DISTRIBUTED AVERAGE CONSENSUS USING PROBABILISTIC QUANTIZATION

Tuncer C. Aysal, Mark Coates, and Michael Rabbat

Department of Electrical and Computer Engineering, McGill University, Montreal, QC
{tuncer.aysal,mark.coates,michael.rabbat}@mcgill.ca

ABSTRACT

In this paper, we develop algorithms for distributed computation of averages of the node data over networks with bandwidth/power constraints or large volumes of data. Distributed averaging algorithms fail to achieve consensus when deterministic uniform quantization is adopted. We propose a distributed algorithm in which the nodes utilize probabilistically quantized information to communicate with each other. The algorithm we develop is a dynamical system that generates sequences achieving a consensus, which is one of the quantization values. In addition, we show that the expected value of the consensus is equal to the average of the original sensor data. We report the results of simulations conducted to evaluate the behavior and the effectiveness of the proposed algorithm in various scenarios.

Index Terms— Distributed algorithms, average consensus, sensor networks

1. INTRODUCTION

Many applications envisioned for sensor networks consist of low–power and low–cost nodes. Foreseen applications such as data fusion and distributed coordination require distributed function computation/parameter estimation under bandwidth and power constraints. Other applications such as camera networks and distributed tracking demands communication of large volumes of data. When the power and bandwidth constraints, or large volume data sets are considered, communication with unquantized values is impractical.

Distributed average consensus, in ad hoc networks, is an important issue in distributed agreement and synchronization problems [1] and is also a central topic for load balancing (with divisible tasks) in parallel computers [2]. More recently, it has also found applications in distributed coordination of mobile autonomous agents [3] and distributed data fusion in sensor networks [4]. In this paper, we focus on a particular class of iterative algorithms for average consensus, widely used in the applications cited above. Each node updates its state by adding a weighted sum of the values at local nodes, *i.e.*,

$$x_i(t+1) = W_{ii}x_i(t) + \sum_{j\in\mathcal{N}_i} W_{ij}x_j(t) \tag{1}$$

for $i = 1, 2, \ldots, N$, and $t \geq 0$ [4, 5]. Here W_{ij} is a weight associated with the edge $\{i, j\}$ and N is the total number of nodes. Furthermore, $\mathcal{N}_i$ denotes the set of nodes that have a direct (bidirectional) communication link with node i. The state at each node in the iteration consists of a single real number, which overwrites the previous value. The algorithm is time–independent, *i.e.*, does not depend on t and is meant to compute the average asymptotically.

One major shortcoming of the majority of the algorithms proposed for distributed average consensus is the reliance on the exchange of analog data [4–6], a clear violation of the common assumption of bandwidth and communication power constraints. There has been some limited investigation of consensus under quantized communication. Recently, Yildiz and Scaglione, in [7], explored the impact of quantization noise through modification of the consensus algorithm proposed by Xaio and Boyd [6]. The algorithm in [6] addresses noisy communication links, modeling the case where each nodes receives noise–corrupted versions of the values at its neighbors, the noise being additive and zero–mean. Yildiz and Scaglione note that the noise component can be considered as the quantization noise and they develop a computationally intensive algorithm for this case [7]. They have shown utilizing simulation studies for small N that, if the increasing correlation among the node states is taken into account, the variance of the quantization noise diminishes and nodes converge to a consensus.

Kashyap *et. al* examine quantization effects from a different point of view, restricting their attention to the transmission of integer values, and introducing the definition of "quantized consensus", which is achieved by a vector $\mathbf{x}$ if it belongs to the set $\mathcal{S}$ defined as:

$$\mathcal{S} \triangleq \left\{ \mathbf{x} : \{x_i\}_{i=1}^N \in \{L, L+1\}, \sum_{i=1}^N x_i = T \right\} \tag{2}$$

with T and L being the sum of initial node values and an integer, respectively [8]. They show that their algorithm, under some constraints and with restricted communications schemes, will converge to a quantized consensus. However, the quantized consensus is clearly not a strict consensus, *i.e.*, all nodes do not have the same value. Furthermore, there is no stopping criterion to determine if the algorithm has achieved a quantized consensus. In addition, since the algorithm requires the use of a single link per iteration, the convergence is very slow [8]. Of note is that both of the consensus algorithms discussed above utilize standard deterministic uniform quantization to quantize the data.

