

Distributed Learning of Distributions via Social Sampling

Anand D. Sarwate, *Member, IEEE*, and Tara Javidi, *Senior Member, IEEE*,

Abstract—A protocol for distributed estimation of discrete distributions is proposed. Each agent begins with a single sample from the distribution, and the goal is to learn the empirical distribution of the samples. The protocol is based on a simple message-passing model motivated by communication in social networks. Agents sample a message randomly from their current estimates of the distribution, resulting in a protocol with quantized messages. Using tools from stochastic approximation, the algorithm is shown to converge almost surely. Examples illustrate three regimes with different consensus phenomena. Simulations demonstrate this convergence and give some insight into the effect of network topology.

I. INTRODUCTION

The emergence of large-network paradigms for communications and the widespread adoption of social networking technologies has resurrected interest in classical models of opinion formation and distributed computation as well as new approaches to distributed learning and inference. In this paper we propose a simple message passing protocol inspired by social communication and show how it allows a network of individuals can learn about global phenomena. In particular, we study a situation wherein each node or agent in a network holds an initial opinion and the agents communicate with each other to infer the distribution of their initial opinions. Our model of messaging is a simple abstraction of social communication in which individuals exchange single opinions. This model is a randomized approximation of consensus procedures. Because agents collect samples of the opinions of their neighbors, we call our model *social sampling*.

In our protocol agents merge the sampled opinions of their neighbors with their own estimates using a weighted average.

Manuscript received May 17, 2013; revised January 9, 2014; second revision . Date of current version June 5, 2014. The work of both authors was funded by National Science Foundation (NSF) Grant Number CCF-1218331. A.D. Sarwate received additional support from the California Institute for Telecommunications and Information Technology (CALIT2) at UC San Diego. T. Javidi received additional support from NSF Grant Numbers CCF-1018722 and CNS-1329819. Simulations were supported by the UCSD FWGrid Project, NSF Research Infrastructure Grant Number EIA-0303622. Parts of this work were presented at the 46th Annual Conference on Information Sciences and Systems (CISS), Princeton, NJ, USA, March 2012 [1], and the 49th Annual Allerton Conference on Communication, Control and Computation, Monticello, IL, USA, September 2011 [2].

Anand D. Sarwate is with the Department of Electrical and Computer Engineering, Rutgers, The State University of New Jersey, 84 Brett Road, Piscataway, NJ 08854, USA. Email: asarwate@ece.rutgers.edu. Tara Javidi is with the Department of Electrical and Computer Engineering, University of California, San Diego, 9500 Gilman Dr. MC 0407, La Jolla CA 92093-0407, USA. Email: tjavidi@ucsd.edu.

Communicated by xxxxxxxxxx.

Digital Object Identifier xxxxxxxxxx

Copyright (c) 2014 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

Averaging has been used to model opinion formation for decades, starting with the early work of French [3], Harary [4], and DeGroot [5]. These works focused on averaging as a means to an end – averaging the local opinions of a group of peers was a simple way to model the process of negotiation and compromise of opinions represented as scalar variables. A natural extension of the above work is that where all agents are interested in the local reconstruction of the empirical distribution of discrete opinions. Such locally constructed empirical distributions not only provide richer information about global network properties (such as the outcome of a vote, the confidence interval around the mean, etc), but from a statistical estimation perspective provide estimates of local sufficient statistics when the agents’ opinions are independent and identically distributed (i.i.d.) samples from a common distribution.

For opinions taking value in a finite, discrete set, we can *compute* the empirical distribution of opinions across a network by running an average consensus algorithm for each possible value of the opinion. This can even be done in parallel so that at each time agents exchange their entire histogram of opinions and compute weighted averages of their neighbors’ histograms to update their estimate. In a social network, this would correspond to modeling the interaction of two agents as a complete exchange of their entire beliefs in every opinion, which is not realistic. In particular, if the number of possible opinions is large (corresponding to a large number of bins or elements in the histogram), communicating information about all opinions may be very inefficient, especially if the true distribution of opinions is far from uniform.

In contrast, this paper considers a novel model in which agents’ information is disseminated through randomly selected samples of locally constructed histograms [2]. The use of random samples results in a much lower overhead because it accounts for the popularity/frequency of histogram bins and naturally enables finite-precision communication among neighboring nodes. It is not hard to guarantee that the expectation of the node estimates converges to the true histogram when the mean of any given (randomized and noisy) shared sample is exactly the local estimate of the histogram. However, to ensure convergence in an almost sure sense we use techniques from stochastic approximation. We identify three interesting regimes of behavior. In the first, studied by Narayanan and Niyogi [6], agents converge to atomic distributions on a common single opinion. In the second, agents converge to a common consensus estimate which in general is not equal to the true histogram. Finally, we demonstrate a randomized protocol which, under mild technical assumptions, ensures convergence of agents’ local estimates to the global

histogram almost surely. The stochastic approximation point of view suggests that a set of decaying weights can control the accumulation of noise along time and still compute the average histogram.

Related work

In addition to the work in mathematical modeling of opinion formation [3]–[5], there has been a large body of work on consensus in terms of decision making initiated by Aumann [7]. Borkar and Varaiya [8] studied distributed agreement protocols in which agents are trying to estimate a common parameter. The agents randomly broadcast conditional expectations based on all of the information they have seen so far, and they find general conditions under which the agents would reach an asymptotic agreement. If the network is sufficiently connected (in a certain sense), the estimates converge to the centralized estimate of the parameter, even when the agents’ memory is limited [9]. In these works the questions are more about whether agreement is possible at all, given the probability structure of the observation and communication.

There is a significant body of work on consensus and information aggregation in sensor networks [10]–[18]. From the protocol perspective, many authors have studied the effect of network topology on the rate of convergence of consensus protocols [10]–[12], [16]–[21]. For communication networks the speed can be accelerated by exploiting network properties [22]–[25] (see surveys in [19], [26] for more references). Others have studied how quantization constraints impact convergence [27]–[34]. However, in all of these works the agents are assumed to be some sort of computational devices like robotic networks or sensor networks. A comprehensive view of this topic is beyond the scope of this paper. Instead, we focus on a few papers most relevant to our model and study: consensus with quantized messages and consensus via stochastic approximation. However, it is important to note that in contrast to all the studies discussed below, our work primarily deals with an extension of the classic consensus (linear combination of private values) in that we are interested in ensuring agreement over the space of distributions (histograms).

Our goal in this paper is to ensure the convergence of each agent’s local estimate to a true and global discrete distribution via a low-overhead algorithm in which messages are chosen in a discrete set. Our work is therefore related to the extensive recent literature on quantized consensus [28], [30], [31], [33]. In these works, as in ours, the communication between nodes is discretized (and in some cases the storage/computation at nodes as well [30]) and the question is how to ensure consensus (within a bin) to the average. This is in sharp contrast to our model which uses discrete messages to convey and ensure consensus on the network-wide histogram of discrete values. As a result, in contrast to the prior work on quantization noise [27], [31], [35], the “noise” is manufactured by our randomized sample selection scheme and hence plays a significantly different role.

Our analysis uses similar tools from stochastic approximation as recent studies of consensus protocols [34], [36], [37].

However, these works use stochastic approximation to address the effect of random noise in network topology, message transmission, and computation for a scalar consensus problem, while our use of standard theorems in stochastic approximation is to handle the impact of the noise that comes from the sampling scheme that generates our random messages. In other words, our noise is introduced by design even though our technique to control its cumulative effect is similar.

II. MODEL AND ALGORITHMS

Let $[n]$ denote the set $\{1, 2, \dots, n\}$ and let $\mathbf{e}_i \in \mathbb{R}^M$ denote the i -th elementary row vector in which the i -th coordinate is 1 and all other coordinates are 0. Let $\mathbf{1}(\cdot)$ denote the indicator function and $\mathbf{1}$ the column vector whose elements are all equal to 1. Let $\|\cdot\|$ denote the Euclidean norm for vectors and the Frobenius norm for matrices. We will represent probability distributions on finite sets as row vectors, and denote the set of probability distributions on a countable set A by $\mathbb{P}(A)$.

A. Problem setup

Time is discrete and indexed by $t \in \{0, 1, 2, \dots\}$. The system contains n agents or “nodes.” At time t the agents can communicate with each other according to an undirected graph $\mathcal{G}(t)$ with vertex set $[n]$ and edge set $\mathcal{E}(t)$. Let $\mathcal{N}_i(t) = \{j : (i, j) \in \mathcal{E}(t)\}$ be the set of neighbors of node i . If $(i, j) \in \mathcal{E}(t)$ then nodes i and j can communicate at time t . Let $G(t)$ denote the adjacency matrix of $\mathcal{G}(t)$ and let $D(t)$ be the diagonal matrix of node degrees. The Laplacian of $\mathcal{G}(t)$ is $L(t) = D(t) - G(t)$.

At time 0 every node starts with a single discrete *sample* $X_i \in \mathcal{X} = [M]$. The goal of the network is for each node to estimate the empirical distribution, or normalized histogram, of the observations $\{X_i : i \in [n]\}$:

$$\Pi(x) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(X_i = x) \cdot \mathbf{e}_x \quad \forall x \in [M].$$

To make it simpler to characterize the overall communication overhead, we assume that

- Agents can exchange messages $Y_i(t)$ lying in a finite set \mathcal{Y} .
- At each time $t = 1, 2, \dots$ agents can transmit a single message to all of their neighbors.

At each time $t = 0, 1, 2, \dots$, each node i maintains an internal estimate $Q_i(t)$ of the distribution Π , and we take $Q_i(0) = \mathbf{e}_{X_i}$. Every node i generates its message $Y_i(t) \in \mathcal{Y}$ as a function of this internal estimate $Q_i(t)$. Furthermore, each node i receives the messages $\{Y_j(t) : j \in \mathcal{N}_i(t)\}$ from its neighbors and use these messages to perform an update of its estimate $Q_i(t+1)$ using $Q_i(t)$, $\{Y_j(t) : j \in \mathcal{N}_i(t)\}$.

