Proceedings of the
47th IEEE Conference on Decision and Control
Cancun, Mexico, Dec. 9-11, 2008

Stochastic Multiscale Approaches to Consensus Problems

Jong-Han Kim, Matthew West, Sanjay Lall, Eelco Scholte, and Andrzej Banaszuk

***Abstract*— While peer-to-peer consensus algorithms have enviable robustness and locality for distributed estimation and computation problems, they have poor scaling behavior with network diameter. We show how deterministic multi-scale consensus algorithms overcome this limitation and provide optimal scaling with network size, but at the cost of requiring global knowledge of network topology. To obtain the benefits of both single- and multi-scale consensus methods we introduce a class of stochastic message-passing schemes that require no topology information and yet transmit information on several scales, achieving scalability. The algorithm is described by a sequence of random Markov chains, allowing us to prove convergence for general topologies.**

I. Introduction

A variety of problems occurring in the area of multi-agent coordination can be solved by consensus algorithms, *e.g.,* formation, alignment, decision making, synchronization, data fusion, and so on [1][2]. These algorithms are robust to network failure and changing topologies, since algorithms depend on local information only, and every node functions identically.

However, existing consensus schemes require a great deal of communication or unrealistically dense network topologies to ensure acceptable convergence in practice when applied to large scale networks. Ironically, this prevents the practical implementation of decentralized estimation techniques for large scale real world problems with limited bandwidth, even though they were originally aimed at exactly such large systems. There have been a variety of efforts to achieve fast convergence of decentralized consensus [3]– [5].

On the other hand, the *multigrid* computational method, which was originally developed for efficiently solving elliptic boundary value problems, is an example of a scalable linear iterative solver and is a well-established technique for solving large-scale problems. The method builds a hierarchy in the state domain and separates solving the various wave-number components on different layers, thus quickly decaying all wave-number components of the residual. This results in accelerated convergence compared to conventional iterative relaxation schemes (*e.g.,* Jacobi or Gauss-Seidel) [6][7], and importantly converges in a number of iterations *independent* of the problem size.

This work was supported in part by DARPA grant number FA9550-07-C-0024.

J.-H. Kim and S. Lall are with the Department of Aeronautics and Astronautics, Stanford University {jonghank,lall}@stanford.edu

M. West is with the Department of Mechanical Science and Engineering, University of Illinois at Urbana-Champaign mwest@illinois.edu

E. Scholte and A. Banaszuk are with the United Technologies Research Center (UTRC) {scholte,banasza}@utrc.utc.com

The multiscale consensus scheme was developed after observing the similarity between the slow convergence of decentralized consensus and the slow convergence encountered in conventional relaxation schemes[8]; high wave-number components diminish quickly in several iterations, but after that the node values do not change very quickly since they have local information only. A multiscale scheme can thus also accelerate the slow convergence of decentralized consensus. The basic idea of this is to construct a virtual multilevel hierarchy, across which the local information is passed to distant nodes. All the nodes undergo two operation modes (*wake* and *sleep*), where the current mode is determined by the hierarchical structure. That is, as the consensus goes up to coarse levels, the number of sleeping nodes increases. This simple concept was demonstrated in a basic consensus problem with poor network connectivity and it was shown that the proposed scheme substantially accelerates the convergence.

However, such deterministic multi-scale schemes require every node to maintain some knowledge of a multilevel hierarchy, which reduces the robustness and expandability of distributed systems. In this paper, a stochastic multiscale consensus scheme is developed to overcome such drawbacks. The wake/sleep behavior of nodes in deterministic multiscale consensus is mimicked by stochastic mode transition control, and therefore the new scheme does not require each node to maintain hierarchy information. This is advantageous because the robustness and expandability of the stochastic multiscale consensus can thus be as good as that of conventional consensus schemes.

Mode transition strategies determine the distribution of wake/sleep modes, and thus control the convergence rate of the consensus scheme. Several strategies are chosen and analyzed to explore the possibilities to recover similar wake-sleep distributions with the deterministic consensus.

For simplicity this paper only considers sensors arranged in a 1D network topology and assumes that the data can pass through sleeping nodes with no delay. These restrictions are not fundamental and will be removed in future works.

Previous studies on robust information fusion addressed convergence in the presence of stochastic communication graphs or switching topologies [9][10]. The current work investigates how the stochastic communication patterns can be used to accelerate consensus.

In the following sections, conventional and multiscale consensus schemes are reviewed first, and then the stochastic consensus scheme is introduced. Several efforts to improve and characterize the scalability are described, and similarities between multiscale schemes and stochastic schemes are analyzed and compared. The proposed scheme is applied to the conventional consensus schemes and the randomized gossip algorithm. Performance and limitation of the proposed scheme are also discussed.

II. Conventional Consensus

A. Consensus Scheme

Consensus is an iterative process to let every node in a networked group of n nodes on a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ asymptotically compute the average value of a specific information using only local communication. Let $s_i(0)$, the initial information to be averaged, where the subscript $i \in \mathcal{V}$ denotes the node number. Each node communicates with locally connected peer nodes to update its information. Note that the node does not require the information from distant (non-adjacent) nodes.

