Distributed Consensus with Link Failures as a Structured Stochastic Uncertainty Problem

Stacy Patterson and Bassam Bamieh

Abstract— **We consider the standard distributed average consensus algorithm under the conditions of random communication link failures, for which we analyze convergence of nodes to average consensus in the mean square sense. We first recast this problem as a discrete-time linear system with multiplicative random coefficients. We then rewrite the system equations as a nominal system in feedback with diagonally structured time-varying stochastic uncertainty; a problem for which necessary and sufficient mean square stability conditions have recently been derived. We investigate the particular instance of these conditions in the case of networked consensus with random link failures. In particular, we show that for circulant graphs, mean square convergence is guaranteed for any probability of link failure other than 1. We anticipate our particular analysis techniques to be applicable to the robust performance problem as well.**

I. Introduction

We consider the distributed average consensus problem over a connected, undirected network under the conditions of random communication link failures. The objective of distributed average consensus is for all nodes to reach consensus at the average of their initial values using a distributed algorithm that relies only on communication between neighboring nodes in the network graph. Distributed average consensus is a well-studied problem that has been investigated under a variety of contexts including vehicle formations [1], [2], [3], aggregation in sensor networks and peer-to-peer networks [4], [5], and load balancing in parallel processors [6], [7].

In a network where the communication structure is fixed, it has been shown that convergence of the consensus algorithm depends on the second smallest eigenvalue of the Laplacian matrix of the graph, see [8], [9] for example. Recently, interest has turned to investigating scenarios where the communication structure is time-varying. One such scenario is where the network evolves over time due to the mobility of nodes, and work has been done to identify necessary and sufficient conditions for convergence in this model [1], [2], [3]. Additionally, it has also been shown that in a network with a time-varying topology, as long as the union of all infinitely occurring graph instances is connected, there is a distributed consensus that will eventually converge [4].

Another scenario, which is the subject of this work, is one where the underlying network graph is fixed, but where communication links are not necessarily reliable. For example, in wireless networks, messages may be lost due to interference, and in wired networks, buffer overflow may result in packet loss. This scenario can be modeled using stochastic communication link failures. For a completely connected network where links fail with uniform probability, it has been shown that the consensus algorithm convergences almost surely [10]. In [11], the authors give necessary and sufficient conditions for almost sure convergence in networks where the algorithm weight matrices are i.i.d, and in [12], the authors analyze a model where the network topology is arbitrary and links may fail with non-uniform probability, and they establish sufficient conditions for mean square convergence based on the mean Laplacian. Additionally, both [13] and [14] study the effects of stochastic communication failures on the convergence rates of the consensus algorithms in directed and undirected graphs respectively.

In this work, we consider a simple distributed average consensus algorithm over a network where links fail with uniform probability. We illustrate how this problem can be recast as a nominal system in feedback with diagonally structured time-varying stochastic uncertainty; a problem for which necessary and sufficient mean square stability conditions have recently been derived [15], [16]. We then present an analysis of this system for circulant graphs and prove mean square stability for any link failure probability other than 1. We anticipate that our analysis techniques will be applicable to the robust performance problem as well.

An outline of this paper is as follows. In Section II, we formally define the system model and consensus algorithm. Section III illustrates how the problem can be reformulated as a structured stochastic uncertainty problem. In Section IV, we give stability analysis of this reformulation for circulant graphs. Finally, we conclude in Section V with a brief discussion of future research directions.

II. Preliminaries

We model the network by a connected, undirected graph $G = (V, E)$ where V is the set of N nodes and E is the set of M communication links between them. Each communication link has independent identical probability p of failing in each round, where $0 \leq p < 1$. If a link fails, no communication takes place across that link in either direction in that round. Each node has an initial value $x_j(0)$, and the objective is for nodes to reach consensus at the average of all values $x_{ave} := \sum_{j=1}^{N} x_j(0)$. We study a simple distributed averaging protocol where, in each round, each node exchanges information with all neighbors with which

This work was funded in part by NSF grant IIS 02-23022.

