Convergence Rates of Distributed Average Consensus With Stochastic Link Failures

Stacy Patterson, Bassam Bamieh, *Fellow, IEEE*, and Amr El Abbadi, *Senior Member, IEEE*

***Abstract*—We consider a distributed average consensus algorithm over a network in which communication links fail with independent probability. In such stochastic networks, convergence is defined in terms of the variance of deviation from average. We first show how the problem can be recast as a linear system with multiplicative random inputs which model link failures. We then use our formulation to derive recursion equations for the second order statistics of the deviation from average in networks with and without additive noise. We give expressions for the convergence behavior in the asymptotic limits of small failure probability and large networks. We also present simulation-free methods for computing the second order statistics in each network model and use these methods to study the behavior of various network examples as a function of link failure probability.**

***Index Terms*—Distributed systems, gossip protocols, multiplicative noise, packet loss, randomized consensus.**

WE study the distributed average consensus problem over a network with stochastic link failures. Each node has some initial value and the goal is for all nodes to reach consensus at the average of all values using only communication between neighbors in the network graph. Distributed average consensus is an important problem that has been studied in contexts such as vehicle formations [1]–[3], aggregation in sensor networks and peer-to-peer networks [4], load balancing in parallel processors [5], [6], and gossip algorithms [7], [8].

Distributed consensus algorithms have been widely investigated in networks with static topologies, where it has been shown that the convergence rate depends on the second smallest eigenvalue of the Laplacian of the communication graph [9], [10]. However, the assumption that a network topology is static, i.e. that communication links are fixed and reliable throughout the execution of the algorithm, is not always realistic. In mobile networks, the network topology changes as the agents change position, and therefore the set of nodes with which each node can communicate may be time-varying. In sensor networks and mobile ad-hoc networks, messages may be lost due to interference, and in wired networks, networks may suffer from packet loss and buffer overflow. In scenarios such as these, it is desirable to quantify the effects that topology changes and communication failures have upon the performance of the averaging algorithm.

In this work, we consider a network with an underlying topology that is an arbitrary, connected, undirected graph where links fails with independent but not necessarily identical probability. In such stochastic networks, we define convergence in terms of the variance of deviation from average. We show that the averaging problem can be formulated as a linear system with multiplicative noise and use our formulation to derive a recursion equation for the second order statistics of the deviation from average. We also give expressions for the mean square convergence rate in the asymptotic limits of small failure probability and large networks.

Additionally, we consider the scenario where node values are perturbed by additive noise. This formulation can be used to model load balancing algorithms in peer-to-peer networks or parallel processing systems, where the additive perturbations represent file insertions and deletions or job creations and completions, with the goal of equilibrating the load amongst the participants. A measure of the performance of the averaging algorithm in this scenario is not how quickly node values converge to the average, but rather how close the node values remain to each other, and therefore to the average of all values as this average changes over time. This problem has been previously studied in networks without communication failures [10], [11], however we are unaware of any existing work that addresses this problem in networks with communication failures. We show how our formulation for static-valued networks can be extended to incorporate the additive perturbations and give an expression for the steady-state deviation from average. Finally, for both problem formulations, we present simulation-free methods for computing the second order statistics of the variance of the deviation from average, and we use these methods to study the behavior of various network examples as a function of link failure probability.

Although there has been work that gives conditions for convergence with communication failures, to our knowledge, this is the first work that quantifies the effects of stochastic communication failures on the performance of the distributed average consensus algorithm. We briefly review some of the related work below.

Related Work: The distributed consensus problem has been studied in switching networks, where convergence is defined in a deterministic sense. The works by Jadbabaie *et al.* [1] and Xiao and Boyd [12] show that in undirected, switching communication networks, convergence is guaranteed if there is an infinitely occurring, contiguous sequence of bounded time intervals

Manuscript received December 19, 2008; revised June 18, 2009. First published February 02, 2010; current version published April 02, 2010. This work was supported in part by NSF Grant IIS 02-23022 and NSF Grant ECCS-0802008. Recommended by Associate Editor M. Egerstedt.

S. Patterson and B. Bamieh are with the Department of Mechanical Engineering, University of California, Santa Barbara, CA 93106 USA (e-mail: sep@engineering.ucsb.edu; bamieh@engineering.ucsb.edu).

A. El Abbadi is with the Department of Computer Science, University of California, Santa Barbara, CA 93106 USA (e-mail: amr@cs.ucsb.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TAC.2010.2041998

in which the network is jointly connected. The same condition also guarantees convergence in directed networks, as shown by Olfati-Saber and Murray [2] and Moreau [3]. Cao *et al.* [13] identify a similar convergence condition for consensus in directed networks based on an infinitely occurring sequence of jointly rooted graphs. Recent works have also studied the convergence rates of averaging algorithms in switching networks. In [14], Olshevsky and Tsitsiklis give upper and lower bounds on the convergence rate in a directed network in terms of the length of the bounded time interval of joint connectivity, and in [13], Cao *et al.* give bounds on the convergence rate in terms of length of the interval of connectivity of the rooted graph.

Convergence conditions for the distributed averaging algorithm have also been investigated in stochastic networks. In [15], Hatano and Mesbahi study the Erdős-Rényi random graph model where each edge fails with identical probability. The authors use analysis of the expected Laplacian to prove that nodes converge to consensus with probability 1. The work by Wu [16] considers a more general directed random graph model where edge failure probabilities are not necessarily identical and proves convergence in probability. In [17], Porfiri and Stilwell study a similar model, a random directed graph where each edge fails with independent non-uniform probability, but additionally where edges are weighted. The authors also use analysis based on the expected Laplacian to show that, in the case where edge weights are non-negative, if the expected graph is strongly connected, the system converges asymptotically to consensus almost surely. For arbitrary weights, the authors show that asymptotic almost-sure convergence is guaranteed if the network topology changes "sufficiently fast enough". Tahbaz-Salehi and Jadbabaie [18] consider directed networks where the weight matrices are stochastic i.i.d and give a necessary and sufficient condition for almost sure convergence based on the second largest eigenvalue of expected weight matrix. In [19], Kar and Moura give sufficient conditions for mean square convergence in undirected networks with non-uniform link failure probabilities based on the second largest eigenvalue of the expected weight matrix. Additionally, our recent work [20] also gives sufficient conditions for mean square convergence in undirected networks where links fail with uniform probability. The analysis depends on reformulating the problem as a structured stochastic uncertainty problem and deriving conditions for convergence based on the nominal component. We also note that in [21], Kar and Moura study averaging algorithms over a network with stochastic communication failures where communication links are also corrupted by additive noise. In order to achieve consensus in such a model, the weight of each edge is decreased as the algorithm executes. This problem is similar to the averaging algorithm with additive noise that is described in this paper. However, in this work, we consider additive perturbations at the nodes as opposed to the communication channels, and we consider algorithms where the edge weights remain constant.

The remainder of this paper is organized as follows. In Section I, we formally define our system model and the distributed consensus algorithm. Section II gives our convergence results for systems with no additive noise, and in Section III, we give an extension of the model and convergence results for networks with additive noise. In Section IV, we describe our computational methods, and in Section V, we present computational results for different network scenarios. Finally, we conclude in Section VI.

I. Problem Formulation

We model the network as a connected, undirected graph $G=(V,E)$ where V is the set of nodes, with $|V|=N$, and E is the set of communication links between them, with $|E|=M$. We assume that each link $(r,s)\in E$ has an independent, but not necessarily identical probability $p_{(r,s)}$ of failing in each round. If a link fails, no communication takes place across the link in either direction in that round. A link that does not fail in round k is *active*. The neighbor set of node r, denoted by $\mathcal{N}_r(k)$ for round k, is the set of nodes with which node r has active communication links in round k.

We consider the following simple distributed consensus algorithm. Every node r has an initial value $x_r(0)$. The objective of the algorithm is to converge to an equilibrium where $x_r(k)=(1/N)\sum_{s=1}^{N}x_s(k)$ for all $r\in V$. In each round, each node sends a fraction β of its current value to each neighbor with which it has an active communication link. Each node's value is updated according to the following rule:

$$x_r(k+1)=\beta\sum_{s\in\mathcal{N}_r(k)}x_s(k)+\left(1-\beta\left|\mathcal{N}_r(k)\right|\right)x_r(k)$$

where β is the parameter that defines an instance of the algorithm. This algorithm can be implemented without any *a priori* knowledge of link failures.

In a network with no communication failures, this algorithm can be expressed as an $N\times N$ matrix, $A:=I-\beta\mathcal{L}$, where $\mathcal{L}$ is the Laplacian matrix[1] of the graph G. The evolution of the system is described by the following recursion equation:

$$x(k+1)=Ax(k). \tag{1}$$

It is a well known result that the system converges to consensus at the average of all node values if and only if the magnitude of the second largest eigenvalue of A, $\lambda_2(A)$, is strictly less than 1, and that if the graph is connected, it is always possible to chose a β that guarantees convergence [1], [9], [10], [22], [23]. In this work, we place no restriction on the choice of β other than that the resulting A matrix is such that $|\lambda_2(A)|<1$. The diagonal entries of A may be negative.

We now demonstrate how (1) can be extended to include stochastic communication failures. We note that a similar model for communication failures in directed graphs is given in [24]. Let $b_{(r,s)}$ be the N-vector with the r'th entry equal to 1, the s'th entry equal to -1 and all other entries equal to 0. $B_{(r,s)}$ is defined as

$$B_{(r,s)}:=\beta b_{(r,s)}b_{(r,s)}^{*}. \tag{2}$$

[1]Let $\mathcal{E}$ be the adjacency matrix of G and $\mathcal{D}$ be the diagonal matrix with the diagonal entry in row j equal to the degree of node j. Then the Laplacian of a graph G is defined as $\mathcal{L}:=\mathcal{D}-\mathcal{E}$.

