

Finite-time Analysis of the Distributed Detection Problem

Shahin Shahrampour, Alexander Rakhlin and Ali Jadbabaie

Abstract—This paper addresses the problem of distributed detection in fixed and switching networks. A network of agents observe partially informative signals about the *unknown* state of the world. Hence, they collaborate with each other to identify the true state. We propose an update rule building on distributed, stochastic optimization methods. Our main focus is on the *finite-time* analysis of the problem. For fixed networks, we bring forward the notion of Kullback-Leibler cost to measure the efficiency of the algorithm versus its centralized analog. We bound the cost in terms of the network size, spectral gap and relative entropy of agents' signal structures. We further consider the problem in random networks where the structure is realized according to a *stationary* distribution. We then prove that the convergence is exponentially fast (with high probability), and the *non-asymptotic* rate scales inversely in the spectral gap of the expected network.

I. INTRODUCTION

Distributed detection, estimation, prediction and optimization has been an interesting subject of study in science and engineering for so many years [1]–[6]. Decentralized algorithms are ubiquitous in numerous scenarios ranging from sensor and robot to social and economic networks [7]–[10]. In these applications, a network of agents aim to accomplish a team task for which they only have *partial* knowledge. Therefore, they must communicate with each other to benefit from local observations. In fact, the global spread of information in the network is sufficient for agents to achieve the network goal. In most of these schemes, *consensus* protocols are employed to converge agents to a common value [11], [12].

We would like to focus on the distributed detection problem in this work. The problem was first considered in the case that each agent sends its private observations to a fusion center [1], [2], [7]. The fusion center then faces a classical hypothesis testing to make a decision over the value of the parameter. While collecting data is decentralized in these cases, the decision making part is done in a centralized fashion.

Distributed detection has also been considered in works where no fusion center is necessary [13]–[15]. These works mostly focus on asymptotic analysis of the problem. Cattivelli et al. [13] develop a fully distributed algorithm based on

This work was supported by ONR BRC Program on Decentralized, Online Optimization.

Shahin Shahrampour and Ali Jadbabaie are with the Department of Electrical and Systems Engineering at the University of Pennsylvania, Philadelphia, PA 19104 USA. (e-mail: shahin@seas.upenn.edu; jadbabai@seas.upenn.edu).

Alexander Rakhlin is with the Department of Statistics at the University of Pennsylvania, Philadelphia, PA 19104 USA. (e-mail: rakhlin@wharton.upenn.edu).

the connection of Neyman-Pearson detection and minimum-variance estimation to tackle the problem. Jakovetić et al. [14] propose a consensus+innovations algorithm for detection in the case of *Gaussian* observations. Their method possesses an *asymptotic* exponential error rate even under noisy communication of agents. In [15], the consensus+innovations method is extended to *generic* (non-Gaussian) observations over random networks.

Another social learning model inspiring several works in the literature is proposed by Jadbabaie et al. [16]. In this model, a fixed true *state* or *hypothesis* is aimed to be recovered by a network of agents. The state might be a decision or opinion of an individual, the correct price of a product or any quantity of interest. The state is assumed to belong to a *finite* collection of states. Agents receive a stream of *private* signals about the true state. However, the signals do *not* provide enough information for agents to detect the underlying state of the world. Hence, agents engage in *local* interactions to compensate for their imperfect knowledge about the states. Numerous works build on this approach to study social learning and distributed detection [17]–[19]. The focus of these works is the asymptotic behavior of the model. Though appealing in certain cases, asymptotic analysis might not unveil all important parameters for learning. Therefore, *finite-time* analysis of the problem is also an interesting complementary direction to study.

Let us first provide more details on the asymptotic analysis of the problem. The authors in [16] propose a non-Bayesian update rule for social learning applications. Each individual updates her Bayesian prior, and averages the result with the opinion of her neighbors. Under mild technical assumptions, the authors prove that agents' beliefs converge to a delta distribution on the true state. The convergence occurs exponentially fast and in almost sure sense. Shahrampour and Jadbabaie [17] consider an optimization-based approach to the problem, inspired by the work of Duchi et al. [20] on distributed dual averaging. Their proposed update rule is essentially a distributed stochastic optimization. They establish that the sequence of beliefs is weakly consistent (convergence in probability) when agents employ a *gossip* communication protocol. A communication-efficient variant of the problem is considered in [19] in which agents switch between Bayesian and non-Bayesian regimes to detect the true state. Furthermore, Lalitha et al. [18] develop a strategy where agents perform a Bayesian update, and geometrically average the posteriors in their local neighborhood. The authors then provide the convergence and rate analysis of their method. In [16]–[19], the convergence occurs exponentially fast, and the *asymptotic* rate is characterized via the *relative*

entropy of individuals' signal structures and their *eigenvector centralities* (in directed networks).