S. Patterson is with the Department of Computer Science, University of California, Santa Barbara, Santa Barbara, CA 93106, USA sep@cs.ucsb.edu

B. Bamieh is with the Department of Mechanical Engineering, University of California, Santa Barbara, Santa Barbara, CA 93106, USA bamieh@engineering.ucsb.edu

it has active (not failed) communication links. We denote this neighbor set by $\mathcal{N}_j(k)$. Each node updates its state according to the following rule,

$$x_j(k+1) = \beta \sum_{i \in N_j(k)} x_i(k) + (1 - |N_j(k)|)\beta x_j(k),$$

where $\beta > 0$ is a constant that is identical for all nodes.

Let L denote the Laplacian matrix of the graph G. Then, in the case where $p = 0$, the state $x(k)$ evolves according to the linear recursion

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

where the weight matrix A is given by $A := I - \beta L$. It is well known that the system converges to consensus at x_{ave} if and only if $|\lambda_2(A)| < 1$ [8], [9], [17], [18], [1]. If the graph is connected, there always exists a β for which this inequality holds, for example $\beta < \frac{1}{D}$ where D is the maximum vertex degree of the graph [9]. In this work, we assume that A is such that $|\lambda_2(A)| < 1$.

As shown in [14], the system (1) can be extended to incorporate link failures. We note that a similar model for communication failures in directed graphs is given in [19]. Let B be an oriented incidence matrix of the graph G. B is an $N \times M$ matrix where each column b_j, $j = 0, \ldots, M-1$ corresponds to an edge in the graph. Each b_j has a 1 in the row corresponding to one vertex incident to edge j and a -1 in the row corresponding to the other incident vertex[1]. The remainder of the entries are 0. In a system where all links have identical probability p of failing in each round, the state vector x evolves according to the following recursion,

$$x(k+1) = Ax(k) + \sum_{j=0}^{M-1} \mu_j(k) \beta b_j b_j^* x(k), \tag{2}$$

where $\mu_j(k)$ are Bernoulli random variables with

$$\mu_j(k) \ := \ \begin{cases} 1 & \text{with probability } p \\ 0 & \text{with probability } 1-p. \end{cases}$$

By choosing $\delta_j(k) := p - \mu_j(k)$ and noting that the $\delta_j(k)$ are zero mean, (2) can be rewritten as a linear recursion with zero mean multiplicative noise,

$$x(k+1) = \overline{A}x(k) + \sum_{j=0}^{M-1} \delta_j(k) \beta b_j b_j^* x(k),$$

where $\overline{A}$ is the expected weight matrix,

$$\overline{A} := A + p \sum_{j=0}^{M-1} \beta b_j b_j^*.$$

We measure how far the system state is from consensus at x_{ave} using the deviation from average vector, with each component defined by $\tilde{x}_i(k) := x_i(k) - x_{ave}$. The vector of deviations is the projection of $x(k)$ onto the subspace orthogonal to span($\mathbf{1}$), where $\mathbf{1}$ is the N-vector with all entries equal to 1. Equivalently,

$$\tilde{x}(k) := Px(k),$$

where $P := I - \frac{1}{N}\mathbf{1}\mathbf{1}^*$. The evolution of the this deviation from average vector is governed by the following recursion, which was first derived in [14],

$$\tilde{x}(k) = \tilde{A}\tilde{x}(k) + \sum_{j=0}^{M-1} \delta_j(k) \beta b_j {b_j}^* \tilde{x}(k), \tag{3}$$

with $\tilde{A} := P\overline{A}$. As (2) and, therefore, (3) are stochastic systems, we study convergence to x_{ave} in the mean square sense. Specifically, we say that the system *converges in mean square* if the total deviation from average converges to 0 as $k \to \infty$, i.e.

$$lim_{k \to \infty} \mathbf{E}\left[\|\tilde{x}(k)\|_2^2\right] = 0.$$

In the next section, we show how the mean square convergence of (3) can be characterized using stochastic structured uncertainty analysis.