In this paper, we propose a simple distributed and iterative scheme to compute the average at each sensor node utilizing only quantized information communication. We adopt the probabilistic quantization (PQ) scheme to quantize the information before transmitting to the neighboring sensors [9]. PQ has been shown to be very effective in decentralized estimation since the quantized state is equal to the analog state in expectation [9]. This makes PQ suitable for average–based algorithms. In the scheme considered here, each node exchanges quantized state information with its neighbors in a simple and bidirectional manner. This scheme does not involve routing messages in the network; instead, it diffuses information across the network by updating each node's data with a weighted average of its neighbors' quantized values. It is shown here that the distributed average computation utilizing probabilistic consensus indeed achieves a consensus and the consensus is one of the quantization levels. Furthermore, the expected value of the achieved consensus is equal to the desired value, *i.e.*, the average of the analog measurements.

The remainder of this paper is organized as follows. Section 2 introduces the distributed average consensus problem and the PQ scheme. The proposed algorithm, along with its properties, is detailed in Section 3. Numerical examples are provided in Section 4. Finally, we conclude with Section 5.

2. PROBLEM FORMULATION

We consider a set of nodes of a network, each with an initial real valued scalar $y_i \in [-U, U]$, where $i = 1, 2, \ldots, N$. We assume strongly connected network topologies, according to which, each node can establish bidirectional noise–free communication with a subset of the nodes in the network. The N–node topology is represented by the $N \times N$ matrix $\mathbf{\Phi} = [\Phi_{ij}]$, where for $i \neq j$, $\Phi_{ij} = 1$ if nodes i and j directly communicate, and $\Phi_{ij} = 0$, otherwise. We assume that transmissions are always successful and that the topology is fixed. Moreover, we define $\mathcal{N}_i \triangleq \{j \in \{1, 2, \ldots, N\} : \Phi_{ij} \neq 0\}$.

Let $\mathbf{1}$ denote the vector of ones. Our goal is to develop a distributed iterative algorithm that computes at every node in the network, the value $\overline{\mathbf{y}} \triangleq (N)^{-1}\mathbf{1}^{\mathrm{T}}\mathbf{y}$ utilizing quantized state information communication. We hence aim to design a system such that the states at all nodes converge to a consensus and the statistical expectation of the consensus achieved, in the limit, is the average of the initial states.

Remark 1. *When the observations follow $y_i = \theta + n_i$ where $i = 1, 2, \ldots, N$ and θ is the scalar to be estimated, and $\{n_i\}_{i=1}^N$ are independent and identically distributed (i.i.d.) zero–mean Gaussian with variance σ^2, the maximum likelihood estimate is given by the average, $\hat{\theta} = (N)^{-1}\mathbf{1}^{\mathrm{T}}\mathbf{y}$ with the associated mean square error σ^2/N.*

In the following, we present a brief review of the quantization scheme adopted in this paper. Suppose that the scalar value $x_i \in \mathbb{R}$ is bounded to a finite interval $[-U, U]$. Furthermore, suppose that we wish to obtain a quantized message q_i with length l bits, where l is application dependent. We therefore have $L = 2^l$ quantization points given by the set $\boldsymbol{\tau} = \{\tau_1, \tau_2, \ldots, \tau_L\}$. The points are uniformly spaced such that $\Delta = \tau_{j+1} - \tau_j$ for $j \in \{1, 2, \ldots, L-1\}$. It follows that $\Delta = 2U/(2^l - 1)$. Now suppose $x_i \in [\tau_j, \tau_{j+1})$ and let $q_i \triangleq \mathcal{Q}(x_i)$ where $\mathcal{Q}(\cdot)$ denotes the PQ operation. Then x_i is quantized in a probabilistic manner:

$$\Pr\{q_i = \tau_{j+1}\} = r, \text{ and } \Pr\{q_i = \tau_j\} = 1 - r \quad (3)$$

where $r = (x_i - \tau_j)/\Delta$. The following lemma, adopted from [9], discusses two important properties of the PQ.

Lemma 1. *[9] Let q_i be an l–bit quantization of $x_i \in [-U, U]$. The message q_i is an unbiased representation of x_i,* i.e.*,*

$$\mathbb{E}\{q_i\} = x_i, \text{ and } \operatorname{var}(q_i) \leq \frac{U^2}{(2^l - 1)^2} \quad (4)$$

where $\operatorname{var}(\cdot)$ denotes the variance of its argument.