The asymptotic analysis presented in the works above only characterizes the dominant factors in the long run. In real world applications, however, the *decision* has to be made in a *finite* time horizon. Hence, it is important to study the finite-time version of the problem to understand the role of *network parameters* in detection quality. Serving this goal, the works of [21]–[23] study the *non-asymptotic* problem. While the network structure in [21] is assumed to be fixed, the works of [22], [23] address switching protocols that are deterministic.

In this paper we build on the results of [21] to extend our setup to random networks. For fixed networks, we define the notion of *Kullback-Leibler (KL)* cost to compare the performance of distributed setting to its centralized counterpart. We provide an upper bound on the cost in terms of the *spectral gap* of the network. Our cost bound is independent of time with *high probability*. We further consider the stochastic communication setting in which the structure is realized randomly at each iteration. We prove that in this case, the rate scales inversely in the spectral gap of the expected network. Our result also guarantees the almost sure learning in random networks.

The rest of the paper is organized as follows: we formalize the notation, problem and the distributed detection scheme in Section II. Section III is devoted to the finite-time analysis of the problem. We consider both fixed and switching network topologies and provide non-asymptotic results. Section IV provides the concluding remarks, and proofs are included in the Appendix.

II. THE PROBLEM DESCRIPTION AND ALGORITHM

A. Notation

We use the following notation throughout the paper:

$[n]$	The set $\{1, 2, \dots, n\}$ for any integer n
x^\top	Transpose of the vector x
$x(k)$	The k -th element of vector x
e_k	Delta distribution on k -th component
$\mathbb{1}$	Vector of all ones
$\ \mu - \pi\ _{\text{TV}}$	Total variation distance between $\mu, \pi \in \Delta_m$
$D_{\text{KL}}(\mu \pi)$	KL-divergence of $\pi \in \Delta_m$ from $\mu \in \Delta_m$
$\sigma_i(W)$	The i -th largest singular value of matrix W

Furthermore, all the vectors are assumed to be in column format.

B. Observation Model

We consider a setting in which $\Theta = \{\theta_1, \theta_2, \dots, \theta_m\}$ denotes a finite set of *states* of the world. A network of n agents seek the *unique*, true state of the world $\theta_1 \in \Theta$ (unknown of the problem). At each time $t \in [T]$, the belief of agent i is represented by $\mu_{i,t} \in \Delta_m$, where Δ_m is a probability distribution over the set Θ . For instance, $\mu_{i,0} \in \Delta_m$ denotes the prior belief of agent $i \in [n]$ about the states of the world, and it is assumed to be uniform¹.

¹The assumption of uniform prior only avoids notational clutter. The analysis in the paper holds for any prior with full support.

The detection model is defined with a conditional likelihood function $\ell(\cdot|\theta_k)$ parametrized by some state of the world $\theta_k \in \Theta$. For each $i \in [n]$, $\ell_i(\cdot|\theta_k)$ denotes the i -th marginal of $\ell(\cdot|\theta_k)$, and we use the vector representation $\ell_i(\cdot|\theta) = [\ell_i(\cdot|\theta_1), \dots, \ell_i(\cdot|\theta_m)]^\top$ to stack all states. At each time $t \in [T]$, the signal $s_t = (s_{1,t}, s_{2,t}, \dots, s_{n,t}) \in \mathcal{S}_1 \times \dots \times \mathcal{S}_n$ is generated based on the true state θ_1 . As a result, for each $i \in [n]$, the signal $s_{i,t} \in \mathcal{S}_i$ is a sample drawn from the likelihood $\ell(\cdot|\theta_1)$ where \mathcal{S}_i is the sample space.

The signals are i.i.d. over time horizon, and also the marginals are independent across agents, i.e., $\ell(\cdot|\theta_k) = \prod_{i=1}^n \ell_i(\cdot|\theta_k)$ for any $k \in [m]$. For simplicity, we define the log-marginal $\psi_{i,t} := \log \ell_i(s_{i,t}|\theta)$ which is a sample corresponding to $\Psi_i := \log \ell_i(\cdot|\theta)$ for any $i \in [n]$.

A1. We assume that all log-marginals are uniformly bounded such that $\|\psi_{i,t}\|_\infty \leq B$ for any $s_{i,t} \in \mathcal{S}_i$, i.e., we have $|\log \ell_i(\cdot|\theta_k)| \leq B$ for any $i \in [n]$ and $k \in [m]$.

Based on assumption **A1**, every private signal has bounded information content [24]. This bound can be found, for instance, when the signal space is discrete and provides a full support for distribution. Let us define $\bar{\Theta}_i$ as the set of states that are observationally equivalent to θ_1 for agent $i \in [n]$; in other words, $\bar{\Theta}_i = \{\theta_k \in \Theta : \ell_i(s_i|\theta_k) = \ell_i(s_i|\theta_1) \forall s_i \in \mathcal{S}_i\}$ almost surely with respect to the signal space. As evident from the definition, any state $\theta_k \neq \theta_1$ in the set $\bar{\Theta}_i$ is not distinguishable from the true state by observation of samples from the i -th marginal. Let $\bar{\Theta} = \bigcap_{i=1}^n \bar{\Theta}_i$ be the set of states that are observationally equivalent to θ_1 from all agents perspective.

