



Distributed Graph Algorithms

Sample Solution Exercise Sheet 1

Background. In this exercise sheet, we study the relationship between coloring and maximal independent set (MIS) problems in the LOCAL model of distributed computing.

Notation. Let $G = (V, E)$ be a simple graph with maximum degree Δ .

Exercise 1: From MIS to Coloring

(10 Points)

Suppose you are given a distributed algorithm \mathcal{B} that computes an MIS in T rounds.

- Describe how to use \mathcal{B} to compute a $(\Delta + 1)$ -coloring in at most T rounds.
- Explain why the number of colors is guaranteed to be at most $\Delta + 1$.

Sample Solution

We define a new graph $H = (V_H, E_H)$, such that an MIS of H can be turned into a $\Delta + 1$ coloring of our input graph G . We then argue, that nodes can simulate running the MIS algorithm on the graph H in T rounds, even though the actual communication network is G .

Constructing H : To create H , we create $\Delta + 1$ copies of every node in G , one for every possible color.

$$V_H = V[H] = \{v_1, v_2, \dots, v_{\Delta+1} \mid v \in V[G]\}$$

For each v , all of the copies are connected to each other, as a result only one of the copies can be in an MIS. To satisfy this for every v , we need the following edges.

$$E_v = \{\{v_i, v_j\} \mid 1 \leq i, j \leq \Delta + 1 \text{ with } i \neq j\}$$

Additionally, we need to connect the copies $v_1, \dots, v_{\Delta+1}$ of v , to the copies $u_1, \dots, u_{\Delta+1}$ of a node u that is a neighbor of v in G . We do so, by connecting each v_i with u_i , for all $1 \leq i \leq \Delta + 1$.

$$E_{u,v} = \{\{u_i, v_i\} \mid 1 \leq i \leq \Delta + 1\}$$

Now our actual set of edges contains E_v , the connections between the copies, for every node $v \in V[G]$ and the connections between neighbors, so the set $E_{u,w}$, for every edge $\{u, w\} \in E[G]$.

$$E_H = V[H] = \bigcup_{v \in V[G]} E_v \cup \bigcup_{\{u,w\} \in E[G]} E_{u,w}$$

Claim: Any valid MIS of H can be turned into a valid $\Delta + 1$ coloring of G .

Proof. Let $I \subset V[H]$ be an MIS of H . Because $E_v \subset E[H]$ all of the copies of a node $v \in V[G]$ are adjacent by construction. As a result we obtain

Fact 1: Only one of the copies of a node v can be in any valid MIS.

Next, we define the set of copies of v 's neighbors

$$N_v = \{u_1, \dots, u_{\Delta+1} \mid u \in \mathcal{N}(v)\}$$

, where $\mathcal{N}(v)$ is the 1 hop neighborhood of v . Notice that v has at most Δ neighbors in G . By Fact 1, for each neighbor u of v only one of the copies $u_1, \dots, u_{\Delta+1}$ in H is in the MIS I , we get the following inequality:

$$|N_v \cap I| \leq \Delta$$

As a result, there must exist at least one index $1 \leq j \leq \Delta + 1$, such that v_j is not adjacent to any neighboring copy that is in the MIS. So if no other copy of v is in I , then v_j must be in I , otherwise I is not maximal. As a result we obtain:

Fact 2: For each node v , exactly one of the copies is in I .

With this we can now define our $(\Delta + 1)$ -coloring $\phi : V[G] \rightarrow [\Delta + 1]$ of G . For every $v \in V[G]$, let j_v be the index of the copy of v that is in I . Note that because of Fact 2 this is well defined, that is for every v there is exactly one copy of v in the MIS and hence, there is a unique value for j_v . We prove that by setting $\phi(v) = j_v$ for all nodes, we ensure that ϕ is a proper $\Delta + 1$ coloring of G .

For contradiction, assume that ϕ is not a proper coloring, then there exists a monochromatic edge $\{v, u\} \in E[G]$. By the definition of ϕ , this means that $j_v = j_u$, let $j := j_u = j_v$. Furthermore, by the definition of j_u, j_v this means that both copies $u_j, v_j \in V[H]$ must be in the independent set I . But since the edge $\{v, u\}$ exists in G , this means that $\{u_j, v_j\} \in E_{u,v}$ exists in H and so u_j and v_j are neighbors, so they cannot both be in the independent set I , which is the desired contradiction. Hence ϕ must be a proper coloring of G , so we turned the MIS I into a proper $(\Delta + 1)$ -coloring, which proves the claim. \square

To complete the argument, we show how to transform the MIS algorithm \mathcal{B} on H into a $(\Delta + 1)$ -coloring algorithm \mathcal{A} on G , without increasing the round complexity. In \mathcal{A} , each node $v \in V[G]$ locally simulates the execution of \mathcal{B} on all its corresponding copies $v_1, \dots, v_{\Delta+1} \in V[H]$.

Observe that in one round of the LOCAL model on H , each copy v_i can only send messages to (i) other copies of the same node v , or (ii) copies of neighbors of v in G . The first type of communication—between copies of the same node—can be simulated internally by v , since it represents all its copies. The second type—communication between copies of neighboring nodes—can be simulated by having v send the corresponding messages to its actual neighbors in G , who in turn simulate the appropriate copies. Hence, one communication round in H can be simulated by a single communication round in G .

Once the simulation of \mathcal{B} on H is done, each node v knows which one of its copies is in the MIS and outputs the appropriate color, as described in the proof of the above Claim. After the termination of \mathcal{B} this can be done without communication and by the guarantees of the above Claim, this results in a proper coloring of G .

It follows that \mathcal{A} runs in exactly T rounds and produces a valid $(\Delta + 1)$ -coloring of G , as claimed.

Exercise 2: Coloring, using more colors.

(10 Points)

Give a *randomized* distributed algorithm that computes an $O(\Delta \log n)$ -coloring in a single communication round.

- Describe the random process each node performs to select its color.
- Explain why, with high probability, the resulting coloring is proper.

Sample Solution

We use a palette of $2c\Delta \log n$ colors, for some arbitrary constant $c > 0$. The palette is divided into $c \log n$ buckets, each containing 2Δ colors. Every node v independently selects one color uniformly at random from each bucket and communicates its choices to all of its neighbors.

We say that a bucket is *successful* for v if none of v 's neighbors selected the same color as v from that bucket. After all buckets are processed, node v checks whether it has at least one successful bucket; if so, v outputs its color choice from any such successful bucket.

If every node has at least one successful bucket, then we obtain a proper coloring, by the definition of successful buckets. So what remains to show is that every node has at least one successful bucket with large enough probability.

Observe that the random experiment in each bucket corresponds to one round of the 2Δ -coloring algorithm discussed in the lecture. Hence, for every bucket, the probability that it is successful for v is at least $\frac{1}{2}$. Consequently, the probability that v has no successful bucket at all is at most

$$\left(\frac{1}{2}\right)^{c \log n} = \frac{1}{n^c}.$$

Applying a union bound over all n nodes, the probability that *any* node has no successful bucket is at most

$$n \cdot \frac{1}{n^c} = \frac{1}{n^{c-1}}.$$

Therefore, with probability at least $1 - \frac{1}{n^{c-1}}$, all nodes have at least one successful bucket, and the algorithm produces a proper coloring, as desired.

Hint for Exercise 1. Try to build a graph H , such that solving MIS on graph H gives a $\Delta + 1$ coloring of G . Then argue that nodes can *pretend* to be in H .

Hint for Exercise 2. Try to do the entire coloring algorithm from the lecture in one round.