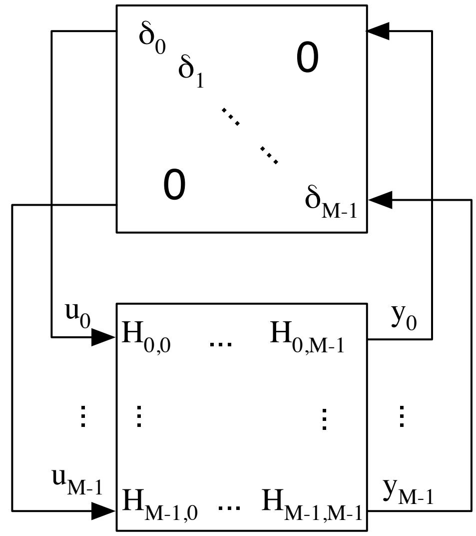

Fig. 1. The feedback system (H, Δ).

III. Problem Reformulation

In this section, we show how (3) can be converted to a form that is a amiable to robust stability analysis by rewriting the dynamics as a feedback between a nominal system H and a diagonally structured stochastic perturbation Δ, as illustrated in Fig. 1. With this formulation, necessary and sufficient conditions for mean square stability are given in terms of only the nominal system.

We first note that (3) can be decomposed into an equivalent representation of M^2 scalar subsystems, one for each pair of edges in G, as follows

$$H_{i,j} : \begin{cases} \tilde{x}(k+1) = \tilde{A}\tilde{x}(k) + \beta b_i u(k) \\ y_j(k) = b_j^* \tilde{x}(k) \end{cases} , i,j = 0, ..., M-1, \tag{4}$$

$$u_l(k) = \delta_l(k) y_l(k), \qquad l = 0, ..., M-1. \tag{5}$$

[1]Note that the choice of orientation for each edge is arbitrary and does not reflect any underlying property of the communication link.

Let $\mathcal{H}$ be the $M \times M$ matrix of H_2 norms of the subsystems

$$\mathcal{H} := \begin{bmatrix} \|H_{0,0}\|_2^2 & \cdots & \|H_{0,N-1}\|_2^2 \\ \vdots & & \vdots \\ \|H_{M-1,0}\|_2^2 & \cdots & \|H_{M-1,M-1}\|_2^2 \end{bmatrix},$$

where the discrete-time H_2 norm of $H_{i,j}$ is given by

$$\begin{aligned} \|H_{i,j}\|_2 &:= \mathbf{tr}\left(b_i{}^* \left(\sum_{l=0}^{\infty} \tilde{A}^l \ \beta b_j \beta b_j{}^* \tilde{A}^l \right) \ b_i \right) \\ &= \beta^2 \ \mathbf{tr}\left(b_i b_i{}^* \left(\sum_{l=0}^{\infty} \tilde{A}^l \ b_j b_j{}^* \tilde{A}^l \right) \right). \end{aligned}$$

In recent work, conditions for the mean square stability of such systems have been derived [15], [16]. We restate the relevant result here.

Theorem 3.1: Let H in (4) be a stable system, and let $\mathbf{var}\,(\delta_j) = \sigma^2$ for all $j = 0, \ldots, M-1$. Then, the system H in (4) in feedback with the uncertainty (5) is mean square stable if and only if

$$\sigma^2 \rho\,(\mathcal{H}) < 1,$$

where $\rho(\cdot)$ is the spectral radius.
This theorem implies that the distributed averaging algorithm over a network where links fail uniformly at random with probability p, as described by (2), converges to average consensus in mean square if and only if the spectral radius of $\mathcal{H}$ is strictly less than $\frac{1}{\sigma^2}$. In the next section, we use this result to prove that in specific network graph structures, namely circulant graphs, mean square convergence is guaranteed for any link failure probability other than 1.

IV. Structured Uncertainty Analysis

In this section, we derive mean square stability conditions for the class of circulant graphs. A circulant graph is any graph for which there exists a circulant adjacency matrix, such as a d-dimensional torus network.