The system can then be described by the following recursion equation

$$x(k+1) = \left(A + \sum_{(r,s)\in E} \delta_{(r,s)} k B_{(r,s)} \right) x(k) \tag{3}$$

where $\delta_{(r,s)}$ is a Bernoulli random variable with

$$\delta_{(r,s)}(k) \quad := \quad \begin{cases} 1 & \text{with probability } p_{(r,s)} \\ 0 & \text{with probability } 1 - p_{(r,s)}. \end{cases}$$

When $\delta_{(r,s)} = 1$, the edge (r,s) has failed. One can interpret (3) as first performing the algorithm on the full underlying network graph G and then simulating the failed edges by undoing the effects of communication over those edges. In essence, each $B_{(r,s)}$ matrix returns the values sent across edge (r,s), yielding the state in which edge (r,s) was not active.

We rewrite (3) in a form that is more convenient for our analysis using zero-mean random variables. Let $\mu_{(r,s)}(k) := \delta_{(r,s)}(k) - p_{(r,s)}$ and observe that they are zero mean. The dynamics can now be rewritten as

$$x(k+1) = \overline{A}x(k) + \sum_{(r,s)\in E} \mu_{(r,s)}(k) B_{(r,s)} x(k) \tag{4}$$

where $\bar{A} := A + \sum_{(r,s)\in E} p_{(r,s)} B_{(r,s)}$.

We measure how far the current state of the system is from the average of all states using the *deviation from average vector* $\tilde{x}$ whose components are

$$\tilde{x}_r(k) := x_r(k) \frac{1}{N} \left(x_1(k) + \ldots + x_N(k) \right).$$

The entire vector $\tilde{x}$ can be written as the projection

$$\tilde{x}(k) = \mathcal{P}x(k)$$

with $\mathcal{P} := (I - (1/N)\mathbf{1}\mathbf{1}^*)$, where $\mathbf{1}$ is the N-vector with all entries equal to 1.

We are primarily interested in characterizing the convergence rate of $\tilde{x}$ to zero. Since the dynamics of x and $\tilde{x}$ are stochastic, we use the decay rate of the worst-case *variance of deviation from average* of each node r, $\mathbf{E}[\tilde{x}_r(k)^2]$, as an indicator of the rate of convergence.

Problem Statement 1: Consider a distributed consensus algorithm over a connected, undirected graph where each link fails with independent probability as modeled by the system with multiplicative noise (3). For a given set of link failure probabilities, determine the worst-case rate (over all initial conditions, over all nodes) at which the deviation from average $\mathbf{E}[\tilde{x}_r(k)^2]$, $r \in N$, converges to 0 as $k \to \infty$.

The key to addressing this problem is to study the equations governing the second order statistics of the states of (4). To this end, we define the autocorrelation matrices of x and $\tilde{x}$ by

$$\begin{aligned} M(k) &:= \mathbf{E}\left[x(k)x^*(k)\right] \\ \tilde{M}(k) &:= \mathbf{E}\left[\tilde{x}(k)\tilde{x}^*(k)\right] \end{aligned}$$

and note that they are related by the projection $\mathcal{P}$

$$\begin{aligned} \tilde{M}(k) &= \mathbf{E}\left[\tilde{x}(k)\tilde{x}^*(k)\right] = \mathbf{E}\left[\mathcal{P}x(k)x^*(k)\mathcal{P}\right] \\ &= \mathcal{P}\mathbf{E}\left[x(k)x^*(k)\right]\mathcal{P} \\ &= \mathcal{P}M(k)\mathcal{P}. \end{aligned}$$

The variance of the deviation from average of each node r, $\tilde{x}_r(k)^2$, is given by the diagonal entry of the r^{th} row of $\tilde{M}$, and the total deviation from average is given by the trace of $\tilde{M}$, $\|\tilde{x}(k)\|_2^2 = \mathbf{tr}(\tilde{M}(k))$.

It is well known that the autocorrelation matrix of a system in the form of (4) with zero-mean multiplicative noise [25] obeys the following recursion equation:

$$M(k+1) = \overline{A}M(k)\overline{A} + \sum_{(r,s)\in E} \sigma_{(r,s)}^2 B_{(r,s)} M(k) B_{(r,s)} \tag{5}$$

where $\sigma_{(r,s)}^2 := \mathbf{var}[\mu_{(r,s)}(k)]$. This is a discrete-time Lyapunov-like matrix difference equation. However, the additional terms multiplying $\sigma_{(r,s)}^2$ make this a nonstandard Lyapunov recursion. The matrix $\tilde{M}(k)$ satisfies a similar recursion relation which we derive in the next section and then study its convergence properties.

II. Characterizing Convergence

In this section, we first derive a recursion equation for $\tilde{M}(k)$, the autocorrelation of $\mathbf{E}[\tilde{x}(k)\tilde{x}^*(k)]$, which has the variance of deviation from average of each node as its diagonal entries. We then characterize the decay rate of these variances in terms of the eigenvalues of a Lyapunov-like matrix-valued operator. An exact computational procedure for these eigenvalues is given in Section IV, while in this section, we give expressions for the asymptotic cases of small, uniform link failure probability p and large network size N.

Lemma 2.1: The matrices $\tilde{M}(k)$ satisfy the recursion

$$\begin{aligned} \tilde{M}(k+1) = (\mathcal{P}\overline{A}\mathcal{P})\tilde{M}(k)(\mathcal{P}\overline{A}\mathcal{P}) \\ + \sum_{(r,s)\in E} \sigma_{(r,s)}^2 B_{(r,s)} \tilde{M}(k) B_{(r,s)}. \end{aligned} \tag{6}$$

Proof: First, we note that the following equalities hold for the action of $\mathcal{P}$ on any of the matrices $B_{(r,s)}$

$$B_{(r,s)}\mathcal{P} = \beta\, b_{(r,s)} b_{(r,s)}^* \left(I - \frac{1}{N}\mathbf{1}\mathbf{1}^* \right) = B_{(r,s)}$$

where the second equality follows from $\mathbf{1}^* b_{(r,s)} = 0$ for any edge (r,s). Similarly, $\mathcal{P}B_{(r,s)} = B_{(r,s)}$. We also note that t $\mathcal{L}$, and consequently A and $\bar{A}$, commute with the projection $\mathcal{P}$. This follows from the fact that $\mathbf{1}$ is both a left and a right eigenvector of $\mathcal{L}$. Using these facts and noting that $\mathcal{P} = \mathcal{P}^2$, (6) follows from multiplying both sides of (5) by $\mathcal{P}$ as follows:

$$\begin{aligned}\tilde{M}(k+1) =& \mathcal{P}M(k+1)\mathcal{P} \\ =& \mathcal{P}\overline{A}M(k)\overline{A}\mathcal{P} \\ &+ \sum_{(r,s)\in E} \sigma_{(r,s)}^2 \mathcal{P}B_{(r,s)}M(k)B_{(r,s)}\mathcal{P} \\ =& \mathcal{P}^2\overline{A}M(k)\overline{A}\mathcal{P}^2 \\ &+ \sum_{(r,s)\in E} \sigma_{(r,s)}^2 B_{(r,s)}\mathcal{P}M(k)\mathcal{P}B_{(r,s)} \\ =& \mathcal{P}\overline{A}\mathcal{P}M(k)\mathcal{P}\overline{A}\mathcal{P} \\ &+ \sum_{(r,s)\in E} \sigma_{(r,s)}^2 B_{(r,s)}\tilde{M}(k)B_{(r,s)} \\ =& \mathcal{P}\overline{A}\mathcal{P}^2M(k)\mathcal{P}^2\overline{A}\mathcal{P} \\ &+ \sum_{(r,s)\in E} \sigma_{(r,s)}^2 B_{(r,s)}\tilde{M}(k)B_{(r,s)} \\ =& \mathcal{P}\overline{A}\mathcal{P}\tilde{M}(k)\mathcal{P}\overline{A}\mathcal{P} \\ &+ \sum_{(r,s)\in E} \sigma_{(r,s)}^2 B_{(r,s)}\tilde{M}(k)B_{(r,s)}.\end{aligned}$$

■

If all edges have an equal probability of failure p in each round, we can derive a simpler form of the recursion equation for $\tilde{M}(k)$.

Corollary 2.2: If each edge fails with uniform probability p, the matrices $\tilde{M}(k)$ satisfy the recursion

$$\begin{aligned}\tilde{M}(k+1) = (\tilde{A}+p\beta\mathcal{L})\tilde{M}(k)(\tilde{A}+p\beta\mathcal{L}) \\ +(p-p^2)\sum_{(r,s)\in E} B_{(r,s)}\tilde{M}(k)B_{(r,s)}\end{aligned} \quad (7)$$

where $\tilde{A} := \mathcal{P}A\mathcal{P}$.

Proof: Note that from the definitions of the matrices $B_{(r,s)}$, their sum is proportional to the graph's Laplacian, i.e. $\sum_{(r,s)\in E} B_{(r,s)} = \beta\mathcal{L}$. $\bar{A}$ is then simply

$$\bar{A} = A + p\sum_{(r,s)\in E} B_{(r,s)} = A + p\beta\mathcal{L}.$$

Additionally, note that $\sigma_{(r,s)}^2 = p - p^2$ for all $(r,s) \in E$.

Therefore, for uniform failure probability p, (6) simplifies as follows:

$$\begin{aligned}\tilde{M}(k+1) =& \mathcal{P}\overline{A}\mathcal{P}\tilde{M}(k)\mathcal{P}\overline{A}\mathcal{P} \\ &+ \sum_{(r,s)\in E} \sigma_{(r,s)}^2 B_{(r,s)}\tilde{M}(k)B_{(r,s)} \\ =& \mathcal{P}(A+p\beta\mathcal{L})\mathcal{P}\tilde{M}(k)\mathcal{P}(A+p\beta\mathcal{L})\mathcal{P} \\ &+ (p-p^2)\sum_{(r,s)\in E} B_{(r,s)}\tilde{M}(k)B_{(r,s)} \\ =& (\tilde{A}+p\beta\mathcal{L})\tilde{M}(k)(\tilde{A}+p\beta\mathcal{L}) \\ &+ (p-p^2)\sum_{(r,s)\in E} B_{(r,s)}\tilde{M}(k)B_{(r,s)}.\end{aligned}$$

■

A. The Decay Rate

To study the decay or growth properties of the matrix sequence $\tilde{M}(k)$, we define the Lyapunov-like operator

$$\mathcal{A}(X) := (\mathcal{P}\overline{A}\mathcal{P})X(\mathcal{P}\overline{A}\mathcal{P}) + \sum_{(r,s)\in E} \sigma_{(r,s)}^2 B_{(r,s)}XB_{(r,s)}. \quad (8)$$

The linear matrix recursion (7) can now be written as

$$\tilde{M}(k+1) = \mathcal{A}\left(\tilde{M}(k)\right). \quad (9)$$

Since this is a linear matrix equation, the condition for asymptotic decay of each entry of $\tilde{M}(k)$ is $\rho(\mathcal{A}) < 1$, where $\rho(\mathcal{A})$ is the spectral radius of $\mathcal{A}$, which we call the *decay factor* of the algorithm instance. Since each entry of $\tilde{M}(k)$ has the asymptotic bound of a constant times $\rho(\mathcal{A})^k$, then so does its trace and consequently $\mathbf{E}[||\tilde{x}(k)||_2^2]$. And, in fact, it can be shown that this upper bound on the decay rate is tight.