Conventional consensus is the update:

$$s_i(k+1) = s_i(k) + \sum_{(i,j)\in\mathcal{E}} w_{ij}\left(s_j(k) - s_i(k)\right)$$

Metropolis weights w_{ij} are given by:

$$w_{ij} = \begin{cases} \min\{1/(1+d_i), 1/(1+d_j)\} & \text{if } (i,j) \in \mathcal{E} \\ 1 - \sum_{(i,l)\in\mathcal{E}} w_{il} & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}$$

where w_{ij} is the ij element of $W \in \mathbb{R}^{n\times n}$, and $d_i = |\mathcal{N}_i|$ represents the degree of node i ($\mathcal{N}_i = \{j \in \mathcal{V} | (i,j) \in \mathcal{E}\}$).

The symmetric weighting matrix W is generally chosen to yield fast consensus. Although numerical techniques to compute the optimal weights (for the fastest mixing) were suggested recently [3], a simple heuristic choice, the Metropolis weight matrix, is still attractive for decentralized consensus, in that it requires only knowledge of the local topology.

B. Convergence

For simple analysis, suppose that the information $s_i(k)$ is a scalar. Assuming time-invariant network topology, the synchronous consensus iteration is simply,

$$\begin{aligned} s(k+1) &= W s(k) \\ &= W^{k+1} s(0) \end{aligned}$$

where $s(k) = [\, s_1(k)\ s_2(k)\ \ldots\ s_n(k)\,]^T$.

Since the row and column sum of the symmetric W sum to 1, this doubly stochastic matrix W has 1 as its largest eigenvalue with corresponding eigenvector $\mathbf{1} \in \mathbb{R}^n$, by the Perron-Frobenius theorem. Hence, multiplying W conserves the average of $s(k)$, and $s(k)$ converges to $\frac{1}{n}\mathbf{1}\mathbf{1}^T s(0)$ as $k \to \infty$. To check this, let the eigenvalues of W be $\lambda_1 = 1 \geq |\lambda_2| \geq \ldots \geq |\lambda_n|$, with corresponding eigenvectors $v_1 = \mathbf{1}$, v_2, v_3, $\ldots$, v_n. Multiple eigenvalues of $\lambda = 1$ (but with distinct eigenvectors) occur when some of nodes are separated from the rest, however no multiplicity is assumed here. Now, let the eigenspace decomposition of $s(k)$,

$$\begin{aligned} s(k) &= \frac{1}{n}\mathbf{1}\mathbf{1}^T s(k) + \sum_{l=2}^{n} c_l v_l \\ &= \frac{1}{n}\mathbf{1}\mathbf{1}^T s(0) + \sum_{l=2}^{n} c_l v_l \end{aligned}$$

for some constant c_l. Defining the averaging error $e(k) = s(k) - \frac{1}{n}\mathbf{1}\mathbf{1}^T s(0) = \sum_{l=2}^{n} c_l v_l$,

$$\begin{aligned} \|e(k+1)\| &= \left\| s(k+1) - \frac{1}{n}\mathbf{1}\mathbf{1}^T s(0) \right\| \\ &= \left\| W s(k) - \frac{1}{n} W \mathbf{1}\mathbf{1}^T s(0) \right\| \\ &= \left\| W \sum_{l=2}^{n} c_l v_l \right\| \\ &= \left\| \sum_{l=2}^{n} c_l \lambda_l v_l \right\| \\ &\leq |\lambda_2|\, \|e(k)\| \end{aligned}$$

As the number of iterations increases, $\|e(k)\|$ shrinks by the second largest eigenvalue modulus (SLEM), $|\lambda_2|$, at least. In case of time-invariant network topology, $|\lambda_2| < 1$ guarantees the convergence:

$$\lim_{k\to\infty} s(k) = \frac{1}{n}\mathbf{1}\mathbf{1}^T s(0)$$

This implies that the convergence is totally dependent on the eigenvalue/eigenvector structures of W. The decay rate of each component is determined by the eigenvalue corresponding to the basis vector. *i.e.,* the $c_l v_l$ component decays quickly for large wave number l. Therefore appropriate design of the eigenstructure of W can improve the convergence.

C. Numerical Example

The following sensor network is considered throughout this paper: 64 sensors in a periodic 1D network (circulant graph Laplacian and weight matrix), whose initial measurements are randomly distributed. We will test the consensus algorithms to track how the estimates converge to the average of initial measurements. Note here that W is circulant and diagonalized by the Fourier matrix $\mathcal{F}$. *i.e.,* $W = \mathcal{F}\Lambda\mathcal{F}^*$.