A2. We assume that no state in the world is observationally equivalent to the true state from the standpoint of the network, i.e., the true state is globally identifiable, and we have $\bar{\Theta} = \{\theta_1\}$.

Assumption **A2** guarantees that agents can benefit from collaboration with each other. In other words, for any false state $\theta_k \neq \theta_1$, there must exist an agent who is able to distinguish θ_1 from θ_k .

Finally, the operator $\mathbb{E}_{\mathcal{S}}[\cdot]$ denotes the expectation with respect to signal space throughout the paper.

C. Network Model

Private signals are not informative enough for agents, so they interact with each other to learn the true state of the world. For any time $t \in [T]$, the graph $G_t = ([n], E_t)$ captures the network structure, i.e. $[n]$ is the set of nodes corresponding to agents ($n > 1$), and E_t is the set of edges for that particular round. Agent i receives information from j only if the pair $(i, j) \in E_t$. We also define the neighborhood of agent i at any time $t \in [T]$ as $\mathcal{N}_{i,t} := \{j \in [n] : (i, j) \in E_t\}$. We represent by $[W(t)]_{ii} > 0$ the *self-reliance* of agent i at time t , and by $[W(t)]_{ij} > 0$ the weight that agent i assigns to information received from agent j in round t . The matrix $W(t)$ is then defined such that $[W(t)]_{ij}$ denotes the entry in its i -th row and j -th column. By construction, $W(t)$ has nonnegative entries, and

$[W(t)]_{ij} > 0$ only if $(i, j) \in E_t$. We further assume that $W(t)$ is doubly stochastic and symmetric; hence,

$$W(t)^\top = W(t) \quad \sum_{j=1}^n [W(t)]_{ij} = \sum_{j \in \mathcal{N}_{i,t}} [W(t)]_{ij} = 1.$$

To form the network structure, $W(t)$ is drawn *independently* over time from a *stationary* distribution \mathcal{W} , i.e. the elements of the sequence $\{W(t)\}_{t=1}^T$ are i.i.d. samples. Note importantly that the source of randomness is independent of signal space. We distinguish any expectation with respect to the network randomness from signal space using $\mathbb{E}_{\mathcal{W}}[\cdot]$.

A3. The network is *connected* in expectation sense. That is, for matrix $W = \mathbb{E}_{\mathcal{W}}[W(t)]$ there exists a bidirectional path from any agent $i \in [n]$ to any agent $j \in [n]$.

The assumption guarantees that (in expectation sense) information can be propagated properly throughout the network. It further implies that $\sigma_1(W) = 1$ is unique, and the other singular values of W are strictly less than one in magnitude [25]. Since the matrix W is symmetric, each agent in the network is *equally* central, and

$$\mathbb{1}^\top W = \mathbb{1}.$$

Therefore, $\mathbb{1}/n$ is the stationary distribution of W . Assumption **A3** entails that the Markov chain W is irreducible and aperiodic [25].

A well-known example of the communication protocol detailed above is *gossip* algorithm. The communication protocol works based on a Poisson clock. Once the clock ticks, an agent from the set $[n]$ is picked uniformly at random. The agent then selects a neighboring agent (with respect to a fixed, predefined structure) at random, and they average their information. For a thorough review of gossip algorithms we refer the interested reader to [26].

D. Centralized vs. Decentralized Detection

Decentralized detection is constructed based on its centralized analog. In the centralized scenario there is only one agent in the world (no network exists), and the agent has *global* information to learn the true state. In other words, the agent has full access to the sequence of signals $\{s_t\}_{t=1}^T$. At any round $t \in [T]$, the agent accumulates an *empirical average* of log-marginals, and forms the *belief* $\mu_t \in \Delta_m$ about the states, where $\Delta_m = \{\mu \in \mathbb{R}^m \mid \mu \succeq 0, \sum_{k=1}^m \mu(k) = 1\}$ denotes the m -dimensional probability simplex. Defining

$$\psi_t := \frac{1}{n} \sum_{i=1}^n \psi_{i,t} = \frac{1}{n} \sum_{i=1}^n \log \ell_i(s_{i,t} | \theta), \quad (1)$$

the following updates capture the centralized detection

$$\phi_t = \phi_{t-1} + \psi_t \quad , \quad \mu_t(k) = \frac{\exp\{\eta \phi_t(k)\}}{\sum_{z=1}^m \exp\{\eta \phi_t(z)\}}. \quad (2)$$

It can be seen from above that the centralized detector aggregates a geometric average of marginals in ϕ_t . The parameter η is called the *learning rate*.