We summarize these results in the following theorem.

Theorem 2.3: Consider a distributed consensus algorithm where links fail with independent probability $p_{(r,s)}$ as modeled by the system with multiplicative noise

$$x(k+1) = \left(A + \sum_{(r,s)\in E} \delta_{(r,s)}(k)B_{(r,s)}\right)x(k)$$

where $\delta_{(r,s)}$ are Bernoulli random variables with

$$\delta(k) \quad := \quad \begin{cases} 1 & \text{with probability } p_{(r,s)} \\ 0 & \text{with probability } 1-p_{(r,s)}. \end{cases}$$

1) The total deviation from average $\mathbf{E}[||\tilde{x}(k)||_2^2]$ converges to 0 as $k \to \infty$ if and only if
$$\rho(\mathcal{A}) < 1$$
where $\mathcal{A}$ is the matrix-valued operator defined in (8).
2) The worst-case asymptotic growth (over all initial conditions) of any $\mathbf{E}[\tilde{x}_r(k)^2]$, $r = 1\ldots N$ is given by
$$\mathbf{E}\left[\tilde{x}_r(k)^2\right] \leq \eta\rho(\mathcal{A})^k$$
where η is a constant. This upper bound is tight.

Proof: As $\mathcal{A}$ is a matrix-valued linear recursion, it is well known that the decay rate of each entry of $\tilde{M}(k)$ is proportional to the spectral radius of $\mathcal{A}$, and this is true for all initial conditions $\tilde{M}(0)$. What remains to be shown is that this worst-case decay rate holds when $\tilde{M}(0)$ is restricted to be a covariance matrix, or equivalently, when $\tilde{M}(0)$ is positive semidefinite. The proof of this is given in the Appendix. ■

Note that in the case that links do not fail, when $p_{(r,s)} = 0$ for all $(r,s) \in E$, we have

$$\mathcal{A}: \quad X \mapsto (\mathcal{P}A\mathcal{P})X(\mathcal{P}A\mathcal{P})$$

and $\rho(\mathcal{A})$ is precisely $(\rho(\mathcal{P}A\mathcal{P}))^2$, the square of the eigenvalue of A with the second largest modulus, as is well known. However, when failures occur with non-zero probability, the additional terms in the operator $\mathcal{A}$ play a role. The operator $\mathcal{A}$ is no longer a pure Lyapunov operator of the form $X \mapsto \tilde{A}X\tilde{A}$ but rather a sum of such terms. Thus, one does not expect a simple relationship between the eigenvalues of $\mathcal{A}$ and those of the constitutive matrices as in the pure Lyapunov operator case.

B. Perturbation Analysis

One important asymptotic case is that of small, uniform link failure probability p. We can analyze this case by doing a first order eigenvalue perturbation analysis of the operator $\mathcal{A}$ in (8) as a function of the parameter p. We first recall the basic setup from analytic perturbation theory for eigenvalues of symmetric operators [26].

Consider a symmetric, matrix-valued function $\mathcal{A}(p, X)$ of a real parameter p and matrix X of the form

$$\mathcal{A}(p, X) = \mathcal{A}_o(X) + p\mathcal{A}_1(X) + p^2\mathcal{A}_2(X).$$

Let $\gamma(p)$ and $W(p)$ be an eigenvalue-eigenmatrix pair of $\mathcal{A}(p, .)$ as p varies, i.e.

$$\mathcal{A}(p, W(p)) = \gamma(p)W(p).$$

It is a standard result of spectral perturbation theory that for isolated eigenvalues of $\mathcal{A}(0, .)$, the functions γ and W are well defined and analytic in some neighborhood $p \in (-\epsilon, \epsilon)$. The power series expansion of γ is

$$\gamma(p) = \lambda + c_1 p + c_2 p^2 + \cdots$$

where λ is an eigenvalue of $\mathcal{A}_o$. The calculation of the coefficient c_1 involves the corresponding eigenmatrix V of λ and is given by

$$c_1 = \frac{\langle V, \mathcal{A}_1(V)\rangle}{\langle V, V\rangle}. \tag{10}$$

Note that we are dealing with matrix-valued operators on matrices, and the inner product on matrices is given by $\langle X, Y\rangle := \mathbf{tr}(X^*Y)$.

In order to apply this procedure to the operator $\mathcal{A}$ in (8), we first note that, when all links have uniform failure probability p, $\sigma^2_{(r,s)} = p - p^2$ for all $(r, s) \in E$. $\mathcal{A}$ can then be written as

$$\mathcal{A} = \mathcal{A}_o + p\mathcal{A}_1 + p^2\mathcal{A}_2$$

where

$$\begin{aligned}
\mathcal{A}_o(X) &= \tilde{A}X\tilde{A}\\
\mathcal{A}_1(X) &= \beta\mathcal{L}X\tilde{A} + \beta\tilde{A}X\mathcal{L} + \sum_{(r,s)\in E} B_{(r,s)}XB_{(r,s)}\\
\mathcal{A}_2(X) &= \beta^2\mathcal{L}X\mathcal{L} - \sum_{(r,s)\in E} B_{(r,s)}XB_{(r,s)}.
\end{aligned}$$

To investigate the first order behavior of the largest eigenvalue, we observe that the eigenmatrix corresponding to the largest eigenvalue of $\mathcal{A}_o$ is $V = ww^*$ where w is an eigenvector, with $\|w\|_2 = 1$, corresponding to the second smallest eigenvalue of the Laplacian $\mathcal{L}$, also called the Fiedler vector. w is also an eigenvector corresponding to the largest eigenvalue of $\tilde{A}$ (equivalently, the second largest eigenvalue of A).

Applying formula (10) to this expression for V yields the first order term in the expansion of the largest eigenvalue of $\mathcal{A}$ to be

$$c_1 = \frac{\langle ww^*, \mathcal{A}_1(ww^*)\rangle}{\langle ww^*, ww^*\rangle}.$$

The denominator can be simplified as follows:

$$\langle ww^*, ww^*\rangle = \mathbf{tr}(ww^*ww^*) = \mathbf{tr}(w^*ww^*w) = \|w\|_2^2 = 1$$

and therefore c_1 is equivalent to

$$\begin{aligned}
c_1 &= \langle ww^*, \mathcal{A}_1(ww^*)\rangle\\
&= \mathbf{tr}\left(ww^*\left(\mathcal{A}_1(ww^*)\right)^*\right)\\
&= \mathbf{tr}(\beta ww^*\mathcal{L}ww^*\tilde{A} + \beta ww^*\tilde{A}ww^*\mathcal{L})\\
&\quad + \mathbf{tr}\left(ww^* \sum_{(r,s)\in E} B_{(r,s)}ww^*B_{(r,s)}\right)\\
&= \mathbf{tr}(\beta ww^*\mathcal{L}ww^*\tilde{A} + \beta ww^*\tilde{A}ww^*\mathcal{L})\\
&\quad + \mathbf{tr}\left(\sum_{(r,s)\in E)} w^*B_{(r,s)}ww^*B_{(r,s)}w\right)\\
&= \mathbf{tr}(\beta ww^*\mathcal{L}ww^*\tilde{A} + \beta ww^*\tilde{A}ww^*\mathcal{L})\\
&\quad + \sum_{(r,s)\in E)} \left(w^*B_{(r,s)}w\right)^2.
\end{aligned} \tag{11}$$

Since w is an eigenvector of $\tilde{A}$, the following equality holds

$$w^*\tilde{A} = \overline{\lambda}(\tilde{A}t)w^* \tag{12}$$

where $\overline{\lambda}(\tilde{A})$ denotes the largest eigenvalue of $\tilde{A}$. w is also an eigenvector of $\mathcal{L}$. Therefore the following equality also holds

$$w^*\mathcal{L} = \underline{\lambda}(\mathcal{L})w^* \tag{13}$$

where $\underline{\lambda}(\mathcal{L})$ denotes the second smallest eigenvalue of $\mathcal{L}$.

