$

Now, we set $A = \left[M\epsilon_n + \frac{4\sigma_s^2 \gamma \ln T}{\Delta_i^2 G(\eta)} \left(1 + \sqrt{1 + \frac{\Delta_i^2 MG(\eta) f(T)}{2\sigma_s^2 \gamma \ln T}}\right)\right]$. It follows from monotonicity of $f(t)$ and $\ln(t)$ and statement (i) of Proposition 1 that event (15) does not hold if $Mn_i^{\text{cent}}(t) > A$.

Therefore, from (12) we see that

$$\begin{aligned} \sum_{k=1}^M \mathbb{E}[n_i^k(T)] &\leq \max\{M, A\} + 2 \sum_{k=1}^M \sum_{t=1}^{t_k^\dagger-1} 1 \\ &\quad + \frac{2}{\ln(1+\eta)} \sum_{k=1}^M \sum_{t=t_k^\dagger}^T \left(\frac{\ln(t)}{t^\gamma} + \frac{\ln((1+\epsilon_n)(1+\eta))}{t^\gamma}\right) \\ &\leq \max\{M, A\} + 2 \sum_{k=1}^M (t_k^\dagger - 1) \\ &\quad + \frac{2M}{\ln(1+\eta)} \sum_{t=1}^T \left(\frac{\ln(t)}{t^\gamma} + \frac{\ln((1+\epsilon_n)(1+\eta))}{t^\gamma}\right) \\ &\leq \max\{M, A\} + 2 \sum_{k=1}^M (t_k^\dagger - 1) + \frac{2M}{\ln(1+\eta)} \left(\frac{1}{(\gamma-1)^2} \right. \\ &\quad \left. + \frac{\gamma \ln(1+\epsilon_n)(1+\eta)}{\gamma-1} + 1\right), \end{aligned}$$

completing the theorem. \square

Remark 1 (Asymptotic Regret for coop-UCB2): In the limit $t \rightarrow +\infty$, $\frac{f(t)}{\ln(t)} \rightarrow 0^+$, $\eta \rightarrow 0$, and

$$\sum_{k=1}^M \mathbb{E}[n_i^k(T)] \leq \left(\frac{8\sigma_s^2 \gamma}{\Delta_i^2} + o(1)\right) \ln T,$$

and we recover the upper bound on regret for a centralized agent as given in (1) within a constant factor. \square

Remark 2 (Performance of Individual Agents):

Theorem 2 provides bounds on the performance of the group as a function of the graph structure, and the logarithmic portion of the bound is independent of agent location. However, the constant factor is dependent on the agent's position in the network since it depends on ϵ_c^k . In this sense, ϵ_c^k can be thought of as a measure of ‘‘explore-exploit centrality,’’ which indicates that agents with a higher ϵ_c^k will contribute more to the group's regret. \square

V. BAYESIAN COOPERATIVE DECISION-MAKING

In this section, we extend the coop-UCB2 algorithm to a Bayesian setting and develop the coop-UCL algorithm. The Bayesian setting allows us to model correlated bandits and incorporate a priori knowledge about reward and correlation structure in the Bayesian prior. We first recall the UCL algorithm proposed in [10], [9] and extend it to the cooperative setting. We then analyze the performance of this algorithm for an uninformative prior.

A. The UCL Algorithm

The UCL algorithm developed in [10] applies the approach of Bayes-UCB [9] to correlated Gaussian bandits. The UCL algorithm at each time computes the posterior distribution of mean rewards at each option and then computes the $(1-1/Kt^a)$ upper-credible-limit for each arm, i.e., an upper bound that holds with probability $(1-1/Kt^a)$ where $K = \sqrt{2\pi e}$, $\gamma > 1$, and $a = 4/3\gamma$. The algorithm chooses the arm with highest upper credible limit. For Gaussian rewards, the $(1-1/Kt^a)$ upper-credible-limit can be written as

$$Q_i(t) = \nu_i(t) + \sigma_i(t) \Phi^{-1}(1-1/Kt^a), \quad (17)$$

where $\nu_i(t)$ is the posterior mean and $\sigma_i(t)$ the posterior standard deviation of mean reward at time t . $\Phi^{-1}(\cdot)$ is the standard Gaussian inverse cumulative distribution function.

