CS:5620 Homework 2, Fall 2016

Due in class on Thu, 9/22

Collaboration: You are welcome to form groups of size 2 and work on your homeworks in groups. Of course, you are not *required* to work in groups. Every group should make one submission and names of both group members should appear on the submission and both students in a group will receive the same score. Other than the TA and the professor, you can only discuss homework problems with your group partner. Collaboration can be positive because talking to someone else about these problems can help to clarify your ideas and you will also (hopefully) get to hear about different ways of thinking about the problem. On the other hand, collaboration can be negative if one member of the group rides on work being done by the other member – please avoid this situation. If your solutions are (even partly) based on material other than what has been posted on the course website, you should clearly acknowledge your outside sources.
Late submissions: No late submissions are permitted. You will receive no points for your submission if your submission is not turned in at the beginning of class on the due date.
Evaluation: Your submissions will be evaluated on correctness *and* clarity. Correctness is of course crucial, but how clearly you communicate your ideas is also quite important.
Who submits what: Undergraduate students are required to solve Problems 1-4, Master's students are required to solve Problems 2-5, and PhD students are required to solve Problems 3-6.

1. The input is a 6-node oriented tree. The root r has two children u and v. Node u has one child w and node v has two children x and y. The ID's of these nodes (in decimal) are as follows: $\texttt{ID}_r = 104$, $\texttt{ID}_u = 110$, $\texttt{ID}_v = 51$, $\texttt{ID}_w = 170$, $\texttt{ID}_x = 35$, and $\texttt{ID}_y = 15$. Show the execution of the Cole-Vishkin 6-coloring algorithm on oriented trees on this input. Make sure to use the correct number of bits in the representation of colors in each round of the Cole-Vishkin algorithm and also make sure to terminate the algorithm after the appropriate number of rounds. Your answer should be a sequence of appropriately labeled illustrations of the given oriented tree.

2. Just before some round i in Luby's algorithm for distributed $(\Delta + 1)$-coloring, the state of the neighborhood of some node u is as follows:

   (i) u has 5 uncolored neighbors v_1, v_2, v_3, v_4, v_5.

   (ii) $P(u) = \{1, 2, 3, 4, 5, 6\}$.

   (iii) $P(v_1) = \{1, 4\}$, $P(v_2) = \{1, 7\}$, $P(v_3) = \{1, 2\}$, $P(v_4) = \{1, 3\}$, and $P(v_5) = \{2, 7\}$.

   (a) What is $Pr(C_u)$? Show all your work.

   (b) Now consider the proof of Proposition 6 in Johansson's paper (posted on the course website). What color from $P(u)$ could play the role of c' in the proof? What color from $P(u)$ could play the role of c'' in the proof? If multiple colors could be c' or c'', identify all of these.

   (c) Continuing with the proof of Proposition 6, pick an appropriate color $c' \in P(u)$, neighbor v of u, and color $c'' \in P(u)$ and replace c' in $P(v)$ by c''. Recalculate $Pr(C_u)$ and compare it with your answer for (a).

3. Describe an algorithm in the CONGEST model, running in $O(\log^* n)$ rounds, to compute the MIS of graph with maximum degree $\Delta \leq 10$. A plain English description of the algorithm will suffice.

4. A *distance-2 coloring* of a graph $G = (V, E)$ is a vertex coloring $c : V \to \mathcal{C}$ such that no two vertices at distance at most two from each other have the same color. For any node v, let $N_2(v)$ (called the *2-neighborhood* of v) denote the set of vertices at distance at most

two from v. Note that $N_2(v)$ contains v and all its neighbors. Let Δ_2 denote the size of the largest 2-neighborhood in the graph.

Describe a randomized algorithm, running in the LOCAL model in $O(\log n)$ rounds (in expectation), that produces a distance-2 coloring of G using Δ_2 colors. Write your algorithm using pseudocode executed by a node v. Make sure that your pseudocode is clear and well-commented.
Hint: Your approach should be to mimic Luby's $(\Delta + 1)$-coloring algorithm.

5. Consider the following (non-distributed) algorithm on an input graph $G = (V, E)$.

while $(V \neq \emptyset)$ **do**
  $S \leftarrow$ all nodes of degree at most 2
  for each edge $e \in E$ **do**
    if e has exactly one endpoint in S **then**
      orient e from its endpoint in S to its endpoint in $V \setminus S$
    if e has both endpoints in S **then**
      orient e arbitrarily
  $V \leftarrow V \setminus S$
  $E \leftarrow$ edges in $G[V]$

(a) Prove that if G is a tree, then this algorithm can be implemented in the CONGEST model to run in $O(\log n)$ rounds, yielding an orientation of edges such that every node has at most 2 out-neighbors.

(b) Show that there is a *deterministic* algorithm in the CONGEST model to compute a 3-coloring of an unoriented tree in $O(\log n)$ rounds.

6. This problem describes a randomized variant of the greedy, distributed, $(\Delta + 1)$-coloring algorithm discussed in class. Recall that each node v has an ID, denoted $\texttt{ID}_v$, represented by $\lceil c \log_2 n \rceil$ bits for some constant $c \geq 1$. Instead of using these ID's to break symmetry, each node v generates a random bit string $r(v)$ of length $c' \cdot \lceil c \log_2 n \rceil$, for some constant c'. Each bit in $r(v)$ is generated independently and uniformly at random from $\{0, 1\}$. Now nodes run the greedy, distributed $(\Delta + 1)$-coloring algorithm, but using r-values for breaking symmetry, rather than using IDs (which are supplied by an adversary). One caveat is that since the r-values are picked at random, there is some chance (though small) that two neighbors have identical r-values. If this happens, nodes can detect it and the algorithm simply aborts.

Consider the execution of this algorithm on a cycle C with n nodes. Prove that with probability at least $1 - 1/n$, the above algorithm terminates with a 3-coloring of C in $O(\log n)$ rounds.