Noting that $A = I - \beta\mathcal{L}$, it follows that [27], [28] for $\beta \leq 1/2\Delta$, where Δ is the maximum node degree of the graph, we have the following relationship between $\overline{\lambda}(\tilde{A})$ and $\underline{\lambda}(\mathcal{L})$

$$\underline{\lambda}(\mathcal{L}) = \frac{1}{\beta}\left(1 - \overline{\lambda}(\tilde{A})\right).$$

This equality allows us to rewrite (13) as

$$w^*\mathcal{L} = \frac{1}{\beta}\left(1 - \overline{\lambda}(\tilde{A})\right)w^*. \tag{14}$$

Using (12) and (14), (11) can be further simplified as follows:

$$\begin{aligned}
c_1 &= 2\overline{\lambda}(\tilde{A})\left(1 - \overline{\lambda}(\tilde{A})\right)\mathbf{tr}\left((ww^*ww^*)\right.\\
&\quad + \sum_{(r,s)\in E)} (w^*B_{(r,s)}w)^2\\
&= 2\overline{\lambda}(\tilde{A}) - 2\left(\overline{\lambda}(\tilde{A})\right)^2\|w\|_2^2 + \sum_{(r,s)\in E)} \left(w^*B_{(r,s)}w\right)^2\\
&= 2\overline{\lambda}(\tilde{A}) - 2\left(\overline{\lambda}(\tilde{A})\right)^2 + \sum_{(r,s)\in E)} \left(w^*B_{(r,s)}w\right)^2.
\end{aligned} \tag{15}$$

Applying this identity and noting that $\rho(\mathcal{A}_o) = (\rho(\tilde{A}))^2$, we arrive at following expression for $\rho(\mathcal{A})$ which is valid up to first order in p:

$$\begin{aligned}\rho(\mathcal{A}) &= \rho(\mathcal{A}_0) + c_1 p \\ &= \overline{\lambda}(\tilde{A})^2 + \Bigg(2\overline{\lambda}(\tilde{A}) - 2\overline{\lambda}(\tilde{A})^2 \\ &\qquad + \sum_{(r,s)\in E} \left(w^* B_{(r,s)} w\right)^2\Bigg) p. \qquad (16)\end{aligned}$$

In the special case of a torus network, $\overline{\lambda}(\tilde{A})$ can be computed analytically [27], [28]. For completeness, we state this result here.

Theorem 2.4: In a d-dimensional torus or d-lattice with N nodes, the asymptotic expression for the second largest eigenvalue of the weight matrix A (equivalently $\overline{\lambda}(\tilde{A})$) is given by

$$\lambda_2(A) = 1 - \beta\frac{8\pi^2}{N^{2/d}} + O\left(\frac{1}{N^{4/d}}\right).$$

Proof: The proof is given in the Appendix. ∎

With this result, we are able to derive an analytic form for the decay factor in tori networks.

Theorem 2.5: For a d-dimensional torus with N nodes, the first order expansion (in p) of the decay factor is given by

$$\rho(\mathcal{A}) = 1 - (1-p)\beta\frac{16\pi^2}{N^{2/d}} + O\left(\frac{1}{N^{3/d}}\right). \qquad (17)$$

Proof: We first note that, by substituting the value for $\overline{\lambda}(\tilde{A})$ given by Theorem 2.4 into (16), we arrive at the following expression for $\rho(\mathcal{A})$

$$\begin{aligned}\rho(\mathcal{A}) = 1 - (1-p)\beta\frac{16\pi^2}{N^{2/d}} &+ O\left(\frac{1}{N^{4/d}}\right) \\ &+ p\sum_{(r,s)\in E}\left(w^* B_{(r,s)} w\right)^2. \qquad (18)\end{aligned}$$

We now prove the theorem by showing that the term containing the summation of $B_{(r,s)}$ matrices is of order $O(1/N^{3/d})$.

Recall that each $B_{(r,s)}$ matrix is of the form $\beta b_{(r,s)} b^*_{(r,s)}$ where $b_{(r,s)}$ is a vector of all zeros, excepting the r^{th} and s^{th} components which are equal to 1 and -1 respectively. Therefore, the following equivalence holds for the summation:

$$\sum_{(r,s)\in E}\left(w^* B_{(r,s)} w\right)^2 = \beta^2 \sum_{(r,s)\in E}(w_r - w_s)^4 \qquad (19)$$

where w_r and w_s are the r^{th} and s^{th} components of w. w is the eigenvector corresponding to the second largest eigenvalue of A, or equivalently, the eigenvector corresponding to the largest eigenvalue of $\tilde{A}$. In the case of a d-dimensional torus, there is an analytical expression for the eigenvectors A. Let n be such that $N = n^d$.

Each eigenvector of A is associated with a multi-dimensional index $(m_1, \ldots, m_d)$, for $0 \leq m_i \leq (n-1)$. The components of such an eigenvector are given by

$$v^{(m_1,\ldots,m_d)}_{(j_1,\ldots,j_d)} = e^{-i\frac{2\pi}{n}(j_1 m_1 + \cdots + j_d m_d)}$$

for $j_i = 0\ldots(n-1)$.

The eigenvector corresponding to the largest eigenvalue of A occurs when $m_1 = m_2 = \ldots = m_d = 0$. The second largest eigenvalue has multiplicity d with d independent eigenvectors; each has one m_i equal to 1 and all other m_i's equal to 0. We compute the asymptotic expression for (19) for the eigenvector with $m_0 = 1$ and $m_2 = \ldots = m_{n-1} = 0$. The computation for the other $d-1$ eigenvectors is similar.

Let w be the eigenvector with multi-index $(1, 0, \ldots, 0)$; its components are given by

$$w^{(1,0,\ldots,0)}_{(j_1,\ldots,j_d)} = e^{-i\frac{2\pi}{n}j_1}$$

for $m_i = 0\ldots(n-1)$. Substituting this expression for the r^{th} and s^{th} components of w in (19), we obtain

$$\begin{aligned}\beta^2 \sum_{(r,s)\in E}&\left(w^{(1,0,\ldots,0))}_{(r_1,\ldots,r_d)} - w^{(1,0,\ldots,0))}_{(s_1,\ldots,s_d)}\right)^4 \\ &= \beta^2\sum_{(r,s)\in E}\left(e^{-i\frac{2\pi}{n}r_1} - e^{-i\frac{2\pi}{n}s_1}\right)^4.\end{aligned}$$

Since (r,s) is an edge in the torus, we know that if nodes r and s share an edge in the first dimension then $|r_1 - s_1| = 1$. Otherwise $|r_1 - s_1| = 0$. Therefore, for all $(r,s) \in E$, we have

$$e^{-i\frac{2\pi}{n}r_1} - e^{-i\frac{2\pi}{n}s_1} \leq \frac{2\pi}{n}|r_1 - s_1|.$$

Applying this bound to (19) and using the fact that in d-dimensional torus with N nodes, there are n edges in each dimension, we get the following bound on the summation term:

$$\sum_{(r,s)\in E}\left(w^* B_{(r,s)} w\right)^2 \leq \frac{(2\pi)^4}{n^3} = \frac{(2\pi)^4}{N^{3/d}}.$$

Therefore the summation term of $B_{(r,s)}$ matrices is of order $O(1/N^{3/d})$ which gives the result in (17). ∎

It is interesting to note that for large N, the leading order behavior of the decay factor is

$$1 - (1-p)\beta\frac{16\pi^2}{N^{2/d}}.$$

Recall that β is the fraction that is sent across each link. Therefore for large N, the failure of links with probability p has the same effect on the convergence rate as decreasing β by a factor of $1-p$.

C. Simulations

In this section, we demonstrate through simulations that the relationship between network size and dimensionality and link failure probability in tori networks stated in Theorem 2.4 appears to hold even for smaller networks and a larger probability of link failure. Specifically, we demonstrate that, for a fixed failure probability, the leading order of the decay factor is related to the network size and dimension as follows:

$$\rho(\mathcal{A}) = 1 - O\left(\frac{1}{N^{2/d}}\right). \qquad (20)$$

In order to evaluate whether this relationship holds for different network sizes, we simulate the algorithm in one-dimensional tori (ring) networks with sizes ranging from 10 to 1000 nodes and in two-dimensional tori networks with sizes ranging from 36 to 1764 nodes. For all simulations we let links fail with a uniform probability of 0.1. In tori networks, the variance of deviation from average is the same at each node, and therefore, by Property 1 of Theorem 2.3

$$\mathbf{E}\left[\frac{\|\tilde{x}(k)\|_2^2}{N}\right] \propto \rho(\mathcal{A})^k$$

or equivalently

$$\log\left(\frac{\|\tilde{x}(k)\|_2^2}{N}\right) \propto k \log\left(\rho(\mathcal{A})\right).$$

To estimate the decay factor, $\rho(\mathcal{A})$, for each network size, we run the algorithm and record the log of the per node variance as a function of time. In order to guarantee that the simulations exhibit the worst case decay behavior, the initial matrix $\tilde{M}(0)$ must be such that it is not orthogonal to the eigenmatrix Y associated with the largest eigenvalue of $\mathcal{A}$, or equivalently we must have $\mathbf{tr}(\tilde{M}(0)Y) \neq 0$. Since Y is positive semidefinite and $Y \neq 0$ (see the proof of Theorem 2.3), any covariance matrix $\tilde{M}(0)$ will satisfy this property so long as $\mathbf{E}[\tilde{x}_r(0)^2] \neq 0$ for all $r \in V$. We achieve this by choosing each $x_r(0)$ uniformly at random from the interval [0,100].

We run each simulation until the plot of $\log(\|\tilde{x}(k)\|_2^2/N)$ is linear, indicating that largest eigenvalue of the $\mathcal{A}$ operator dominates the decay rate. We then find of the slope of this linear plot which gives us an estimate of $\log(\rho(\mathcal{A}))$. If the relationship between the decay rate, the network dimension, and the number of nodes as described in (20) holds, then a plot of $\log(1-\rho(\mathcal{A}))$ as a function of $\log(N)$ should have a slope of -2 for the ring networks and -1 for the 2-dimensional torus networks. Fig. 1 shows $\log(1-\rho(\mathcal{A}))$ versus $\log(N)$ using estimates of $\rho(\mathcal{A})$ generated by the procedure described above. For each type of network, the slope of the linear fit is very close to what is predicted by (20), -1.9792 for the 1-dimensional networks and -1.0011 for the 2-D networks. These results indicate that the relationship in (20) holds even for smaller network sizes.

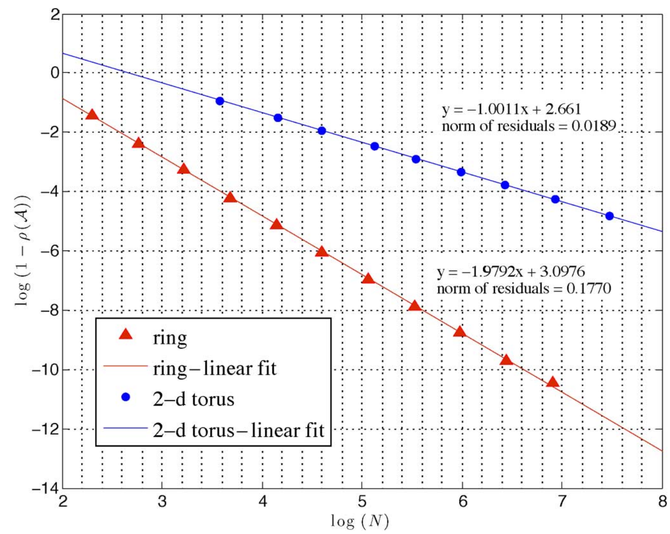

Fig. 1. $\log(1-\rho(\mathcal{A}))$ as function of the logarithm of the network size.

