Consensus with Random Link Failures

Linear Systems with Multiplicative Noise
Structured Stochastic Uncertainty Analysis

Stacy Patterson, Bassam Bamieh and Amr El-Abbadi

Computer Science and Mechanical Engineering, UCSB

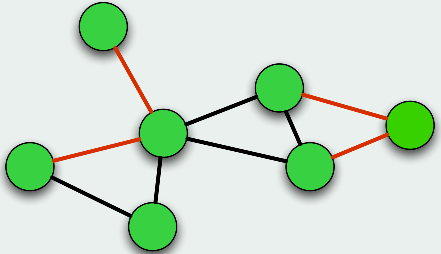

Outline

Linear Systems with Multiplicative Noise

Structured Stochastic Uncertainty

Consensus with Random Link Failures

Linear Systems with Random Multiplicative Inputs

- System model

$$x(k+1) \;=\; A\,x(k) \;+\; \Big(\delta_1(k)\,B_1 \;+\; \cdots \;+\; \delta_n(k)\,B_n\Big)\; x(k)$$

$\{\delta_1, \ldots, \delta_n\}$ uncorrelated, zero-mean white processes

- Covariance of the state, $P(k) \;=\; \mathcal{E}\left\{x(k)x^*(k)\right\}$

$$P(k+1) \;=\; A\,P(k)\,A^* + \Big(\sigma_1 B_1 P(k) B_1^* \;+\; \cdots \;+\; \sigma_n B_n P(k) B_n^*\Big)$$

- Def: System is *Mean Square Stable* (MSS) if

$$\lim_{k\to\infty} P(k) \;=\; 0$$

Linear Systems with Random Multiplicative Inputs

Lyapunov-like, matrix recursion

$$P(k+1) = A\,P(k)\,A^* + \left(\sigma_1 B_1 P(k) B_1^* + \cdots + \sigma_n B_n P(k) B_n^*\right)$$
$$P(k+1) = \mathcal{A}\left(P(k)\right)$$

$\mathcal{A}$ is a matrix-valued operator on matrices

$$X \stackrel{\mathcal{A}}{\longmapsto} A\,X\,A^* + \left(\sigma_1 B_1\,X\,B_1^* + \cdots + \sigma_n B_n\,X\,B_n^*\right)$$

SYSTEM IS MSS IFF $\rho(\mathcal{A}) < 1$

Structured Uncertainty

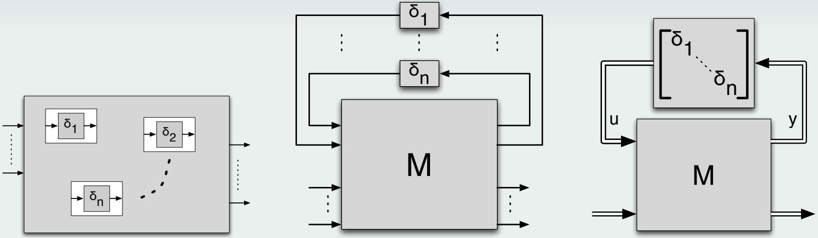

STRUCTURED UNCERTAINTY ANALYSIS:

Give stability conditions for uncertain system

in terms of M and bounds on δ's

Stochastic Structured Uncertainty

- $n = 1$ (Single uncertainty)

$$\text{MSS} \quad \Leftrightarrow \quad \|M\|_2^2 < \frac{1}{\sigma_\delta}$$

S. Boyd (80's)

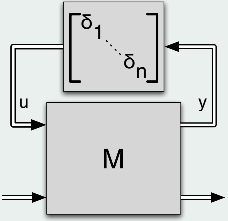

- Finite n

$$\rho\left(\begin{bmatrix} \|M_{11}\|_2^2 & \cdots & \|M_{1n}\|_2^2 \\ \vdots & & \vdots \\ \|M_{n1}\|_2^2 & \cdots & \|M_{nn}\|_2^2 \end{bmatrix}\right) < \frac{1}{\sigma_\delta}$$

Lu & Skelton '02, Elia '05

Consensus with Random Link Failures

Will formulate the problem of *consensus with random link failures* so as to use these tools

Distributed Average Consensus

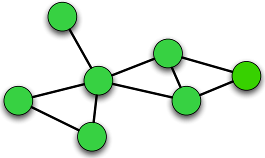

- Network Model
  - Undirected, connected graph $G = (V, E)$ with N nodes and M edges.
  - Each link has independent probability p of failing in each round.