3. DISTRIBUTED CONSENSUS USING PQ

In the following, we propose a quantized distributed average consensus algorithm and incorporate PQ into the consensus framework for networks. Furthermore, we analyze the effect of PQ on the consensus algorithm.

At $t = 0$ (after all sensors have taken the measurement), each node initializes its state as $x_i(0) = y_i$, *i.e.*, $\mathbf{x}(0) = \mathbf{y}$ where $\mathbf{x}(0)$ denotes the initial states at the nodes. It then quantizes its state to generate $q_i(0) = \mathcal{Q}(x_i(0))$. At each following step, each node updates its state with a linear combination of its own quantized state and the quantized states at its neighbors

$$x_i(t+1) = W_{ii}q_i(t) + \sum_{j \in \mathcal{N}_i} W_{ij}q_j(t) \quad (5)$$

for $i = 1, 2, \ldots, N$, where $q_j(t) = \mathcal{Q}(x_j(t))$, and t denotes the time step. Also, W_{ij} is the weight on $x_j(t)$ at node i. Moreover, setting $W_{ij} = 0$ whenever $\Phi_{ij} = 0$, the distributed iterative process reduces to the following recursion

$$\mathbf{x}(t+1) = \mathbf{W}\mathbf{q}(t) \quad (6)$$

where $\mathbf{q}(t)$ denotes the quantized state vector. In the sequel, we assume that $\mathbf{W}$, the weight matrix, satisfies the conditions required for asymptotic average consensus without quantization [6]:

$$\mathbf{W}\mathbf{1} = \mathbf{1},\ \mathbf{1}^{\mathrm{T}}\mathbf{W} = \mathbf{1}^{\mathrm{T}},\ \lambda_1 = 1, \text{ and } |\lambda_i| < 1 \quad (7)$$

where $\{\lambda_i\}_{i=1}^N$ denote the eigenvalues of $\mathbf{W}$ in non–increasing order. Weight matrices satisfying the above conditions are easy to find if the underlying graph is connected and non–bipartite, *e.g.*, Maximum–degree and Metropolis weights [6].

The following theorem considers the convergence of the probabilistically quantized distributed average computation.

Theorem 1. *The distributed iterative consensus utilizing PQ achieves a consensus, almost surely, given by*

$$\lim_{t\to\infty} \mathbf{x}(t) = c\mathbf{1} \tag{8}$$

where $c \in \boldsymbol{\tau}$*.*

Proof. Without loss of generality, we focus on integer quantization in the range $[1, m]$. Define $\mathcal{M}$ as the discrete Markov chain with initial state $\mathbf{q}(0)$ and transition matrix defined by the combination of the deterministic transformation $\mathbf{x}(t+1) = \mathbf{W}\mathbf{q}(t)$ and the probabilistic quantizer $\mathbf{q}(t+1) \sim \Pr\{\mathbf{q}(t+1)|\mathbf{x}(t+1)\}$.

Let $\mathcal{S}_0$ be the set of quantization points that can be represented in the form $q\mathbf{1}$ for some integer q and denote by $\mathcal{S}_k$ the set of quantization points with minimum Manhattan distance k from $\mathcal{S}_0$. Moreover, let $\mathcal{C}(\mathbf{q})$ be the open hypercube centered at $\mathbf{q}$ and defined as $(q_1-1, q_1+1) \times (q_2-1, q_2+1) \times \ldots \times (q_N-1, q_N+1)$. Here q_k denotes the k–th coefficient of $\mathbf{q}$. Note that any point in $\mathcal{C}(\mathbf{q})$ has a non-zero probability of being quantized to $\mathbf{q}$. Let

$$\mathcal{A}_k = \bigcup_{\mathbf{q}\in\mathcal{S}_k} \mathcal{C}(\mathbf{q}). \tag{9}$$

The consensus operator has the important property that $|\mathbf{W}\mathbf{q} - \mu_{\mathbf{W}\mathbf{q}}| < |\mathbf{q} - \mu_{\mathbf{q}}|$ for $|\mathbf{q} - \mu_{\mathbf{q}}| > 0$, where $\mu_{\mathbf{q}}$ denotes the projection of $\mathbf{q}$ onto the $\mathbf{1}$–vector. Moreover, $c\mathbf{W}\mathbf{1} = c\mathbf{1}$. The latter property implies that $\mathbf{q} \in \mathcal{S}_0$ is an absorbing state, since $\mathcal{Q}(\mathbf{x}(t+1)) = \mathcal{Q}(\mathbf{W}\mathbf{q}(t)) = \mathcal{Q}(\mathbf{q}(t)) = \mathbf{q}(t)$. The former property implies that there are no other absorbing states, since $\mathbf{x}(t+1)$ cannot equal $\mathbf{q}(t)$ (it must be closer to the $\mathbf{1}$–vector). This implies, from the properties of the quantizer $\mathcal{Q}$, that there is a non-zero probability that $\mathbf{q}(t+1) \neq \mathbf{q}(t)$.