One can prove that (see e.g. [21]) the following inequality holds

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^n \Psi_i(k) \right] &= \mathbb{E} \left[\sum_{i=1}^n \log \ell_i(\cdot | \theta_k) \right] \\ &< \mathbb{E} \left[\sum_{i=1}^n \log \ell_i(\cdot | \theta_1) \right] = \mathbb{E} \left[\sum_{i=1}^n \Psi_i(1) \right], \end{aligned}$$

for any $k \neq 1$, due to uniqueness of the true state θ_1 (assumption **A2**). In what follows, without loss of generality, we assume the following descending order, i.e.

$$\mathbb{E} \left[\sum_{i=1}^n \Psi_i(1) \right] > \mathbb{E} \left[\sum_{i=1}^n \Psi_i(2) \right] \geq \dots \geq \mathbb{E} \left[\sum_{i=1}^n \Psi_i(m) \right], \quad (3)$$

The assumption will only help us to simplify the derivation of technical results.

We now describe the distributed setting which involves a network of agents. In this scenario, each agent $i \in [n]$ only receives a stream of private signals $\{s_{i,t}\}_{t=1}^T$ generated based on the parametrized likelihood $\ell_i(\cdot | \theta_1)$. Therefore, agent $i \in [n]$ does not directly observe $s_{j,t}$ for any $j \neq i$. In other words, the agent collects local information by calculating a weighted average of log-likelihoods in its neighborhood. Then, the agent forms the belief $\mu_{i,t} \in \Delta_m$ at round $t \in [T]$ as follows:

$$\begin{aligned} \phi_{i,t} &= \sum_{j \in \mathcal{N}_{i,t}} [W(t)]_{ij} \phi_{j,t-1} + \psi_{i,t} \\ \mu_{i,t}(k) &= \frac{\exp\{\eta \phi_{i,t}(k)\}}{\sum_{z=1}^m \exp\{\eta \phi_{i,t}(z)\}}. \end{aligned} \quad (4)$$

As depicted above, each agent updates its belief using purely local diffusion. Let us distinguish the centralized and decentralized detector more specifically. The centralized detector collects *all* log-marginals, whereas the decentralized detector receives *private* log-marginals, and collects a weighted average of *local* information. Regardless of the information collection part, both algorithms are special cases of the well-known *Exponential Weights* algorithm.

It can be verified (see e.g. [17]) that the closed-form solution of $\phi_{i,t}$ in the decentralized update (4) is as follows

$$\phi_{i,t} = \sum_{\tau=1}^t \sum_{j=1}^n \left[\prod_{\rho=0}^{t-1-\tau} W(t-\rho) \right]_{ij} \psi_{j,\tau},$$

for any $i \in [n]$. One can also combine above with (1) and (2) to observe that

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \phi_{i,t} &= \frac{1}{n} \sum_{\tau=1}^t \sum_{j=1}^n \sum_{i=1}^n \left[\prod_{\rho=0}^{t-1-\tau} W(t-\rho) \right]_{ij} \psi_{j,\tau} \\ &= \frac{1}{n} \sum_{\tau=1}^t \sum_{j=1}^n \psi_{j,\tau} = \phi_t, \end{aligned}$$

since product of doubly stochastic matrices remains doubly stochastic. The identity above draws the connection between centralized and decentralized updates.

III. FINITE-TIME ANALYSIS

In this section, we provide our technical results. We first specialize the problem to *fixed* network structures, and then consider *switching* topologies, and prove non-asymptotic results for convergence of beliefs.

A. Fixed Network Topology

One can investigate the convergence of beliefs in fixed networks; however, we present a more general result in this section. We measure the efficiency of the distributed algorithm versus its centralized counterpart using the notion of *decentralization* cost. Throughout this section we assume that $W(t) = W$ with probability one for all $t \in [T]$.

At any round $t \in [T]$, we quantify the cost which agent $i \in [n]$ needs to pay to have the same opinion as a centralized agent with $D_{KL}(\mu_{i,t} \parallel \mu_t)$; then, the agent suffers a total decentralization cost of

$$\text{Cost}_{i,T} := \sum_{t=1}^T D_{KL}(\mu_{i,t} \parallel \mu_t), \quad (5)$$

after T rounds. In general, the KL-divergence measures the dissimilarity of two probability distributions; hence, it could be a reasonable metric to capture the difference between two algorithms as they both output a probability distribution over state space.

More formally, the cost quantifies the difference between a decentralized agent that observes private signals $\{s_{i,t}\}_{t=1}^T$ and a centralized agent that has $\{s_t\}_{t=1}^T$ available. In other words, it shows how well the decentralized algorithm copes with the partial information. It is important to note that $\text{Cost}_{i,T}$ is still a random quantity as it depends on signals. Our goal is to find a bound on the cost in the *high probability* sense.

The connectivity of network plays an important role in the learning as $W^t \rightarrow \frac{1}{n} \mathbf{1}\mathbf{1}^\top$ as $t \rightarrow \infty$. We shall see that the cost bound is governed by the mixture behavior of Markov chain W . We now present the main result of this section in the following theorem. The proof can be found in [21].