III. Incorporating Additive Noise

In this section, we extend our analysis to a network model where node values are perturbed by a zero-mean additive noise in each round. Let $u(k)$ be a zero-mean stochastic process with the autocorrelation matrix R defined by

$$R := \mathbf{E}\left[u(k)u(k)^*\right].$$

We assume that the additive noise processes are not correlated with the state nor with the stochastic processes governing communication failures. This type of noise can be used to model random insertions and deletions from the participating nodes in a distributed file system or data center.

The dynamics of this system are governed by an extension of the recursion equation in (3) that includes both multiplicative and additive noise

$$x(k+1) = \left(A + \sum_{(r,s)\in E} \delta_{(r,s)}(k) B_{(r,s)}\right) x(k) + u(k). \quad (21)$$

As in the first problem formulation, we are interested in the second order statistics of the deviation from average, $\mathbf{E}[\|\tilde{x}(k)\|_2^2]$. However in a system with additive noise, the average of all node values at time k, $(1/N)\sum_{r=1}^N x_r(k)$, drifts in a random walk about the average of the initial values $(1/N)\sum_{r=1}^N x_r(0)$. Additionally, since node values are perturbed in each round, one can no longer expect the nodes to converge to consensus at the current average, or equivalently, each $\tilde{x}_r$ does not converge to 0. In this extended model with additive noise, we do not measure the algorithm performance in terms of the convergence rate. Instead, performance is measured using the *steady-state total variance* of the deviation from average

$$TV_{ss} := \lim_{k\to\infty} \mathbf{E}\left[\|\tilde{x}(k)\|_2^2\right]$$

which is the sum of the variances of the deviation from the current average at each node. We are interested in the network conditions under which TV_{ss} is bounded as well as in quantifying that bound.

Problem Statement 2: Consider a distributed consensus algorithm on a network where each link $(r,s) \in E$ fails with independent probability $p_{(r,s)}$ and where node values are perturbed by a zero-mean stochastic process, as modeled by the system with additive and multiplicative noise (21). For a given input noise covariance R, determine the steady-state total variance of the deviation from average, $\lim_{k\to\infty} \mathbf{E}[\|\tilde{x}(k)\|_2^2]$.

Again, we study $\mathbf{E}[\|\tilde{x}(k)\|_2^2]$ by analyzing the recursion equation for the matrices $\tilde{M}(k) := \mathbf{E}[\tilde{x}(k)\tilde{x}(k)^*]$, noting that TV_{ss} is related to $\tilde{M}(k)$ as follows:

$$TV_{ss} = \lim_{k\to\infty} \mathbf{tr}\left(\mathbf{E}\left[\tilde{x}(k)\tilde{x}(k)^*\right]\right) = \lim_{k\to\infty} \mathbf{tr}\left(\tilde{M}(k)\right).$$

Using the same approach by which we derived the recursion (7), we can derive a recursion equation for the system with additive noise.

Lemma 3.1: The matrices $\tilde{M}(k)$ for the system (21) satisfy the recursion

$$\begin{aligned}\tilde{M}(k+1) &= \mathcal{P}\overline{A}\mathcal{P}\tilde{M}(k)\mathcal{P}\overline{A}\mathcal{P} \\ &\quad + \sum_{(r,s)\in E} \sigma_{(r,s)}^2 B_{(r,s)}\tilde{M}(k)B_{(r,s)} + \mathcal{P}R\mathcal{P} \\ &= \mathcal{A}\left(\tilde{M}(k)\right) + \mathcal{P}R\mathcal{P}\end{aligned}$$

where $\mathcal{A}$ is the matrix-valued operator defined in (8).

If the operator $\mathcal{A}$ is asymptotically stable, this recursion has a limit

$$\tilde{M}_{ss} := \lim_{k\to\infty} \tilde{M}(k)$$

and the limit $\tilde{M}_{ss}$ satisfies the following Lyapunov-like equation

$$\tilde{M}_{ss} = \mathcal{A}(\tilde{M}_{ss}) + \mathcal{P}R\mathcal{P}.$$

These facts lead to the following theorem relating to the second order statistics of the system (21).

Theorem 3.2: Consider the distributed consensus algorithm with random link failures as modeled by the system with multiplicative and additive noise (21).

1) The total variance of the deviation from average $\mathbf{E}[\|\tilde{x}(k)\|_2^2]$ has a steady-state limit if and only if

$$\rho(\mathcal{A}) < 1.$$

2) This limit is equal to the trace of $\tilde{M}_{ss}$, $TV_{ss} = \mathbf{tr}(\tilde{M}_{ss})$, where $\tilde{M}_{ss}$ satisfies the equation

$$\tilde{M}_{ss} = \mathcal{A}(\tilde{M}_{ss}) + \mathcal{P}R\mathcal{P}.$$

This theorem implies that if the consensus algorithm results in convergence to the average in a network with random link failures, the same algorithm executed on the same network with link failures *and additive noise* has a finite steady-state limit for the total deviation from average.

IV. Computational Procedures

We present computational methods for calculating the exact second order statistics of the deviation from average for systems with random communication failures. For the static-valued system model, the procedure involves computing the largest eigenvalue of a matrix-valued operator. For systems with additive noise, one must compute the trace of a solution of a Lyapunov-like equation.

A. Computing the Decay Factor

The decay factor of the static-valued system (3) is the spectral radius of the linear operator $\mathcal{A}$ defined in (8). Therefore, it is not necessary to perform Monte Carlo simulations of the original system (4) to compute decay factors. However, $\mathcal{A}$ is not in a form to which standard eigenvalue computation routines can be immediately applied. We present a simple procedure to obtain a matrix representation of $\mathcal{A}$ which can then be readily used in eigenvalue computation routines.

Recall that the Kronecker product of any two $m\times n$ and $r\times s$ matrices C and D respectively is the $mr \times ns$ matrix

$$C\otimes D := \begin{bmatrix} c_{11}D & \cdots & c_{1n}D \\ \vdots & \ddots & \vdots \\ c_{m1}D & \cdots & c_{mn}D \end{bmatrix}.$$

Let $\mathbf{vec}(X)$ denote the "vectorization" of any $m \times n$ matrix X constructed by stacking the matrix columns on top of one another to form an $mn \times 1$ vector. It then follows that a matrix equation of the form $Y = CXD$ can rewritten using matrix-vector products as

$$\mathbf{vec}(Y) = (C\otimes D)\mathbf{vec}(X).$$

Thus, using Kronecker products, $\mathcal{A}$ in (8) has a matrix representation of the form

$$\mathcal{A} = \left(\mathcal{P}\overline{A}\mathcal{P}\right)\otimes\left(\mathcal{P}\overline{A}\mathcal{P}\right) + \sum_{(r,s)\in E} \sigma_{(r,s)}^2 B_{(r,s)} \otimes B_{(r,s)}.$$

For a graph with N nodes, $\mathcal{A}$ is an $N^2 \times N^2$ matrix. This matrix representation can be used to find $\rho(\mathcal{A})$ via readily available eigenvalue routines in MATLAB. Due to the structure of $\mathcal{A}$, it is also possible to compute the eigenvalues in a more efficient manner. We briefly outline this procedure here. For large values of N, one can use an Arnoldi eigensolver to determine the eigenvalues of $\mathcal{A}$ in a constant number of matrix-vector multiplications that depends on the structure of $\mathcal{A}$. Since $\mathcal{A}$ is the sum of $M+1$ terms, this matrix-vector multiplication can also be computed by multiplying each of the terms by the N^2 vector and summing the result. The product of $(P\overline{A}P)\otimes(P\overline{A}P)$ and an N^2-vector can be computed in $O(N^3)$. Each $B_{(r,s)}\otimes B_{(r,s)}$ contains exactly 16 non-zero elements, and thus the product of each $B_{(r,s)} \otimes B_{(r,s)}$ times an N^2-vector can be computed in $O(1)$. Therefore, it is possible to find the eigenvalues of $\mathcal{A}$ in $O(N^3)$.

B. Computing the Steady-State Total Variance

Recall that the steady-state total variance of the deviation from average TV_{ss} is the trace of $\tilde{M}_{ss}$ where $\tilde{M}_{ss}$ satisfies the Lyapunov-like equation

$$\begin{aligned}\tilde{M}_{ss} &= (\mathcal{P}\overline{A}\mathcal{P})\tilde{M}_{ss}(\mathcal{P}\overline{A}\mathcal{P}) \\ &\quad + \sum_{(r,s)\in E} \sigma_{(r,s)}^2 B_{(r,s)}\tilde{M}_{ss}B_{(r,s)} + \mathcal{P}R\mathcal{P}\end{aligned}$$

where R is the covariance matrix of the additive noise process $u(k)$.

We again use Kronecker products to find an expression for $\tilde{M}_{ss}$

$$\begin{aligned}\mathbf{vec}(\tilde{M}_{ss}) &= \left((\mathcal{P}\overline{A}\mathcal{P})\otimes(\mathcal{P}\overline{A}\mathcal{P})\right)\mathbf{vec}(\tilde{M}_{ss}) \\ &\quad + \left(\sum_{(i,j)\in E} \sigma_{(r,s)}^2 B_{(r,s)}\otimes B_{(r,s)}\right)\mathbf{vec}(\tilde{M}_{ss}) \\ &\quad + \mathbf{vec}(\mathcal{P}R\mathcal{P}).\end{aligned}$$

Using this expression, $\mathbf{vec}(\tilde{M}_{ss})$ can be computed directly for any given algorithm instance and covariance matrix R. One can then reassemble $\tilde{M}_{ss}$ from $\mathbf{vec}(\tilde{M}_{ss})$ and find its trace.