In order to prove that $\mathcal{M}$ is an absorbing Markov chain, it remains to show that it is possible to reach an absorbing state from any other state. We prove this by induction, demonstrating first that

$$\Pr\{\mathbf{q}(t+1) \in \mathcal{S}_0 | \mathbf{q}(t) \in \mathcal{S}_1\} > 0 \tag{10}$$

and subsequently that

$$\Pr\{\mathbf{q}(t+1) \in \bigcup_{i=0}^{i=k-1} S_i | \mathbf{q}(t) \in \mathcal{S}_k\} > 0. \tag{11}$$

Define the open set $\mathcal{V}_k$ as

$$\mathcal{V}_k = \left\{\mathbf{x} : |\mathbf{x} - \mu_{\mathbf{x}}| < k\sqrt{N-1}/\sqrt{N}\right\}. \tag{12}$$

To commence, observe that $\mathcal{V}_1 \subset \mathcal{A}_0$. The distance $|\mathbf{q} - \mu_{\mathbf{q}}| = \sqrt{N-1}/\sqrt{N}$ for $\mathbf{q} \in \mathcal{S}_1$. Hence, if $\mathbf{q}(t) \in \mathcal{S}_1$, $\mathbf{x}(t+1) = \mathbf{W}\mathbf{q}(t) \in \mathcal{V}_1 \in \mathcal{A}_0$ and $\Pr\{\mathbf{q}(t+1) \in \mathcal{S}_0\} > 0$. Similarly, the set

$$\mathcal{V}_k \subset \bigcup_{i=0}^{i=k-1} \mathcal{A}_i. \tag{13}$$

The maximum distance $|\mathbf{q} - \mu_{\mathbf{q}}|$ for any point $\mathbf{q} \in \mathcal{S}_k$ is $k\sqrt{N-1}/\sqrt{N}$. This implies that

$$\mathbf{x}(t+1) = \mathbf{W}\mathbf{q}(t) \in \mathcal{V}_k \in \bigcup_{i=0}^{i=k-1} \mathcal{A}_i. \tag{14}$$

There is thus some $i < k$ and some $\mathbf{q} \in \mathcal{S}_i$ such that $\Pr\{\mathcal{Q}(\mathbf{x}(t+1)) = \mathbf{q}\} > 0$. This argument implies that for any starting state $\mathbf{q}(0)$ such that $\mathbf{q}(0) \in \mathcal{S}_k$ for some k, there exists a sequence of transitions with non–zero probability whose application results in absorption. □

The following theorem discusses the expectation of the limiting random vector, *i.e.*, the expected value of $\mathbf{x}(t)$ as t tends to infinity.

Theorem 2. *The expectation of the limiting random vector is given by*

$$\mathbb{E}\left\{\lim_{t\to\infty} \mathbf{x}(t)\right\} = (N)^{-1}\mathbf{1}\mathbf{1}^{\mathrm{T}}\mathbf{x}(0). \tag{15}$$

Proof. Note that $||\mathbf{x}(t)|| \leq \sqrt{N}U$, for $t \geq 0$, and, $\{x_i(t) : i = 1, 2, \ldots, N\}$ are bounded for all t. Moreover, from Theorem 1, we know that the random vector sequence $\mathbf{x}(t)$ converges in the limit, *i.e.*, $\lim_{t\to\infty} \mathbf{x}(t) = c\mathbf{1}$ for some $c \in \boldsymbol{\tau}$. Thus, by dominated convergence theorem [10], we have

$$\mathbb{E}\left\{\lim_{t\to\infty} \mathbf{x}(t)\right\} = \lim_{t\to\infty} \mathbb{E}\{\mathbf{x}(t)\}. \tag{16}$$

In the following, we derive $\lim_{t\to\infty} \mathbb{E}\{\mathbf{x}(t)\}$ and utilize the above relationship to arrive at the desired result.