Since in a single time t there are potentially $2|\mathcal{E}(t)|$ messages transmitted in the network, a first approximation for the communication overhead of this class of schemes is simply proportional to the number of edges in the network multiplied by the logarithm of the cardinality of set \mathcal{Y} .

We are interested in the properties of the estimates $Q_i(t)$ as $t \rightarrow \infty$. In particular, we are interested in the case where every

element in the set $\{Q_i(t) : i \in n\}$ converges almost surely to a common random variable \mathbf{q}^* . In this case, we call the random vector \mathbf{q}^* the consensus variable. Different algorithms that we consider will result in different properties of this consensus variable. For example, the we will consider the support of the distribution of \mathbf{q}^* as well as its expectation.

B. Social sampling and linear updates

In this paper we assume that $\mathcal{Y} = \{\mathbf{0}, \mathbf{e}_1, \mathbf{e}_1, \dots, \mathbf{e}_M\}$, with the convention that node i transmits nothing (or remains silent) when $Y_i(t) = \mathbf{0}$. Furthermore, we consider the class of schemes where the random message $Y_i(t) \in \mathcal{Y}$ of node i at time t is generated according to a distribution $P_i(t) \in \mathbb{P}(\mathcal{Y})$ which itself is a function of the estimate $Q_i(t)$. In other words, $P_i(t)$ is a row vector of length M where $\mathbb{P}(Y_i(t) = \mathbf{e}_m) = P_{i,m}(t)$. We frequently refer to the random messages $Y_i(t) \in \mathcal{Y}, i \in [n], t = 0, 1, \dots$ as *social samples* because they correspond to nodes obtaining random samples of their neighbor's opinions. Note that although the random variable $Y_i(t)$ takes values in \mathbb{R}^M , it is supported only on the finite set \mathcal{Y} and hence requires communicating $\log |\mathcal{Y}|$ information bits.

For simplicity in this paper, we often rely on matrix representation across the network. Accordingly, let $\mathbf{Y}(t)$ be the $n \times M$ matrix whose i -th row is $Y_i(t)$. Then we have $\mathbb{E}[\mathbf{Y}(t)] = \mathbf{P}(t)$.

Let $\{W(t) : t = 0, 1, 2, \dots\}$ be a sequence of $n \times n$ matrices with nonnegative entries, such that $W_{ij}(t) = 0$ for all $(i, j) \notin \mathcal{E}(t)$. We study linear updates of the form

$$Q_i(t+1) = (1 - \delta(t)A_{ii}(t))Q_i(t) - \delta(t)B_{ii}(t)Y_i(t) + \sum_{j \in \mathcal{N}_i(t)} \delta(t)W_{ij}(t)Y_j(t). \quad (1)$$

Here the parameter $\delta(t)$ is a step size for the algorithm. Let $\mathbf{Q}(t)$ be the $n \times M$ matrix whose i -th row is $Q_i(t)$. We can write the iterates more compactly as

$$\mathbf{Q}(t+1) = (I - \delta(t)A(t))\mathbf{Q}(t) - \delta(t)B(t)\mathbf{Y}(t) + \delta(t)W(t)\mathbf{Y}(t),$$

where $A(t)$ and $B(t)$ are diagonal matrices.

In the next section we will analyze this update and identify conditions under which the estimates $Q_i(t)$ converge to a common $q^* \in \mathbb{P}(\mathcal{Y})$ and additional conditions under which $q^* = \Pi$. To provide a unified analysis of these different algorithms, we transform the update equation into a stochastic iteration

$$\mathbf{Q}(t+1) = \mathbf{Q}(t) + \delta(t) [\bar{H}(t)\mathbf{Q}(t) + \mathbf{C}(t) + \mathbf{M}(t)]. \quad (2)$$

In this form of the update, the matrix $\bar{H}(t)$ represents the mean effect of the network connectivity, $\mathbf{M}(t)$ is a martingale difference term related to the randomness in the network topology and social sampling, and $\mathbf{C}(t)$ is a correction term associated with the difference between the estimate $\mathbf{Q}(t)$ and the sampling distribution $\mathbf{P}(t)$.

Lemma 1. *The iteration in (1) can be rewritten as (2), where*

$$\bar{H}(t) \triangleq \bar{W}(t) - \bar{B}(t) - \bar{A}(t) \quad (3)$$

$$\mathbf{C}(t) \triangleq (W(t) - B(t))(\mathbf{P}(t) - \mathbf{Q}(t)) + (W(t) - B(t) - \bar{W}(t) + \bar{B}(t))\mathbf{Y}(t) \quad (4)$$

$$\mathbf{M}(t) \triangleq (W(t) - B(t) - A(t) - \bar{W}(t) + \bar{B}(t) + \bar{A}(t))\mathbf{Q}(t) + (\bar{W}(t) - \bar{B}(t))(\mathbf{Y}(t) - \mathbf{P}(t)) + (\bar{W}(t) - \bar{B}(t) - W(t) + B(t))\mathbf{P}(t). \quad (5)$$

and the term $\mathbf{M}(t)$ is a martingale difference sequence:

$$\mathbb{E}[\mathbf{M}(t)|\mathcal{F}_t] = \mathbf{0}.$$

Proof. Rewriting the iterates, we see that

$$\mathbf{Q}(t+1) = \mathbf{Q}(t) + \delta(t) \left[-A(t)\mathbf{Q}(t) + (W(t) - B(t))\mathbf{Y}(t) \right], \quad (6)$$

and the term multiplied by $\delta(t)$ can be expanded:

$$\begin{aligned} & -A(t)\mathbf{Q}(t) + (W(t) - B(t))\mathbf{Y}(t) \\ &= (\bar{W}(t) - \bar{B}(t) - \bar{A}(t))\mathbf{Q}(t) + (W(t) - B(t) - A(t) - \bar{W}(t) + \bar{B}(t) + \bar{A}(t))\mathbf{Q}(t) \\ & \quad + (W(t) - B(t))(\mathbf{Y}(t) - \mathbf{Q}(t)) \\ &= (\bar{W}(t) - \bar{B}(t) - \bar{A}(t))\mathbf{Q}(t) + (W(t) - B(t) - A(t) - \bar{W}(t) + \bar{B}(t) + \bar{A}(t))\mathbf{Q}(t) \\ & \quad + (W(t) - B(t))(\mathbf{P}(t) - \mathbf{Q}(t)) \\ & \quad + (W(t) - B(t))(\mathbf{Y}(t) - \mathbf{P}(t)) \\ &= (\bar{W}(t) - \bar{B}(t) - \bar{A}(t))\mathbf{Q}(t) + (W(t) - B(t) - A(t) - \bar{W}(t) + \bar{B}(t) + \bar{A}(t))\mathbf{Q}(t) \\ & \quad + (W(t) - B(t))(\mathbf{P}(t) - \mathbf{Q}(t)) \\ & \quad + (\bar{W}(t) - \bar{B}(t))(\mathbf{Y}(t) - \mathbf{P}(t)) \\ & \quad + (\bar{W}(t) - \bar{B}(t) - W(t) + B(t))\mathbf{P}(t) \\ & \quad + (W(t) - B(t) - \bar{W}(t) + \bar{B}(t))\mathbf{Y}(t)]. \quad (7) \end{aligned}$$

Define \bar{H} , \mathbf{C} , and $\mathbf{M}(t)$ as in (3), (4) and (5). Furthermore, define

$$\mathbf{M}(t) = (W(t) - B(t) - A(t) - \bar{W}(t) + \bar{B}(t) + \bar{A}(t))\mathbf{Q}(t) + (\bar{W}(t) - \bar{B}(t))(\mathbf{Y}(t) - \mathbf{P}(t)) + (\bar{W}(t) - \bar{B}(t) - W(t) + B(t))\mathbf{P}(t).$$

Taking conditional expectation of both sides of (6) and noting that $\mathbb{E}[\mathbf{Y}(t)|\mathcal{F}_t] = \mathbf{P}(t)$, we have the result. \square

Loosely speaking, the term $\mathbf{C}(t)$ will be asymptotically vanishing if $\mathbf{P}(t) \rightarrow \mathbf{Q}(t)$ and the matrices $W(t)$ and $B(t)$ are asymptotically independent of $\mathbf{Y}(t)$. In the next section, we show that this stochastic approximation scheme converges under certain conditions on update rule.

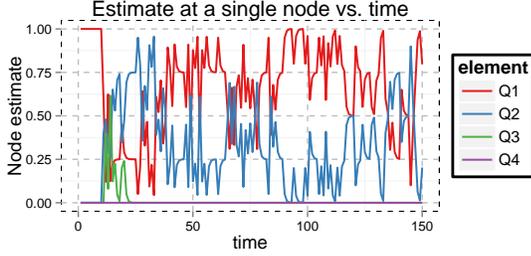


Fig. 1. Trace of estimate of $Q_i(t)$ for a single node i and $M = 4$, using the algorithm in (8) on a 5×5 grid graph. The four lines correspond to the 4 entries of the vector $Q_i(t)$. The estimates of all nodes converge to a random elementary vector $\mathbf{q}^* \in \mathcal{J}$; furthermore, $\mathbb{E}[\mathbf{q}^*] = \Pi$.

C. Example Algorithms

There are many algorithms which have the general form of the update rule (1). Before we proceed with the analysis of (1), however, we look at three examples in which we can see interesting regimes of consensus-like behavior stemming from the update rule. In these numerical examples, the graph $\mathcal{G}(t)$ is a 5 grid, so the maximum degree of any node is $d_{\max} = 4$. The initial values in the grid were drawn i.i.d. from the distribution $(0.4, 0.3, 0.2, 0.1)$ on $M = 4$ elements.