Theorem 1: Let the sequence of beliefs $\{\mu_{i,t}\}_{t=1}^T$ for each agent $i \in [n]$ be generated by the Distributed Detection algorithm with the choice of learning rate $\eta = \frac{1-\sigma_2(W)}{16B \log n}$. Given bounded log-marginals (assumption **A1**), global identifiability of the true state (assumption **A2**), and connectivity of the network (assumption **A3** with $W(t) = W$ a.e.), we have

$$\begin{aligned} \text{Cost}_{i,T} \leq & \frac{18B^2}{\mathcal{I}^2(\theta_1, \theta_2)} \max \left\{ \log \frac{6m}{\delta}, \frac{3B\sqrt{2}}{\mathcal{I}(\theta_1, \theta_2)} \right\} \\ & + \frac{48B \log n}{\mathcal{I}(\theta_1, \theta_2)} \frac{\log m + 2}{1 - \lambda_{\max}(W)}, \end{aligned}$$

with probability at least $1 - \delta$.

We remark that the special choice of learning rate only simplifies the bound. We do not tune the learning rate for optimization purposes; otherwise, using the same η for both algorithms would not provide a fair comparison. One can also work with $\eta = 1$ for both algorithms, and derive a bound

which looks slightly more complicated than the bound in Theorem 1.

The dependence of bound to the inverse of $\mathcal{I}(\theta_1, \theta_2)$ is quite natural since it can be seen as the asymptotic convergence rate of beliefs (see e.g. [17]–[19]). It simply means that when observations under θ_2 (the second likeliest state) are as likely as observations under θ_1 (the true state), the cost of the algorithm increases. Intuitively, when signals hardly reveal the difference between the best two candidates for the true state, agents must spend more effort to discern the two. Hence, this results in suffering a larger cost caused by slower learning rate.

The cost scales logarithmically with the network size n and the number of states m . It further scales inversely in

$$\gamma(W) := 1 - \sigma_2(W),$$

defined as the *spectral gap* of the network. Interestingly, in a fixed network, the detection cost is time-independent (with high probability), showing the best behavior with respect to time. Therefore, the average expected cost (per iteration cost) asymptotically tends to zero. We finally indicate that dependence of the bound to $\sigma_2(W)$ is important from network design perspective [21].

B. Switching Network Topology

In this section, we investigate the convergence of agents' beliefs in *time-varying* networks. As described in Section II-C, at every time t the network structure is realized with a random matrix $W(t)$. Agents then communicate with each other following $W(t)$. We assume that the matrix is generated from a *stationary* distribution, and therefore, we have $\mathbb{E}_{\mathcal{W}}[W(t)] = W$ for all $t \in [T]$. Assumption **A3** guarantees that the network is connected in *expectation* sense. We show that even in *time-varying* networks the mixture behavior of W (expected network) affects convergence of beliefs with high probability.

We now establish that agents have arbitrarily close opinions in a network that is connected in expectation sense. The following proposition proves that the convergence rate is governed by relative entropy of signals and network characteristics.

Proposition 2: Let the sequence of beliefs $\{\mu_{i,t}\}_{t=1}^T$ for each agent $i \in [n]$ be generated by the Distributed Detection algorithm (4) with the learning rate $\eta = 1$. Given bounded log-marginals (assumption **A1**), global identifiability of the true state (assumption **A2**), and connectivity of the expected network (assumption **A3**), for each individual agent $i \in [n]$ it holds that

$$\begin{aligned} \log \|\mu_{i,t} - \mathbf{e}_1\|_{\text{TV}} \leq & -\mathcal{I}(\theta_1, \theta_2)t + \sqrt{2B^2t \log \frac{m}{\delta}} \\ & + \frac{8B \log n}{1 - \sigma_2(W)} + \log m, \end{aligned}$$

with probability at least $1 - \delta$, where for $k \geq 2$

$$\mathcal{I}(\theta_1, \theta_k) := \frac{1}{n} \sum_{i=1}^n D_{KL}(\ell_i(\cdot | \theta_1) \parallel \ell_i(\cdot | \theta_k)).$$

The proposition provides an *anytime* bound on the log-distance in the high probability sense. It also verifies that the belief $\mu_{i,t}$ of each agent $i \in [n]$ is *strongly* consistent, i.e., it converges almost surely to a delta distribution on the true state. It is evident that the dominant term (asymptotic rate) depends on relative entropy of signals through $\mathcal{I}(\theta_1, \theta_2)$. This is consistent with previous asymptotic results found in [17]–[19] for other updates similar to our update. Proposition 2 complements those results by providing a non-asymptotic convergence rate. Finally, the inverse scaling with spectral gap also remains in effect (for expected network) similar to the case of fixed topologies.

