NCS Lecture 8
A Primer on Graph Theory

Richard M. Murray
Control and Dynamical Systems
California Institute of Technology

Goals

- Introduce some motivating cooperative control problems
- Describe basic concepts in graph theory (review)
- Introduce matrices associated with graphs and related properties (spectra)

Based on CDS 270 notes by Reza Olfati-Saber (Dartmouth) and PhD thesis of Alex Fax (Northrop Grumman).

References

- R. Diestel, Graph Theory. Springer-Verlag, 2000.
- C. Godsil and G. Royle, Algebraic Graph Theory. Springer, 2001.
- R. A. Horn and C. R. Johnson, Matrix Analysis. Cambridge Univ Press,1987.
- R. Olfati-Saber and M, "Consensus Problems in Networks of Agents", *IEEE Transactions on Automatic Control*, 2004.

Cooperative Control Applications


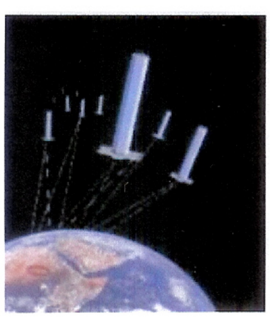


Transportation

- Air traffic control
- Intelligent transportation systems (ala California PATH project)

Military

- Distributed aperture imaging
- Battlespace management

Scientific

- Distributed aperture imaging
- Adaptive sensor networks (eg, Adaptive Ocean Sampling Network)

Commercial

- Building sensor networks (related)

Non-vehicle based applications

- Communication networks (routing, ...)
- Power grid, supply chain mgmt

Example: RoboFlag (D'Andrea, Cornell)

D'Andrea & M
ACC 2003

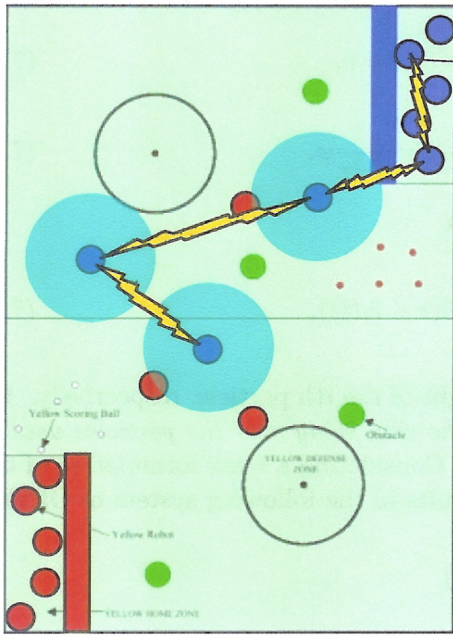

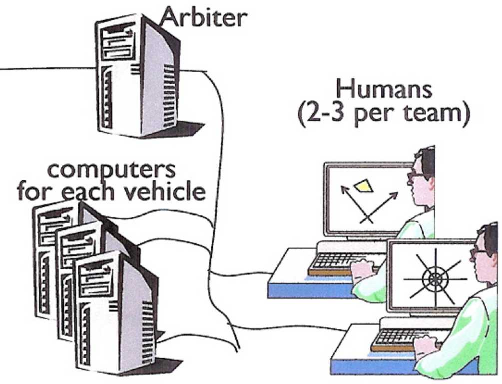

Robot version of "Capture the Flag"

- Teams try to capture flag of opposing team without getting tagged
- Mixed initiative system: two humans controlling up to 6-10 robots
- Limited BW comms + limited sensing

Problem Framework

Agent dynamics

$$\dot{x}^i = f^i(x^i, u^i) \quad x^i \in \mathbb{R}^n, u^i \in \mathbb{R}^m$$
$$y^i = h^i(x^i) \quad y^i \in \mathrm{SE}(3),$$

Vehicle "role"

- $\alpha \in \mathcal{A}$ encodes internal state + relationship to current task
- Transition $\alpha' = r(x, \alpha)$

Communications graph $\mathcal{G}$

- Encodes the system information flow
- Neighbor set $\mathcal{N}^i(x, \alpha)$

Task

- Encode as finite horizon optimal control

$$J = \int_0^T L(x, \alpha, u)\, dt + V(x(T), \alpha(T)),$$

- Assume task is *coupled*

Strategy

- Control action for individual agents

$$u^i = \gamma(x, \alpha) \qquad \{g^i_j(x, \alpha) : r^i_j(x, \alpha)\}$$

$$\alpha^{i\,\prime} = \begin{cases} r^i_j(x, \alpha) & g(x, \alpha) = \text{true} \\ \text{unchanged} & \text{otherwise.} \end{cases}$$

