



Principles of Distributed Computing

Exercise 7: Sample Solution

1 Deterministic Maximal Independent Set

- a) Consider the graph consisting of a connected chain of k nodes v_1, \dots, v_k . We add $i - 1$ additional edges leading to $i - 1$ additional nodes at each node v_i for all $i \in \{1, \dots, k - 1\}$ and k additional nodes and edges to node v_k . The degree $\delta(v_1)$ of v_1 is 1 and for all other nodes $v_i \in \{2, \dots, k\}$ we have that $\delta(v_i) = i + 1$. All additional nodes have degree 1.

In the first round, all nodes except v_k have a neighbor with a larger degree, thus only v_k joins the MIS. Afterwards, v_{k-1} can decide, then v_{k-2} and so on. Thus, after k time all nodes v_1, \dots, v_k and also the additional nodes have decided to join or not to join the MIS.

The number of nodes in this graph is

$$n = k + \sum_{i=1}^k (i - 1) + 1 = 1 + \sum_{i=1}^k i = 1 + \frac{k(k + 1)}{2} < \frac{(k + 1)^2}{2}.$$

The time complexity is thus $k \geq \sqrt{2n} - 1 \in \Omega(\sqrt{n})$.

- b) We first show, that the above lower bound is tight. Assume a node v_0 of degree $\delta(v_0)$ that is still undecided at time \sqrt{n} . This implies that there was a neighbor v_1 at time $\sqrt{n} - 1$ that had a higher degree than v_0 , that is $\delta(v_0) < \delta(v_1)$. However, v_1 might have been removed from the set of undecided nodes due to having a neighbor that joined the MIS. We conclude that v_1 had an undecided neighbor of higher degree at time $\sqrt{n} - 1$. Using the same argument v_2 had an undecided neighbor at time $\sqrt{n} - 2$ of higher degree, which in turn had an undecided higher-degree neighbor v_3 at time $\sqrt{n} - 3$. By induction it follows that there are nodes $v_0, \dots, v_{2\sqrt{n}}$ such that nodes v_i and v_{i-1} are neighbors and $\delta(v_{i-1}) < \delta(v_i)$. Since we consider a tree, it is not possible that any two nodes among $v_0, \dots, v_{2\sqrt{n}}$ have the same neighbors (else there was a cycle and the graph not a tree). Thus there must be at least $\sum_{i=0}^{2\sqrt{n}} \delta(v_i) - 1$ nodes in the tree. We are interested in minimizing this sum and it is minimal for $\delta(v_0) = 1$ and $\delta(v_{i-1}) = \delta(v_i) + 1$. This yields that there are

$$\sum_{i=0}^{2\sqrt{n}} \delta(v_i) - 1 = \sum_{i=0}^{2\sqrt{n}} i = \frac{2\sqrt{n}(2\sqrt{n} + 1)}{2} > 2n.$$

Since our graph should have only n nodes, this is a contradiction to the assumption that there is a graph such that the algorithm does not remove all nodes within time \sqrt{n} .

To construct a lower bound for general graphs we now consider a ring of k nodes v_1, \dots, v_k instead of a chain. We use $k - 1$ additional nodes u_1, \dots, u_{k-1} to increase the degrees of the nodes v_i : There is an edge $\{v_i, u_j\}$ from all nodes v_i to all nodes u_j for which $j \in \{1, \dots, k - i\}$. It is easy to see that the degree $\delta(v_i)$ of node v_i is $k + 2 - i$, and that $\delta(u_j) = k - j$.

In the first round, only v_1 joins the MIS. This means that all nodes u_1, \dots, u_{k-1} and also v_2 and v_k can no longer join the MIS. Thus, in the second round, all these nodes broadcast to all their neighbors that they will not join the MIS. In the third round, only v_3 decides to join the MIS because all other undecided nodes have an undecided neighbor with a larger degree. Subsequently, only v_4 decides (not to join the MIS) in round 4. Repeating this argument, we get that the last node v_{k-1} makes its decision not before round $k - 1$. Since $n = k + (k - 1)$, the time complexity is thus $k - 1 = \frac{n-1}{2} \in \Omega(n)$.

2 (Local) Reductions

- a) We use one of the fast MIS algorithms from the script on the *line graph* of G , i.e., the graph $L(G) = (E, F)$ having the edges of G as nodes, where two nodes (edges of G) are connected exactly if they are adjacent to each other in G (formally: $F = \{\{e, f\} \in \binom{E}{2} \mid e \cap f \neq \emptyset\}$). Thus, a (node) coloring of $L(G)$ is an edge coloring of G . Since $L(G)$ has $m \in O(n^2)$ nodes and maximum degree $2(\Delta - 1)$ (an edge may be adjacent to $\Delta - 1$ others at each of its nodes), the algorithm will need $O(\log n^2) = O(\log n)$ time and $2(\Delta - 1) + 1 = 2\Delta - 1$ colors.

The line graph can be simulated locally, where nodes of the line graph (i.e., edges of G) are simulated by one of their incident nodes. The nodes simulating adjacent edges are connected by them and therefore at most 2 hops away from each other. Thus, two rounds and (at most) two messages are required to simulate one round of communication and one message on the line graph, respectively. Hence, the time complexity is doubled, but still in $O(\log n)$.

If we do not have edge orientations or identifiers, the decision which of the nodes plays the part of the edge can, e.g., be made w.h.p. in a single round by exchanging random bit strings of size $O(\log n)$ between neighbors.

- b) First, we 3-color the ring by means of Algorithm Six-2-Three (or its uniform variant from the first exercise sheet). Next, all nodes with color 0 join the dominating set and inform their neighbors. Then, all nodes with color 1 having no neighbor of color 0 join the set and inform their neighbors. Finally, still uncovered nodes with color 2 join the dominating set.

Obviously, the resulting set is a dominating set and the algorithm has a time complexity of $O(\log^* n)$. However, the constructed set is also a (maximal) independent set, as no two neighbors join. An independent set in a ring cannot have more than $n/2$ nodes, while a dominating set must contain at least $n/3$ nodes (each node covers itself and its two neighbors). In other words, the computed set is at most a factor of $3/2$ larger than any dominating set and hence also than a minimum dominating set.

- c) Again we use one of the fast MIS algorithms to compute a maximal independent set I within $O(\log n)$ time. It is also a dominating set (because a node without a neighbor in I could be added) and we claim that it is at most C times larger than a minimum dominating set M .

To prove this, consider a node $v \in I$. Since M is a dominating set, there must be at least one node in $(N(v) \cup \{v\}) \cap M$, i.e., a node from the optimal solution is in v 's neighborhood. For each $v \in I$, we count such a node. Because the graph is of bounded independence, no node $m \in M$ is counted more than C times, because there cannot be more than that many independent neighbors of m . Therefore, $|I| \leq C|M|$.