If G is a circulant graph, then H is circulant, and therefore, so is the matrix $\mathcal{H}$. One can find the eigenvalues of $\mathcal{H}$ by taking the Discrete Fourier Transform (DFT) over any column of $\mathcal{H}$, $(\|h_0\|_2^2, \|h_1\|_2^2, ..., \|h_{M-1}\|_2^2)$. Here, we use the DFT over the first column of $\mathcal{H}$, which is

$$\begin{aligned} \widehat{h}_r &:= \sum_{j=0}^{M-1} \|h_j\|_2^2 e^{-i\frac{2\pi}{M}jr} \\ &= \beta^2 \ \mathbf{tr}\left(b_0 b_0{}^* \left(\sum_{l=0}^{\infty} \tilde{A}^l \sum_{j=0}^{M-1} b_j b_j{}^* e^{-i\frac{2\pi}{M}jr} \ \tilde{A}^l \right) \right), \end{aligned}$$

for $r = 0, \ldots, M-1$. As the stability of H depends on the spectral radius of $\mathcal{H}$, we must identify the Fourier coefficient with maximal absolute magnitude. This identification can be obtained from the following lemma.

Lemma 4.1: Let $\widehat{c}_r$ be the Discrete Fourier Transform over a sequence of non-negative reals, $(c_0, c_1, \ldots c_{M-1})$. The Fourier coefficient with maximal modulus occurs at $r = 0$, i.e.

$$|\widehat{c}_r| \le \widehat{c}_0 = c_1 + c_2 + \cdots c_{M-1}, \qquad r = 0, \ldots, M-1.$$

Proof: Assume, without loss of generality, that

$$c_0 + c_1 + \ldots + c_{M-1} = 1,$$

which also implies that $\widehat{c}_0 = 1$. Each Fourier coefficient $\widehat{c}_r, r = 0 \ldots M-1$, is a convex combination of a subset of the M roots of unity $e^{\frac{-2\pi i}{M}j}$, $j = 0 \ldots M-1$. Therefore each $\widehat{c}_r$ lies in the convex hull of those roots and thus also lies in the unit disk. This implies that $|\widehat{c}_r| \le 1$ for $r = 0 \ldots M-1$, and as $\widehat{c}_0 = 1$, we also have $|\widehat{c}_r| \le \widehat{c}_0$ for $r = 0 \ldots M-1$. ■

This result proves that $\rho(\mathcal{H})$ is the modulus of the Fourier coefficent with $r = 0$,

$$\widehat{h}_0 = \beta^2 \ \mathbf{tr}\left(b_0 b_0{}^* \left(\sum_{l=0}^{\infty} \tilde{A}^l \sum_{j=0}^{M-1} b_j b_j{}^* \ \tilde{A}^l \right) \right) \quad (6)$$

$$= \beta^2 \ \mathbf{tr}\left(b_0 b_0{}^* \sum_{l=0}^{\infty} \tilde{A}^l L \tilde{A}^l \right) \quad (7)$$

$$= \beta^2 \ \mathbf{tr}\left(b_0 b_0{}^* \ L \ (I - \tilde{A}^2)^{-1} \right). \quad (8)$$

Equation (7) is derived from (6) by observing that, by definition of b_j, $\sum_{j=0}^{M-1} b_j b_j^* = L$. $\tilde{A}$ and L are both circulant matrices, and therefore, they commute, which allows us to derive (8) from (7).

In order to simplify the expression for $\widehat{h}_0$ further, we observe that the choice of the column of $\mathcal{H}$ over which the DFT is taken is arbitrary; $\widehat{h}_r$ are equivalent for each column. Therefore, the following relationship holds for all $i, j = 0, \ldots, M-1$

$$\mathbf{tr}\left(b_i b_i{}^* \ L \ (I - \tilde{A}^2)^{-1} \right) = \mathbf{tr}\left(b_j b_j{}^* \ L \ (I - \tilde{A}^2)^{-1} \right),$$

and we can rewrite $\widehat{h}_0$ as the average of the 0^{th} Fourier coefficients for each column of $\mathcal{H}$,

$$\widehat{h}_0 = \frac{1}{M} \sum_{j=0}^{M-1} \beta^2 \ \mathbf{tr}\left(b_j b_j{}^* \ L \ (I - \tilde{A}^2)^{-1} \right) \quad (9)$$

$$= \frac{\beta^2}{M} \ \mathbf{tr}\left(L^2 (I - \tilde{A}^2)^{-1} \right). \quad (10)$$