Simple Averaging Protocol

In each round, each node sends equal fraction to each neighbor and keeps remaining fraction for self

$$x_i(k+1) = \beta \sum_{j \in N_i(k)} x_j(k) + (1 - \beta |N_i(k)|) x_i(k)$$

In a **static network**, dynamics can be represented by recursion equation

$$x(k+1) = Ax(k)$$

where $A = I - \beta L$. L is the Laplacian matrix of the graph.

e.g. 4 node ring network with $\beta := \frac{1}{3}$,

$$A := \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & 0 & \frac{1}{3} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 \\ 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & 0 & \frac{1}{3} & \frac{1}{3} \end{bmatrix}$$

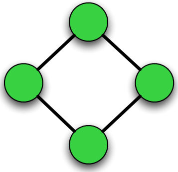

Well known that convergence to x_{ave} is guaranteed if $|\lambda_2(A)| < 1$.

System with Random A Matrix

Problem equivalent to

$$x(k+1) \;=\; A(k)\, x(k)$$

$A(k)$ is a matrix-valued random variable

Convergence in Stochastic Networks

- Some Related Work
  - Convergence is guaranteed almost surely in random graphs [Hatano and Mesbahi 2005, Porfiri and Stilwell 2007].
  - $|\lambda_2(\mathbf{E}[A(k)])| < 1$ is both a necessary and sufficient condition for almost sure convergence in random networks [Tahbaz-Salehi and Jadbabaie 2008].
    - Analysis based on ergodicity of sequence of A matrices.
  - $|\lambda_2(\mathbf{E}[A(k)])| < 1$ is sufficient condition for mean square convergence [Kar and Moura 2007].

Distributed Average Consensus in a Static Network

- The dynamics of the system can be represented by recursion equation

$$x(k+1) = Ax(k)$$

In a 4 node ring network with $\beta := \frac{1}{3}$,

$$A := \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & 0 & \frac{1}{3} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 \\ 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & 0 & \frac{1}{3} & \frac{1}{3} \end{bmatrix}$$

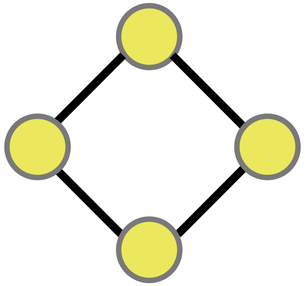

- Well known that convergence depends on $\lambda_2(A)$
  - Convergence guaranteed if $|\lambda_2(A)| < 1$
  - In d-dimensional torus or d-lattice with N nodes
    [Kranakis et al. 1994, Patterson et al. 2006, Carli et al. 2007]

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

Incorporating Communication Failures

- Consider the failure of edge $(1,2)$ in a ring with $\beta := \frac{1}{3}$

$$\begin{bmatrix} x(k+1) \end{bmatrix} = \begin{bmatrix} \frac{2}{3} & 0 & 0 & \frac{1}{3} \\ 0 & \frac{2}{3} & \frac{1}{3} & 0 \\ 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & 0 & \frac{1}{3} & \frac{1}{3} \end{bmatrix} \begin{bmatrix} x(k) \end{bmatrix}$$

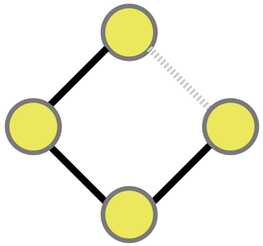

- Intuition: Perform protocol as if no failure occurred, then undo effects across failed links.