1) *Averaging with social samples*: Suppose $\mathbf{P}(t) = \mathbf{Q}(t)$ for all t and consider the update

$$Q_i(t+1) = \frac{d_{\max} + 1 - d_i}{d_{\max} + 1} Q_i(t) + \sum_{j \in \mathcal{N}_i(t)} \frac{1}{d_{\max} + 1} Y_j(t). \quad (8)$$

This corresponds to $\delta(t) = 1$, $A_{ii}(t) = \frac{d_i}{d_{\max} + 1}$, $B_{ii}(t) = 0$, and $W_{ij}(t) = \frac{1}{d_{\max} + 1}$. A trace of a single node's estimates for $M = 4$ is shown in Figure 1. The four lines correspond to the 4 elements of $Q_i(t)$. As shown in [6], this procedure results in all $Q_i(t)$ converging to a consensus value that is a random singleton \mathbf{q}^* in $\{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_M\}$ such that $\mathbb{E}[\mathbf{q}^*] = \Pi$.

2) *Averaging with social samples and decaying step size*: Suppose $\mathbf{P}(t) = \mathbf{Q}(t)$ for all t and consider the update

$$Q_i(t+1) = Q_i(t) - \frac{d_i}{d_{\max} + 1} \delta(t) Q_i(t) + \delta(t) \sum_{j \in \mathcal{N}_i(t)} \frac{1}{d_{\max} + 1} Y_j(t), \quad (9)$$

with $\delta(t) = 10/(t+1)$. This corresponds to $A_{ii}(t) = \frac{d_i}{d_{\max} + 1}$, $B_{ii}(t) = 0$, and $W_{ij}(t) = \frac{1}{d_{\max} + 1}$. Figure 2 shows the estimates of all agents in a 5×5 grid tracking the estimate of $\Pi(m)$ for a single m . The estimates of the agents converge to a consensus \mathbf{q}^* which is not equal to the true value Π but $\mathbb{E}[\mathbf{q}^*] = \Pi$. More generally, we will show that under standard assumptions on the step size and weights, if $\mathbf{P}(t) = \mathbf{Q}(t)$ then the iteration (1) converges to a consensus state whose expectation is Π .

3) *Exchange with social samples and censoring*: Let $\delta(t) = 1/t$ and suppose that for each (i, m) , $m \neq 0$, we set

$$P_{i,m}(t) = \begin{cases} 0 & Q_{i,m}(t) < d_{\max} \delta(t) \\ Q_{i,m}(t) & \text{otherwise} \end{cases}$$

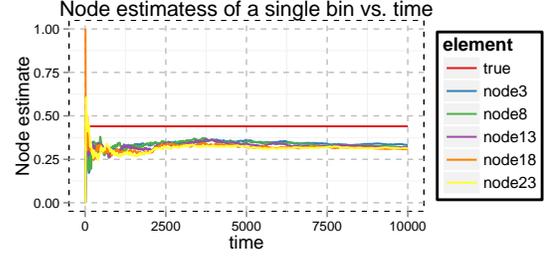


Fig. 2. Trace of estimate of $Q_{i,m}(t)$ over all $i \in [n]$ for a single $m \in [M]$ with $M = 4$, using the algorithm in (9) on a 5×5 grid graph. The estimates of all nodes converge to a common random distribution $\mathbf{q}^* \in \mathbb{P}(\mathcal{J})$ whose expectation is equal to the true distribution Π .

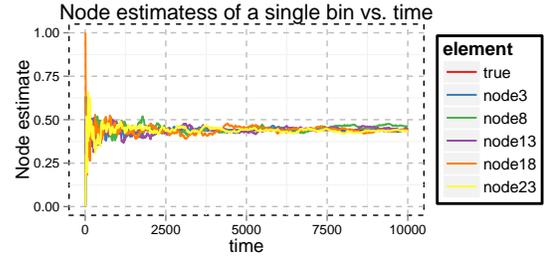


Fig. 3. Trace of $Q_{i,m}(t)$ for a few different i and a single $m \in M$ with $M = 4$, using the algorithm in (10) on a 5×5 grid graph. For this update rule, the estimates of all nodes converge almost surely to Π .

That is, under $\mathbf{P}(t)$ node i sends random messages corresponding to those elements of $\mathbf{Q}(t)$ that are larger than $d_{\max} \delta(t)$ while it remains silent with probability $P_{i,0}(t) = 1 - \sum_{m \neq 0} P_{i,m}(t)$. Let $N_i(t) = \sum_{j \in \mathcal{N}_i} \mathbf{1}(Y_j(t) \neq 0)$ denote the total number of neighbors of node i which did not remain silent at time t . Node i updates its local estimate of the global histogram according to the following rule:

$$Q_i(t+1) = \begin{cases} Q_i(t) & Y_i(t) = 0 \\ Q_i(t) - \delta(t) N_i(t) Y_i(t) + \delta(t) \sum_{j \in \mathcal{N}_i} Y_j(t) & Y_i(t) \neq 0 \end{cases}. \quad (10)$$

The behavior of (10), when $\delta(t) = \frac{10}{t+1}$ is illustrated in Figure 3. The estimates $Q_i(t)$ of all agents converge to Π almost surely under this update rule. More generally, we show that given certain technical conditions, the mean across the agents of the sample paths $\mathbf{Q}(t)$ is always equal to Π and the estimate of each agent converges almost surely to Π . In this case we show that the rate of convergence is on the order of $1/t$.

These examples illustrate that the algorithms can display different qualitative behaviors depending on the choice of the step size $\delta(t)$ as well as the choice of \mathbf{P} . In the first case, all agent estimates converge to a common random singleton, whereas under the second scenario they seem to converge to a common estimated histogram, even though this common estimated histogram might be far from the true histogram of the given initial values. Finally, in the case where agents ‘‘censor’’ their low values and follow update rule (10), Figure 3 suggests an almost sure convergence to Π . In the next section,

we analytically confirm these empirical findings in with a unified analysis.

III. ANALYSIS

We now turn to the analysis of the general protocol in (1). We will need a number of additional conditions on the iterates in order to guarantee convergence. Condition 1 is that the agents compute convex combinations of their current estimates and the messages at time t . This guarantees that the estimated distributions of the agents $Q_i(t)$ are proper probability distributions on $[M]$.

A. Mean preservation

Let \mathcal{F}_t be the σ -algebra generated by

$$\{\mathbf{Q}_i(s) : s < t\} \cup \{\mathcal{G}(s) : s < t\},$$

so $\mathbf{Q}(t+1)$ is measurable with respect to \mathcal{F}_t .

Condition 1 (Mixing coefficients). *For all $t > 0$ and all $i \in [n]$,*

$$\sum_{j \in \mathcal{N}_i} W_{ij}(t) - A_{ii}(t) - B_{ii}(t) = 0. \quad (11)$$

Let

$$\begin{aligned} \bar{A}(t) &= \mathbb{E}[A(t)|\mathcal{F}_t] \\ \bar{B}(t) &= \mathbb{E}[B(t)|\mathcal{F}_t] \\ \bar{W}(t) &= \mathbb{E}[W(t)|\mathcal{F}_t]. \end{aligned}$$

Note that the coefficient $W(t)$ and $A(t)$ can, in general, depend on the messages $\mathbf{Y}(t)$ as well as the graph $\mathcal{G}(t)$.

Our first result is a trivial consequence of the linearity of the update rule, and does not require any conditions beyond the fact that the estimate \mathbf{Q} is itself a distribution.

Lemma 2. *Suppose Condition 1 holds. If $\mathbf{P}(t) = \mathbf{Q}(t)$ and $\mathbf{W}(t)$ and $\mathbf{B}(t)$ are independent of $\mathbf{Y}(t)$, then for all t ,*

$$\mathbb{E}[\mathbf{1}^\top \mathbf{Q}(t)] = \mathbf{1}^\top \mathbf{Q}(0) = n\Pi.$$

Proof. Given Condition 1,

$$\begin{aligned} \mathbb{E}[\mathbf{Q}(t+1)|\mathcal{F}_t] &= (I - \delta(t)\bar{A}(t))\mathbf{Q}(t) - \delta(t)\bar{B}(t)\mathbf{Q}(t) \\ &\quad + \delta(t)\bar{W}(t)\mathbf{Q}(t). \end{aligned}$$

And therefore

$$\mathbb{E}[\mathbf{1}^\top \mathbf{Q}(t+1)|\mathcal{F}_t] = \mathbf{1}^\top \mathbf{Q}(t). \quad (12)$$

On the other hand, since $\mathbf{1}^\top \mathbf{Q}(0) = n\Pi$, the proof is complete. \square

This result is simple to see – if the expected message $Y_i(t)$ is equal to $Q_i(t)$, then the mean of the dynamics are just those of average consensus. However, it is not necessarily the case that the nodes converge to a consensus state, and if they do converge to a consensus state, that state may not be equal to Π on almost every sample path. The expected average of the node estimates will be equal to Π , and if they do reach a consensus the expected consensus value will be Π . In this latter case it is sometimes possible to characterize the consensus

value more explicitly. For example, Narayanan and Niyogi [6] show that in one version of this update rule, for all $i \in [n]$, $Q_i(t) \rightarrow \mathbf{q}^* \in \mathcal{Y} - \{\mathbf{0}\}$ and $\mathbb{E}[\mathbf{q}^*] = \Pi$.

Lemma 3 (Singleton convergence [6]). *Suppose Condition 1 holds. If $\mathbf{P}(t) = \mathbf{Q}(t)$ and $W(t) = W$ and $B(t) = B$ are independent of time t and (random) social samples $\mathbf{Y}(t)$, then*

$$\mathbb{P}\left\{\lim_{t \rightarrow \infty} \mathbf{Q}(t) = \mathbf{1}\mathbf{q}^*, \mathbf{q}^* \in \mathcal{Y} - \{\mathbf{0}\}\right\} = 1.$$

B. Almost sure convergence

The main result of this paper is obtaining sufficient conditions under which the update rule in (1) converges to a state in which all agents have the same estimate of the histogram Π . In general, the limiting state need not equal Π , but in some cases the process does converge almost surely to Π . To show almost sure convergence we will need some additional conditions. Condition 2 is a standard assumption from stochastic approximation on the step size used in the iteration. A typical choice of step size is $\delta(t) = \Theta(1/t)$ which we used in the examples earlier.