Decentralized strategy

$$u^i(x, \alpha) = u^i(x^i, \alpha^i, x^{-i}, \alpha^{-i})$$
$$x^{-i} = \{x^{j_1}, \ldots, x^{j_{m_i}}\}$$
$$j_k \in \mathcal{N}^i \qquad m_i = |\mathcal{N}^i|$$

- Similar structure for role update

Information Flow in Vehicle Formations

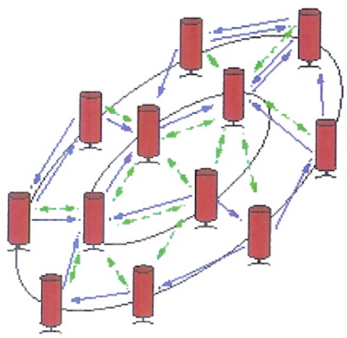

Example: satellite formation

- Blue links represent *sensed* information
- Green links represent *communicated* information

Sensed information

- Local sensors can see some subset of nearby vehicles
- Assume small time delays, pos'n/vel info only

Communicated information

- Point to point communications (routing OK)
- Assume limited bandwidth, some time delay
- Advantage: can send more complex information

Topological features

- Information flow (sensed or communicated) represents a directed graph
- Cycles in graph ⇒ information feedback loops

Question: How does topological structure of information flow affect stability of the overall formation?

Basic Definitions (1 of 2)

Definition 1. A *graph* is a pair $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ that consists of a set of vertices $\mathcal{V}$ and a set of edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$:

- Vertices: $v_i \in \mathcal{V}$
- Edges: $e_{ij} = (v_i, v_j) \in \mathcal{E}$

Example:

$\mathcal{V} = \{1, 2, 3, 4, 5, 6\}$

$\mathcal{E} = \{(1,6), (2,1), (2,3), (2,6), (6,2), (3,4), (3,6), (4,3), (4,5), (5,1), (6,1), (6,2), (6,4)\}$

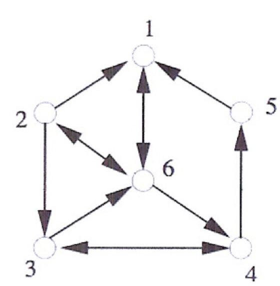

Notation:

- Order of a graph = number of nodes: $|\mathcal{V}|$
- v_i and v_j are *adjacent* if there exists $e = (v_i, v_j)$
- An adjacent node v_j for a node v_i is called a *neighbor* of v_i
- $\mathcal{N}_i$ = set of all neighbors of v_i
- $\mathcal{G}$ is *complete* if all nodes are adjacent

Basic Definitions (2 of 2)

Undirected graphs

- A graph is *undirected* if $e_{ij} \in \mathcal{E} \implies e_{ji} \in \mathcal{E}$
- Degree of a node: $\deg(v_i) := |\mathcal{N}_i|$
- A graph is *regular* (or *k-regular*) if all vertices of a graph have the same degree k

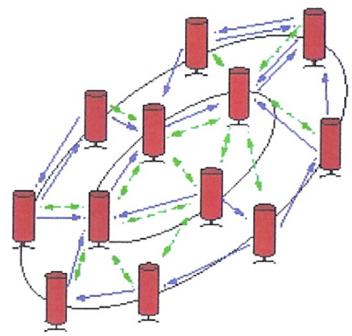

Directed graphs (digraph)

- Out-degree of v_i: $\deg_{\text{out}}$ = number of edges $e_{ij} = (v_i, v_j)$
- In-degree of v_i: $\deg_{\text{in}}$ = number of edges $e_{ki} = (v_k, v_i)$

Balanced graphs

- A graph is *balanced* if out-degree = in-degree at each node

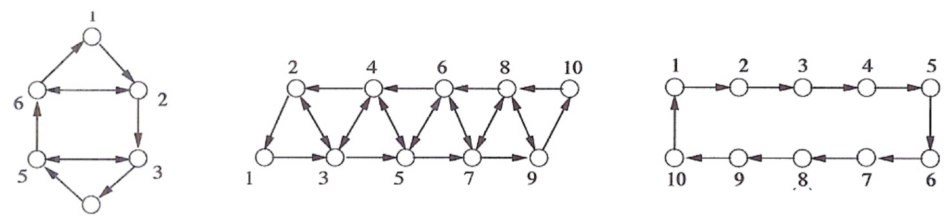

Connectedness of Graphs (1 of 2)