$$\begin{bmatrix} x(k+1) \end{bmatrix} = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & 0 & \frac{1}{3} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 \\ 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & 0 & \frac{1}{3} & \frac{1}{3} \end{bmatrix} \begin{bmatrix} x(k) \end{bmatrix} + \left(\frac{1}{3}\right) \begin{bmatrix} 1 & -1 & 0 & 0 \\ -1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x(k) \end{bmatrix}$$

General Recursion Equation

General equation that includes failure of all links is

$$x(k+1) = Ax(k) + \sum_{(i,j)\in E} \delta_{(i,j)}(k) B_{(i,j)} x(k)$$

where $\delta_{(i,j)}$ is a Bernoulli random variable

$$\delta_{(i,j)}(k) \; := \; \begin{cases} 1 & \text{with probability } p_{(i,j)} \; : \; \text{link has failed} \\ 0 & \text{with probability } (1 - p_{(i,j)}) \; : \; \text{link is active} \end{cases}$$

Rewrite recursion equation with zero-mean multiplicative noise

$$x(k+1) = \left(A + \sum_{(i,j)\in E} p_{(i,j)} B_{(i,j)} \right) x(k) + \sum_{(i,j)\in E(k)} \mu_{(i,j)}(k) B_{(i,j)} x(k)$$

where $\mu_{(i,j)}(k) := \delta_{(i,j)}(k) - p_{(i,j)}$ is zero-mean

$\bar{A} := A + \sum_{(i,j)\in E} p_{(i,j)} B_{(i,j)}$ is the *mean* (or *expected*) protocol matrx

Defining Convergence

We measure how far the system is from consensus at x_{ave} with **deviation from average vector**

$$\begin{aligned} \tilde{x}_i(k) &= x_i(k) - \frac{1}{n}(x_1(k) + ... + x_n(k)) \\ \tilde{x}(k) &= \left(I - \frac{1}{n}\mathbf{1}\mathbf{1}^*\right) x(k) \end{aligned}$$

Autocorrelation of deviation from average

$$\tilde{M}(k) := \mathbf{E}\{\tilde{x}(k)\tilde{x}^*(k)\} = \left(I - \frac{1}{n}\mathbf{1}\mathbf{1}^*\right) \mathbf{E}\{x(k)x^*(k)\} \left(I - \frac{1}{n}\mathbf{1}\mathbf{1}^*\right)$$

Goals:

- Determine conditions under which entries of $\tilde{M}(k)$ converge to 0 as $k \to \infty$
- Determine rate of convergence

Analyzing Convergence

- Dynamics obey Lyapunov-like recursion

$$\tilde{M}(k+1) = \left(I - \frac{1}{n}\mathbf{1}\mathbf{1}^*\right)\overline{A}\,\tilde{M}(k)\,\left(I - \frac{1}{n}\mathbf{1}\mathbf{1}^*\right)\overline{A} + \sum_{(i,j)\in E} \sigma^2_{(i,j)} B_{(i,j)} \tilde{M}(k) B_{(i,j)}$$

- **Decay Factor** - the factor by which entries of $\tilde{M}(k)$ decay in each round
- Decay factor is largest eigenvalue of matrix-valued operator $\mathcal{A}$

$$X \;\overset{\mathcal{A}}{\longmapsto} \left(I - \frac{1}{n}\mathbf{1}\mathbf{1}^*\right)\overline{A}\; X\; \left(I - \frac{1}{n}\mathbf{1}\mathbf{1}^*\right)\overline{A} + \sum_{(i,j)\in E} \sigma^2_{(i,j)} B_{(i,j)}\; X\; B_{(i,j)}$$

- If $p_{(i,j)} = 0$ for all $(i,j) \in E$

$$X \longmapsto \left(I - \frac{1}{n}\mathbf{1}\mathbf{1}^*\right) A\; X \left(I - \frac{1}{n}\mathbf{1}\mathbf{1}^*\right) A$$