Equation (9) indicates that one can determine $\rho(\mathcal{H})$ from the eigenvalues of L and $\tilde{A}$, and, as both are circulant operators, the eigenvalues can be obtained analytically using the DFT. However, in order to show that our original system (2) converges in mean square, we can employ a simpler approach using well known bounds on the eigenvalues of the Laplacian. First we note that the eigenvalues of $\tilde{A}$, denoted $\lambda_i(\tilde{A}), i = 1, \ldots, N$ are related to the eigenvalues of L, denoted $\lambda_i(L), i = 1, \ldots, N$, where the eigenvalues of L are given in order of decreasing magnitude, by the following,

$$\lambda_i(\tilde{A}) = \begin{cases} 1 - (1-p)\beta \ \lambda_i(L) & \text{for } i = 1, \ldots, N-1 \\ 0 & \text{for } i = N. \end{cases}$$

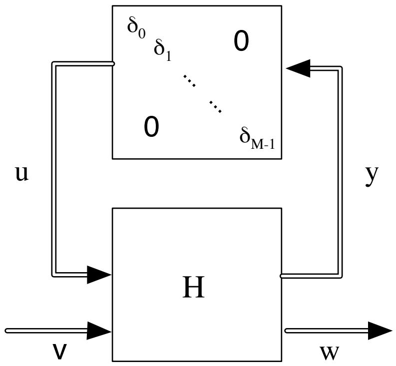

Fig. 2. The feedback system (H, Δ) with input v and output w.

Therefore it is possible to rewrite (9) in terms of just the eigenvalues of L,

$$\widehat{h}_0 = \frac{\beta^2}{M} \sum_{i=1}^{N-1} \frac{\lambda_i(L)^2}{1 - (1-(1-p)\beta\ \lambda_i(L))^2}. \tag{11}$$

It has been shown that the eigenvalues of the Laplacian of a connected graph are such that $0 = \lambda_N(L) < \lambda_{N-1}(L) \leq \lambda_{N-2}(L) \leq \ldots \leq \lambda_1(L) = 2D$, where D is the maximum vertex degree of the graph [20]. Since $\widehat{h}_0$ is monotonically increasing over $\lambda_i(L)$, the summation in (11) can be bounded by using the upper bound on $\lambda_N(L)$, giving,

$$\rho(\mathcal{H}) \leq \beta^2 \left(\frac{N-1}{M}\right) \left(\frac{(2D)^2}{1-(1-(1-p)\beta\ 2D)^2}\right). \tag{12}$$

Recall that by Theorem 3.1, the system H is mean square stable if and only if $\sigma^2 \rho(\mathcal{H}) < 1$. For uniform link failure probability p, $\sigma^2 = p - p^2$. Using this fact and the bound in (12) gives conditions that guarantee the mean square stability of the system. We state this result in the following theorem.

Theorem 4.2: For a circulant network with N nodes and M edges with a maximum vertex degree of D, the system that obeys the dynamics in (2) converges in mean square if β and p satisfy the following inequality

$$\left(\frac{N-1}{M}\right) \left(\frac{p\beta D}{1-(1-p)\beta D}\right) < 1.$$

As a direct consequence of this theorem, it is possible to select a β that guarantees mean square convergence for any probability of link failure other than 1.

Corollary 4.3: Consider the dynamics (2) for a circulant network. The system converges in mean square for any link failure probability $0 \leq p < 1$ if β is chosen such that $0 < \beta < \frac{1}{D}$.

We note that this choice of β also guarantees convergence for the algorithm over the network with no communication failures [9].