Since every node is connected to two other nodes ($d_i = d_j = 2$), every edge's weight is set to $w_{ij} = 1/3$ if $(i,j) \in \mathcal{E}$. In case $(i,j) \in \mathcal{E}$ and $d_i = d_j = 1$, double message paths are established, and the effective weights becomes $w_{ij} = 2/3$. The latter happens in the highest layer of some multiscale schemes and some stochastic schemes.

The consensus history by a conventional scheme is shown in Fig. 1. It is observed that the high wave-number components diminish rapidly in several consensus sweeps, resulting in a smooth profile and slow convergence thereafter. This premature stagnation occurs because the low wave-number components correspond to the slowly decaying modes (eigenval-

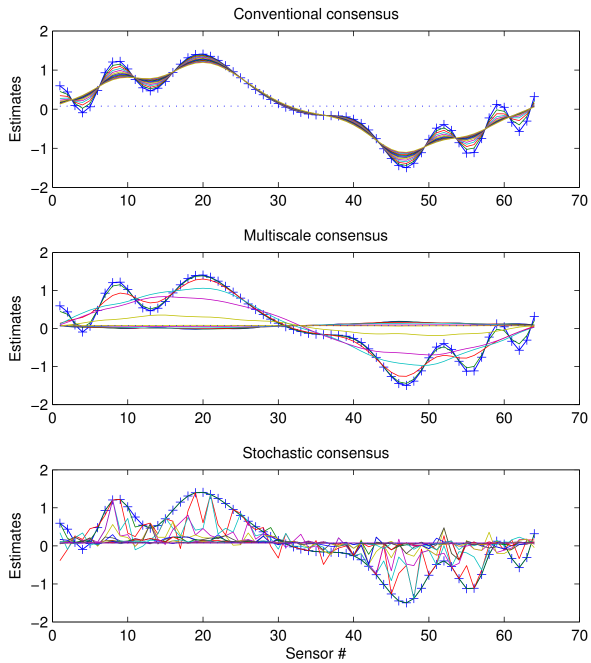

Fig. 1. Consensus history by conventional scheme (top), multiscale scheme (middle), and stochastic scheme (bottom). Initial measurements are indicated by cross marks. The lines represent the consensus results after successive iterations - the same color indicate equivalent numbers of node updates.

ues close to 1) of the weight matrix, whereas the high wave-number components are associated with rapidly decaying small eigenvalues (See the red line in Fig. 4). Unfortunately, this frequently happens in conventional consensus schemes which are based on local diffusion mechanisms.

III. Deterministic Multiscale Consensus

A. Deterministic Multiscale Consensus

In the previous example, we have observed that the nodes start to reduce the update amount, as soon as the spatial profile smooths. This is because each node, to its (local) knowledge, believes it has achieved a satisfactory approximation of the true average, even though it still has a large deviation from a global-scale view. This sort of problem is frequently encountered in iterative methods for solving systems of linear equations.

To resolve this problem we propose a multiscale consensus scheme which transfers the information between distant nodes, so that the nodes can obtain the global information on multiple scales. In principle, the basic concept of this approach is analogous to the fundamental multigrid computation idea.

A virtual hierarchy of nodes is constructed, in each level of which the consensus scheme is executed on a different scale. We call the series of consensus sweeps along the different levels the *cycle*, compared to the term from the multigrid computation field, *v-cycle*.

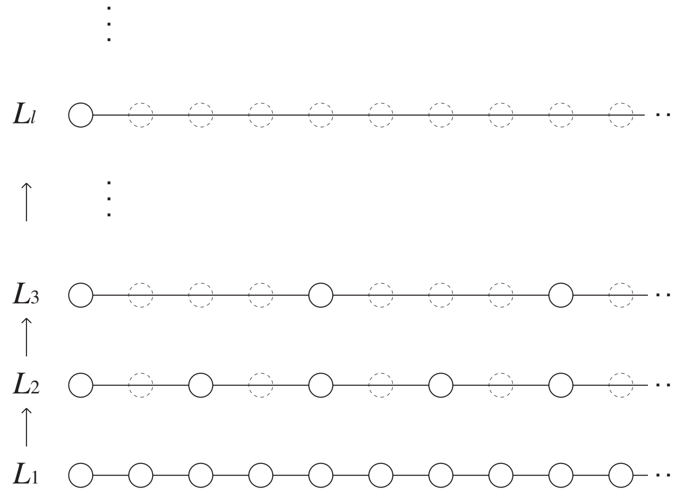

Fig. 2. Basic concept of multiscale consensus (Note that the graph reduces by half at the next upper level).

Designing the hierarchy and the cycle (the number of layers, the number of consensus sweeps in each layer, and the sequence of levels in which the consecutive consensus sweeps occur) is not trivial and could be formulated as another complicated problem. However, here we present one of the simplest choices for such a scheme.