Condition 2 (Decreasing step size). *The sequence $\delta(t) \rightarrow 0$ with $\delta(t) \geq 0$ satisfies*

$$\sum_{t=1}^{\infty} \delta(t) = \infty \quad \text{and} \quad \sum_{t=1}^{\infty} \delta(t)^2 < \infty.$$

Condition 3 states that the expected weight matrices $\bar{H}(t)$ at each time are perturbations around a fixed time-invariant contraction matrix \bar{H} . This condition is satisfied in all three examples of interest above. Furthermore, it allows us to simplify the analysis. Note that it seems to us that this condition is rather technical and can be relaxed at the cost of more cumbersome notation and complicated analysis. As a result, relaxing this assumption remains an area of future work.

Condition 3 (Limiting dynamics). *There exists a symmetric matrix \bar{H} such that*

$$\bar{H}(t) = \bar{H} + \tilde{H}(t),$$

where $\tilde{H}_{ij}(t) = O(\delta(t))$. Furthermore, if λ is an eigenvalue of \bar{H} then we have $|\lambda| < 1$ and in particular $\bar{H}\mathbf{1} = \mathbf{0}$. That is, \bar{H} is a contraction.

Condition 4 implies that the perturbation term $\mathbf{C}(t)$ in (2) vanishes as the step size decreases. This condition guarantees that the mean dynamics given by \bar{H} govern the convergence to the final consensus state.

Condition 4 (Bounded perturbation). *We have*

$$\|\mathbb{E}[\mathbf{C}(t)|\mathcal{F}_t]\| = O(\delta(t)).$$

Given these three conditions and the conditions on the coefficients we can show that the agent estimates in (1) converge almost surely to a random common consensus state whose expected value, by Lemma 2, is equal to Π . Thus for almost every sample path of the update rule, the estimates converge to a common value, but that value may differ across sample paths. The expectation of the random consensus state is the true average.

Theorem 1. *Suppose Conditions 1, 2, 3, and 4 hold. Then the estimate of each node i governed by the update rule (1) converges almost surely to a random variable $\mathbf{q}^* \in \mathbb{P}(\mathcal{Y})$ which is a consensus state. That is, $Q(t) \rightarrow \mathbf{Q}^* = \mathbf{1}\mathbf{q}^*$ where $\mathbb{E}[\mathbf{Q}^*] = \mathbf{1}\Pi$.*

Proof. The result follows from a general convergence theorem for stochastic approximation algorithms [38, Theorem 5.2.1]. Define

$$\mathbf{V}(t) = \bar{H}(t)\mathbf{Q}(t) + \mathbf{C}(t) + \mathbf{M}(t).$$

Several additional conditions are needed to guarantee properties of $\mathbf{V}(t)$ which ensure convergence of the update in (2). Condition 2 guarantees that the step sizes decay slowly enough to take advantage of the almost sure convergence of stochastic approximation procedures. The limit point is a fixed point of the matrix map \bar{H} and Conditions 1 and 3 show that this map is a contraction so the limit points are consensus states. The final condition is that the noise in the updates can be ‘‘averaged out.’’ This follows in part because the process is bounded, and in part because Condition 4 shows that the perturbation must be decaying sufficiently fast.

We must verify a number of conditions [38, p.126] to use this theorem.

- 1) Condition 2 shows that the step sizes not summable but are square summable [38, (5.1.1) and (A2.4)].
- 2) The iterates are bounded in the sense that $\sup_t \mathbb{E}[\|\mathbf{V}(t)\|^2] < \infty$. This follows because $Q_i(t)$ is a probability distribution for all t , so the updates must also be bounded [38, (A2.1)].
- 3) If we take the expected update

$$\begin{aligned} \mathbb{E}[\mathbf{V}(t)|\mathcal{F}_t] &= \bar{H}(t)\mathbf{Q}(t) + \mathbb{E}[\mathbf{C}(t)|\mathcal{F}_t] \\ &= \bar{H}\mathbf{Q}(t) + \tilde{H}(t)\mathbf{Q}(t) + \mathbb{E}[\mathbf{C}(t)|\mathcal{F}_t], \end{aligned}$$

so we can write this conditional expectation as the sum of a measurable function $\bar{H}\mathbf{Q}(t)$ and a random yet diminishing perturbation $\tilde{H}(t)\mathbf{Q}(t) + \mathbb{E}[\mathbf{C}(t)|\mathcal{F}_t]$. Furthermore, from Condition 3, the map \bar{H} is continuous [38, (A2.2)-(A2.3)].

- 4) The final thing to check [38, (A2.5)] is that the random perturbation in the expected update decays sufficiently quickly:

$$\begin{aligned} &\sum_{t=1}^{\infty} \delta(t) \left\| \tilde{H}(t)\mathbf{Q}(t) + \mathbb{E}[\mathbf{C}(t)|\mathcal{F}_t] \right\| \\ &\leq \sum_{t=1}^{\infty} \delta(t) \left\| \tilde{H}(t) \right\| \|\mathbf{Q}(t)\| + \sum_{t=1}^{\infty} \delta(t) \|\mathbb{E}[\mathbf{C}(t)|\mathcal{F}_t]\| \\ &< \infty. \end{aligned}$$

The last step follows from Conditions 2 and 4 as well as boundedness of $\mathbf{Q}(t)$.

Applying Theorem 5.2.1 of Kushner and Yin [38] shows that the estimates converge to a limit set of the linear map \bar{H} . Furthermore, from Condition 3 we know \bar{H} is a contraction with a single eigenvector at $\mathbf{1}$. In other words, the limit points are of the form $\mathbf{Q}^* = \mathbf{1}\mathbf{q}^*$ where every row is identical. \square

The preceding theorem shows that the updates converge almost surely to a limit when the step sizes are decreasing,

even though as shown in [6], we know that decreasing step size is not necessary for almost sure convergence. So far the algorithm has no provable advantage to that of [6], in that each node’s estimate converges to a consensus state \mathbf{q}^* , but \mathbf{q}^* need not equal Π . However, by ensuring that the sample path of the algorithm is ‘‘mean preserving’’ (the sum of the j -th components of all $Q_i(t)$ ’s is equal to Π_j), this consensus limit becomes equal to Π .

Condition 5 (Mean preservation). *The average of the node estimates is Π*

$$\mathbf{1}^\top \mathbf{Q}(t) = \Pi \quad \forall t.$$

Corollary 1. *Suppose Conditions 1, 2, 3, 4, and 5 hold. Then $Q(t) \rightarrow \mathbf{Q}^*$ almost surely, where $\mathbf{Q}^* = \mathbf{1}\Pi$ almost surely.*

C. Rate of convergence

We now turn to bounds on the expected squared error of in the case where $\mathbf{Q}(t) \rightarrow \mathbf{Q}^*$ almost surely.

Theorem 2 (Rate of convergence). *Suppose that Conditions 1, 2, 3, and 4 also hold. Then there exists a constant C such that*

$$\mathbb{E} \left[\|\mathbf{Q}(t) - \mathbf{Q}^*\|^2 \right] \leq C\delta(t).$$

Proof. First note that in the process (1), $Q_i(t)$ is a probability distribution, so the entire process lies in a bounded compact set, and under Conditions 3 and 4 we can write the iteration as

$$\mathbf{Q}(t+1) = \mathbf{Q}(t) + \delta(t) [\bar{H}\mathbf{Q}(t) + \mathbf{D}(t) + \mathbf{M}(t)],$$

where the perturbation term $\|\mathbb{E}[\mathbf{D}(t)|\mathcal{F}_t]\| = O(\delta(t))$. We can now apply Theorem 24 of Benveniste et al. [39, p. 246], which requires checking similar conditions as the previous Theorem.

- 1) Condition 2 shows the step sizes are not summable [39, (A.1)].
- 2) Treat the tuple of random variables $(\mathbf{A}(t), \mathbf{B}(t), \mathbf{W}(t), \mathbf{Y}(t))$ as a state variable $\mathbf{S}(t)$. This state is measurable with respect to \mathcal{F}_t , and there exists a conditional probability distribution corresponding to the update [39, (A.2)].
- 3) Let

$$\mathbf{N}(t) = \mathbf{M}(t) + \mathbf{C}(t) - \mathbb{E}[\mathbf{D}(t)|\mathcal{F}_t],$$

so $\mathbf{N}(t)$ is still a martingale difference. If we define $\mathbf{J}(t) = \frac{1}{\delta(t)}\mathbb{E}[\mathbf{C}(t)|\mathcal{F}_t]$ we can rewrite the iterates as

$$\mathbf{Q}(t+1) = \mathbf{Q}(t) + \delta(t) [\bar{H}\mathbf{Q}(t) + \mathbf{N}(t)] + \delta(t)^2\mathbf{J}(t).$$

Again, the terms $\|\bar{H}\mathbf{Q}(t) + \mathbf{N}(t)\|$ and $\|\mathbf{J}(t)\|$ are bounded by a constant [39, (A.3) and (A.5)].

- 4) Since \bar{H} is a linear contraction, it is also Lipschitz, and the martingale difference $\mathbf{N}(t)$ is bounded, which implies condition (A.4) of Benveniste et al. [39, p. 216].

With the validity of the above conditions, the assertion of the theorem follows directly [39, p. 216]:

$$\mathbb{E} \left[\|\mathbf{Q}(t) - \mathbf{Q}^*\|^2 \right] = O(\delta(t)).$$

\square

D. Example Algorithms Revisited

We can now describe how the results apply to the algorithms described in section II-C.