Paths

- A *path* is a subgraph $\pi = (\mathcal{V}, \mathcal{E}_\pi) \subset \mathcal{G}$ with distinct nodes $\mathcal{V} = \{v_1, v_2, \ldots, v_m\}$ and

$$\mathcal{E}_\pi := \{(v_1, v_2), (v_2, v_3), \ldots, (v_{m-1}, v_m)\}.$$

- The *length* of π is defined as $|\mathcal{E}_\pi| = m - 1$.
- A *cycle* (or m-cycle) $C = (\mathcal{V}, \mathcal{E}_C)$ is a path (of length m) with an extra edge $(v_m, v_1) \in \mathcal{E}$.

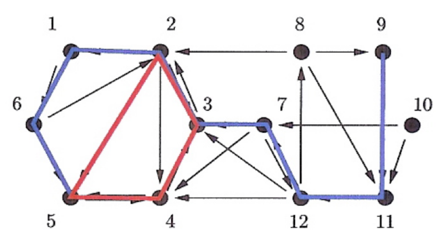

Connectivity of undirected graphs

- An undirected graph $\mathcal{G}$ is called *connected* if there exists a path π between any two distinct nodes of $\mathcal{G}$.
- For a connected graph $\mathcal{G}$, the length of the longest path is called the *diameter* of $\mathcal{G}$.
- A graph with no cycles is called *acyclic*
- A *tree* is a connected acyclic graph

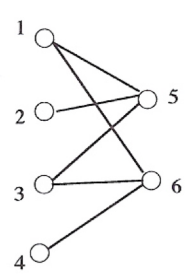

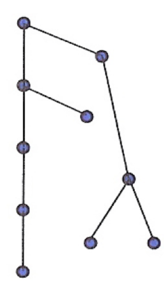

Connectedness of Graphs (2 of 2)

Connectivity of directed graphs

- A digraph is called *strongly connected* if there exists a directed path π between any two distinct nodes of $\mathcal{G}$.
- A digraph is called *weakly connected* if there exists an undirected path between any two distinct nodes of $\mathcal{G}$.

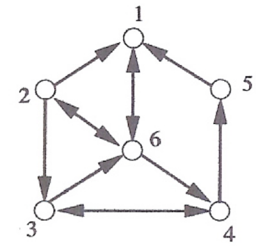

Strongly connected

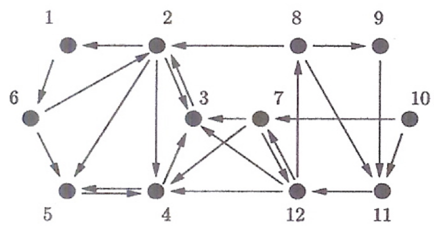

Weakly connected

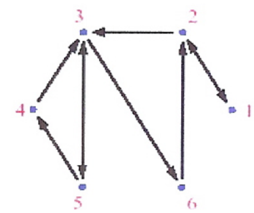

???

- Q: how do we check connectivity of a graph?

Matrices Associated with a Graph

Capturing properties of a graph using matrices

- The *adjacency matrix* $A = [a_{ij}] \in \mathbb{R}^{n\times n}$ of a graph $\mathcal{G}$ of order n is given by:

$$a_{ij} := \begin{cases} 1 & \text{if } (v_i, v_j) \in \mathcal{E} \\ 0 & \text{otherwise} \end{cases}$$

- The *degree matrix* of a graph as a diagonal $n \times n$ ($n = |\mathcal{V}|$) matrix $\Delta = \text{diag}\{\deg_{out}(v_i)\}$ with diagonal elements equal to the out-degree of each node and zero everywhere else.
- The *Laplacian matrix* L of a graph is defined as

$$L = \Delta - A$$

- The row sums of the Laplacian are all 0.

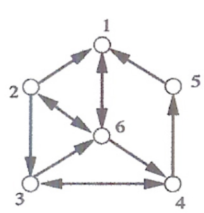

$$A = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 0 \end{bmatrix} \qquad L = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & -1 \\ -1 & 3 & -1 & 0 & 0 & -1 \\ 0 & 0 & 2 & -1 & 0 & -1 \\ 0 & 0 & -1 & 2 & -1 & 0 \\ -1 & 0 & 0 & 0 & 1 & 0 \\ -1 & -1 & 0 & -1 & 0 & 3 \end{bmatrix}$$

Periodic Graphs and Weighted Graphs

Periodic and acyclic graphs

- A graph with the property that the set of all cycle lengths has a common divisor $k > 1$ is called *k-periodic*.
- A graph without cycles is said to be *acyclic*.