V. Discussion

Our aim in this paper has been to demonstrate how structured uncertainty analysis in the spirit of robust control can be used to address consensus problems with link failures. In this particular case, the uncertainty is of the time-varying stochastic type. While, in this initial work, we have shown mean square convergence for circulant networks, our ultimate goal is to quantify performance in the presence of link failures using the technique of equivalence between robust performance and robust stability. It may be possible to derive analytical bounds for the robust performance by studying the system described by Figure 2 and quantifying the variance of the output w in response to the input v. This technique has been previously studied in a deterministic context [21], [22], and we anticipate it will be applicable to the stochastic setting proposed in this work.

References

[1] A. Jadbabaie, J. Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Transactions on Automatic Control*, vol. 48, pp. 988–1001, June 2003.

[2] R. Olfati-Saber and R. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Transactions on Automatic Control*, vol. 49, pp. 1520–1533, September 2004.

[3] L. Moreau, "Stability of multiagent systems with time-dependent communication links," *IEEE Transactions on Automatic Control*, vol. 50, pp. 169–182, February 2005.

[4] L. Xiao, S. Boyd, and S. Lall, "A scheme for robust distributed sensor fusion based on average consensus," in *Proceedings of the 4th International Conference on Information Processing in Sensor Networks*, 2005, pp. 63–70.

[5] ——, "A space-time diffusion scheme for peer-to-peer least-squares estimation," in *Proceedings of the 5th International Conference on Information Processing in Sensor Networks*, 2006, pp. 168–176.

[6] G. Cybenko, "Dynamic load balancing for distributed memory multiprocessors," *Journal of Parallel and Distributed Computing*, vol. 7, no. 2, pp. 279–301, 1989.

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

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

[9] L. Xiao and S. Boyd, "Fast linear iterations for distributed averaging," in *Systems and Control Letters*, vol. 53, September 2004, pp. 65–78.

[10] Y. Hatano and M. Mesbahi, "Agreement over random networks," *IEEE Transactions on Automatic Control*, vol. 50, pp. 1867–1872, November 2005.

[11] A. Tahbaz-Salehi and A. Jadbabaie, "A necessary and sufficient condition for consensus over random networks," *IEEE Transactions on Automatic Control*, vol. 53, pp. 791–795, April 2008.

[12] S. Kar and J. Moura, "Sensor networks with random links: Topology design for distributed consensus," *IEEE Transactions on Signal Processing*, vol. 56, pp. 3315–3326, July 2008.

[13] F. Fagnani and S. Zampieri, "Average consensus with packet drop communication," *SIAM Journal on Control and Optimization (to appear)*, 2008.

[14] S. Patterson, B. Bamieh, and A. El Abbadi, "Distributed average consensus with stochastic communication failures," in *Proceedings of the 46th IEEE Conference on Decision and Control*, December 2007, pp. 490–495.

[15] N. Elia, "Remote stabilization over fading channels," *Systems and Control Letters*, vol. 54, pp. 237–249, March 2005.

[16] J. Lu and R. E. Skelton, "Mean-square small gain theorem for stochastic control: discrete-time case," *IEEE Transactions on Automatic Control*, vol. 47, pp. 490–494, March 2002.

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

[18] J. N. Tsitsiklis, "Problems in decentralized decision making and computation," Ph.D. dissertation, Massachusetts Institute of Technology, 1985.

[19] N. Elia, "Emergence of power laws in networked control systems," in *Proceedings of the 45th IEEE Conference on Decision and Control*, December 2006, pp. 490–495.

[20] B. Mohar, "The laplacian spectrum of graphs," in *Graph Theory, Combinatorics, and Applications*, E. Y. Alavi, G. Chartrand, O. R. Oellermann, and A. J. Schwenk, Eds., vol. 2, 1991, pp. 871–898.

[21] M. Khammash and J. B. Pearson, “Performance robustness of discrete-time systems with structured uncertainty,” *IEEE Transactions on Automatic Control*, vol. 36, pp. 398–412, April 1991.

[22] J. Shamma, “Robust stability with time-varying structured uncertainty,” *Automatic Control, IEEE Transactions on*, vol. 39, no. 4, pp. 714–724, 1994.