Decay factor is $\lambda_2(A)^2$

Computing the Decay Factor

Kronecker product of matrices C and D

$$C_{m\times n}\otimes D_{r\times s} \;:=\; \begin{bmatrix} c_{11}D & \cdots & c_{1n}D \\ \vdots & \ddots & \vdots \\ c_{m1}D & \cdots & c_{mn}D \end{bmatrix}_{mr\times ns}$$

Matrix equation of the form $Y = CXD$ can be rewritten as

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

$$\tilde{M}(k+1) = \left(I - \frac{1}{n}\mathbf{1}\mathbf{1}^*\right)\overline{A}\;\tilde{M}(k)\;\left(I - \frac{1}{n}\mathbf{1}\mathbf{1}^*\right)\overline{A} + \sum_{(i,j)\in E}\sigma^2_{(i,j)}B_{(i,j)}\tilde{M}(k)B_{(i,j)}$$

becomes

$$\mathbf{vec}(\tilde{M}(k{+}1)) = \left(\left(\left(I - \frac{1}{n}\mathbf{1}\mathbf{1}^*\right)\overline{A}\right)\otimes\left(\left(I - \frac{1}{n}\mathbf{1}\mathbf{1}^*\right)\overline{A}\right) + \sum_{(i,j)\in E}\sigma^2_{(i,j)}B_{(i,j)}\otimes B_{(i,j)}\right)\mathbf{vec}(\tilde{M}(k))$$