Weighted graphs

- A *weighted graph* is graph $(\mathcal{V}, \mathcal{E})$ together with a map $\varphi : \mathcal{E} \to \mathbb{R}$ that assigns a real number $w_{ij} = \varphi(e_{ij})$ called a *weight* to an edge $e_{ij} = (v_i, v_j) \in \mathcal{E}$.
- The set of all weights associated with $\mathcal{E}$ is denoted by $\mathcal{W}$.
- A weighted graph can be represented as a triplet $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{W})$.

Weighted Laplacian

- In some applications it is natural to "normalize" the Laplacian by the outdegree
- $\tilde{L} := \Delta^{-1} L = I - P$, where $P = \Delta^{-1} A$ (weighted adjacency matrix).

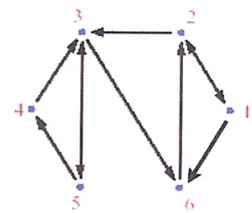

$$\begin{bmatrix} 1 & -\frac{1}{2} & 0 & 0 & 0 & -\frac{1}{2} \\ -\frac{1}{2} & 1 & -\frac{1}{2} & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & -\frac{1}{2} & -\frac{1}{2} \\ 0 & 0 & -1 & 1 & 0 & 0 \\ 0 & 0 & -\frac{1}{2} & -\frac{1}{2} & 1 & 0 \\ 0 & -1 & 0 & 0 & 0 & 1 \end{bmatrix}$$

Weighted Laplacian

Application: Consensus Protocols

Consider a collection of N agents that communicate along a set of undirected links described by a graph $\mathcal{G}$. Each agent has a state x_i with initial value $x_i(0)$ and together they wish to determine the average of the initial states $\text{Ave}(x(0)) = 1/N \sum x_i(0)$.

The agents implement the following *consensus protocol*:

$$\dot{x}_i = \sum_{j \in \mathcal{N}_i} (x_j - x_i) = -|\mathcal{N}_i|(x_i - \text{Ave}(x_{\mathcal{N}_i}))$$

which is equivalent to the dynamical system

$$\dot{x} = u \qquad u = -Lx.$$

Proposition 1. *If the graph is connected, the state of the agents converges to* $x_i^* = \text{Ave}(x(0))$ *exponentially fast.*

- Proposition 1 implies that the spectra of L controls the stability (and convergence) of the consensus protocol.
- To (partially) prove this theorem, we need to show that the eigenvalues of L are all positive.

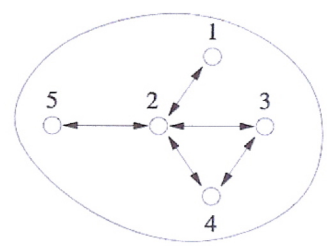

Sensor Network

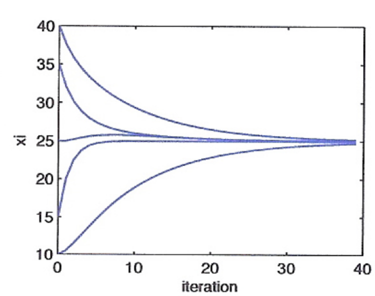

Gershgorin Disk Theorem

Theorem 2 (Gershgorin Disk Theorem). *Let $A = [a_{ij}] \in \mathbb{R}^{n \times n}$ and define the deleted absolute row sums of A as*

$$r_i := \sum_{j=1, j \neq i}^{n} |a_{ij}| \tag{1}$$

Then all the eigenvalues of A are located in the union of n disks

$$G(A) := \bigcup_{i=1}^{n} G_i(A), \text{ with } G_i(A) := \{z \in \mathbb{C} : |z - a_{ii}| \leq r_i\} \tag{2}$$

Furthermore, if a union of k of these n disks forms a connected region that is disjoint from all the remaining $n-k$ disks, then there are precisely k eigenvalues of A in this region.