Let the prior on rewards from each arm be multivariate Gaussian with mean vector $\nu_0 \in \mathbb{R}^N$ and covariance matrix $\Sigma_0 \in \mathbb{R}^{N \times N}$. Then, the posterior mean and covariance of mean reward at time t can be computed using the following recursive update rule [19]:

$$\begin{aligned} \mathbf{q}(t) &= \frac{r(t)\phi(t)}{\sigma_s^2} + \Lambda(t-1)\nu(t-1) \\ \Lambda(t) &= \frac{\phi(t)^\top \phi(t)}{\sigma_s^2} + \Lambda(t-1), \quad \Sigma(t) = \Lambda(t)^{-1} \\ \nu(t) &= \Sigma(t)\mathbf{q}(t), \end{aligned} \quad (18)$$

where $\phi(t)$ and $\nu(t)$ are column vectors of $\phi_i(t)$ and $\nu_i(t)$, respectively, and $\phi_i(t)$ is the indicator function of selecting arm i at time t . The update equation (18) can be reduced to

$$\begin{aligned} \nu(t) &= (\Lambda_0 + \Gamma(t)^{-1})^{-1}(\Gamma(t)^{-1}\boldsymbol{\mu}(t) + \Lambda_0\nu_0) \\ \Lambda(t) &= \Lambda_0 + \Gamma(t)^{-1}, \quad \Sigma(t) = (\Lambda(t))^{-1}, \end{aligned} \quad (19)$$

where $\Lambda_0 = \Sigma_0^{-1}$, $\Gamma(t)$ is a diagonal matrix with entries $\frac{\sigma_s^2}{n_i(t)}$, and $\boldsymbol{\mu}(t)$ is the vector of $\mu_i(t)$, which is the empirical mean of rewards from arm $i \in \{1, \dots, N\}$ until time t . Note that diagonal entries of $\Sigma(t)$ are $(\sigma_i(t))^2$, $i \in \{1, \dots, N\}$.

B. The coop-UCL Algorithm

We now extend the UCL algorithm to the distributed cooperative setting and propose the coop-UCL algorithm. In the coop-UCL algorithm, each agent first computes an approximate posterior distribution of mean rewards conditioned on rewards obtained by all the agents. To this end, each agent uses the approximate frequentist estimator $\hat{\mu}_i^k$ from Section III in update equation (19).

Let the prior of agent k be a multivariate Gaussian distribution with mean ν_0^k and covariance Σ_0^k . Let $\hat{\Sigma}^k(t)$ and $\hat{\nu}^k(t)$ be the estimated covariance matrix and posterior mean at time t , respectively. Then, the coop-UCL algorithm performs cooperative approximate Bayesian estimation:

$$\begin{aligned} \hat{\nu}^k(t) &= (\Lambda_0^k + \Gamma^k(t)^{-1})^{-1}(\Gamma^k(t)^{-1}\hat{\boldsymbol{\mu}}^k(t) + \Lambda_0^k\nu_0^k) \\ \hat{\Lambda}^k(t) &= \Lambda_0^k + \Gamma^k(t)^{-1}, \quad \hat{\Sigma}^k(t) = (\hat{\Lambda}^k(t))^{-1}, \end{aligned} \quad (20)$$

where $\Gamma^k(t)$ is a diagonal matrix with diagonal entries $\sigma_s^2/M\hat{n}_i^k(t)$, $i \in \{1, \dots, N\}$, and $\Lambda_0^k = (\Sigma_0^k)^{-1}$.

After computing $\hat{\nu}^k(t-1)$ and $\hat{\Sigma}^k(t-1)$, the coop-UCL algorithm at time t requires each agent k to choose the option with maximum $(1 - \alpha(t))$ -upper-credible-limit given by

$$Q_i^k(t-1) = \hat{\nu}_i^k(t-1) + \hat{\sigma}_i^k(t-1)\Phi^{-1}(1 - \alpha(t-1)), \quad (21)$$

where $\alpha(t)$ is defined such that

$$\Phi^{-1}(1 - \alpha(t)) = \sqrt{\frac{\hat{n}_i^k(t) + f(t)}{G(\eta)\hat{n}_i^k(t)}}\Phi^{-1}\left(1 - \frac{1}{Kt^a}\right),$$

where $\hat{\nu}_i^k(t)$ is the i -th entry of $\hat{\nu}^k(t)$, $(\hat{\sigma}_i^k(t))^2$ is the i -th diagonal entry of $\hat{\Sigma}^k(t)$, $K = \sqrt{2\pi e}$, $\gamma > 1$, and $a = 4/3\gamma$.