1) *Averaging with social samples*: Our first algorithm in (8) was one in which the nodes perform a weighted average of the distribution of the messages they receive with their current estimate. The general form of the algorithm was

$$\mathbf{Q}(t+1) = (I - A)\mathbf{Q}(t) + W\mathbf{Y}(t),$$

which corresponds to choosing $\delta(t) = 1$ and $B(t) = 0$. For the specific example in (8), $A = \frac{1}{d_{\max}+1}D$ and $W = \frac{1}{d_{\max}+1}G$, where G is the adjacency matrix of the graph and D is the diagonal matrix of degrees. Furthermore, $\mathbf{P}(t) = \mathbf{Q}(t)$ for all t .

$$\mathbf{Q}(t+1) = \mathbf{Q}(t) + (W - A)\mathbf{Q}(t) + W(\mathbf{Y}(t) - \mathbf{Q}(t)). \quad (13)$$

The term $W - A$ is the graph Laplacian of the graph with edge weights given by W . The following is a corollary of Lemma 2 and Lemma 3.

Corollary 2. *For the update given in (13), the estimates $\mathbf{Q} \rightarrow \mathbf{Q}^*$ almost surely, where \mathbf{Q}^* is a random matrix taking values in the set $\{\mathbf{1}\mathbf{q}^* : \mathbf{q}^* \in \mathcal{Y}\}$ such that $\mathbb{E}[\mathbf{Q}^*] = \mathbf{1}\Pi$.*

Examining (13), we see that the the Laplacian term drives the iteration to a consensus state, but the only stable consensus states are those for which $\mathbf{Y}(t) - \mathbf{Q}(t) = \mathbf{0}$, which means $\mathbf{Y}(t)$ must be equal to $\mathbf{Q}(t)$ almost surely. This means each row of $\mathbf{Q}(t)$ must correspond to a degenerate distribution of the form \mathbf{e}_m .

2) *Averaging with social samples and decaying step size*: The second class of algorithms, exemplified by (9), has the following general form:

$$\mathbf{Q}(t+1) = (1 - \delta(t)A)\mathbf{Q}(t) + \delta(t)W\mathbf{Y}(t),$$

and again $\mathbf{P}(t) = \mathbf{Q}(t)$ for all t . This is really the same as (13) but with a decaying step size $\delta(t)$:

$$\mathbf{Q}(t+1) = \mathbf{Q}(t) + \delta(t)(W - A)\mathbf{Q}(t) + \delta(t)W(\mathbf{Y}(t) - \mathbf{Q}(t)). \quad (14)$$

However, the existence of a decreasing step size means that the iterates under this update behave significantly differently than those governed by (13). The convergence of this algorithm is characterized by Theorems 1 and 2.

Corollary 3. *For the update given in (14) with $\delta(t) = 1/t$, the estimates $\mathbf{Q}(t) \rightarrow \mathbf{Q}^*$ almost surely, where \mathbf{Q}^* is a random matrix in the set $\{\mathbf{1}\mathbf{q}^\top : \|\mathbf{q}\|_1 = 1, q_m > 0\}$ and $\mathbb{E}[\mathbf{Q}^*] = \mathbf{1}\Pi^\top$. Furthermore, $\mathbb{E}[\|\mathbf{Q}(t) - \mathbf{Q}^*\|^2] = O(1/t)$.*

3) *Exchange with social samples and censoring*: The last algorithm in (10) has a more complex update rule, but it is a special case of the generic update

$$\mathbf{Q}(t+1) = \mathbf{Q}(t) - \delta(t)B(t)\mathbf{Y}(t) + \delta(t)W(t)\mathbf{Y}(t). \quad (15)$$

For a fixed weight matrix W , define the the thresholds $\Delta(t) = \delta(t) \sum_{j \in \mathcal{N}_i} W_{ij}$ and the sampling distribution $P_{i,m}(t) = \mathbf{1}(Q_{i,m}(t) > \Delta(t))$. The social samples $\mathbf{Y}(t)$ are sampled

according to this distribution and the weight matrices are defined by

$$W_{ij}(t) = \begin{cases} 0 & i \neq j, Y_i = \mathbf{0} \text{ or } Y_i \neq \mathbf{0} \\ W_{ij} & i \neq j, Y_i = \mathbf{0} \text{ and } Y_i \neq \mathbf{0} \end{cases}$$

$$B_{ii}(t) = \sum_{j \in \mathcal{N}_i} W_{ij}(t)$$

In this algorithm the iterates keep $\sum_i Q_{i,m}(t)$ constant over time by changing the sampling distribution $\mathbf{P}(t)$ over time and by using the weight matrix $B(t)$ to implement a ‘‘mass exchange’’ policy between nodes. At each time, agent i samples a opinion $Z_i(t)$. If $Q_{i,Z_i(t)}(t)$ is large enough, the agent sends $Y_i(t) = Z_i(t)$, giving $\delta(t)W_{ij}$ mass to each neighbor j and subtracting the corresponding mass from its own opinion. If $Q_{i,Z_i(t)}(t)$ is not large enough it exchanges nothing with its neighbors. The distribution $\mathbf{P}(t)$ implements this ‘‘censoring’’ operation. By keeping the total sum on each opinion fixed, Corollary 1 shows that the estimates converge almost surely to Π .

Corollary 4. *For the update given in (15) with $\delta = 1/t$, the estimates $\mathbf{Q}(t) \rightarrow \mathbf{1}\Pi$ almost surely, and $\mathbb{E}[\|\mathbf{Q}(t) - \mathbf{1}\Pi\|^2] = O(1/t)$.*

IV. EMPIRICAL RESULTS

The preceding analysis shows the almost-sure convergence of all node estimates for some social sampling strategies, and in some cases characterizes the rate of convergence. However, the analysis does not capture the effect of problem parameters such as the initial distribution of node values and the network topology. These factors are well known to affect the rate of convergence of many distributed estimation and consensus procedures – in this section we provide some empirical results about these effects.

We considered a number of different topologies for our simulations:

- The $\sqrt{n} \times \sqrt{n}$ grid has vertex set $[\sqrt{n}] \times [\sqrt{n}]$ and edge exists between two nodes whose L_1 distance is 1.
- A star network has a single central vertex which is connected by a single edge to $n - 1$ other vertices.
- An Erdős-Rényi graph [40] on vertex set $[n]$ contains each edge (i, j) independently with probability p . We choose $p = 0.6$.
- A preferential attachment graph [41], [42] is constructed by adding vertices one at a time. A new vertex is connected to an existing vertex with a probability that is a function of the current degree of the vertices. We allowed each new vertex to be connected to 3 preceding vertices.
- A 2-dimensional Watts-Strogatz graph is a grid with randomly ‘‘rewired’’ edges [43]. We chose rewiring probability 0.1.

Details on the random graph models can be found in the `igraph` package for the R statistical programming language [44]. In the simulations we calculated the average of $\frac{1}{n} \|\mathbf{Q}(t) - \mathbf{1}\Pi\|^2$ across runs of the simulation, which is the average mean-squared-error (MSE) per node of the estimates.

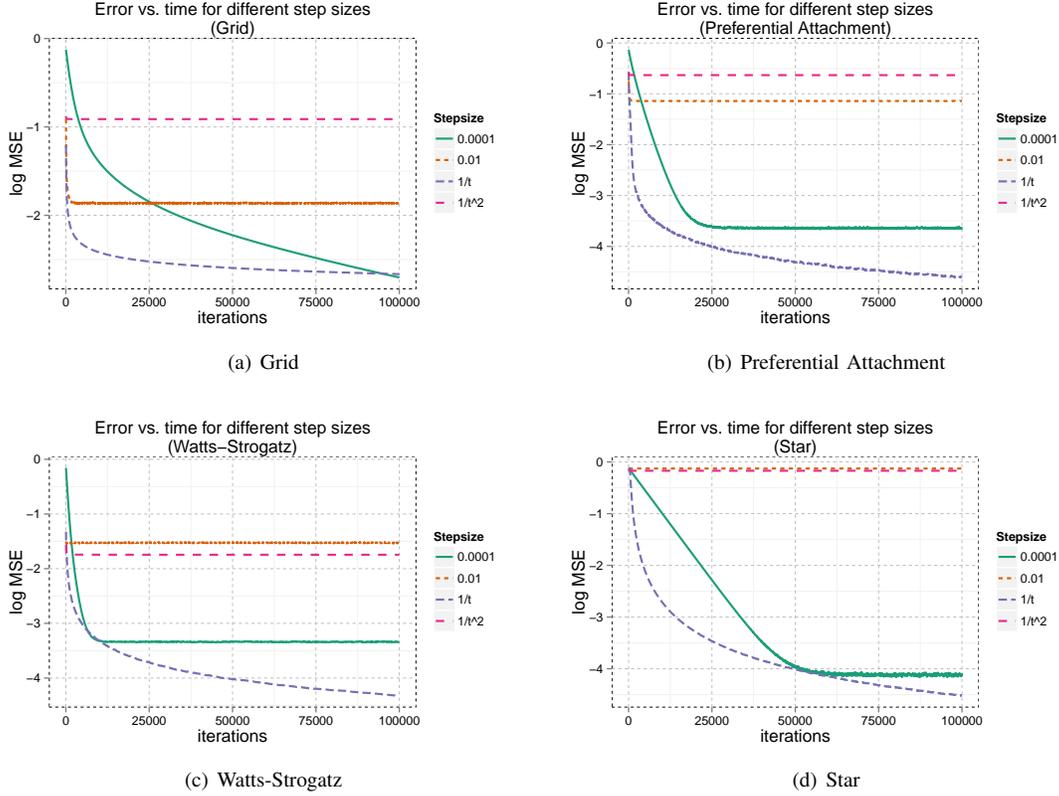


Fig. 4. Average MSE between the agent estimates and the true histogram initial node values versus time for 4 different 100-node graphs : a 10×10 grid, preferential attachment graphs with three edges generated per new node, Watts-Strogatz graphs with rewiring probability 0.1, and a star with one central node and 99 peripheral nodes.