In the next section, we use our computational procedures to calculate the decay factor and steady-state total deviation from average for various network examples.

V. Examples

We examine the second order statistics of the deviation from average for the consensus algorithm as a function of uniform link failure probability. For static-valued networks, we give computational results for different network topologies and values of β to illustrate the relationship between the probability of failure, the structure of the network, and the choice of β. For networks with additive noise, we give results that consider all three of these factors, and we also explore the effects of the size of the variance of the additive noise process on the variance of the deviation from average. For each class of problems, MATLAB was used to produce results according to the computational procedures described in the previous section.

A. Decay Factors

We first investigate the behavior of the decay factor $\rho(\mathcal{A})$ in systems with no additive noise. For each network topology, we compute the decay factor for several values of β including the value that is optimal for each graph when there are no communication failures. This optimal β is the edge weight that yields the smallest decay factor in networks with reliable communication links. The value is given by the following [10]

$$\beta^* = \frac{2}{\underline{\lambda}(\mathcal{L}) + \overline{\lambda}(\mathcal{L})}$$

where $\underline{\lambda}(\mathcal{L})$ and $\overline{\lambda}(\mathcal{L})$ are the second smallest and the largest eigenvalues of the Laplacian matrix of the graph, respectively.

Figs. 2 and 3 give the decay factors for a ring network with 9 nodes and a 2-dimensional discrete torus with 25 nodes. For each topology, we show the decay factors for the optimal β, a β that is larger than optimal, $\beta = 1/\Delta$, where Δ is the degree of each node in the network, and a β that is smaller than optimal, $\beta = 1/2\Delta$. For the ring network, the larger β is 0.5, the optimal β is approximately 0.4601, and the smaller β is 0.25. For the 2-dimensional torus, the larger β is 0.25, the optimal β is approximately 0.2321 and the smaller β is 0.125.

As expected, in both networks, when there are no link failures, the decay factor is smallest for the optimal β. Surprisingly, for the maximum β, the decay factors decrease for small probabilities of failure, and this edge weight yields better performance than the optimal weight. The decay factor continues to decrease until the failure probability reaches approximately 0.1 and then steadily increases. For the case where $\beta = 1/2\Delta$, the decay factor is consistently larger than that for the optimal β. Similar trends can be observed in the decay factors larger networks, however the difference for the various choices of β is not as pronounced.

We also compute the decay factors for an Erdős-Rényi (ER) random graph [29] of 50 nodes where each pair of nodes is connected with probability 0.25. The graph has 319 edges and a maximum node degree of 20. The decay factors are given in Fig. 4. The optimal β is approximately 0.071. We also show decay factors for values of β that are larger and smaller than optimal, $\beta = 0.081$ and $\beta = 0.061$, respectively. As in the results for the torus networks, the optimal β yields the smallest decay factor when there is zero probability of edge failure. When failures are introduced, the decay factor initially decreases for the larger value of β, where it actually results in faster convergence than the optimal β.

We conjecture that link failures reduce the effective weight of the values that are sent across each edge over a large number of rounds. In the case where β is larger than the optimal choice, the introduction of failures decreases the effective weight to approach the optimal β, and thus the algorithm performance actually improves. These results demonstrate that there is a relationship between the failure probability and the choice of β, and therefore it seems possible to select a β that optimizes performance for a given failure probability.

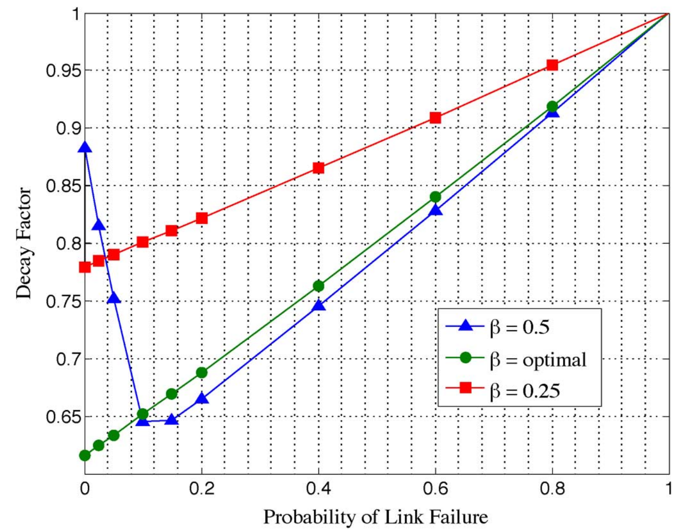

Fig. 2. Decay factor for various link failure probabilities in a 9 node ring network.

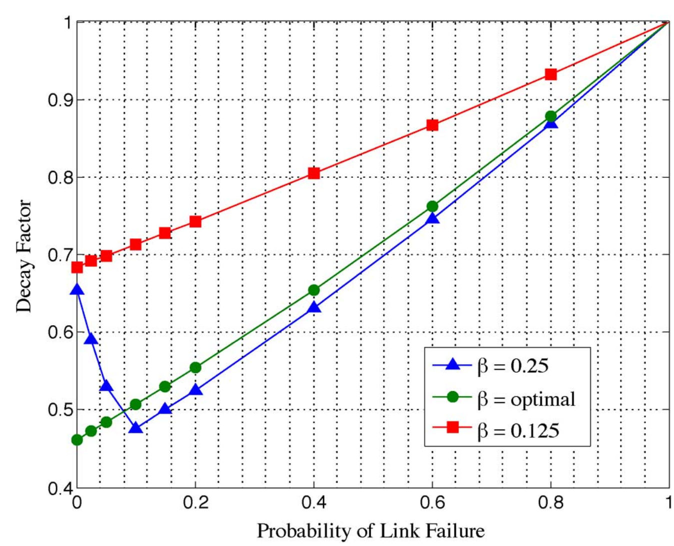

Fig. 3. Decay factor for various link failure probabilities in a 25 node 2-D torus.

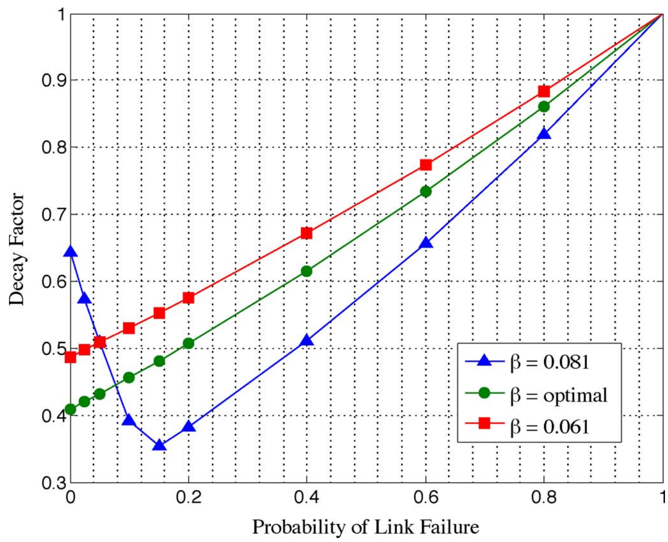

Fig. 4. Decay factor for various link failure probabilities in a 50 node ER random graph.

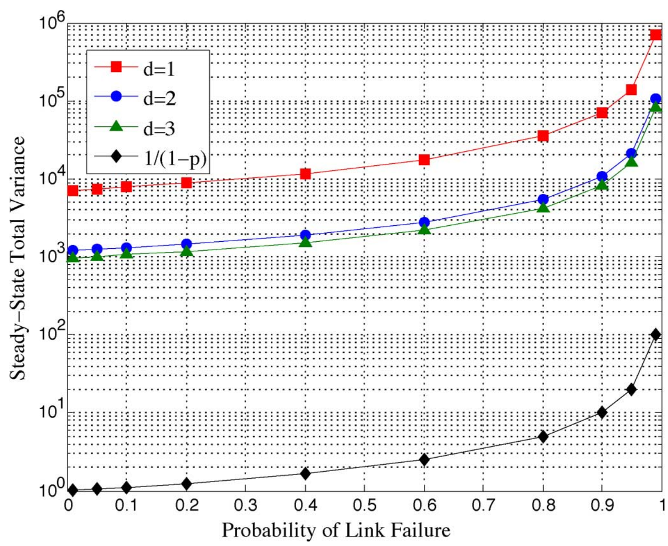

Fig. 5. Steady-state total variance of the deviation from average in 64 node torus networks of dimensions 1, 2, and 3.

B. Steady-State Total Variance

We next examine the steady-state total variance for systems with communications failures where the state values are perturbed by additive noise. While we do not know of any analytical result for the optimal choice of β for these systems when there are no communication failures, it has been shown that the optimal edge weight can be bounded above and below as $1/\overline{\lambda}(\mathcal{L}) \leq \beta^* < 2/\overline{\lambda}(\mathcal{L})$ [30].

Fig. 5 shows the results for 64 node tori networks of dimension 1, $(d=1)$, 2 $(d=2)$, and 3 $(d=3)$. For all networks, the variance of the additive noise is 10. For each network, we select β to be the lower bound of the optimal value, $\beta = 1/\overline{\lambda}(\mathcal{L})$. In a torus, this value corresponds to $1/2\Delta$ where Δ is the degree of each node in the graph. So, for $d=1$, we have $\beta = 1/4$, for $d=2$ we have $\beta = 1/8$, and for $d=3$, we have $\beta = 1/12$. While the magnitude of TV_{ss} is different for each of the three networks, the effect of increasing the probability of communication failure is appears to be same regardless of the dimension of the torus. In fact, for all three networks, the seems to grow as $1/(1-p)$, which is also shown in the figure.

In Fig. 6, we show the steady-state total variance for a 9 node ring network. The node values are perturbed by a zero-mean additive noise with a variance of 10. We use both the upper bound on the optimal value of β, β_{UB}, which is approximately 0.2578, and the lower bound on the optimal value, β_{LB}, which is approximately 0.5155. We observe that for β_{UB}, introducing a small probability of communication failure decreases the steady-state total variance. Just as the introduction of communication failures can decrease the decay factor in systems with no additive noise, this result demonstrates that communication failures can also improve performance by decreasing variance in systems with additive noise.