C. Regret of the coop-UCL Algorithm

We now derive bounds on the expected cumulative regret for each agent using the coop-UCL algorithm with uninformative priors for each agent. For an uninformative prior, $\Lambda_0^k = 0$, for each $k \in \{1, \dots, M\}$, and consequently, $\hat{\nu}^k(t) = \hat{\mu}^k(t)$ and $\hat{\Sigma}^k(t) = \Gamma^k(t)$. In addition, we first present a bound on $\Phi^{-1}(\cdot)$.

Lemma 1 (Inverse Gaussian CDF Bounds): For the standard normal random variable z and the associated inverse cumulative distribution function $\Phi^{-1}(\cdot)$, the following hold for any $\alpha \in [0, 0.5]$, $t \in \mathbb{N}$ and $a > 1$:

$$\Phi^{-1}(1 - \alpha) \leq \sqrt{-2\log(\alpha)}$$

$$\Phi^{-1}(1 - \alpha) > \sqrt{-\log(2\pi\alpha^2(1 - \log(2\pi\alpha^2)))}$$

$$\Phi^{-1}\left(1 - \frac{1}{\sqrt{2\pi e t^a}}\right) > \sqrt{\nu \log t^a},$$

for $0 < \nu \leq 1.59$.

Proof: The first inequality can be found in [20]. The second inequality was established in [10]. To establish the last inequality, it suffices to show that

$$-\log\left(\frac{1}{e t^2}\left(1 - \log\left(\frac{1}{e t^2}\right)\right)\right) - \nu \log t \geq 0,$$

for $0 < \nu \leq 1.59$. The left hand side of the above inequality is

$$g(t) := 1 - \log 2 + (2 - \nu) \log t - \log(1 + \log t).$$

It can be verified that g admits a unique minimum at $t = e^{(\nu-1)/(2-\nu)}$ and the minimum value is $\nu - \log 2 + \log(2-\nu)$, which is positive for $0 < \nu \leq 1.59$. \square

In the following, we select $\nu = 3/2$.

Theorem 3 (Regret of the coop-UCL Algorithm): For the Gaussian MAB problem and the coop-UCL algorithm with uninformative priors for each agent, the number of times a suboptimal arm i is selected by all agents until time T satisfies

$$\begin{aligned} \sum_{k=1}^M \mathbb{E}[n_i^k(T)] &\leq 2 \sum_{k=1}^M (t_k^\dagger - 1) + \max\left\{M, \left\lceil M\epsilon_n \right. \right. \\ &\quad \left. \left. + \left(\frac{4\sigma_s^2 \ln KT^a}{\Delta_i^2 MG(\eta)} \left(1 + \sqrt{1 + \frac{\Delta_i^2 MG(\eta)}{2\sigma_s^2} \frac{f(T)}{\ln KT^a}} \right) \right) \right\rceil \right\} \\ &\quad + \frac{2M}{\ln(1+\eta)} \left(\frac{1}{(\gamma-1)^2} + \frac{\gamma \ln(1+\epsilon_n)(1+\eta)}{\gamma-1} + 1 \right) \end{aligned}$$

where $t_k^\dagger = f^{-1}(\epsilon_c^k)$.

Proof: For uninformative priors, coop-UCL is analogous to coop-UCB2 with $C_i^k(t) = \hat{\sigma}_i^k(t)\Phi^{-1}(1 - \alpha(t))$, where $\hat{\sigma}_i^k(t) = \sigma_s/\sqrt{M\hat{n}_i^k(t)}$. Similar to the proof of Theorem 2,

we first note that for (14) simple manipulations lead to

$$\begin{aligned} \frac{\hat{s}_i^k(t) - m_i \hat{n}_i^k(t)}{\sqrt{\hat{n}_i^k(t) + f(t)}} &\geq \frac{\sigma_s}{\sqrt{MG(\eta)}} \Phi^{-1} \left(1 - \frac{1}{Kt^a} \right) \\ &> \frac{\sigma_s}{\sqrt{MG(\eta)}} \sqrt{\frac{3a}{2} \ln t} \\ &= \frac{\sigma_s}{\sqrt{MG(\eta)}} \sqrt{2 \ln t^\gamma} \end{aligned} \quad (22)$$

where (22) follows from Lemma 1 for $K = \sqrt{2\pi e}$.