A. Network size and topology

We were interested in finding how the convergence time depends on the step size and network topology. To investigate this we simulated the grid, preferential attachment, Watts-Strogatz, and star topologies described above on networks of $n = 100$ nodes with $M = 5$. The initial node values were drawn i.i.d. from a distribution $(0.1, 0.25, 0.15, 0.3, 0.2)$ on 5 elements. Simulations were averaged over 100 instances of the network and initial values.

Figure 3 shows that the estimates converge almost surely to the true histogram Π when $\delta(t) = 1/t$. While our theoretical analysis was for this case, in practice stochastic optimization is often used with a constant step size because the algorithm converges faster to a neighborhood of the optimal solution when $\delta(t)$ is appropriately small. In order to assess if this is the case in our model, we simulated variants of the algorithm in (10) with different settings for $\delta(t)$.

Figure 4 shows the error between the local estimates and the true histogram of initial node values under four different topologies and four different choices for $\delta(t)$. For $\delta(t) = 1/t$ the algorithm satisfies the conditions of the theorem and we can see the rather rapid convergence to the mean. If $\delta(t)$ is constant, then the error does not converge to 0 but can still be quite small if the step size is small. This is similar to the fixed-weight algorithm with a weight matrix that has very small off-diagonal entries. By contrast, the weight sequence $\delta(t) = 1/t^2$ decays too quickly and there is a large residual MSE.

We see a greater effect on the convergence time by looking at different graph topologies for the same number of nodes. For graphs with fast mixing time such as the preferential attachment and Watts-Strogatz model, the error decreases much more rapidly than for the grid or star. This suggests that the mixing time of the random walk associated with the weight matrix of the algorithm should affect the rate of convergence, as is the case in other consensus algorithms. The effect of choosing weight sequences that do not guarantee almost sure convergence also varies depending on the network topology. For sparsely connected networks like the star, the performance is quite poor unless the weight sequence is chosen appropriately. However, for denser networks like the Watts-Strogatz model, the difference may be more modest.

B. The effect of the initial distribution

The size and shape of the histogram to be estimated also affects the rate of convergence. To illustrate this, we sampled initial values from a uniform distribution on M items for different values of M . Figure 5 shows the average time to get to an MSE of 10^{-2} versus M for this scenario. Here the effect of the network topology is quite pronounced; topological features such as the network diameter seem to have a significant impact on the time to convergence.

To see the effect of the number of nonzero elements in a fixed example we simulated the Erdős-Rényi, grid, preferential attachment, and Watts-Strogatz models for $n = 100$ with

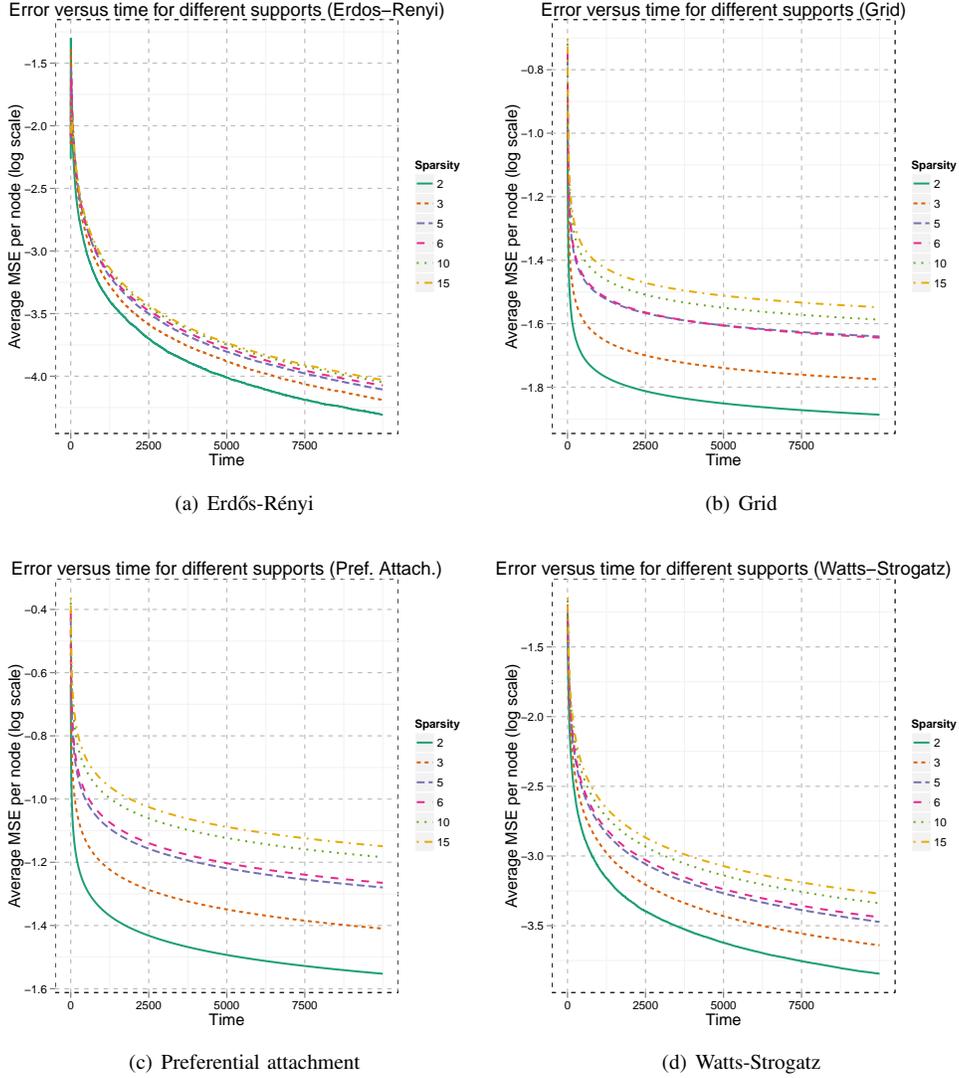


Fig. 6. Average MSE per node versus time (on a $\log_{10}(\cdot)$ scale) versus time for different support sizes for four different network topologies. The distribution is uniform on a subset of size $M^* = 2$ to $M^* = 15$ (the Sparsity level) out of $M = 150$. The MSE decays much more quickly for sparser distributions for the same alphabet size M .

initial values for the nodes sampled from a set of sparse distributions with different values of M . More precisely, we considered a sparse distribution over $M = 150$ bins where the actual distribution of opinions is uniform on an (unknown) sparse subset of $M^* \ll M$ bins where in our simulations, M^* ranges from $M^* = 2$ to $M^* = 15$. The effect of this “sparsity” is shown in Figure 6, where the log MSE per node is plotted against the number of time steps of the algorithm. Here the difference in the network topologies is more stark – for the Erdős-Rényi graph the effect of changing the number of elements is negligible, but the average MSE per node increases in the other three graph models. The difference is greatest in the preferential attachment model, where the increase in M corresponds to a nearly linear increase in the log MSE.

Next we consider a closely related question regarding the shape of the histogram to be estimated. In particular, we considered initial distributions which are heavily concentrated on a few elements but still contain many elements with relatively

low popularity. Specifically, in our simulations we chose the initial values to be drawn from the following distribution

$$\Pi = \left(0.38, 0.38, \frac{0.24}{M-2}, \dots, \frac{0.24}{M-2} \right) \quad (16)$$

for values of M ranging from $M = 5$ to $M = 26$. We simulated each network 50 times, uniformly assigning the initial values to the nodes. The average error is shown in Figure 7. Here we see that when the distribution is biased such that most of the weight is on the first two elements, the support size M does not have an appreciable effect on the convergence time. What Figure 7 suggests is that the shape of the distribution is more important than the support. This is not surprising – because we are measuring squared error, elements in Π which are small will contribute relatively little to the overall error and so in this sense the uniform distribution is the “worst case” for convergence. In these scenarios, different measures of convergence may be important, such as the Kullback-Leibler

