

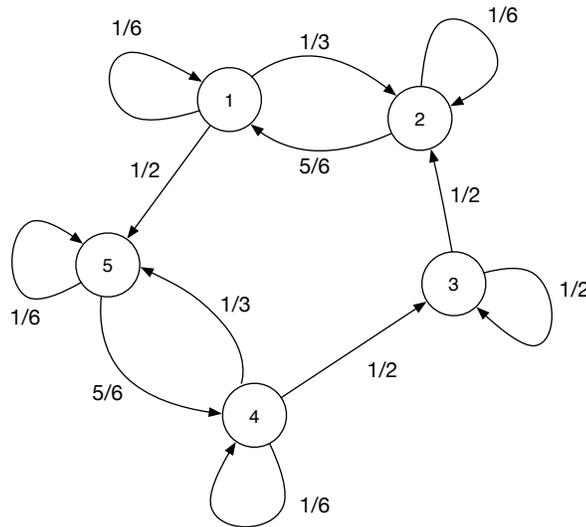
Theory of Networked Systems

Homework 3 (100 points)

Due Thursday, October 9, 2014 at 12:00 pm, in class

1. 15 points

Let A be the weighted adjacency matrix corresponding to the weighted digraph below.



- What is the per-step convergence factor of A ?
- What is the asymptotic convergence factor of A ?

2. 50 points

Consider an undirected, connected network of n nodes. Let Δ be the diameter of the network graph; if we consider the shortest paths between every pair of nodes, the diameter of a graph is the length of the longest shortest path.

- Each node i has a value $y_i(0)$. The goal is for all nodes to learn the maximal value over all n values. Design a discrete-time (possibly non-linear) algorithm by which the nodes can learn this maximal value. The nodes can only exchange information with their neighbors in the graph. (Note that once you have a max-finding algorithm, you can also implement a min-finding algorithm that has the same complexity.)
- How many time steps are needed by your algorithm for every node to learn the maximal value?
- Suppose that every node i has an initial value $x_i(0)$, and the goal is for all nodes to learn the average of these values $\frac{1}{n} \sum_{i=1}^n x_i(0)$. The nodes run a discrete-time averaging algorithm that is guaranteed to converge to the average value (for example, the algorithm defined by A_β). Design a “local convergence test” that a node can use to determine when its value $x(\ell)$ is within ϵ of the average, where $\epsilon > 0$ is a parameter of the test algorithm. Assume that the nodes know Δ and ϵ . You may want to use the max-finding and min-finding algorithms in your solution.
- Explain why your local convergence test is correct.

3. 35 points

In the pairwise-symmetric-exchange averaging algorithm, a single edge (i, j) is chosen in each time step. Nodes i and j average their values with one another,

$$x_i(\ell + 1) = x_j(\ell + 1) = \frac{1}{2} (x_i(\ell) + x_j(\ell)).$$

Every other node k keeps the same value,

$$x_k(\ell + 1) = x_k(\ell).$$

Let the network be a ring graph (an undirected cycle) of 100 nodes, where the nodes are numbered $1, 2, \dots, 100$. Consider the following two edge selection policies:

- (i) *Round-robin*: The edges are selected in the following sequence, $(1, 2), (2, 3), (3, 4), \dots, (99, 100), (100, 1)$, and the sequence is repeated.
- (ii) *Uniform symmetric gossip*: In each iteration, one edge (i, j) is selected uniformly at random.

Initialize each $x(0)$ to a random vector where each component is drawn from a uniform distribution over $[0, 10]$.

- (a) Write code to simulate the averaging algorithm for each edge selection policy and turn in your code.
- (b) Run one simulation for each edge selection policy and plot the norm of the deviation vector $\|\tilde{x}(\ell)\|_2$ as a function of the time step ℓ . You should generate a single graph that shows both policies.
- (c) From your graph, what can you say about the convergence behavior of the two edge selection policies? Does one converge more quickly than the other? To answer this question, you may need to run the simulations for thousands of time steps and plot the norms on log scale.