Using Theorem 1 we get that

$$\begin{aligned} &\mathbb{P} \left((14) \text{ holds } | t \geq t_k^\dagger \right) \\ &\leq \mathbb{P} \left(\frac{\hat{s}_i^k(t) - m_i \hat{n}_i^k(t)}{\sqrt{\frac{1}{M} (\hat{n}_i^k(t) + f(t))}} \geq \sigma_s \sqrt{\frac{2\gamma \ln(t)}{G(\eta)}} \mid t \geq t_k^\dagger \right) \\ &\leq \left(\frac{\ln(t)}{\ln(1+\eta)} + \frac{\ln(1+\epsilon_n)}{\ln(1+\eta)} + 1 \right) \frac{1}{t^\gamma} \end{aligned}$$

resulting in sub-logarithmic regret as in Theorem 2.

We now examine the event (15). Following the argument in the proof of Theorem 2 and using the upper bound on $\Phi^{-1}(\cdot)$ from Lemma 1, we obtain that the event (15) does not hold if

$$\hat{n}_i^k(t) \geq \frac{4\sigma_s^2 \ln Kt^a}{\Delta_i^2 MG(\eta)} \left(1 + \sqrt{1 + \frac{\Delta_i^2 MG(\eta)}{2\sigma_s^2} \frac{f(t)}{\ln Kt^a}} \right).$$

We set $A = \left[M\epsilon_n + \frac{4\sigma_s^2 \ln Kt^a}{\Delta_i^2 MG(\eta)} \left(1 + \sqrt{1 + \frac{\Delta_i^2 MG(\eta)}{2\sigma_s^2} \frac{f(T)}{\ln Kt^a}} \right) \right]$ and the theorem follows by proceeding similarly to the proof of Theorem 2. \square

VI. NUMERICAL ILLUSTRATIONS

In this section, we elucidate our theoretical analyses from the previous sections with numerical examples. We first demonstrate that the ordering of the performance of nodes obtained through numerical simulations is identical to the ordering predicted by the nodal explore-exploit centrality measure: the larger the ϵ_c^k the lower the performance. We then investigate the effect of the graph connectivity on the performance of agents in random graphs.

For all simulations we consider a 10-armed bandit problem with mean rewards drawn from a normal random distribution with mean 75 and standard deviation 25. The sampling standard deviation is $\sigma_s = 30$. These parameters were selected to give illustrative results within the displayed time horizon, but the relevant conclusions hold across a wide variation of parameters. The simulations used $f(t) = \sqrt{\ln t}$.

Example 1 (Regret on Fixed Graphs): Consider the set of agents communicating according to the graph in Fig. 1(a) and using the coop-UCB2 algorithm to handle the explore-exploit tradeoff in the distributed cooperative MAB problem. The values of ϵ_c^k for nodes 1, 2, 3, and 4 are 2.31, 2.31, 0, and 5.43, respectively. As noted in Remark 2, agent 3 should have the lowest regret, agents 1 and 2 should have equal and intermediate regret, and agent 4 should have the highest regret. These predictions are validated in our simulations






Fig. 1: (a) Simulation results comparing expected cumulative regret for different agents in the fixed network shown, using P as in (3) and $\kappa = \frac{d_{\max}}{d_{\max}-1}$. Note that agents 1 and 2 have nearly identical expected regret. (b) Simulation results of expected cumulative regret for several different MAB algorithms using the fixed network shown in Fig. 1(a). (c) Simulation results of expected cumulative regret as a function of normalized ϵ_c^k for nodes in ER graphs at $T = 500$. Also shown in red is the best linear fit.

shown in Fig. 1(a). The expected cumulative regret in our simulations is computed using 5000 Monte-Carlo runs.

Fig. 1(b) demonstrates the relative performance differences between coop-UCB, coop-UCB2, coop-UCL, and single agent UCB with the same run parameters. Here the coop-UCL algorithm is shown with an informative prior and no correlation structure. Each agent in the coop-UCL simulation shown here has $\Sigma_0 = 625 \cdot \mathcal{I}_M$ and $\nu_0 = 75 \cdot \mathbf{1}_M$. The use of priors markedly improves performance.

We now explore the effect of ϵ_c^k on the performance of an agent in an Erdős-Rényi (ER) random graph. ER graphs are a widely used class of random graphs where any two agents are connected with a given probability ρ [21].

Example 2 (Regret on Random Graphs): Consider a set of 10 agents communicating according to an ER graph and using the coop-UCB2 algorithm to handle the explore-exploit tradeoff in the aforementioned MAB problem. In our simulations, we consider 100 connected ER graphs, and for each ER graph we compute the expected cumulative regret of agents using 1000 Monte-Carlo simulations with $\rho = \ln(10)/10$, P as in (3), and $\kappa = d_{\max}/(d_{\max} - 1)$. We show the behavior of the expected cumulative regret of each agent as a function of the normalized ϵ_c^k in Fig. 1(c). It is evident that increased ϵ_c^k results in a sharp decrease in performance. Conversely, low ϵ_c^k is indicative of better performance. This disparity is due to the relative scarcity of information at nodes that are in general less “central.”