Sketch of proof Let λ be an eigenvalue of A and let v be a corresponding eigenvector. Choose i such that $|v_i| = \max_j |v_j > 0$. Since v is an eigenvector,

$$\lambda v_i = \sum_i A_{ij} v_j \quad \implies \quad (\lambda - a_{ii}) v_i = \sum_{i \neq j} A_{ij} v_j$$

Now divide by $v_i \neq 0$ and take the absolute value to obtain

$$|\lambda - a_{ii}| = |\sum_{j \neq i} a_{ij} v_j| \leq \sum_{j \neq i} |a_{ij}| = r_i$$

Properties of the Laplacian (1)

Proposition 3. *Let L be the Laplacian matrix of a digraph $\mathcal{G}$ with maximum node out–degree of $d_{max} > 0$. Then all the eigenvalues of $A = -L$ are located in a disk*

$$B(\mathcal{G}) := \{s \in \mathbb{C} : |s + d_{max}| \leq d_{max}\} \tag{3}$$

that is located in the closed LHP of s-plane and is tangent to the imaginary axis at $s = 0$.

Proposition 4. *Let $\tilde{L}$ be the weighted Laplacian matrix of a digraph $\mathcal{G}$. Then all the eigenvalues of $A = -L$ are located inside a disk of radius 1 that is located in the closed LHP of s-plane and is tangent to the imaginary axis at $s = 0$.*

Theorem 5 (Olfati-Saber). *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ be a weighted digraph of order n with Laplacian L. If $\mathcal{G}$ is strongly connected, then rank$(L) = n - 1$.*

Remarks:

- Proof for the directed case is standard
- Proof for undirected case is available in Olfati-Saber & M, 2004 (IEEE TAC)
- For directed graphs, need $\mathcal{G}$ to be strongly connected and converse is not true.

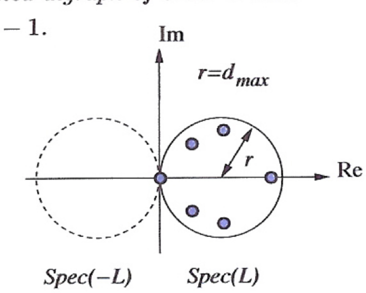

Proof of the Consensus Protocol

$$\dot{x} = -Lx \qquad L = \Delta - A$$

Note first that the subspaced spanned by $\mathbf{1} = (1, 1, \ldots, 1)^T$ is an invariant subspace since $L \cdot \mathbf{1} = 0$ Assume that there are no other eigenvectors with eigenvalue 0. Hence it suffices to look at the convergence on the complementary subspace $\mathbf{1}^{\perp}$.

Let δ be the disagreement vector

$$\delta = x - \text{Ave}(x(0))\,\mathbf{1}$$

and take the square of the norm of δ as a Lyapunov function candidate, i.e. define

$$V(\delta) = \|\delta\|^2 = \delta^T \delta \tag{4}$$

Differentiating $V(\delta)$ along the solution of $\dot{\delta} = -L\delta$, we obtain

$$\dot{V}(\delta) = -2\delta^T L \delta < 0, \quad \forall \delta \neq 0, \tag{5}$$

where we have used the fact that $\mathcal{G}$ is connected and hence has only 1 zero eigenvalue (along $\mathbf{1}$). Thus, $\delta = 0$ is globally asymptotically stable and $\delta \to 0$ as $t \to +\infty$, i.e. $x^* = \lim_{t \to +\infty} x(t) = \alpha_0 \mathbf{1}$ because $\alpha(t) = \alpha_0 = \text{Ave}(x(0)), \forall t > 0$. In other words, the average–consensus is globally asymptotically achieved. □

Perron-Frobenius Theory

Spectral radius:

- $\text{spec}(L) = \{\lambda_1, \ldots, \lambda_n\}$ is called the *spectrum* of L.
- $\rho(L) = |\lambda_n| = \max_k |\lambda_k|$ is called the *spectral radius* of L

Theorem 6 (Perron's Theorem, 1907). *If $A \in \mathbb{R}^{n \times n}$ is a positive matrix ($A > 0$), then*

1. *$\rho(A) > 0$;*
2. *$r = \rho(A)$ is an eigenvalue of A;*
3. *There exists a positive vector $x > 0$ such that $Ax = \rho(A)x$;*
4. *$|\lambda| < \rho(A)$ for every eigenvalue $\lambda \neq \rho(A)$ of A, i.e. $\rho(A)$ is the unique eigenvalue of maximum modulus; and*
5. *$[\rho(A)^{-1}A]^m \to R$ as $m \to +\infty$ where $R = xy^T$, $Ax = \rho(A)x$, $A^T y = \rho(A)y$, $x > 0$, $y > 0$, and $x^T y = 1$.*

Theorem 7 (Perron's Theorem for Non–Negative Matrices). *If $A \in \mathbb{R}^{n \times n}$ is a non–negative matrix ($A \geq 0$), then $\rho(A)$ is an eigenvalue of A and there is a non–negative vector $x \geq 0$, $x \neq 0$, such that $Ax = \rho(A)x$.*

Irreducible Graphs and Matrices

Irreducibility

- A directed graph is irreducible if, given any two vertices, there exists a path from the first vertex to the second. (Irreducible = strongly connected)
- A matrix is irreducible if it is not similar to a block upper triangular matrix via a permutation.
- A digraph is irreducible if and only if its adjacency matrix is irreducible.