IV. CONCLUSION

We considered the distributed detection problem in fixed and switching network topologies. A network of agents observe a stream of private signals which do not provide enough information about the true state of the world. Therefore, agents must communicate with each other to augment their imperfect knowledge with local observations. Iteratively forming a belief about the states, each agent uses purely local diffusion to update itself. We analyzed the detection problem in *finite-time* domain. We first specialized to the case of fixed networks, and study the efficiency of our algorithm versus its centralized analog. We introduced a KL cost to measure the dissimilarity of the two algorithms, and bounded the cost in terms of relative entropy of signals, network size and spectral gap. We further extended our results to *switching* network topologies. We investigated convergence of beliefs, and provided an anytime bound on the detection error in the high probability sense. In this case, the spectral gap of the expected network proves to be crucial in convergence rate. As a future direction, we would like to consider the scenario where the signal distribution is not stationary. Studying drifting distributions allows us to examine the robustness of detection in dynamic framework.

APPENDIX : OMITTED PROOFS

We use McDiarmid’s inequality in Lemma 3 for the proof of Proposition 2.

Proof of Proposition 2. Letting $\eta = 1$ in (4), we write

$$\begin{aligned} \mu_{i,t}(1) &= \frac{\exp\{\phi_{i,t}(1)\}}{\sum_{k=1}^m \exp\{\phi_{i,t}(k)\}} \\ &= \left(1 + \sum_{k=2}^m \exp\{\phi_{i,t}(k) - \phi_{i,t}(1)\}\right)^{-1} \\ &\geq 1 - \sum_{k=2}^m \exp\{\phi_{i,t}(k) - \phi_{i,t}(1)\}, \end{aligned} \quad (6)$$

using the simple inequality that $(1+x)^{-1} \geq 1-x$ for any $x \geq 0$. Since we know

$$\begin{aligned} \|\mu_{i,t} - \mathbf{e}_1\|_{\text{TV}} &= \frac{1}{2} \left(1 - \mu_{i,t}(1) + \sum_{k=2}^m \mu_{i,t}(k)\right) \\ &= 1 - \mu_{i,t}(1), \end{aligned}$$

we can combine above with (6) to get

$$\|\mu_{i,t} - \mathbf{e}_1\|_{\text{TV}} \leq \sum_{k=2}^m \exp\{\phi_{i,t}(k) - \phi_{i,t}(1)\}. \quad (7)$$

For any $k \in [m]$, define

$$\Phi_{i,t}(k) := \sum_{\tau=1}^t \sum_{j=1}^n \left[\prod_{\rho=0}^{t-1-\tau} W(t-\rho) \right]_{ij} \log \ell_j(\cdot|\theta_k),$$

and note that given $\{W(\tau)\}_{\tau=1}^t$, $\Phi_{i,t}(k)$ is a function of nt random variables $\{s_{i,\tau}\}_{\tau=1}^t$ for all $i \in [n]$. To use McDiarmid’s inequality in Lemma 3, set $H = \Phi_{i,t}(k) - \Phi_{i,t}(1)$, fix the samples for $nt - 1$ random variables, and draw two different samples $s_{j,\tau}$ and $s'_{j,\tau}$ for some $j \in [n]$ and some $\tau \in [t]$. The fixed samples are simply cancelled in the subtraction, and we derive

$$\begin{aligned} &|H(\dots, s_{j,\tau}, \dots) - H(\dots, s'_{j,\tau}, \dots)| \\ &= \left[\prod_{\rho=0}^{t-1-\tau} W(t-\rho) \right]_{ij} \left| \log \frac{\ell_j(s_{j,\tau}|\theta_k)}{\ell_j(s_{j,\tau}|\theta_1)} - \log \frac{\ell_j(s'_{j,\tau}|\theta_k)}{\ell_j(s'_{j,\tau}|\theta_1)} \right| \\ &\leq \left[\prod_{\rho=0}^{t-1-\tau} W(t-\rho) \right]_{ij} 2B, \end{aligned}$$

where we applied assumption **A1**. Since the product of doubly stochastic matrices remains doubly stochastic, summing over $j \in [n]$ and $\tau \in [t]$, we obtain

$$\sum_{\tau=1}^t \sum_{j=1}^n \left(\left[\prod_{\rho=0}^{t-1-\tau} W(t-\rho) \right]_{ij} 2B \right)^2 \leq 4B^2 t.$$