Multiscale consensus scheme:

1) *For the finest level (L_1, $\mathcal{G} = (\mathcal{V}, \mathcal{E})$),*
   a) *Execute a consensus sweep for $\mathcal{G}$*
2) *For the next coarse level (L_2),*
   a) *Form 2 subgroups $\mathcal{V}_0^{(2)} = \{i \in \mathcal{V} | (i \mod 2) = 0\}$, $\mathcal{V}_1^{(2)} = \{i \in \mathcal{V} | (i \mod 2) = 1\}$, each of them with isomorphic topologies*
   b) *Execute consensus sweeps for each group $\mathcal{G}_0^{(2)}$, and $\mathcal{G}_1^{(2)}$*

   $\vdots$
3) *For the l-th level (L_l)*
   a) *Form 2^{l-1} subgroups $\mathcal{V}_k^{(l)} = \{i \in \mathcal{V} | (i \mod 2^{l-1}) = k\}$, for $0 \leq k \leq 2^{l-1} - 1$, each of them with isomorphic topologies*
   b) *Execute consensus sweeps for each group $\mathcal{G}_0^{(l)}, \ldots, \mathcal{G}_{2^{l-1}-1}^{(l)}$*

   $\vdots$
4) *Finishing the coarsest level ($L_{\lfloor \log_2 n \rfloor}$), go back to [1]*
   ** Number of active subgroups in each level depends on cycle design*

Note that most of nodes undergo periodical wake-sleep transitions, which is governed by the cycle structures. When a node is sleeping, it does not mix the incoming message but passes it to the other side. We assume that the communication cost is ignorable, so instantaneous message passing between distant nodes through sleeping nodes are made. *i.e.,* we presume that this is equivalent to periodical change of graph topology.

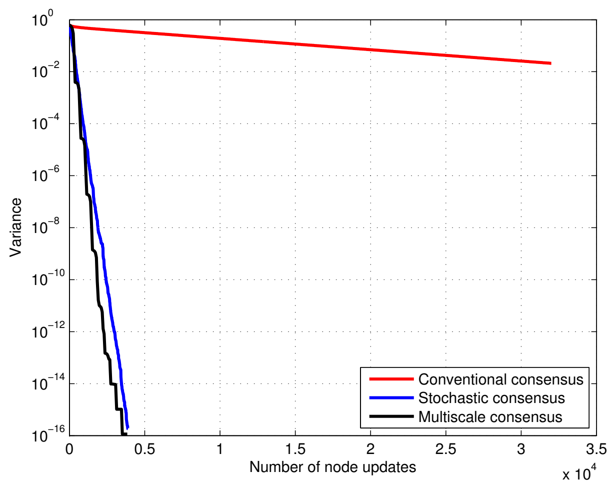

Fig. 3. Convergence rates. Stochastic scheme is with $p = 0.1$. Comparisons are based on the equal number of node updates.

The consensus history and convergence of the proposed multiscale scheme is plotted in Fig. 1 and Fig. 3. It is obvious from Fig. 4 that the proposed scheme substantially accelerates the convergence, eliminating low wave-number components efficiently. Comparisons in these plots are based on the equal number of node updates. In Fig. 4, the eigenvalues are scaled to the equivalent number of node updates. ($|\lambda|^{1/N}$, where N is the number of layers)

B. Convergence

A cycle of multi-level consensus is expressed by a series of conventional consensus iterations.

$$\begin{aligned} s_i(k+1) &= W^{(N)}W^{(N-1)}\cdots W^{(1)}s(k) \\ &= W_M s(k) \end{aligned}$$

where W_M is the composite weight matrix, $W_M = W^{(N)}W^{(N-1)}\cdots W^{(1)}$, and $N = \lfloor \log_2 n \rfloor$ is the number of levels.

$W^{(l)}$ matrices are determined by the same procedure as $W^{(1)}$ is generated, but with the l-th level topology, $\mathcal{G}^{(l)} = (\mathcal{V}^{(l)}, \mathcal{E}^{(l)})$. Suppose that $W^{(1)}, W^{(2)}, \ldots, W^{(N)}$ are all Fourier diagonalizable, then the eigenvalue of W_M is just the product of the eigenvalues of each matrix.

$$\begin{aligned} \mathcal{F}^*W^{(l)}\mathcal{F} &= \Lambda^{(l)} \quad \text{for } l = 1, 2, \ldots, N \\ \mathcal{F}^*W_M\mathcal{F} &= \Lambda^{(N)}\Lambda^{(N-1)}\cdots\Lambda^{(1)} \\ &= \operatorname{diag}\left(\left[\prod_{i=1}^{N}\lambda_1^{(i)}, \prod_{i=1}^{N}\lambda_2^{(i)}, \ldots, \prod_{i=1}^{N}\lambda_n^{(i)}\right]\right) \end{aligned}$$

We can define the scaled eigenvalue $\bar{\lambda}_k$ as an index of average convergence rate. See Fig. 4 for eigenvalues for various multiscale structures.

$$\bar{\lambda}_2 = \lambda_2^{1/N} = \left(\prod_{i=1}^{N}\lambda_2^{(i)}\right)^{1/N}$$

C. Scalability

The spectral gap for growing network size is plotted in Fig. 8. It is observed that the spectral gap of the multiscale scheme decreases moderately as the network (n) grows, compared to the conventional consensus.