Finally, in Fig. 7, we show the steady-state total variance for an ER random graph with 30 nodes, where an edge exists between each pair of nodes with probability 0.25. The graph has 132 edges and a maximum node degree of 15. We use both the upper and lower bounds on the optimal β, $\beta_{UB} \approx 0.1167$ and $\beta_{LB} \approx 0.0584$. We show results for systems with zero-mean additive noise with variance of 1, 10, and 100. As in the previous scenario, a small probability of communication failure decreases the total variance for β_{UB} in all cases. An interesting observation is that the variance of the additive noise process does not affect the relationship between the probability of communication failure and the steady-state total variance. For all three additive noise processes, the behavior of the steady-state total variance is the same with respect to the probability of failure. Additionally, after the initial decrease, the variance appears to grow as $1/(1-p)$ for all network instances.

VI. Conclusion

We have presented an analysis of the distributed average consensus algorithm in networks with stochastic communication failures and shown that the problem can be formulated as a linear system with multiplicative noise. For systems with no additive noise, we have shown that the convergence rate of the consensus algorithm can be characterized by the spectral radius of a Lyapunov-like matrix recursion, and we have developed expressions for the multiplicative decay factor in the asymptotic limits of small failure probability and large networks. For systems with additive noise, we have shown that the steady-state total deviation from average is given by the solution of a Lyapunov-like equation. For both models, we have presented simulation-free methods for computing the second order statistics of the deviation from average. Using these methods, we have computed these second order statistics for various network topologies as a function of link failure probability. These computations indicate that there is a relationship between the network topology, the algorithm parameter β, and the probability of failure that is more complex than intuition would suggest. In particular, we show that for certain choices of β, communication failures can actually improve algorithm performance.

As the subject of current work, we are investigating the extension of our model and analysis to incorporate communication

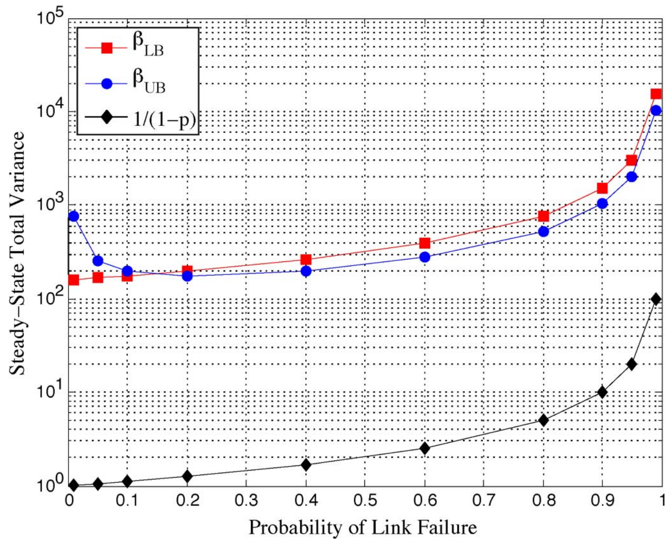

Fig. 6. Steady-state total variance of the deviation from average for various link failure probabilities in a nine-node ring network.

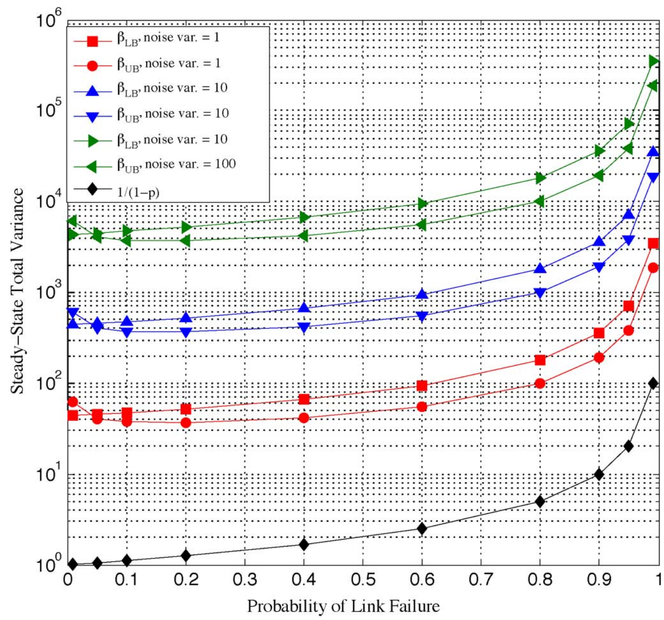

Fig. 7. Steady-state total variance of the deviation from average for various link failure probabilities in a 30 node ER random graph.

failures that are spatially and temporally correlated. Such extensions will enable the study of other network conditions such as network partitions and node failures.

Appendix A
Proof of Theorem 2.3

Proof: In order to prove the existence of a covariance matrix $\tilde{M}(0)$ for which the decay factor of the linear recursion (6) is precisely $\rho(\mathcal{A})$, we show that every eigenvalue of $\mathcal{A}$ has an associated positive semidefinite eigenmatrix. By setting $\tilde{M}(0)$ to be the eigenmatrix associated with the largest eigenvalue of $\mathcal{A}$, the worst case decay rate is achieved.

We first show that every for every eigenvalue-eigenmatrix pair (γ, W) of $\mathcal{A}$, there exists a symmetric matrix Z such that (γ, Z) is also an eigenvalue-eigenmatrix pair of $\mathcal{A}$. Let Z be the symmetric matrix $Z = W + W^*$. Then, we have

$$\begin{aligned}\mathcal{A}(Z) =& \mathcal{P}\overline{A}\mathcal{P}(W+W^*)\mathcal{P}\overline{A}\mathcal{P} \\ &+ \sum_{(r,s)\in E} \sigma^2_{(r,s)} B_{(r,s)}(W+W^*)B_{(r,s)} \\ =& \mathcal{P}\overline{A}\mathcal{P}W\mathcal{P}\overline{A}\mathcal{P} + \sum_{(r,s)\in E} \sigma^2_{(r,s)} B_{(r,s)} W B_{(r,s)} \\ &+ \mathcal{P}\overline{A}\mathcal{P}W^*\mathcal{P}\overline{A}\mathcal{P} \\ &+ \sum_{(r,s)\in E} \sigma^2_{(r,s)} B_{(r,s)} W^* B_{(r,s)} \\ =& \gamma W + \gamma^* W^*.\end{aligned}$$

Since $\mathcal{A}$ is self-adjoint, all of its eigenvalues are real, and so $\gamma = \gamma^*$, giving

$$\mathcal{A}(Z) = \gamma(W + W^*) = \gamma Z.$$

Let $\overline{\gamma}$ be the largest eigenvalue of $\mathcal{A}$, and let $\overline{Z}$ be the corresponding symmetric eigenmatrix. Then the decay factor of the $\mathcal{A}$ operator acting on an initial state of $\overline{Z}$ is precisely $|\overline{\gamma}|$. We note that as $\overline{Z}$ is symmetric, it can be decomposed as $\overline{Z} = \overline{Z}_+ + \overline{Z}_-$, where $\overline{Z}_+$ and $\overline{Z}_+$ are positive and negative semidefinite respectively. By the linearly of $\mathcal{A}$, we have

$$\mathcal{A}^k(\overline{Z}) = \mathcal{A}^k(\overline{Z}_+ + \overline{Z}_-) = \mathcal{A}^k(\overline{Z}_+) + \mathcal{A}^k(\overline{Z}).$$

Therefore the decay rate of $\mathcal{A}$ with the initial conditions $\overline{Z}$ is equivalent to the maximum of the decay rates of $\mathcal{A}$ with the initial condition $\overline{Z}_+$. and $\mathcal{A}$ with the initial condition $\overline{Z}_-$. This implies that there exists a covariance (positive semi-definite) matrix $\tilde{M}(0)$ such that the decay factor of the $\mathcal{A}$ operator acting on the initial $\tilde{M}(0)$ is the spectral radius of $\mathcal{A}$. ■

Appendix B
Proof of Theorem 2.4

Proof: We consider an N node tori network of dimension d as a d-dimensional n-array where $n^d = N$. The distributed average consensus algorithm is given by the following recursion equation:

$$\begin{aligned}x_{(j_1,\ldots,j_d)}(k+1) =& \alpha x_{(j_1,\ldots,j_d)}(k) + \beta\big(x_{(j_1-1,\ldots,j_d)}(k) \\ &+ x_{(j_1+1,\ldots,j_d)}(k) + \cdots + x_{(j_1,\ldots,j_d-1)}(k) + x_{(i_j,\ldots,j_d+1)}(k)\big)\end{aligned}$$

where each j_h ranges from 0 to $n-1$.

