Problem Set 2 (due Tuesday, March 1)

Remarks:

- Since the background of the students in the class varies, some may find these problems easy while others may find some problems challenging. The main purpose of these problem sets is for you to learn the course material better by applying ideas learnt in class and/or exploring related problems.
- Please work on these problems on your own, or in collaboration with fellow students in class. Please do not seek out solutions or solution approaches from the web. If you need hints, ask the course instructor.
- Please typeset your solution. Latex, plain text, pdf, and Word are all acceptable formats.
- The total points for the problem set is 75, worth 7.5% of the grade.

Problem 1. (5 + 5 = 10 points) Condition for decoding in gossip protocol

In our study of gossiping using network coding, we considered the following scenario. We have a network of n nodes, each having an L-bit message that they want to share with all the other nodes. Each message $\vec{m}_i$ is an element of $\mathcal{F}_2^L$. We studied a network coding algorithm in which each node broadcast a packet consisting of two vectors, an n-bit vector $\vec{\mu}$ from $\mathcal{F}_2^n$, and an L-bit message vector $\vec{m}$ equaling the linear combination $\sum_i \mu(i)\vec{m}_i$, where $\mu(i)$ represents the ith component of the $\vec{\mu}$ and the sum is the component-wise addition mod 2. The particular message vector that a node transmits is drawn uniformly at random from the subspace spanned by the messages that it has received thus far.

We further showed the following statement: if S_v denotes the coefficient vectors received by v, and there is no vector in $\mathcal{F}_2^n$ that is orthogonal to the subspace spanned by S_v, then v can decode all of the n original messages.

Let B_n denote the set of n vectors in F_2^n:

$$\{(1,0,0,\ldots,0),(0,1,0,\ldots,0),\ldots,(0,0,\ldots,1)\}.$$

(a) Prove or disprove: B_n is a basis for F_2^n in the sense that every vector in F_2^n can be written as a linear combination of vectors in B_n.

(b) Show why the following condition is not sufficient for decoding: none of the n vectors in B_n is orthogonal to the subspace spanned by S_v.

Problem 2. (10 + 5 + 15 = 30 points) SIS model on the star network

Let S_n denote the star graph with n vertices: a root node 0 with edges to the remaining $n-1$ nodes labeled 1 through $n-1$.

(a) Show that the largest eigenvalue of the adjacency matrix of S_n is $\sqrt{n-1}$.

(b) Show that the edge expansion of S_n is 1.

From the results that we have shown in class it follows that for S_n (i) if the infection rate $\beta < 1/\sqrt{n}$, then the epidemic die-out occurs in expected $O(\log n)$ time, assuming a recovery rate of 1 (per unit time); and (ii) if the infection rate $\beta > 1$, then the epidemic die-out occurs in expected exponential time.

We want to understand what happens when $1/\sqrt{n} \leq \beta \leq 1$.

(c) In this exercise you need to derive an upper bound on the expected epidemic die-out time when $\beta = C/\sqrt{n}$ for a constant $C > 1$.

We will consider a simplified discrete variant of the process. Assume that every node is infected at time 0. Let $S_j(t)$ be 1 if j is infected at t, and 0 otherwise. Let $H_j(t)$ denote a random variable that takes 1 with probability 1/2, and 0 otherwise (this represents healing of node j, if infected). Let $R_j(t)$ denote a random variable that is 1 with probability β, and 0 otherwise (this represents receipt of an infection from 0). Finally, let $T_j(t)$ denote a random variable that is 1 with probability β, and 0 otherwise (this represents transmitting an infection from 0). Define for $0 < j < n$ and $t > 0$:

$$S_j(t+1) = \max\{0, S_j(t) - H_j(t), S_0(t)R_j(t)\}.$$

Finally, define for $t > 0$:

$$S_0(t+1) = \max\{0, S_0(t) - H_0(t), \max_{j>0}(S_j(t)T_j(t))\}.$$

Show that in $\tau = O(\log n)$ steps, $S_j(\tau)$ is 0 for all j with probability $1 - 1/n^{\alpha}$ for some constant $\alpha > 0$.

Problem 3. (15 + 5 + 15 = 35 points) Myopic routing on a network with randomly chosen long-range links

Consider n nodes labeled 0 through $n-1$ organized in a ring network; so every node has exactly two neighbors. Now suppose each node has two additional *long-range links*, each of which is to a node in the ring drawn uniformly at random, with replacement. Let G denote the resulting random graph. The total number of edges is exactly $3n$. (Note that the graph is undirected; even though the long-range links are selected by each node in a directional manner, the edges formed are undirected.)

(a) Show that the diameter of G is $O(\log n)$ with high probability, i.e., with probability at least $1 - 1/n^c$ for some constant $c > 0$, where the hidden constant in the big-Oh notation bound for the diameter may depend on c.

Consider the myopic search algorithm for routing a message from a source s to destination t. Every intermediate node u forwards the message to its neighbor v (on the ring or via a long-range edge) such that v is the node nearest to t along the ring among all neighbors of v, breaking ties arbitrarily; formally, $(v - t) \bmod n$ is minimum among all neighbors v of u.

(b) Show that the myopic search algorithm always terminates.

(c) Show that there exist s and t such that the expected time for myopic search to complete is $\Omega(\sqrt{n})$.