However, we should notice that every node requires some knowledge of network hierarchy, which breaks the "decentralization" policy, and thus loses the robustness and expandability of decentralized systems. Also, implementing the multiscale consensus in a decentralized way is not trivial. This particular issue is addressed in the following section.

IV. Stochastic Multiscale Consensus

In the multiscale consensus, each node undergoes periodical wake-sleep mode transitions, which is governed by the multilevel hierarchy construction. It accelerates the convergence and makes the algorithm scalable. However it deteriorates the robustness and expandability of distributed network systems, and thus is disagreeable for decentralized algorithms.

In order to resolve this, we develop a stochastic multiscale consensus scheme, in which the mode transition is controlled stochastically. Because this operation is independent of the network hierarchy, the resulting algorithm is fully decentralized, with the stochastic wake-sleep mode transition still providing the accelerated convergence.

Several strategies to control the wake-sleep mode distribution are discussed in this section.

A. Naive Stochastic Consensus

A simple idea is to let the transition be controlled by the predetermined probability; we call it the *wake probability*. When a message arrives, the node decides whether to be awake (receive the message and compute average) or to be asleep (just pass the message to the other direction), based on the wake probability. This forms some distribution of wake/sleep mode transitions, and we may want to control the distribution to have a similar behaviors of the deterministic multiscale algorithm. Note that implementing this does not require any knowledge of hierarchy information.

Fig. 1 and Fig. 3 display the consensus result for a specific instance with $p = 0.1$. Note that these plots show the result for only a single sample instance, however the expected results or other sample results do not differ qualitatively.

B. Convergence

The network topology changes randomly, thus $W(0)$, $W(1)$, $\ldots$, $W(k)$ are now a series of random matrices. Applying $\mathbf{E}[Ws] = \mathbf{E}[W]s$ recursively, we get the expectation of $s(k)$ as follows.

$$\begin{aligned} \mathbf{E}[s(k+1)] &= \mathbf{E}[W]s(k) \\ &= \mathbf{E}[W]^{k+1}s(0) \end{aligned}$$

$\mathbf{E}[W]$ satisfies $\mathbf{E}[W]\mathbf{1} = \mathbf{1}$ and $\mathbf{1}^T\mathbf{E}[W] = \mathbf{1}^T$.

For the given numerical example with 64-sensor ring networks, the expectation of the weight matrix with the wake probability p is obtained as follows. Note that $\mathbf{E}[W]$

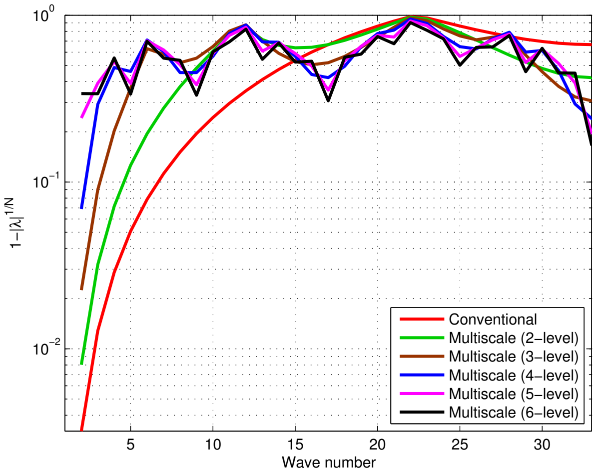

Fig. 4. Fourier coefficients of multiscale schemes W_M, scaled to equivalent number of node updates ($1 - |\lambda|^{1/N}$ is plotted).

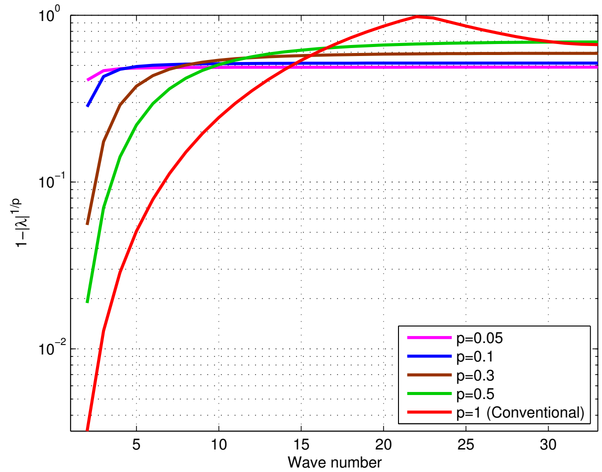

Fig. 5. Fourier coefficients of stochastic schemes $\mathbf{E}[W]$, scaled to equivalent number of node updates ($1 - |\lambda|^{1/p}$ is plotted).