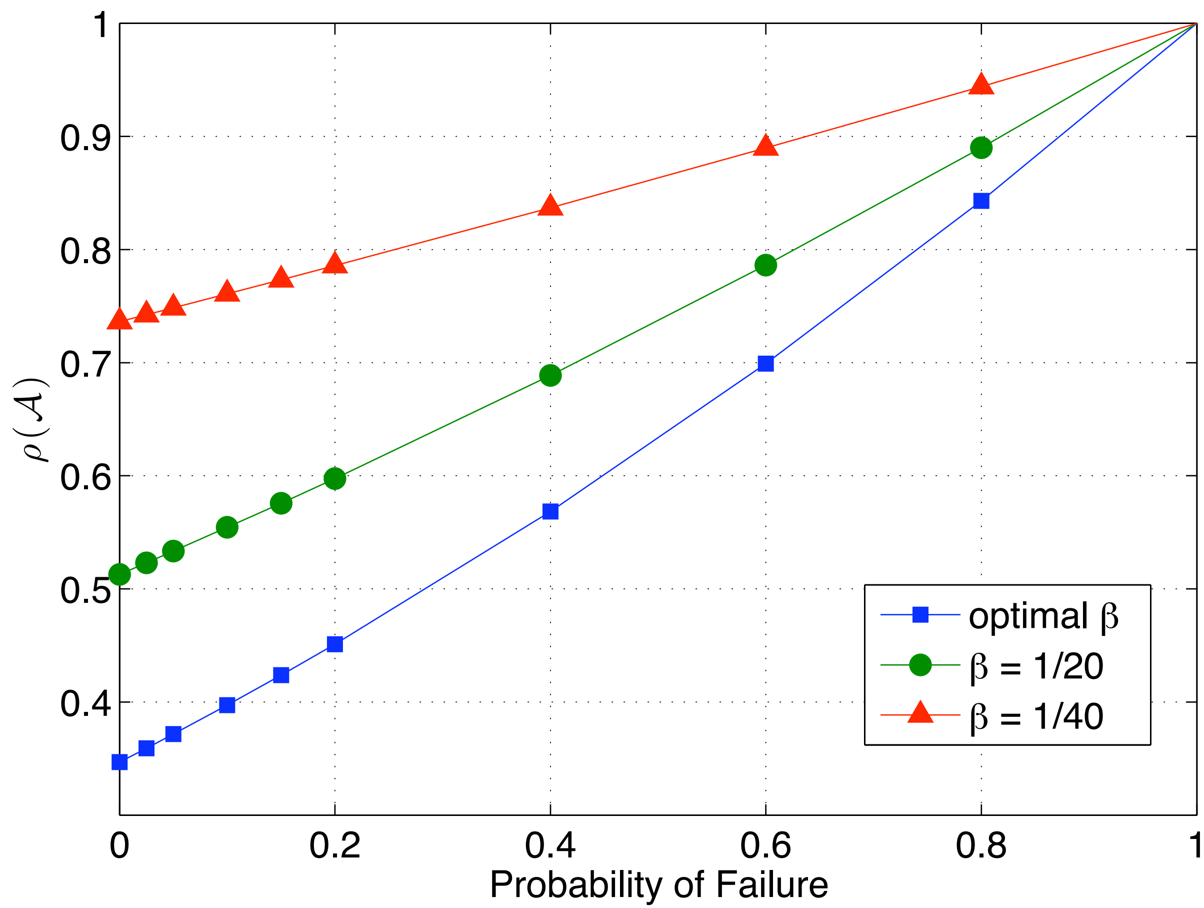

50 node ER graph, each node connected with probability 0.25

More Computational Results

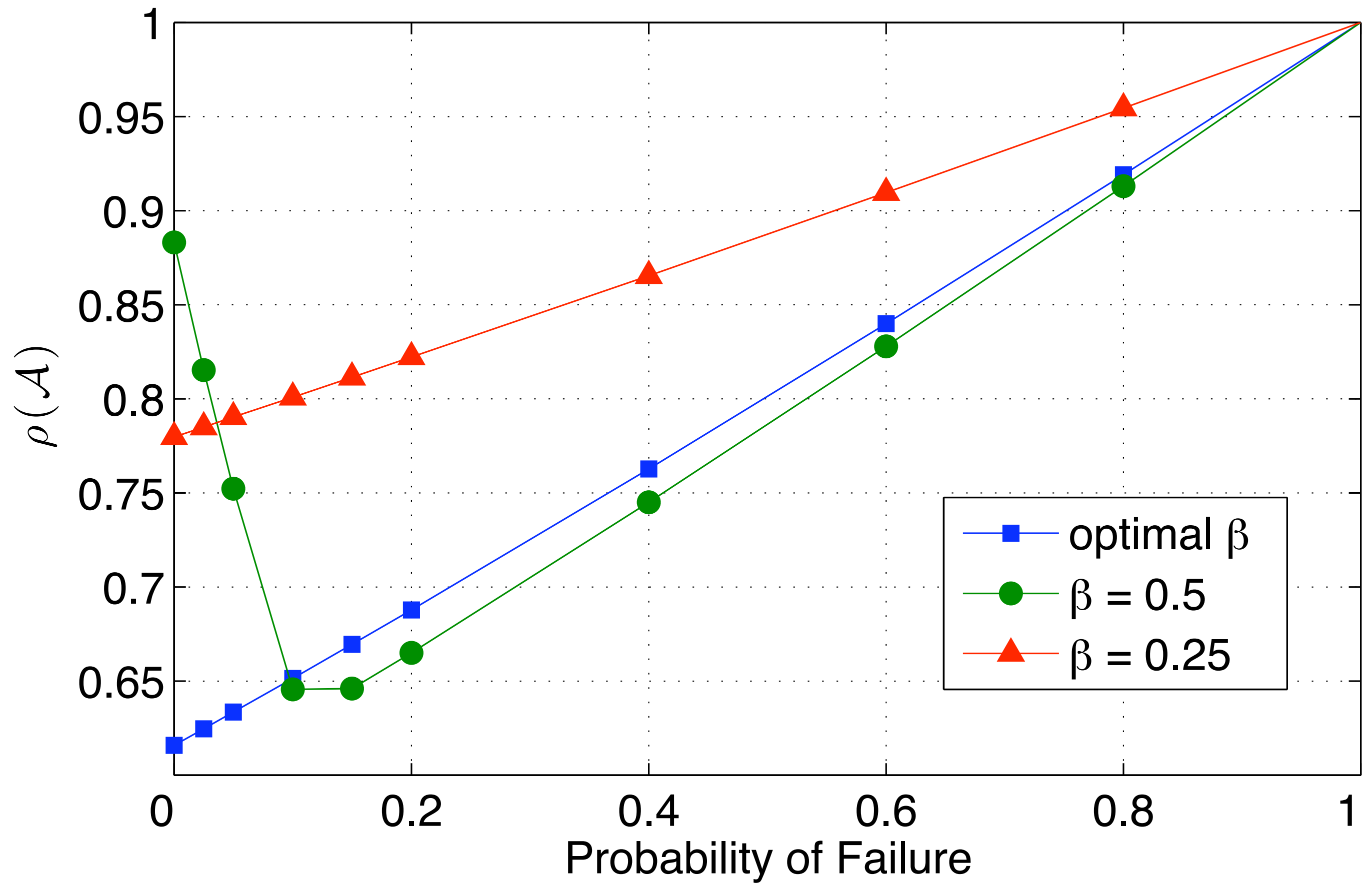

9 node ring network

Spectral Perturbation Analysis

For uniform link failure probability p

$$\mathcal{A}(X) := \left(\tilde{A} + p\beta\mathcal{L}\right) X \left(\tilde{A} + p\beta\mathcal{L}\right) + \ (p - p^2) \sum_{(i,j)\in E} B_{(i,j)} \ X \ B_{(i,j)}$$

where $\tilde{A} = \left(I - \frac{1}{n}\mathbf{1}\mathbf{1}^*\right) A$

$$\Updownarrow$$

$$\mathcal{A}(p, X) = \mathcal{A}_0(X) + p\mathcal{A}_1(X) + p^2\mathcal{A}_2(X)$$

where

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

Perturbation Analysis

Let γ be an eigenvalue of $\mathcal{A}(0, \cdot)$

Power series expansion of γ is

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

where λ is eigenvalue of $\mathcal{A}_0$ with eigenmatrix V and

$$c_1 = \frac{< V, A_1(V) >}{< V, V >}$$

Interested in largest eigenvalue of $\mathcal{A}$ up to first order in p

$$\begin{aligned}
\rho(\mathcal{A}) &= \rho(\mathcal{A}_0) + c_1 p \\
&= \rho(\tilde{A})^2 + \frac{< ww^*, \mathcal{A}_1(ww^*) >}{< ww^*, ww^* >} p \\
