CS:5620 Homework 1, Fall 2016

Due in class on Thu, 9/8

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.

1. This problem is on the *Distributed Greedy Graph Coloring* algorithm discussed in class on Thursday, Aug 25th.

   (a) What is the size (in bits) of each message sent during this algorithm? Express your answer in asymptotic notation as a function of n, the number of nodes in the network. (Show your work to receive credit for this problem.)

   (b) Suppose that the `ID`'s of nodes are not all distinct, but they satisfy the constraint:

   $$\text{for all nodes } v\text{: } \mathtt{ID}_v \neq \mathtt{ID}_w \text{ for all } w \in N(v).$$

   In other words, the `ID` of a node is distinct from the `ID`'s of its neighbors, but is not guaranteed to be distinct from `ID`'s of non-neighbors. Will the Distributed Greedy Graph Coloring algorithm still work correctly? You can answer this question in one of two ways: (i) present an argument showing the algorithm is still correct or (ii) present a counter-example showing that the algorithm no longer works, as described.

2. An *independent set* in a graph $G = (V, E)$ is a subset $V' \subseteq V$ of vertices such that no two vertices in V' are adjacent. A *maximal independent set (MIS)* in a graph $G = (V, E)$ is a subset $V' \subseteq V$ such that (i) V' is an independent set and (ii) for any $v \in V \setminus V'$: $V' \cup \{v\}$ is *not* an independent set. In other words, an MIS of G is an independent set to which no more nodes can be added without violating the independence property.

   Design and present a distributed greedy algorithm in the CONGEST model, with round complexity $\Theta(n)$, to compute an MIS of a given graph. There is a distributed algorithm for the MIS problem that is quite similar to the Distributed Greedy Graph Coloring algorithm and this is what I have in mind for this problem. Your answer should be in the form of pseudocode (for each node v for each round i), similar to what I presented in class for the Distributed Greedy Graph Coloring problem. Note that the MIS can be represented by a boolean variable $I(v)$ at each node v such that $I(v)$ is True iff v is in the MIS.

3. The input is a graph $G = (V, E)$ with diameter at most 10. Recall that the diameter of a graph is the maximum (shortest path) distance between any pair of vertices in the graph. Show that there is distributed algorithm that runs in at most 25 rounds in the LOCAL model and computes a proper vertex coloring of G using a palette with $\chi(G)$ colors. Recall that $\chi(G)$ denotes the *chromatic number* of graph G and it is the fewest number of colors needed for a proper vertex coloring of G[1]. You may describe your algorithm using pseudocode (as in Problem 3) or in plain English. The main thing is to aim for precision and clarity.

[1] Computing a proper vertex coloring of a graph G using $\chi(G)$ colors is NP-complete and in general this problem is notoriously difficult in the sequential setting.

4. The input is a graph $G = (V, E)$ whose vertices have already been assigned a proper vertex coloring $c : V \to \{1, 2, 3, 4\}$. In other words, each node knows its color, denoted by the local variable $c(v)$. Furthermore, $c(v) \in \{1, 2, 3, 4\}$ for all $v \in V$ and $c(u) \neq c(v)$ for all $\{u, v\} \in E$. Design and present a distributed algorithm in the CONGEST model running in 4 rounds for computing an MIS of G. Use pseudocode to describe your algorithm.