Theorem 8 (Frobenius). *Let $A \in \mathbb{R}^{n\times n}$ and suppose that A is irreducible and non-negative. Then*

1. *$\rho(A) > 0$;*
2. *$r = \rho(A)$ is an eigenvalue of A;*
3. *There is a positive vector $x > 0$ such that $Ax = \rho(A)x$;*
4. *$r = \rho(A)$ is an algebraically simple eigenvalue of A; and*
5. *If A has h eigenvalues of modulus r, then these eigenvalues are all distinct roots of $\lambda^h - r^h = 0$.*

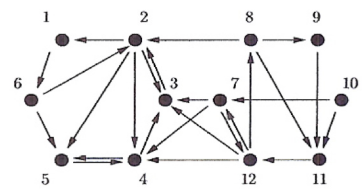

Figure 4.1: Sample Graph $\mathcal{G}$.

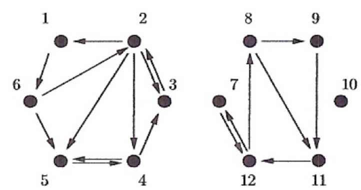

Figure 4.2: Induced Subgraphs of Components of $\mathcal{G}$.

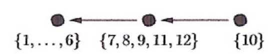

Figure 4.3: Graph of Components of $\mathcal{G}$.

Properties of Laplacians (2)

Properties of L

- If $\mathcal{G}$ is strongly connected, the zero eigenvalue of L is simple.
- If $\mathcal{G}$ is aperiodic, all nonzero eigenvalues lie in the interior of the Gershgorin disk.
- If $\mathcal{G}$ is k-periodic, L has k evenly spaced eigenvalues on the boundary of the Gershgorin disk.

Unidirectional tree

$\lambda = 0,1$

Undirected graph

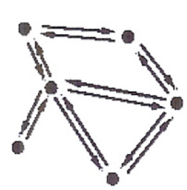

$\lambda \in [0,2]$

Cycle

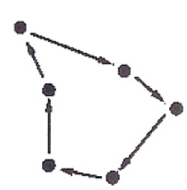

$\lambda_i = 1 - e^{2\pi(i-1)j/N}$

Periodic graph

$\lambda_1 = 0, \lambda_N = 2$

Algebraic Connectivity

Theorem 9 (Variant of Courant-Fischer). *Let $A \in \mathbb{R}^{n \times n}$ be a Hermitian matrix with eigenvalues $\lambda_1 \le \lambda_2 \le \cdots \le \lambda_n$ and let w_1 be the eigenvector of A associated with the eigenvalue λ_1. Then*

$$\lambda_2 = \min_{\substack{x \neq 0, x \in \mathbb{C}^n, \\ x \perp w_1}} \frac{x^* A x}{x^* x} = \min_{\substack{x^* x = 1, \\ x \perp w_1}} x^* A x \qquad (6)$$

Remarks:

- λ_2 is called the *algebraic connectivity* of L
- For an undirected graph with Laplacian L, the rate of convergence for the consensus protocol is bounded by the second smallest eigenvalue λ_2

Example

- Directed graph with a bottleneck

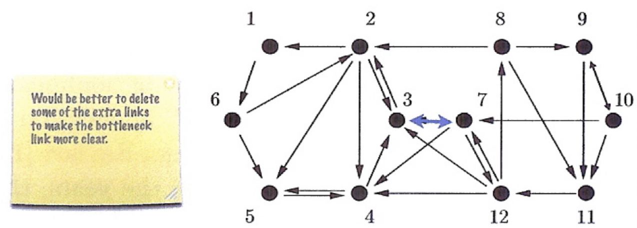

Cyclically Separable Graphs

Definition (Cyclic separability). A digraph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is *cyclically separable* if and only if there exists a partition of the set of edges $\mathcal{E} = \cup_{k=1}^{n_c} \mathcal{E}_k$ such that each partition $\mathcal{E}_k$ corresponds to either the edges of a cycle of the graph, or a pair of directed edges ij and ji that constitute an undirected edge. A graph that is not cyclically separable is called *cyclically inseparable.*