Note that the quantization noise, v_i, can be viewed as a Bernoulli random variable taking values at $r_i\Delta$ and $(r_i-1)\Delta$ with probabilities $1 - r_i$, and r_i, respectively. In terms of quantization noise $\mathbf{v}(t)$, $\mathbf{q}(t)$ is written as

$$\mathbf{q}(t) = \mathbf{x}(t) + \mathbf{v}(t). \tag{17}$$

The distributed iterative process reduces to the following recursion

$$\mathbf{x}(t+1) = \mathbf{W}\mathbf{x}(t) + \mathbf{W}\mathbf{v}(t). \tag{18}$$

Repeatedly utilizing the state recursion gives

$$\mathbf{x}(t) = \mathbf{W}^t\mathbf{x}(0) + \sum_{j=0}^{t-1} \mathbf{W}^{t-j}\mathbf{v}(j) \tag{19}$$

where $\mathbf{W}^t = \mathbf{W} \cdot \mathbf{W} \cdot \ldots \mathbf{W}$, t times. Taking the statistical expectation of $\mathbf{x}(t)$ as $t \to \infty$ and noting that the only random variables are $\mathbf{v}(j)$ for $j = 0, 1, \ldots, t-1$, yields

$$\lim_{t\to\infty} \mathbb{E}\{\mathbf{x}(t)\} = \lim_{t\to\infty} \mathbf{W}^t \mathbf{x}(0) + \sum_{j=0}^{t-1} \mathbf{W}^{\mathbf{t-j}} \mathbb{E}\{\mathbf{v}(j)\} \quad (20)$$

$$= \lim_{t\to\infty} \mathbf{W}^t \mathbf{x}(0) \quad (21)$$

since $\mathbb{E}\{\mathbf{v}(j)\} = \mathbf{0}$ for $j = 0, 1, \ldots, t-1$; a corollary of Lemma 1. Furthermore, noting that

$$\lim_{t\to\infty} \mathbf{W}^t = (N)^{-1} \mathbf{1}\mathbf{1}^{\mathrm{T}} \quad (22)$$

gives

$$\lim_{t\to\infty} \mathbb{E}\{\mathbf{x}(t)\} = (N)^{-1} \mathbf{1}\mathbf{1}^{\mathrm{T}} \mathbf{x}(0). \quad (23)$$

Recalling (16) gives the desired result. □

This result indicates that the statistical mean of the limiting random vector is indeed equal to the initial state average. Furthermore, this theorem, combined with the previous one, indicates that the consensus value, c, is a discrete random variable with support defined by $\boldsymbol{\tau}$ whose expectation is equal to the average of the initial states.

4. NUMERICAL EXAMPLES

This section details numerical examples evaluating the performance of the distributed average computation using probabilistic quantization. Throughout the simulations we utilize the Metropolis weight matrix defined for a graph [6]. The Metropolis weights on a graph are defined as follows:

$$W_{ij} = \begin{cases} (1 + \max\{\mathcal{K}_i, \mathcal{K}_j\})^{-1}, & i \neq j, \text{ and } \Phi_{ij} = 1 \\ 1 - \sum_{k\in\mathcal{N}_i} W_{ik}, & i = j \\ 0, & \text{otherwise} \end{cases} \quad (24)$$

where $\mathcal{K}_i = |\mathcal{N}_i|$ with $|\cdot|$ denoting cardinality of its argument. This method of choosing weights is adapted from the Metropolis algorithm in the literature of Markov chain Monte Carlo. The Metropolis weights are very simple to compute and are well suited for distributed implementation. Furthermore, the nodes do not need any global knowledge of the communication graph or even the total number of nodes.

We simulate a network with $N = 50$ nodes randomly dispersed on the unit square $[0, 1] \times [0, 1]$, connecting two nodes by an edge if the distance between them is less than the connectivity radius (a link exists between any two nodes that are at a range less than d, where d is the connectivity radius [5, 6]) where $d = \sqrt{\log(N)/N}$. The initial states $\mathbf{x}(0)$ are drawn i.i.d. from a Gaussian distribution with mean 0.85 and unit variance. The initial states are then regularized such that $\overline{\mathbf{x}}(0) = 0.85$. The quantization resolution is taken as $\Delta = 0.1$. Figure 1 depicts the trajectories taken by the distributed consensus computation using PQ at each of the 50 nodes in the network overlayed on one plot. The figure indicates that the algorithm does indeed achieve consensus as all