Let us define the event A as

$$A := \left\{ \phi_{i,t}(k) - \phi_{i,t}(1) > \mathbb{E}[\Phi_{i,t}(k) - \Phi_{i,t}(1)] + \varepsilon \right\},$$

where $\mathbb{E}[\cdot]$ is expectation over all sources of randomness (signals and network structures). We then have

$$\mathbb{P}(A) = \mathbb{E}_{\mathcal{W}} [\mathbb{P}(A | \{W(\tau)\}_{\tau=1}^t)],$$

where the expectation $\mathbb{E}_{\mathcal{W}}[\cdot]$ is taken with respect to network randomness. We now apply McDiarmid’s inequality in Lemma 3 to obtain

$$\mathbb{P}(A | \{W(\tau)\}_{\tau=1}^t) \leq \exp \left\{ \frac{-\varepsilon^2}{2B^2 t} \right\},$$

for each fixed k . Letting the probability above equal to δ/m and taking a union bound over $k \in [m]$, the following event (A^C) holds (given $\{W(\tau)\}_{\tau=1}^t$)

$$\phi_{i,t}(k) - \phi_{i,t}(1) \leq \mathbb{E}[\Phi_{i,t}(k) - \Phi_{i,t}(1)] + \sqrt{2B^2 t \log \frac{m}{\delta}}, \quad (8)$$

simultaneously for all $k = 2, \dots, m$, with probability at least $1 - \delta$, which implies

$$\mathbb{P}(A^C) = \mathbb{E}_{\mathcal{W}} [\mathbb{P}(A^C | \{W(\tau)\}_{\tau=1}^t)] \geq 1 - \delta.$$

On the other hand, in view of assumption **A1** and the fact

that $\mathbb{E}_{\mathcal{W}}[W(t)] = W$ for all t , we have

$$\begin{aligned}
& \mathbb{E} [\Phi_{i,t}(k) - \Phi_{i,t}(1)] \\
&= \sum_{\tau=1}^t \sum_{j=1}^n [W^{t-\tau}]_{ij} \mathbb{E}_{\mathcal{S}} [\log \ell_j(\cdot|\theta_k) - \log \ell_j(\cdot|\theta_1)] \\
&= \sum_{\tau=1}^t \sum_{j=1}^n \left([W^{t-\tau}]_{ij} - \frac{1}{n} \right) \mathbb{E}_{\mathcal{S}} [\log \ell_j(\cdot|\theta_k) - \log \ell_j(\cdot|\theta_1)] \\
&+ \frac{1}{n} \sum_{\tau=1}^t \sum_{j=1}^n \mathbb{E}_{\mathcal{S}} [\log \ell_j(\cdot|\theta_k) - \log \ell_j(\cdot|\theta_1)] \\
&\leq 2B \sum_{\tau=1}^t \sum_{j=1}^n \left| [W^{t-\tau}]_{ij} - \frac{1}{n} \right| \\
&- \frac{t}{n} \sum_{j=1}^n D_{KL}(\ell_j(\cdot|\theta_1) \|\ell_j(\cdot|\theta_k)) \\
&= 2B \sum_{\tau=1}^t \sum_{j=1}^n \left| [W^{t-\tau}]_{ij} - \frac{1}{n} \right| - \mathcal{I}(\theta_1, \theta_k)t \\
&\leq \frac{8B \log n}{1 - \sigma_2(W)} - \mathcal{I}(\theta_1, \theta_k)t,
\end{aligned}$$

where we applied Lemma 2 in [21] to derive the last step. Using (3), we simplify above to get

$$\mathbb{E} [\Phi_{i,t}(k) - \Phi_{i,t}(1)] \leq \frac{8B \log n}{1 - \sigma_2(W)} - \mathcal{I}(\theta_1, \theta_2)t, \quad (9)$$

for any $k = 2, \dots, m$. Substituting (9) into (8) and combining with (7), we have

$$\begin{aligned}
& \|\mu_{i,t} - \mathbf{e}_1\|_{\text{TV}} \\
&\leq \sum_{k=2}^m \exp \left\{ -\mathcal{I}(\theta_1, \theta_2)t + \sqrt{2B^2t \log \frac{m}{\delta}} + \frac{8B \log n}{1 - \sigma_2(W)} \right\} \\
&\leq m \exp \left\{ -\mathcal{I}(\theta_1, \theta_2)t + \sqrt{2B^2t \log \frac{m}{\delta}} + \frac{8B \log n}{1 - \sigma_2(W)} \right\},
\end{aligned}$$

with probability at least $1 - \delta$, and the proof is completed. ■

Lemma 3: (McDiarmid's Inequality) Let $X_1, \dots, X_N \in \chi$ be independent random variables and consider the mapping $H : \chi^N \mapsto \mathbb{R}$. If for $i \in \{1, \dots, N\}$, and every sample $x_1, \dots, x_N, x'_i \in \chi$, the function H satisfies

$$|H(x_1, \dots, x_i, \dots, x_N) - H(x_1, \dots, x'_i, \dots, x_N)| \leq c_i,$$

then for all $\varepsilon > 0$,

$$\mathbb{P} \left\{ H(x_1, \dots, x_N) - \mathbb{E} [H(X_1, \dots, X_N)] \geq \varepsilon \right\} \leq \exp \left\{ \frac{-2\varepsilon^2}{\sum_{i=1}^N c_i^2} \right\}.$$