&= \overline{\lambda}(\tilde{A})^2 + \left(2\overline{\lambda}(\tilde{A}) - 2\overline{\lambda}(\tilde{A})^2 + \frac{1}{||w||^2} \sum_{(i,j)\in E} \left(w^* B_{(i,j)} w \right)^2 \right) p
\end{aligned}$$

where w is eigenvector corresponding to $\overline{\lambda}(\tilde{A})$ (Fiedler vector)

Perturbation Analysis in Tori Networks

In a d-dimensional torus with N nodes

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

For d-dimensional tori with N nodes, the first order expansion (in p) of decay factor

$$\rho(\mathcal{A}) \;=\; \overline{\lambda}(\tilde{A})^2 + \left(2\overline{\lambda}(\tilde{A}) - 2\overline{\lambda}(\tilde{A})^2 \;+\; \frac{1}{||w||^2}\sum_{(i,j)\in E}\left(w^* B_{(i,j)} w\right)^2\right) p$$

$$\Updownarrow$$

$$\rho(\mathcal{A}) \;=\; 1 \;-\; (1-p)\beta\,\frac{16\pi^2}{N^{2/d}} \;+\; O\left(\frac{1}{N^{4/d}}\right)$$

For large network size, link failures reduce decay factor by $(1-p)$

Problem Reformulation

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

- Idea: decompose system into 2 components.

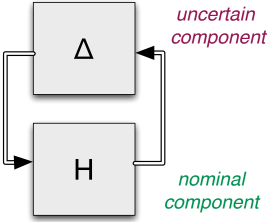

- Mean square stability conditions can be given in terms of only the nominal system.