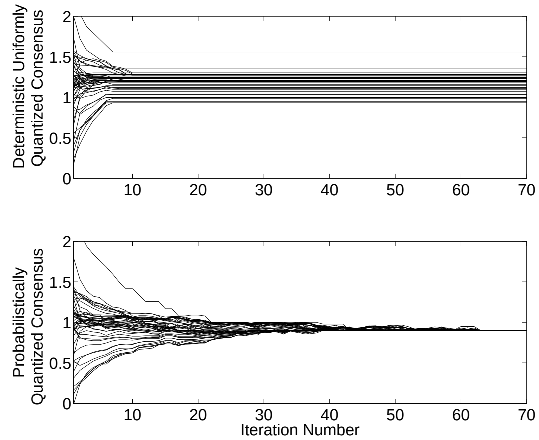

Fig. 1. Individual node trajectories using (top:) uniform quantization and (bottom:) probabilistic quantization.

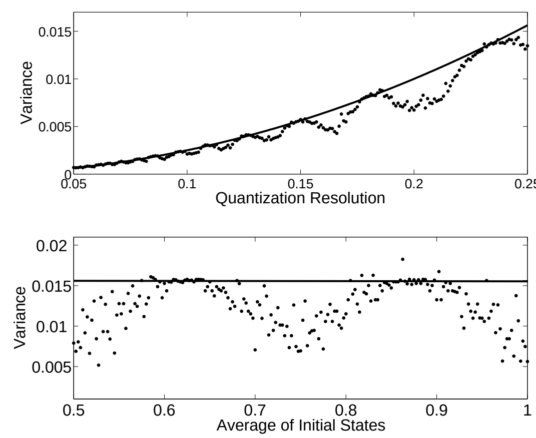

Fig. 2. The variance of the consensus with respect to (a) the quantization resolution and (b) the initial state average.

trajectories are converging to the same point. In this case, the consensus is $\lim_{t\to\infty} \mathbf{x}(t) = (0.9)\mathbf{1}$, which is in agreement with the theoretical results indicating that the consensus is at one of the quantization levels. Also note that distributed average computation utilizing deterministic uniform quantization (to closest quantization level) also converges, but, there is no consensus due to the bias induced by uniform quantization.

We next investigate the effect of the quantization resolution and the location of the initial state average on the consensus variance. Figure 2 plots the variance of the consensus for varying $\Delta \in [0.05, 0.25]$ when $\overline{\mathbf{x}}(0) = 0.85$ and for varying $\overline{\mathbf{x}}(0) \in [0.5, 1]$ when $\Delta = 0.25$. Note that each data point in the plots is an ensemble average of 100 trials. Also plotted is the curve $\Delta^2/4$. The variance, as expected, tends to increase as Δ increases and exhibits a harmonic behavior as the location of the average changes. This is due to the effect induced by the distance of the average to the quantization levels. The plots also suggest that $\Delta^2/4$ is an upper bound for consensus variance, which is a topic of current exploration.

Figure 3 shows the behavior of the average mean square

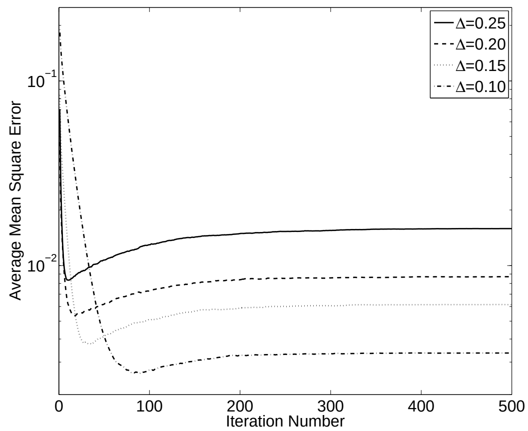

Fig. 3. The average MSE of the consensus utilizing PQ for varying quantization resolution.