VII. FINAL REMARKS

In this paper we used the distributed multi-agent MAB problem to explore cooperative decision-making in networks. We designed the coop-UCB2 and coop-UCL algorithms, which are frequentist and Bayesian distributed algorithms, respectively, in which agents do not need to know the graph structure. We proved bounds on performance, showing order-optimal performance for the group. Additionally, we investigated the performance of individual agents in the network as a function of the graph topology, using a proposed measure of nodal explore-exploit centrality.

Future research directions include rigorously exploring other communications schemes, which may offer better performance or be more suitable for modeling certain networked systems. It will be important to consider the tradeoff between communication frequency and performance as well as the presence of noisy communications.

REFERENCES

- [1] V. Srivastava, P. Reverdy, and N. E. Leonard. Surveillance in an abruptly changing world via multiarmed bandits. In *IEEE CDC*, pages 692–697, 2014.
- [2] M. Y. Cheung, J. Leighton, and F. S. Hover. Autonomous mobile acoustic relay positioning as a multi-armed bandit with switching costs. In *IEEE/RSS Int. Conf. on Intelligent Robots and Systems*, pages 3368–3373.
- [3] J. R. Krebs, A. Kacelnik, and P. Taylor. Test of optimal sampling by foraging great tits. *Nature*, 275(5675):27–31, 1978.
- [4] V. Srivastava, P. Reverdy, and N. E. Leonard. On optimal foraging and multi-armed bandits. In *Allerton Conference on Comm., Control, and Computing*, pages 494–499, 2013.

- [5] A. Anandkumar, N. Michael, A. K. Tang, and A. Swami. Distributed algorithms for learning and cognitive medium access with logarithmic regret. *IEEE Journal on Selected Areas in Communications*, 29(4):731–745, 2011.
- [6] S. Bubeck and N. Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Machine Learning*, 5(1):1–122, 2012.
- [7] T. L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- [8] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2):235–256, 2002.
- [9] E. Kaufmann, O. Cappé, and A. Garivier. On Bayesian upper confidence bounds for bandit problems. In *Int. Conf. on Artificial Intelligence and Statistics*, pages 592–600, 2012.
- [10] P. B. Reverdy, V. Srivastava, and N. E. Leonard. Modeling human decision making in generalized Gaussian multiarmed bandits. *Proceedings of the IEEE*, 102(4):544–571, 2014.
- [11] V. Anantharam, P. Varaiya, and J. Walrand. Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays-part I: I.I.D. rewards. *IEEE Transactions on Automatic Control*, 32(11):968–976, Nov 1987.
- [12] D. Kalathil, N. Nayyar, and R. Jain. Decentralized learning for multiplayer multiarmed bandits. *IEEE Transactions on Information Theory*, 60(4):2331–2345, 2014.
- [13] Y. Gai and B. Krishnamachari. Distributed stochastic online learning policies for opportunistic spectrum access. *IEEE Transactions on Signal Processing*, 62(23):6184–6193, 2014.
- [14] S. Kar, H. V. Poor, and S. Cui. Bandit problems in networks: Asymptotically efficient distributed allocation rules. In *IEEE CDC and ECC*, pages 1771–1778, 2011.
- [15] P. Braca, S. Marano, and V. Matta. Enforcing consensus while monitoring the environment in wireless sensor networks. *IEEE Transactions on Signal Processing*, 56(7):3375–3380, 2008.
- [16] V. Srivastava and N. E. Leonard. Collective decision-making in ideal networks: The speed-accuracy trade-off. *IEEE Transactions on Control of Network Systems*, 1(1):121–132, 2014.
- [17] P. Landgren, V. Srivastava, and N. E. Leonard. On distributed cooperative decision-making in multiarmed bandits. In *European Control Conference*, Aalborg, Denmark, 2016.
- [18] R. Olfati-Saber and R. M. Murray. Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on Automatic Control*, 49(9):1520–1533, 2004.
- [19] S. M. Kay. *Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory*. Prentice Hall, 1993.
- [20] M. Abramowitz and I. A. Stegun, editors. *Handbook of Mathematical Functions: with Formulas, Graphs, and Mathematical Tables*. Dover Publications, 1964.
- [21] B. Bollobás. *Random Graphs*. Springer, 1998.