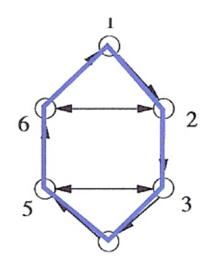

Lemma 10. *Let L be the Laplacian matrix of a cyclically separable digraph $\mathcal{G}$ and set $u = -Lx, x \in \mathbb{R}^n$. Then $\sum_{i=1}^n u_i = 0, \forall x \in \mathbb{R}^n$ and $\mathbf{1} = (1, \ldots, 1)^T$ is the left eigenvector of L.*

Proof. The proof follows from the fact that by definition of cyclic separability. We have

$$-\sum_{i=1}^n u_i = \sum_{ij \in \mathcal{E}} (x_j - x_i) = \sum_{k=1}^{n_c} \sum_{ij \in \mathcal{E}_k} (x_j - x_i) = 0$$

because the inner sum is zero over the edges of cycles and undirected edges of the graph. □

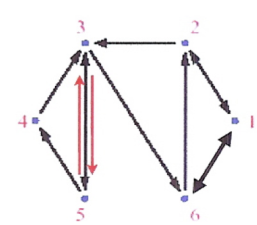

- Provides a "conservation" principle for average consensus

Balanced Graphs

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a digraph. We say $\mathcal{G}$ is *balanced* if and only if the in–degree and out–degree of all nodes of $\mathcal{G}$ are equal, i.e.

$$\deg_{\text{out}}(v_i) = \deg_{\text{in}}(v_i), \quad \forall v_i \in \mathcal{V} \tag{7}$$

Theorem 11. *A digraph is cyclically separable if and only if it is balanced.*

Corollary 11.1. *Consider a network of integrators with a directed information flow $\mathcal{G}$ and nodes that apply the consensus protocol. Then, $\alpha = \text{Ave}(x)$ is an invariant quantity if and only if $\mathcal{G}$ is balanced.*

Remarks

- Balanced graphs generalized undirected graphs and retain many key properties

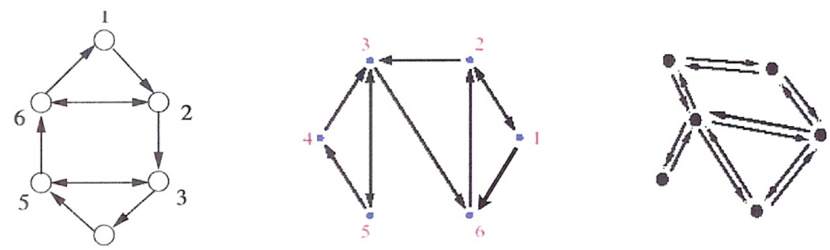

Consensus Protocols for Balanced Graphs

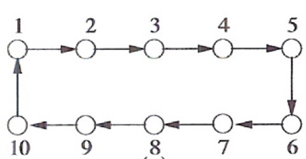

(a)

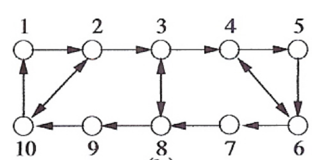

(b)

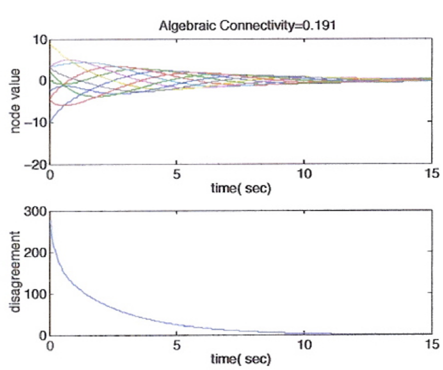

(a)

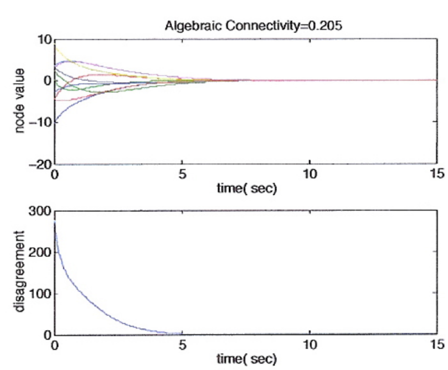

(b)