is a positive symmetric circulant matrix, thus by the Perron-Frobenius theorem for positive matrices, the unique largest eigenvalue (λ_1) of $\mathbf{E}[W]$ is 1, and the following eigenvalues are strictly less than 1. *i.e.,* $1 > |\lambda_2| \geq \ldots \geq |\lambda_n|$.

$$
\begin{aligned}
\mathbf{E}[w_{ij}] &= \frac{2}{3}\Pr((i,j) \in \mathcal{E}, d_i = d_j = 1)) \\
&\quad + \frac{1}{3}\Pr((i,j) \in \mathcal{E}, d_i = d_j = 2) \\
&= \frac{2}{3}p^2(1-p)^{n-2} \\
&\quad + \frac{1}{3}p^2\{(1-p)^{j-i-1} - (1-p)^{n-2} \\
&\qquad + (1-p)^{n+i-j-1} - (1-p)^{n-2}\} \\
&= \frac{1}{3}p^2\{(1-p)^{j-i-1} + (1-p)^{n+i-j-1}\} \\
\mathbf{E}[w_{ji}] &= \mathbf{E}[w_{ij}] \quad (i < j) \\
\mathbf{E}[w_{ii}] &= 1 - \sum_{(i,j)\in\mathcal{E}} \mathbf{E}[w_{ij}]
\end{aligned}
$$

Similarly with the deterministic case, the averaging error $\|e(k)\|$ contracts as k increases.

$$
\begin{aligned}
\mathbf{E}\,\|e(k+1)\| &= \mathbf{E}\left\| s(k+1) - \frac{1}{n}\mathbf{1}\mathbf{1}^T s(0) \right\| \\
&= \mathbf{E}\left\| W(k)s(k) - \frac{1}{n}W(k)\mathbf{1}\mathbf{1}^T s(0) \right\| \\
&= \mathbf{E}\left\| W(k)\sum_{l=2}^{n} c_l(k)v_l(k) \right\| \\
&= \mathbf{E}\left\| \sum_{l=2}^{n} c_l(k)\lambda_l(k)v_l(k) \right\| \\
&\leq \mathbf{E}|\lambda_2|\,\|e(k)\|
\end{aligned}
$$

Since $\mathbf{E}|\lambda_2(k)|$ is strictly less than 1, the stochastic multiscale consensus converges to the correct average.

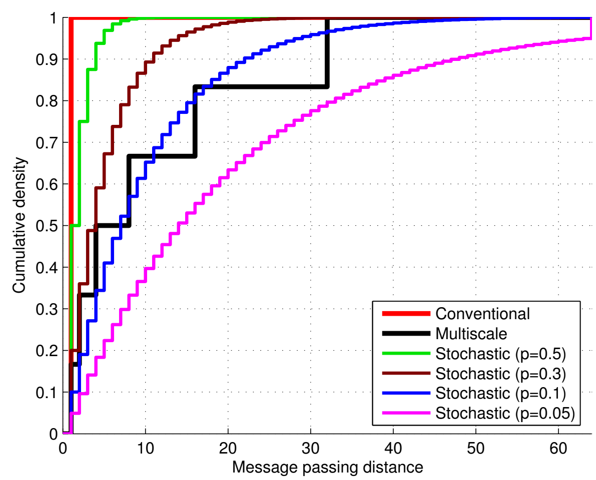

Fig. 6. Message passing distance distributions of various consensus schemes. Observe that the naive stochastic scheme with proper choice of p imitates the multiscale scheme.

Compare the eigenvalue structures of the naive stochastic scheme in Fig. 5 with that of the full multiscale scheme in Fig. 4. This simple approach significantly improves the spectral gap. Also, the changed distribution resembles that of the multiscale scheme, in that the low wave-number regions are improved much and the high wave-number regions are deteriorated. Still, this is acceptable since the minimum value is increased and it is what affects the long run convergence behaviors most.

The improved convergence can be interpreted in a different way, by the *message passing distance distribution*. For instances, conventional scheme relies only on the length-1 message passes, while multiscale scheme evenly distributes to length-1,2,4,8,... message passes. In the distributions for various schemes presented in Fig. 6, the naive stochastic

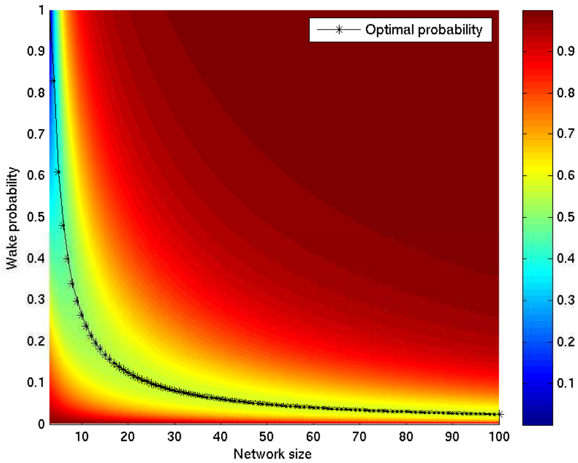

Fig. 7. Scaled SLEM of $\mathbf{E}[W]$ for various wake probabilities and network sizes. Star marks are the optimal p for given network size.

