Distributed Algorithms and Lower Bounds – Exercise 4

Question 1: Luby's MIS Algorithm

In this question we will prove that Luby's algorithm requires $O(\log n)$ rounds with high probability (a correct proof this time!).

Consider step i of Luby's algorithm, and let E be the set of remaining edges that were not removed in previous steps.

We'll use the following notation:

- $N(u) \subseteq V$: the set of neighbors of node u in E
- $d(u) = |N(u)|$ is the remaining degree of u
- $x(u) \in [0,1]$: the random priority chosen by node u in step i

Let $Z_{u \to v}$ be an indicator for the event that u got the highest priority out of all the nodes in u's neighborhood *and also* in v's neighborhood: that is, $Z_{u \to v} = 1$ iff for all $w \in N(u) \cup N(v)$ we have $x(u) > x(w)$.

(1) What is $\Pr[Z_{u \to v} = 1]$?

(2) Can there be two neighbors $u, w \in N(v)$ (where $u \neq w$) such that $Z_{u \to v} = 1$ and also $Z_{w \to v} = 1$?

(3) When $Z_{u \to v} = 1$, we say that node u killed all the directed edges (v, w) for $w \in N(v)$. Note that the real edges of the graph are undirected; a real edge $\{v, w\}$ gets removed if either some node u killed (v, w) or some node u killed (w, v).
Find a lower bound for the number of directed edges killed in step i, in terms of the indicators $Z_{u \to v}$ for the edges $\{u, v\} \in E$ and the remaining degrees $d(u)$ of the nodes in the graph.

(4) Prove that the expected number of edges removed from the graph in step i of Luby's algorithm is at least $|E|/2$. (Careful not to over-count!)

(5) Bonus: prove that Luby's algorithm terminates in $O(\log n)$ rounds with probability at least $2/3$. (Hint: you need to show concentration around the expectation.)

Question 2

Let $\epsilon \in (0,1)$. Consider the following partial function for two players:

$$GapDisj(X, Y) = \begin{cases} 0, if\ |X \cap Y| \geq \epsilon \cdot n, \\ 1, if\ |X \cap Y| = 0 \end{cases}$$

For inputs X, Y such that $0 < |X \cap Y| < \epsilon \cdot n$, the output is undefined, and the protocol is allowed to output whatever it wants.

What is the randomized communication complexity of $GapDisj(X, Y)$ in terms of the input size n and the gap ϵ? Prove an upper and lower bound.

Bonus question: for those who like extremal combinatorics

What is the deterministic communication complexity of $GapDisj(X, Y)$ in terms of n and ϵ, assuming that $\epsilon < \frac{1}{3}$?