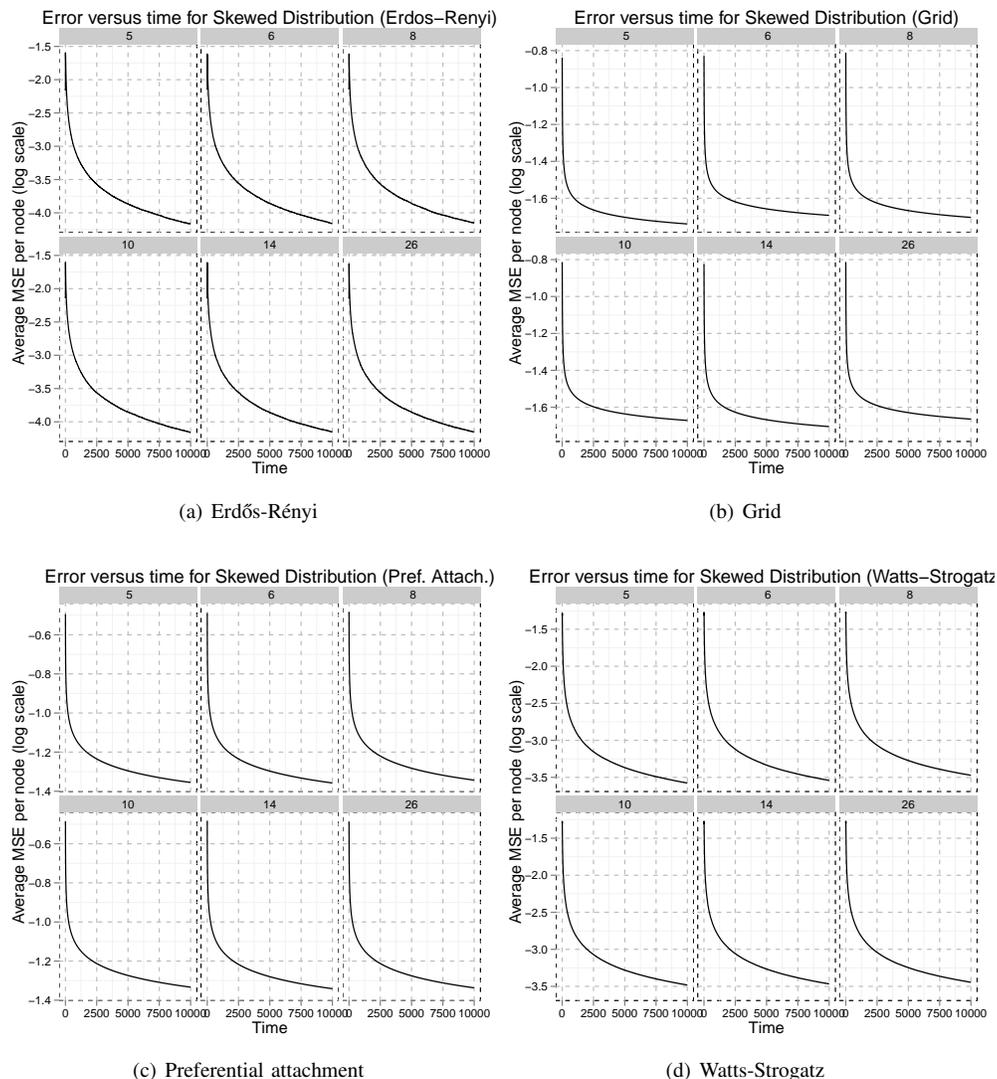


Fig. 7. Average MSE per node versus time (on a $\log_{10}(\cdot)$ scale) versus time for different M for a distribution in (16) that is skewed with larger M . In general, the error is dominated by the convergence on the larger elements of the histogram.

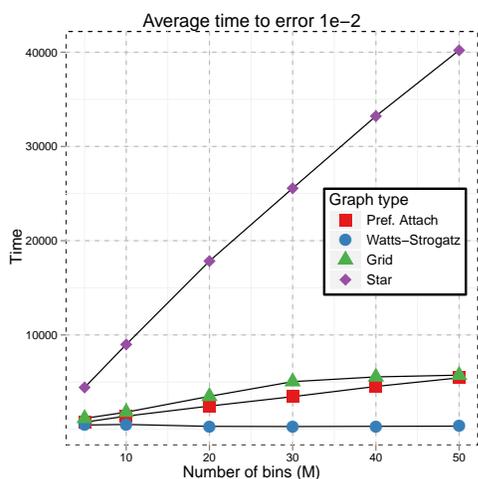


Fig. 5. Time to get to MSE of 10^{-2} , averaged across nodes, versus M for a uniform distribution on the four different network topologies with $n = 100$ nodes.

divergence between the estimated distributions and Π . Other quantities related to Π may impact the rate of convergence of the algorithm.

V. DISCUSSION

In this paper we studied a simple model of message passing in which nodes or agents communicate random messages generated from their current estimates of a global distribution. The message model is inspired by models of social messaging in which agents communicate only a part of their current beliefs at each time. This family of processes contains several interesting instances, including a recent consensus-based model for language formation and an exchange-based algorithm that results in agents learning the true distribution of initial opinions in the network [6]. To analyze this latter process we found a stochastic optimization procedure corresponding to the algorithm. The simulation results confirm the theory and also show that while the topology of the

network affects the rate of convergence, the shape of the overall histogram Π may play larger role than its support size when considering L_2 convergence.

One interesting theoretical question is whether the error $\sqrt{t}(\mathbf{Q}(t) - \mathbf{Q}^*)$ converges to a normal distribution when $\delta(t) = 1/t$. Such a result was obtained by Rajagopal and Wainwright [37] for certain cases of noisy communication in consensus schemes for scalars. They showed a connection between the network topology and the covariance of the normalized asymptotic error. Such a result will not transfer immediately to our scenario because of the additional perturbation term $\mathbf{C}(t)$. However, because this term decays rapidly, we do not believe it will impact the covariance matrix. Characterizing the asymptotic distribution of the error in terms of the graph topology, M , and Π may yield additional insights into the convergence rates in terms of measures other than L_2 norm of error vector.

The results in this paper also apply to the “gossip” scenario wherein only one pair of nodes exchanges messages at a time. This corresponds to selecting a random graph $\mathcal{G}(t)$ which contains only a single edge. In terms of time, the convergence in this setting will be slower because only one pair of messages is exchanged in a single time slot. The analysis framework is fairly general – to get the almost-sure convergence we need mild assumptions on the message distributions. Both finding other interesting instances of the algorithm and extending the analysis for metrics such as divergence and other statistical measures are interesting directions for future work. Solving the latter problem may yield some new techniques for analyzing other statistical procedures which can be cast as stochastic optimization, such as empirical risk minimization.

This model of random message passing may be useful in other contexts such as inference and optimization. Stochastic coordinate ascent is used in convex optimization over large data sets; extending this framework to the distributed optimization setting is a promising future direction, especially for high-dimensional problems. In belief propagation, stochastic generation of beliefs can ensure convergence even when the state space is very large [45]. Finally, the framework here can also be applied to a model for distributed parametric inference in social networks [46]–[48] in which agents both observe and communicate over time. In these applications and in others, the same ideas behind the social sampling model in this paper appear to be useful in reducing the message complexity while allowing consistent inference in distributed setting.

ACKNOWLEDGEMENTS

The authors would like to thank Vivek Borkar for helpful suggestions regarding stochastic approximation techniques, and the anonymous reviewers, whose detailed feedback helped improve the organization and exposition of the paper.

REFERENCES

[1] A. D. Sarwate and T. Javidi, “Distributed learning from social sampling,” in *Proceedings of the 46th Annual Conference on Information Sciences and Systems (CISS)*, Princeton, NJ, USA, March 2012. [Online]. Available: <http://dx.doi.org/10.1109/CISS.2012.6310767>

[2] —, “Opinion dynamics and distributed learning of distributions,” in *Proceedings of the 49th Annual Allerton Conference on Communication, Control and Computation*, Monticello, IL, USA, September 2011. [Online]. Available: <http://dx.doi.org/10.1109/Allerton.2011.6120297>

[3] J. French, Jr., “A formal theory of social power,” *Psychological Review*, vol. 63, no. 3, pp. 181–194, May 1956. [Online]. Available: <http://dx.doi.org/10.1037/h0046123>

[4] F. Harary, *Studies in Social Power*. Ann Arbor, MI: Institute for Social Research, 1959, ch. A criterion for unanimity in French’s theory of social power, pp. 168–182.

[5] M. H. DeGroot, “Reaching a consensus,” *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974. [Online]. Available: <http://www.jstor.org/stable/2285509>

[6] H. Narayanan and P. Niyogi, “Language evolution, coalescent processes, and the consensus problem on a social network,” October 2011, under revision.

[7] R. J. Aumann, “Agreeing to disagree,” *The Annals of Statistics*, vol. 4, no. 6, pp. 1236–1239, 1976. [Online]. Available: <http://www.jstor.org/stable/2958591>

[8] V. Borkar and P. Varaiya, “Asymptotic agreement in distributed estimation,” *IEEE Transactions on Automatic Control*, vol. AC-27, no. 3, pp. 650–655, June 1982. [Online]. Available: <http://dx.doi.org/10.1109/TAC.1982.1102982>

[9] J. N. Tsitsiklis and M. Athans, “Convergence and asymptotic agreement in distributed decision problems,” *IEEE Transactions on Automatic Control*, vol. AC-29, no. 1, pp. 42–50, January 1984. [Online]. Available: <http://dx.doi.org/10.1109/TAC.1984.1103385>

[10] J. Fax, “Optimal and cooperative control of vehicle formation,” Ph.D. dissertation, California Institute of Technology, Pasadena, CA, 2001.

[11] J. Fax and R. Murray, “Information flow and cooperative control of vehicle formations,” *IEEE Transactions on Automatic Control*, vol. 49, no. 9, pp. 1465–1476, September 2004. [Online]. Available: <http://dx.doi.org/10.1109/TAC.2004.834433>

[12] R. Olfati-Saber and R. 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, September 2004. [Online]. Available: <http://dx.doi.org/10.1109/TAC.2004.834113>

[13] R. Agaev and P. Chebotarev, “The matrix of maximum out forests of a digraph and its applications,” *Automation and Remote Control*, vol. 61, pp. 1424–1450, September 2000. [Online]. Available: <http://dx.doi.org/10.1023/A:1002862312617>

[14] R. Agaev and P. Y. Chebotarev, “Spanning forests of a digraph and their applications,” *Automation and Remote Control*, vol. 62, pp. 443–466, March 2001. [Online]. Available: <http://dx.doi.org/10.1023/A:1002862312617>

[15] P. Chebotarev and R. Agaev, “Forest matrices around the Laplacian matrix,” *Linear Algebra and its Applications*, vol. 356, pp. 253–274, 2002. [Online]. Available: [http://dx.doi.org/10.1016/S0024-3795\(02\)00388-9](http://dx.doi.org/10.1016/S0024-3795(02)00388-9)

[16] P. Chebotarev, “Comments on ‘Consensus and Cooperation in Networked Multi-Agent Systems’,” *Proceedings of the IEEE*, vol. 98, no. 7, pp. 1353–1354, July 2010. [Online]. Available: <http://dx.doi.org/10.1109/JPROC.2010.2049911>

[17] R. Olfati-Saber, J. Fax, and R. Murray, “Reply to ‘Comments on ‘Consensus and Cooperation in Networked Multi-Agent Systems’ ’,” *Proceedings of the IEEE*, vol. 98, no. 7, pp. 1354–1355, July 2010. [Online]. Available: <http://dx.doi.org/10.1109/JPROC.2010.2049912>

[18] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, “Randomized gossip algorithms,” *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2508–2530, June 2006. [Online]. Available: <http://dx.doi.org/10.1109/TIT.2006.874516>