Stochastic Structured Uncertainty Problem

Rewrite

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

as M^2 scalar subsystems, one for each pair of edges.

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

$$u_l(k) = \delta_l(k) y_l(k) \qquad l = 0, \ldots, M-1.$$

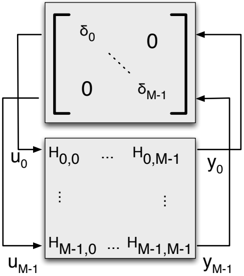

General Convergence Results

Can derive convergence condition in terms of H_2 norms of the subsystems of H

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

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

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

The system is mean square stable if and only if $\sigma^2 \rho(\mathcal{H}) < 1$.

[Elia 2005, Lu and Skelton 2002]

$$\Updownarrow$$

The consensus algorithm converges in mean square if and only if $(p - p^2) \ \rho(\mathcal{H}) < 1$.

Analysis for Circulant Graphs

- G is circulant graph (e.g. torus) $\Rightarrow H$ is circulant $\Rightarrow \mathcal{H}$ is circulant.
- Can find eigenvalues of $\mathcal{H}$ by taking DFT over any column $(\|h_0\|_2^2, \|h_1\|_2^2, ..., \|h_{M-1}\|_2^2)$.

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

- Fourier coefficient with maximal modulus occurs at $r = 0$.

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

Analysis for Circulant Graphs

- Note that $\sum_{j=0}^{M-1} b_j b_j{}^* = L$

$$
\begin{aligned}
\widehat{h}_0 &= \beta^2 \,\mathbf{tr}\left(b_0 b_0{}^* \left(\sum_{l=0}^{\infty} \tilde{A}^l \sum_{j=0}^{M-1} b_j b_j{}^* \; \tilde{A}^l \right)\right) \\
&= \beta^2 \,\mathbf{tr}\left(b_0 b_0{}^* \; L \; (I - \tilde{A}^2)^{-1} \right) \\
&= \frac{1}{M} \sum_{j=0}^{M-1} \beta^2 \,\mathbf{tr}\left(b_j b_j{}^* \; L \; (I - \tilde{A}^2)^{-1} \right) \\
&= \frac{\beta^2}{M} \,\mathbf{tr}\left(L^2 (I - \tilde{A}^2)^{-1} \right)
\end{aligned}
$$

The trace can be determined from the eigenvalues of L and $\tilde{A}$.

Analysis for Circulant Graphs

- $\tilde{A}$ and L are related by the following: $\tilde{A} = I - (1-p)\beta L$.
- Can write $\rho(\mathcal{H})$ in terms of eigenvalues of L.

$$\rho(\mathcal{H}) = \frac{\beta^2}{M} \sum_{i=0}^{N-1} \frac{\lambda_i(L)^2}{1 - (1 - (1-p)\beta\ \lambda_i(L))^2}$$

- Well known that $0 \leq \lambda_i(L) \leq 2(max\_degree)$ for $i = 0 \ldots N-1$.

  Therefore, we can bound $\rho(\mathcal{H})$.

$$\rho(\mathcal{H}) \leq \beta^2 \left(\frac{N-1}{M}\right) \left(\frac{(2(max\_degree))^2}{1 - (1 - (1-p)\beta\ 2(max\_degree))^2}\right)$$

Analysis for Circulant Graphs

- Recall that for a general graph, the system converges in mean square if and only if $(p - p^2)\rho(\mathcal{H}) < 1$.
- In a tori, the system converges in mean square if

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

- For any circulant network, there is a β such that the system converges in mean square for any link failure probability $0 \leq p < 1$.

Summary

- Shown how consensus problem with stochastic communication failures can be recast as stochastic structured uncertainty problem.
- Given mean square convergence conditions for this formulation.
- Demonstrated that for circulant networks, mean square convergence is guaranteed.
- Future work - investigation of performance robustness and convergence rates.