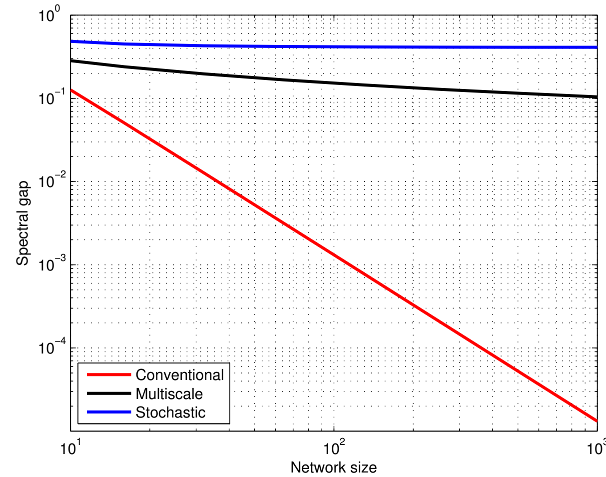

Fig. 8. Spectral gap of consensus networks with increasing sizes, scaled to equal number of node updates.

scheme with a clever choice of p imitates the distribution of multiscale scheme very well. Hence we can conclude that the naive stochastic consensus scheme is a proper approximation of the multiscale scheme.

C. Optimal Wake Probability

Given the network structures (the number of nodes in this numerical example), we can find the optimal p^* that minimizes the scaled SLEM of $\mathbf{E}[W]$, which results in the fastest convergence of the naive stochastic consensus.

Fourier decomposition of $\mathbf{E}[W]$ follows,

$$\mathcal{F}^*\mathbf{E}[W]\mathcal{F} = \mathrm{diag}([\lambda_1\ \lambda_2\ \ldots\ \lambda_n])$$

Here, the scaled SLEM (for fair comparison based on the equal number of node updates) is defined as $\bar{\lambda}_i = |\lambda_i|^{1/p}$.

$$\bar{\lambda}_i = \left|\sqrt{n} f_i^T \mathbf{E}[w_1]\right|^{1/p}$$

where, $f_i \in \mathbb{R}^n$ is the i-th column of $\mathcal{F}$, and $\mathbf{E}[w_1] \in \mathbb{R}^n$ is the first column of $\mathbf{E}[W]$. Note that the vector $\mathbf{E}[w_1]$ is a function of p.

Fig. 7 shows the scaled SLEM of $\mathbf{E}[W]$ for varying p and n. Note that p^* decreases inverse-proportionally as the network grows.

D. Scalability

The scale performances of these schemes are presented in Fig. 8. For fair comparison, the eigenvalues plotted here are scaled to the equal number of node updates; this also implicates the instantaneous message passing assumption. Compared by that metric, it is observed that the naive stochastic scheme with the optimal wake probability is superior to the others in scalability.

V. Application to Gossip Algorithms

In the previous sections, we demonstrated that stochastic message passing can significantly improve the convergence, but under the instantaneous message passing assumption. To investigate this in a more realistic situation where the communication cost matters, we applied a similar concept to the randomized gossip algorithm in a way that does not require such a tough assumption, thus the performance is fairly compared without any significant assumptions.

The gossip algorithm can be interpreted as a variant of consensus algorithms. A node, i, in the given graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is randomly activated. Then the activated node chooses one of its neighboring nodes, $j \in \mathcal{N}_i$, randomly according to some probability distribution p_{ij}, finally the two nodes communicate and average their values to update. A sequence of such actions is repeated until convergence. Optimizing the distribution p_{ij} for fast convergence can be posed as a simple convex problem and thus solved efficiently[9].

For instance, the 64-sensor circular networks, as in the previous numerical examples, has the optimal probability distribution as, $p_{ij} = \frac{1}{2}$ if $(i,j) \in \mathcal{E}$, and $p_{ij} = 0$ otherwise.

A. Gossip Algorithm with Stochastic Message Swap

Now we apply the stochastic message drop/exchange concept to this randomized gossip algorithm. We will denote the following steps as one update.

1) Activate a node randomly. Denote it by i.
2) Choose a neighboring node $j \in \mathcal{N}_i$ according to the probability distribution p_{ij}.
3) With probability q_{ij}, two nodes average their values to update (*wake mode*). Otherwise, two nodes swap their values to update (*sleep mode*).
4) Repeat until convergence.