Each node communicates with its two neighbors along each of the d axes in each round. The numbers α and β then must satisfy $\alpha + 2d\beta = 1$. The sum in the equation above can be written as a multidimensional convolution by defining the array

$$\mathbf{a}_{(j_1,\ldots,j_d)} = \begin{cases} \alpha & j_1 = \cdots = j_d = 0, \\ \beta & j_l = \pm 1, \text{ and } j_m = 0 \text{ for } l \neq m, \\ 0 & \text{otherwise.} \end{cases}$$

We can then express the averaging operation defined above as

$$x(k+1) = M_{\mathbf{a}} x(k)$$

where $M_{\mathbf{a}}$ is the circulant operator associated with the array $\mathbf{a}$. The $N = n^d$ eigenvalues of $M_{\mathbf{a}}$ can be determined using the Discrete Fourier Transform, with $m_h := 0 \dots n-1$, for $j = 1 \dots d$

$$\begin{aligned}
\lambda_{(m_1,\dots,m_d)} &= \sum_{k_1,\dots,k_d \in \mathbb{Z}_n} a_{(k_1,\dots,k_d)} e^{-i\frac{2\pi}{n}(k_1 m_1 + \cdots + k_d m_d)} \\
&= \alpha + \beta\Big(e^{-i\frac{2\pi}{n}m_1} + e^{i\frac{2\pi}{n}m_1} + \cdots + e^{-i\frac{2\pi}{n}m_d} + e^{i\frac{2\pi}{n}m_d}\Big) \\
&= \alpha + 2\beta\left(\cos\left(\frac{2\pi}{n}m_1\right) + \cdots + \cos\left(\frac{2\pi}{n}m_d\right)\right).
\end{aligned}$$

The largest eigenvalue occurs when all m_j's are zero, and this eigenvalue is 1. The next largest eigenvalue occurs when all but one of the m_j's are zero and the non-zero m_j is 1. This eigenvalue corresponds to

$$\alpha + 2\beta(d-1) + 2\beta\cos\left(\frac{2\pi}{n}\right) = 1 - 2\beta\left(1 - \cos\left(\frac{2\pi}{n}\right)\right). \tag{22}$$

When N and consequently n are large, $\cos(2\pi/n)$ can be expressed as

$$\cos\left(\frac{2\pi}{n}\right) = 1 - \left(\frac{2\pi}{n}\right)^2 + O\left(\frac{1}{n^4}\right).$$

Substituting this equivalence into (22) and using the fact that $n = N^{1/d}$, we obtain the following expression for the second largest eigenvalue of $M_{\mathbf{a}}$

$$\lambda_2(M_{\mathbf{a}}) = 1 - \beta\frac{8\pi^2}{N^{2/d}} + O\left(\frac{1}{N^{4/d}}\right).$$

In the case of a lattice network, the A matrix is Toeplitz rather than circulant. However, the spectrum of the A matrix for a d-lattice and d-dimensional torus are equivalent in the limit of large N [31], [32]. Therefore, the convergence results can be applied to lattice networks as well as tori. ■

References

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

[2] R. Olfati-Saber and R. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Trans. Autom. Control*, vol. 49, no. 9, pp. 1520–1533, Sep. 2004.

[3] L. Moreau, "Stability of multiagent systems with time-dependent communication links," *IEEE Trans. Autom. Control*, vol. 50, no. 2, pp. 169–182, Feb. 2005.

[4] L. Xiao, S. Boyd, and S. Lall, "A space-time diffusion scheme for peer-to-peer least-squares estimation," in *Inform. Processing Sensor Networks*, 2006, pp. 168–176.

[5] G. Cybenko, "Dynamic load balancing for distributed memory multiprocessors," *J. Parallel Dist. Comput.*, vol. 7, no. 2, pp. 279–301, 1989.

[6] J. E. Boillat, "Load balancing and poisson equation in a graph," *Concurrency: Practice Exper.*, vol. 2, no. 4, pp. 289–313, 1990.

[7] D. Kempe, A. Dobra, and J. Gehrke, "Gossip-based computation of aggregate information," in *Proc. 44th Annu. IEEE Symp. Found. Comput. Sci.*, 2003, pp. 482–491.

[8] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized gossip algorithms," *IEEE Trans. Inform. Theory*, vol. 52, no. 6, pp. 2508–2530, Jun. 2006.

[9] D. P. Bertsekas and J. N. Tsitsiklis*, Parallel and Distributed Computation: Numerical Methods*. Nashua, NH: Athena Scientific, 1997.

[10] L. Xiao and S. Boyd, "Fast linear iterations for distributed averaging," *Syst. Control Lett.*, vol. 52, pp. 65–78, 2004.

[11] B. Bamieh, M. Jovanovic, P. Mitra, and S. Patterson, "Effect of topological dimension on rigidity of vehicle formations: Fundamental limitations of local feedback," in *Proc. 47th IEEE Conf. Decision Control*, 2008, pp. 369–374.

[12] L. Xiao, S. Boyd, and S. Lall, "A scheme for robust distributed sensor fusion based on average consensus," in *Proc. Inform. Processing Sensor Networks*, 2005, pp. 63–70.

[13] M. Cao, A. S. Morse, and B. D. O. Anderson, "Reaching a consensus in a dynamically changing environment: Convergence rates, measurement delays, and asynchronous events," *SIAM J. Control Optim.*, vol. 47, no. 2, pp. 601–623, 2008.

[14] A. Olshevsky and J. N. Tsitsiklis, "Convergence speed in distributed consensus and averaging," *SIAM J. Control Optim.*, vol. 48, no. 1, pp. 33–55, 2009.

[15] Y. Hatano and M. Mesbahi, "Agreement over random networks," *IEEE Trans. Autom. Control*, vol. 50, no. 11, pp. 1867–1872, Nov. 2005.

[16] C. Wu, "Synchronization and convergence of linear dynamics in random directed networks," *IEEE Trans. Autom. Control*, vol. 51, no. 7, pp. 1207–1210, Jul. 2006.

[17] M. Porfiri and D. J. Stilwell, "Consensus seeking over random weighted directed graphs," *IEEE Trans. Autom. Control*, vol. 52, no. 9, pp. 1767–1773, Sep. 2007.

[18] A. Tahbaz-Salehi and A. Jadbabaie, "A necessary and sufficient condition for consensus over random networks," *IEEE Trans. Autom. Control*, vol. 53, no. 3, pp. 791–795, Apr. 2008.

[19] S. Kar and J. Moura, "Distributed average consensus in sensor networks with random link failures," in *Proc. Int. Conf. Acoust., Speech, Signal Processing*, 2007, pp. 1013–1016.

[20] S. Patterson and B. Bamieh, "Distributed consensus with link failures as a structured stochastic uncertainty problem," in *Proc. 46'th Allerton Conf. Commun., Control, Comput.*, 2008, pp. 623–627.

[21] S. Kar and J. Moura, "Distributed average consensus in sensor networks with random link failures and communication channel noise," *IEEE Trans. Signal Process.*, to be published.

[22] E. Kranakis, D. Krizanc, and J. van den Berg, "Computing boolean functions on anonymous networks," *Inform. Computat.*, vol. 114, no. 2, pp. 214–236, 1994.

[23] J. N. Tsitsiklis, "Problems in Decentralized Decision Making and Computation," Ph.D. dissertation, MIT, Cambridge, MA, 1985.

[24] N. Elia, "Emergence of power laws in networked control systems," in *Proc. 45th IEEE Conf. Decision Control*, Dec. 2006, pp. 490–495.

[25] S. Boyd, L. El Ghauoi, E. Feron, and V. Balakrishnan*, Linear Matrix Inequalities in System and Control Theory*, ser. Studies in Applied Mathematics. Philadelphia, PA: SIAM, 1994, vol. 15.

[26] H. Baumgartel*, Analytic Perturbation Theory for Matrices and Operators*. Basel, Switzerland: Birkhauser Verlag, 1985.

[27] S. Patterson, B. Bamieh, and A. El Abbadi, "Brief announcement: Convergence analysis of scalable gossip protocols," in *Proc. 20th Int. Symp. Distrib. Comput.*, 2006, pp. 540–542.

[28] S. Patterson, B. Bamieh, and A. El Abbadi, "Convergence Analysis of Scalable Gossip Protocols," Tech. Rep. 2006–09, Jul. 2006.

[29] P. Erdos and A. Renyi, "On the evolution of random graphs," *Publ. Math. Inst. Hungarian Acad. Sci.*, vol. 5, pp. 17–61, 1960.

[30] L. Xiao, S. Boyd, and S.-J. Kim, "Distributed average consensus with least-mean-square deviation," *J. Parallel Dist. Comput.*, vol. 67, pp. 33–46, 2007.

[31] E. E. Tyrtyshnikov, "A unifying approach to some old and new theorems on distribution and clustering," *Linear Algebra Appl.*, vol. 232, pp. 1–43, 1996.

[32] S. Salapaka, A. Peirce, and M. Dahleh, "Analysis of a circulant based preconditioner for a class of lower rank extracted systems," *Numer. Lin. Algebra Appl.*, vol. 12, pp. 9–32, Feb. 2005.

Stacy Patterson received the B.S. degree in mathematics and computer science from Rutgers University, Piscataway, NJ, in 1998 and the M.S. and Ph.D. degrees in computer science from the University of California, Santa Barbara, in 2003 and 2009, respectively.

She is currently a Postdoctoral Scholar in the Department of Mechanical Engineering, University of California, Santa Barbara. Her research areas include distributed systems, sensor networks, and pervasive computing.

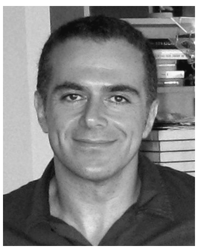

Bassam Bamieh (S'88–M'90–SM'02–F'08) received the Electrical Engineering and Physics degree from Valparaiso University, Valparaiso, IN, in 1983. and the M.Sc. and Ph.D. degrees from Rice University, Houston, TX, in 1986 and 1992, respectively.

From 1991 to 1998, he was with the Department of Electrical and Computer Engineering and the Coordinated Science Laboratory, University of Illinois at Urbana-Champaign. He is currently a Professor of mechanical engineering at the University of California, Santa Barbara. He is currently an Associate Editor of *Systems and Control Letters*. His current research interests are in distributed systems, shear flow turbulence modeling and control, quantum control, and thermo-acoustic energy conversion devices.

Dr. Bamieh received the AACC Hugo Schuck Best Paper Award, the IEEE CSS Axelby Outstanding Paper Award, and the NSF CAREER Award. He is a Control Systems Society Distinguished Lecturer.

Amr El Abbadi (SM'00) received the Ph.D. degree in computer science from Cornell University, Ithaca, NY.

In August 1987, he joined the Department of Computer Science, University of California, Santa Barbara (UC Santa Barbara), where he is currently a Professor. He has served as Area Editor for *Information Systems: An International Journal*, an Editor of *Information Processing Letters* (IPL), and an Associate Editor of the *Bulleten of the Technical Committee on Data Engineering*. He is currently the Chair of the Computer Science Department, UC Santa Barbara. His main research interests and accomplishments have been in understanding and developing basic mechanisms for supporting distributed information management systems, including databases, digital libraries, peer-to-peer systems, and spatial databases.

Dr. El Abbadi received the UCSB Senate Oustanding Mentorship Award in 2007. He served as a Board Member of the VLDB Endowment from 2002 to 2008.