REFERENCES

- [1] R. R. Tenney and N. R. Sandell Jr, "Detection with distributed sensors," *IEEE Transactions on Aerospace Electronic Systems*, vol. 17, pp. 501–510, 1981.
- [2] J. N. Tsitsiklis *et al.*, "Decentralized detection," *Advances in Statistical Signal Processing*, vol. 2, no. 2, pp. 297–344, 1993.
- [3] V. Borkar and P. P. Varaiya, "Asymptotic agreement in distributed estimation," *IEEE Transactions on Automatic Control*, vol. 27, no. 3, pp. 650–655, 1982.
- [4] A. Nedic and A. Ozdaglar, "Distributed subgradient methods for multi-agent optimization," *IEEE Transactions on Automatic Control*, vol. 54, no. 1, pp. 48–61, 2009.
- [5] S. Kar, J. M. Moura, and K. Ramanan, "Distributed parameter estimation in sensor networks: Nonlinear observation models and imperfect communication," *IEEE Transactions on Information Theory*, vol. 58, no. 6, pp. 3575–3605, 2012.
- [6] C. Eksin and A. Ribeiro, "Distributed network optimization with heuristic rational agents," *IEEE Transactions on Signal Processing*, vol. 60, no. 10, pp. 5396–5411, 2012.
- [7] J.-F. Chamberland and V. V. Veeravalli, "Decentralized detection in sensor networks," *IEEE Transactions on Signal Processing*, vol. 51, no. 2, pp. 407–416, 2003.
- [8] F. Bullo, J. Cortés, and S. Martinez, *Distributed control of robotic networks: a mathematical approach to motion coordination algorithms*. Princeton University Press, 2009.
- [9] N. A. Atanasov, J. Le Ny, and G. J. Pappas, "Distributed algorithms for stochastic source seeking with mobile robot networks," *Journal of Dynamic Systems, Measurement, and Control*, 2014.
- [10] S. Shahrampour, S. Rakhlin, and A. Jadbabaie, "Online learning of dynamic parameters in social networks," in *Advances in Neural Information Processing Systems*, 2013.
- [11] A. Jadbabaie, J. Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Transactions on Automatic Control*, vol. 48, no. 6, pp. 988–1001, 2003.
- [12] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Transactions on Automatic Control*, vol. 49, no. 9, pp. 1520–1533, 2004.
- [13] F. S. Cattivelli and A. H. Sayed, "Distributed detection over adaptive networks using diffusion adaptation," *IEEE Transactions on Signal Processing*, vol. 59, no. 5, pp. 1917–1932, 2011.
- [14] D. Jakovetic, J. M. Moura, and J. Xavier, "Distributed detection over noisy networks: Large deviations analysis," *IEEE Transactions on Signal Processing*, vol. 60, no. 8, pp. 4306–4320, 2012.
- [15] D. Bajovic, D. Jakovetic, J. M. Moura, J. Xavier, and B. Sinopoli, "Large deviations performance of consensus+ innovations distributed detection with non-gaussian observations," *IEEE Transactions on Signal Processing*, vol. 60, no. 11, pp. 5987–6002, 2012.
- [16] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, "Non-bayesian social learning," *Games and Economic Behavior*, vol. 76, no. 1, pp. 210–225, 2012.
- [17] S. Shahrampour and A. Jadbabaie, "Exponentially fast parameter estimation in networks using distributed dual averaging," in *IEEE Conference on Decision and Control (CDC)*, 2013, pp. 6196–6201.
- [18] A. Lalitha, A. Sarwate, and T. Javidi, "Social learning and distributed hypothesis testing," in *International Symposium on Information Theory (ISIT)*, 2014, pp. 551–555.
- [19] S. Shahrampour, M. A. Rahimian, and A. Jadbabaie, "Switching to learn," in *American Control Conference (ACC)*, July 2015, pp. 2918–2923.
- [20] J. C. Duchi, A. Agarwal, and M. J. Wainwright, "Dual averaging for distributed optimization: convergence analysis and network scaling," *IEEE Transactions on Automatic Control*, vol. 57, no. 3, pp. 592–606, 2012.
- [21] S. Shahrampour, A. Rakhlin, and A. Jadbabaie, "Distributed detection: Finite-time analysis and impact of network topology," *arXiv preprint arXiv:1409.8606*, 2014.
- [22] A. Nedic, A. Olshevsky, and C. Uribe, "Nonasymptotic convergence rates for cooperative learning over time-varying directed graphs," in *American Control Conference (ACC)*, July 2015, pp. 5884–5889.
- [23] A. Nedic, A. Olshevsky, and C. A. Uribe, "Fast convergence rates for distributed non-bayesian learning," *arXiv preprint arXiv:1508.05161*, 2015.
- [24] K. Drakopoulos, A. Ozdaglar, and J. N. Tsitsiklis, "On learning with finite memory," *IEEE Transactions on Information Theory*, vol. 59, no. 10, pp. 6859–6872, 2013.
- [25] J. S. Rosenthal, "Convergence rates for markov chains," *SIAM Review*, vol. 37, no. 3, pp. 387–405, 1995.
- [26] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized gossip algorithms," *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2508–2530, 2006.