Because this simple modification does not increase the communication or computational load, we can make a fair

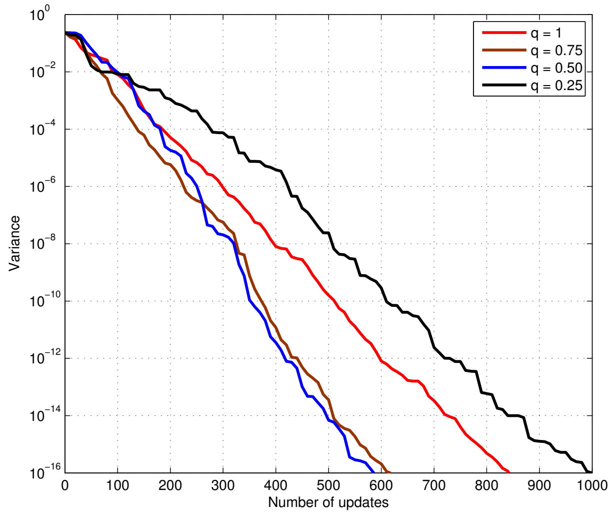

Fig. 9. Convergence histories of gossip algorithms for 10-sensor circular networks. q represents the uniform wake probability for edges. $q = 1$ corresponds to the original randomized gossip algorithm in [9].

comparison between this and the original scheme.

Fig. 9 displays the convergence histories of gossip algorithms with several wake probabilities, q. Note the improved convergence for some range of wake probability. Although this example confines itself to a single random instance, it has the following nontrivial implications: stochastic message swapping increases the chances of messages traveling farther, thus improving the convergence; furthermore, the plot claims the existence of the optimal wake probability for the fastest convergence, since the convergence rate is increased as the wake probability decreases from 1 to some critical point, and then starts to decrease.

This demonstrates a simple application of the proposed stochastic message passing idea to the gossip protocol. In order to improve more, we may try to jointly optimize p_{ij} and q_{ij}, or let q_{ij} be a function of other information which is locally available to node i and j.

VI. Conclusion and Further Research

A multiscale consensus scheme was developed to improve the unbalanced slow convergence of the conventional peer-to-peer, diffusion-based consensus. Although it is demonstrated to achieve substantial acceleration in eliminating large wave-number error components, this scheme requires every node to maintain some level of hierarchy information of the whole network, which is not desirable for robustness or scalability issues.

To overcome such drawbacks, a stochastic multiscale scheme is proposed. Stochastic mode transition control simulates the wake-sleep behaviors of the deterministic multiscale consensus, and therefore the new scheme does not require each node to maintain the hierarchy information. It is shown that the proposed stochastic scheme provides equal convergence acceleration to the multiscale scheme, while recovering the robustness and scalability of standard consensus methods. Similarities between the stochastic scheme and the multiscale scheme were also analyzed in several different aspects.

The results presented in this paper assume instantaneous message passing between distant nodes. *i.e.,* when connected nodes are in the sleep mode, the incoming message is not mixed but passed to the other side, and this happens immediately without latency. However, this assumption may not be practically acceptable for some applications where the communication process is presumably expensive. Analysis and algorithmic development for those situations remains to be done.

Asynchronous peer-to-peer message passing or shuffling/swapping can provide potential solutions, and will allow easy practical implementation. As an example of this, we combined the proposed scheme with the randomized gossip algorithm, and demonstrated that the combination with a proper choice of wake probability can improve the convergence properties. This numerical experiment demonstrates only a simple case and further investigation is required.

References

[1] A. Jadbabaie, J. Lin, and A. S. Morse, “Coordination of groups of mobile autonomous agents using nearest neighbor rules,” *IEEE Trans. on Automatic Control,* vol. 48, no. 6, pp. 988–1001, June 2003.

[2] W. Ren, R. W. Beard, and E. M. Atkins, “A survey of consensus problems in multi-agent coordination,” in *Proc. 2005 American Control Conference,* pp. 1859–1864, June 2005.

[3] S. Boyd, P. Diaconis, and L. Xiao, “Fastest mixing Markov chain on a graph,” *SIAM Review,* vol. 46, no. 4., pp. 667–689, Oct. 2004.

[4] L. Xiao and S. Boyd, “Fast linear iterations for distributed averaging,” *Systems and Control Letters,* vol. 53, pp. 65–78, 2004.

[5] R. Olfati-Saber, “Ultrafast consensus in small-world networks,” in *Proc. 2005 American Control Conference,* pp. 2371–2378, June 2005.

[6] I. Yavneh, “Why multigrid methods are so efficient,” *Computing in Science & Engineering,* vol. 8, no. 6, pp. 12–22, Nov. 2006.

[7] W. L. Briggs, V. E. Henson, and S. F. McCormick, *A Multigrid Tutorial.* SIAM, 2000.

[8] J.-H. Kim, M. West, E. Scholte, and S. Narayanan, “Multiscale consensus for decentralized estimation and its application to building systems,” in *Proc. 2008 American Control Conference,* pp. 888–893, June 2008.

[9] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, “Randomized gossip algorithms,” *IEEE Trans. on Information Theory,* vol. 52, no. 6, pp. 2508–2530, June 2006.

[10] R. Olfati-Saber and R. M. Murray “Consensus problems in networks of agents with switching topology and time-delays,” *IEEE Trans. on Automatic Control,* vol. 49, no. 9, pp. 1520–1533, Sep. 2004.