[19] R. Olfati-Saber, J. Fax, and R. M. Murray, “Consensus and cooperation in networked multi-agent systems,” *Proceedings of the IEEE*, vol. 95, no. 1, pp. 215–233, January 2007. [Online]. Available: <http://dx.doi.org/10.1109/JPROC.2006.887293>

[20] F. Fagnani and S. Zampieri, “Randomized consensus algorithms over large scale networks,” *IEEE Journal on Selected Areas in Communication*, vol. 26, no. 4, pp. 634–649, May 2008. [Online]. Available: <http://dx.doi.org/10.1109/JSAC.2008.080506>

[21] A. Olshevsky and J. N. Tsitsiklis, “Convergence speed in distributed consensus and averaging,” *SIAM J. Control and Optimization*, vol. 48, no. 1, pp. 33–55, 2009. [Online]. Available: <http://dx.doi.org/10.1137/060678324>

[22] A. Dimakis, A. Sarwate, and M. Wainwright, “Geographic gossip: Efficient averaging for sensor networks,” *IEEE Transactions on Signal Processing*, vol. 56, no. 3, pp. 1205–1216, 2008. [Online]. Available: <http://dx.doi.org/10.1109/TSP.2007.908946>

- [23] F. Benezit, A. G. Dimakis, P. Thiran, and M. Vetterli, "Order-optimal consensus through randomized path averaging," *IEEE Transactions on Information Theory*, vol. 56, no. 10, pp. 5150–5167, October 2010. [Online]. Available: <http://dx.doi.org/10.1109/TIT.2010.2060050>
- [24] T. C. Aysal, M. E. Yildiz, A. D. Sarwate, and A. Scaglione, "Broadcast gossip algorithms for consensus," *IEEE Transactions on Signal Processing*, vol. 57, no. 7, pp. 2748–2761, Jul. 2009. [Online]. Available: <http://dx.doi.org/10.1109/TSP.2009.2016247>
- [25] A. D. Sarwate and A. G. Dimakis, "The impact of mobility on gossip algorithms," *IEEE Transactions on Information Theory*, vol. 58, no. 3, pp. 1731–1742, March 2012. [Online]. Available: <http://dx.doi.org/10.1109/TIT.2011.2177753>
- [26] A. G. Dimakis, S. Kar, J. M. Moura, M. G. Rabbat, and A. Scaglione, "Gossip algorithms for distributed signal processing," *Proceedings of the IEEE*, vol. 98, no. 11, pp. 1847–1864, November 2010. [Online]. Available: <http://dx.doi.org/10.1109/JPROC.2010.2052531>
- [27] T. C. Aysal, M. J. Coates, and M. G. Rabbat, "Distributed average consensus with dithered quantization," *IEEE Transactions on Signal Processing*, vol. 56, no. 10, pp. 4905–4918, 2008. [Online]. Available: <http://dx.doi.org/10.1109/TSP.2008.927071>
- [28] A. Nedic, A. Olshevsky, A. Ozdaglar, and J. Tsitsiklis, "On distributed averaging algorithms and quantization effects," *IEEE Transactions on Automatic Control*, vol. 54, no. 11, pp. 2506–2517, November 2009. [Online]. Available: <http://dx.doi.org/10.1109/TAC.2009.2031203>
- [29] R. Carli, F. Bullo, and S. Zampieri, "Quantized average consensus via dynamic coding/decoding schemes," *International Journal of Robust and Nonlinear Control*, vol. 20, no. 2, pp. 156–175, 2010. [Online]. Available: <http://dx.doi.org/10.1002/rnc.1463>
- [30] A. Kashyap, T. Başar, and R. Srikant, "Quantized consensus," *Automatica*, vol. 43, no. 7, pp. 1192–1203, 2007. [Online]. Available: <http://dx.doi.org/10.1016/j.automatica.2007.01.002>
- [31] R. Carli, F. Fagnani, P. Frasca, and S. Zampieri, "Gossip consensus algorithms via quantized communication," *Automatica*, vol. 46, no. 1, pp. 70–80, January 2010. [Online]. Available: <http://dx.doi.org/10.1016/j.automatica.2009.10.032>
- [32] M. Zhu and S. Martínez, "On the convergence time of asynchronous distributed quantized averaging algorithms," *IEEE Transactions on Automatic Control*, vol. 56, no. 2, pp. 386–390, February 2011. [Online]. Available: <http://dx.doi.org/10.1109/TAC.2010.2093276>
- [33] J. Lavaei and R. M. Murray, "Quantized consensus by means of gossip algorithm," *IEEE Transactions on Automatic Control*, vol. 57, no. 1, pp. 19–32, January 2012. [Online]. Available: <http://dx.doi.org/10.1109/TAC.2011.2160593>
- [34] K. Srivastava and A. Nedic, "Distributed asynchronous constrained stochastic optimization," *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, no. 4, pp. 772–790, August 2011. [Online]. Available: <http://dx.doi.org/10.1109/JSTSP.2011.2118740>
- [35] M. Yildiz and A. Scaglione, "Coding with side information for rate-constrained consensus," *IEEE Transactions on Signal Processing*, vol. 56, no. 8, pp. 3753–3764, August 2008. [Online]. Available: <http://dx.doi.org/10.1109/TSP.2008.919636>
- [36] S. Kar and J. Moura, "Distributed consensus algorithms in sensor networks: Quantized data and random link failures," *IEEE Transactions on Signal Processing*, vol. 58, no. 3, pp. 1383–1400, March 2010. [Online]. Available: <http://dx.doi.org/10.1109/TSP.2009.2036046>
- [37] R. Rajagopal and M. Wainwright, "Network-based consensus averaging with general noisy channels," *IEEE Transactions on Signal Processing*, vol. 59, no. 1, pp. 373–385, January 2011. [Online]. Available: <http://dx.doi.org/10.1109/TSP.2010.2077282>
- [38] H. J. Kushner and G. G. Yin, *Stochastic Approximation and Recursive Algorithms and Applications*, 2nd ed. Springer, 2010.
- [39] A. Benveniste, M. Métivier, and P. Priouret, *Adaptive Algorithms and Stochastic Approximations*, ser. Applications of Mathematics. Berlin: Springer-Verlag, 1990, no. 22.
- [40] P. Erdős and A. Rényi, "On the evolution of random graphs," in *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, vol. 5, 1960, pp. 17–61. [Online]. Available: https://www.renyi.hu/~p_erdos/1960-10.pdf
- [41] A.-L. Barabási and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, no. 5439, p. 509, October 1999. [Online]. Available: <http://dx.doi.org/10.1126/science.286.5439.509>
- [42] B. Bollobás and O. Riordan, "The diameter of a scale-free random graph," *Combinatorica*, vol. 24, no. 1, pp. 5–34, 2004. [Online]. Available: <http://dx.doi.org/10.1007/s00493-004-0002-2>
- [43] D. J. Watts and S. H. Strogatz, "Collective dynamics of 'small-world' networks," *Nature*, vol. 393, no. 6684, pp. 440–442, 1998.
- [44] G. Csárdi and T. Nepusz, "The igraph software package for complex network research," *InterJournal Complex Systems*, no. 1695, 2006.
- [45] N. Noorshams and M. J. Wainwright, "Stochastic belief propagation: A low-complexity alternative to the sum-product algorithm," *IEEE Transactions on Information Theory*, vol. 59, no. 4, pp. 1981–2000, April 2013. [Online]. Available: <http://dx.doi.org/10.1109/TIT.2012.2231464>
- [46] P. Molavi and A. Jadbabaie, "Network structure and efficiency of observational social learning," in *Proceedings of the 51st IEEE Conference on Decision and Control*, Maui, HI, USA, 2012. [Online]. Available: <http://dx.doi.org/10.1109/CDC.2012.6426454>
- [47] 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. [Online]. Available: <http://dx.doi.org/10.1016/j.geb.2012.06.001>
- [48] P. Molavi, K. R. Rad, A. Tahbaz-Salehi, and A. Jadbabaie, "On consensus and exponentially fast social learning," in *Proceedings of the 2012 American Control Conference*, 2012.



Anand D. Sarwate (S'99–M'09) received B.S. degrees in electrical engineering and computer science and mathematics from the Massachusetts Institute of Technology (MIT), Cambridge in 2002 and the M.S. and Ph.D. degrees in electrical engineering from the Department of Electrical Engineering and Computer Sciences (EECS) at the University of California, Berkeley (U.C. Berkeley).

He is currently an Assistant Professor in the Department of Electrical and Computer Engineering at Rutgers, The State University of New Jersey, since January 2014. He was previously a Research Assistant Professor from 2011–2013 at the Toyota Technological Institute at Chicago; prior to this he was a Postdoctoral Researcher from 2008–2011 at the University of California, San Diego. His research interests include information theory, machine learning, distributed signal processing and optimization, and privacy and security.

Dr. Sarwate received the Samuel Silver Memorial Scholarship Award and the Demetri Angelakos Memorial Award from the EECS Department at U.C. Berkeley. He was awarded an NDSEG Fellowship from 2002 to 2005. He is a member of Phi Beta Kappa and Eta Kappa Nu.



Tara Javidi (S'96–M'02–SM'12) studied electrical engineering at the Sharif University of Technology, Tehran, Iran, from 1992 to 1996. She received the M.S. degrees in electrical engineering (systems) and applied mathematics (stochastics) and Ph.D. degree in electrical engineering and computer science from the University of Michigan, Ann Arbor, in 1998, 1999, and 2002, respectively.

From 2002 to 2004, she was an Assistant Professor with the Electrical Engineering Department, University of Washington, Seattle. She joined the University of California, San Diego, in 2005, where she is currently an Associate Professor of electrical and computer engineering. Her research interests are in communication networks, stochastic resource allocation, and wireless communications.

Dr. Javidi was a Barbour Scholar during the 1999–2000 academic year and received an NSF CAREER Award in 2004.