Formation Operations: Graph Switching

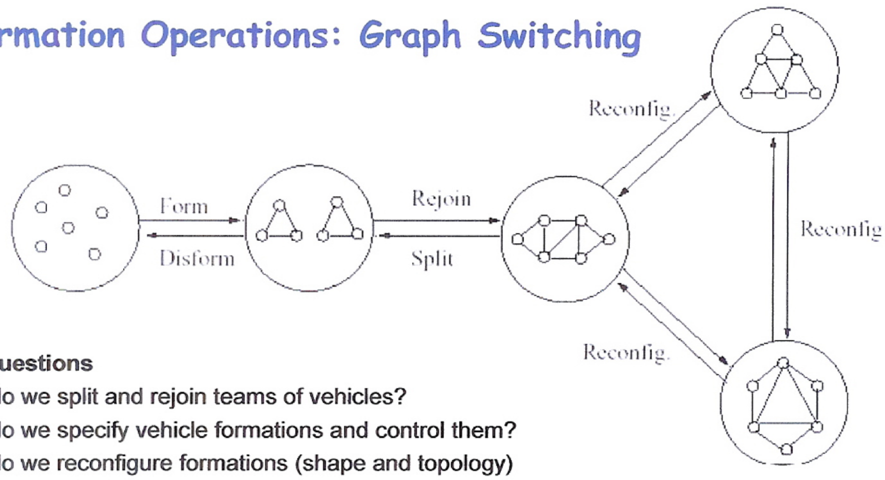

Control questions

- How do we split and rejoin teams of vehicles?
- How do we specify vehicle formations and control them?
- How do we reconfigure formations (shape and topology)

Consensus-based approach using balanced graphs

- If each subgraph is balanced, disagreement vector provides common Lyapunov fcn
- By separately keeping track of the flow in and out of nodes, can preserve center of mass of of subgraphs after a split manuever

Other Uses of Consensus Protocols

Computation of other functions besides the average

- Can adopt the basic approach to compute max, min, etc
- Chandy/Charpentier: can compute *superidempotent* functions:

$$f(X \cup Y) = f(f(X) \cup Y)$$

Basic idea: local conservation implies global conservation

- Can extend these cases to handle splitting and rejoining as well

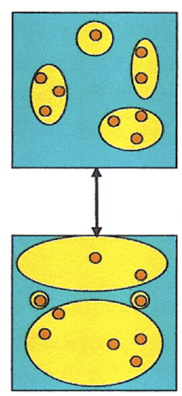

Distributed Kalman filtering

- Maintain local estimates of global average and covariance
- Need to be careful about choosing rates of convergence

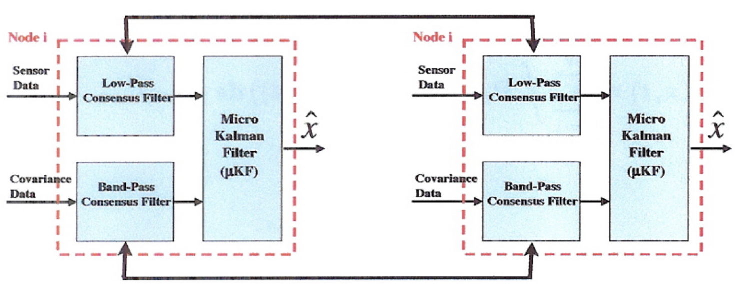

Summary

Graphs

- Directed, undirected, connected, strongly complete
- Cyclic versus acyclic; irreducible, balanced

Graph Laplacian

- L = Δ - A
- Spectral properties related to connectivity of graph
- # zero eigenvalues = # strongly connected components
- Second largest nonzero eigenvalue ~ weak links

Spectral Properties of Graphs

- Gershgorin's disk theorem
- Perron-Frobenius theory

Examples

- Consensus problems, distributed computing
- Cooperative control (coming up next...)

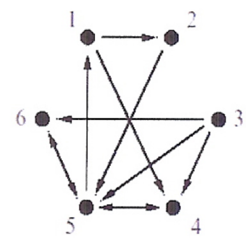

$$L = \begin{bmatrix} 1 & -\frac{1}{2} & 0 & -\frac{1}{2} & 0 & 0 \\ 0 & 1 & 0 & 0 & -1 & 0 \\ 0 & 0 & 1 & -\frac{1}{3} & -\frac{1}{3} & -\frac{1}{3} \\ 0 & 0 & 0 & 1 & -1 & 0 \\ -\frac{1}{3} & 0 & 0 & -\frac{1}{3} & 1 & -\frac{1}{3} \\ 0 & -\frac{1}{2} & 0 & 0 & -\frac{1}{2} & 1 \end{bmatrix}$$