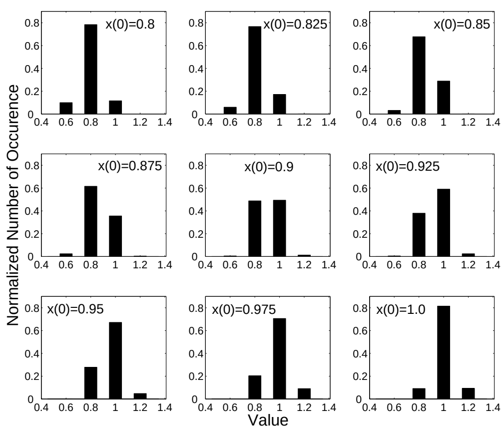

Fig. 4. The consensus value histograms for varying initial state average, $\overline{\mathbf{x}}(0) \in \{0.80, 0.825, \ldots, 1.00\}$ and $\Delta = 0.2$.

error (MSE) per iteration defined as: $\mathrm{MSE}(t) = (N)^{-1}||\mathbf{x}(t) - \overline{\mathbf{x}}(0)||^2$, for $\Delta \in \{0.10, 0.15, 0.20, 0.25\}$. In other words, $\mathrm{MSE}(t)$ is the average mean squared distance of the states at iteration t from the initial mean. Each curve is an ensemble average of 1000 experiments. The smaller quantization fidelity yields a smaller steady state MSE, with a slightly larger convergence time. The quasi–convex shape of the MSE curves are due to the fact that the algorithm, after all the state values converge into a quantization range $[\tau_i, \tau_{i+1})$ for some $i \in \{1, 2, \ldots, L-1\}$, drifts to a quantization value.

Considered next is the consensus value of the probabilistically quantized distributed average consensus. Figure 4 plots the histograms of the consensus value for varying initial state average, *i.e.*, $\overline{\mathbf{x}}(0) \in \{0.80, 0.825, \ldots, 1.00\}$ for $\Delta = 0.2$. Note that the consensus values shifts as the initial average value shifts from 0.80 to 1.00 and its support is always $\boldsymbol{\tau}$ corroborating the theoretical results. Moreover, the shift of consensus values is due to the fact that the consensus, in expectation, is equal to the average of initial states.

5. CONCLUDING REMARKS

We have proposed a framework for distributed computation of averages of the node data over networks with bandwidth/power constraints or large volumes of data. The proposed method unites the distributed average consensus and the probabilistic quantization algorithms. Theoretical and simulation results demonstrate that the proposed algorithm indeed achieves a consensus. Furthermore, the consensus is a discrete random variable whose support is the quantization values and expectation is equal to the average of the initial states. Provided numerical examples shows the effectiveness of the proposed algorithm under varying conditions.

6. REFERENCES

[1] N.A. Lynch, *Distributed Algorithms*, Morgan Kaufmann Publishers, Inc., San Francisco, CA, 1996.

[2] Y. Rabani, A. Sinclair, and R. Wanka, "Local divergence of Markov chains and the analysis of iterative load-balancing schemes," in *Proceedings of the IEEE Symp. on Found. of Comp. Sci.*, Palo Alto, CA, Nov. 1998.

[3] W. Ren and R.W. Beard, "Consensus seeking in multiagent systems under dynamically changing interaction topologies," *IEEE Trans. Autom. Control*, vol. 50, no. 5, pp. 655–661, 2005.

[4] C. C. Moallemi and B. V. Roy, "Consensus propagation," *IEEE Trans. Inf. Theory*, vol. 52, no. 11, pp. 4753–4766, Nov. 2006.

[5] D. S. Scherber and H. C. Papadopoulos, "Distributed computation of averages over ad hoc networks," *IEEE Trans. Wireless Commun.*, vol. 23, no. 4, pp. 776–787, Apr. 2005.

[6] L. Xiao, S. Boyd, and S.-J. Kim, "Distributed average consensus with least–mean–square deviation," *Journal of Parallel and Distributed Computing*, vol. 67, no. 1, pp. 33–46, Jan. 2007.

[7] M. E. Yildiz and A. Scaglione, "Differential nested lattice encoding for consensus problems," in *Proceedings of the Information Processing in Sensor Networks*, Cambridge, MA, Apr. 2007.

[8] A. Kashyap, T. Basar, and R.Srikant, "Quantized consensus," *Automatica*, vol. 43, pp. 1192–1203, July 2007.

[9] J.-J. Xiao and Z.-Q. Luo, "Decentralized estimation in an inhomogeneous sensing environment," *IEEE Trans. Inf. Theory*, vol. 51, no. 10, pp. 3564–3575, Oct. 2005.

[10] O. Kallenberg, *Foundations of Modern Probability*, Springer–Verlag, Second Ed., 2002.