University of Freiburg
Dept. of Computer Science
Prof. Dr. F. Kuhn
S. Faour, M. Fuchs,
Z. Parsaeian, G. Schmit

Theory of Distributed Systems
Sample Solution Exercise Sheet 8

Due: Wednesday, 28th of June 2023, 12:00 noon

Exercise 1: Matching *(5 Points)*

A *matching* of a graph $G = (V, E)$ is a subset of edges $M \subseteq E$ such that no two edges in M are adjacent. A matching is maximal if no edge can be added without violating this property.

Give an algorithm that computes a maximal matching in $O(\log n)$ rounds w.h.p. in the synchronous message passing model. That is, after the algorithm terminates each node needs to know which of its adjacent edges are part of the maximal matching.

Sample Solution

Let $G = (V, E)$ be the graph for which we want to construct the matching. The so-called line graph G' is defined as follows: for every edge in G there is a node in G'; two nodes in G' are connected by an edge if their respective edges in G are adjacent. G' can be simulated on G with constant overhead. A (maximal) independent set in the line graph G is a (maximal) matching in the original graph G, and vice versa. If G has n nodes, G' has at most n^2 nodes and so we can compute an MIS on G' in time $O(\log n^2) = O(\log n)$.

Exercise 2: Dominating Set *(8 Points)*

A *dominating set* of a graph $G = (V, E)$ is a subset of the nodes $D \subseteq V$ such that each node is in D or adjacent to a node in D. A minimum dominating set is a dominating set containing the least possible number of nodes. $G = (V, E)$ has *neighborhood independence* β if for every node $v \in V$ the largest independent set of the neighborhood $N(v) := \{u \in V \mid \{v, u\} \in E\}$ of v is of size at most β.

a) Show that for an MIS M and a minimum dominating set D of a graph it holds $|D| \leq |M|$.

b) Give a class of graphs each containing an independent set I and a dominating set D with $\frac{|I|}{|D|} = O(n)$.

c) Show that for graphs with neighborhood independence $\beta \geq 1$, a β-approximation to a minimum dominating set (that is a dominating set which is at most β times larger than a minimum dominating set) can be found in time $O(\log n)$ w.h.p. in the synchronous message passing model.

d) A unit disc graph is a graph (V, E) with $V \subset \mathbb{R}^2$ and $E = \{\{u, v\} \mid \|u - v\|_2 \leq 1\}$. Show that one can compute a 5-approximation to a minimum dominating set in disc graphs in time $O(\log n)$ w.h.p. in the synchronous message passing model.$\mathbb{R}$

Sample Solution

a) By definition, every MIS is a dominating set so any MIS must be at least as large as the *minimum* dominating set.

b) For every n, take a star graph with n nodes which has a dominating set of size 1 and an independent set of size $n-1$.

c) We compute an MIS I in time $O(\log n)$ which is a dominating set. We have to show that for a minimum dominating set D we have $|I| \leq \beta|D|$. Each node in I is either in D or neighbor of a node in D. So for counting the nodes in I, we can iterate through the nodes in D and count in each step the number of nodes in I which are covered by the corresponding node in D. Each node in D has at most β neighbors in I. Therefore we will count at most $\beta|D|$ nodes.

d) We show that disc graphs have neighborhood independence 5. Consider a node v. We need to show that in the unit circle around v (i.e., v's neighborhood), there can fit at most five nodes with pairwise distance > 1. Assume we try to place six nodes with pairwise distance > 1. Consider one of those nodes u inside the circle. We partition the circle into six parts of equal size (as in the picture) such that u is located on one of the partitioning lines.

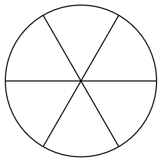

Roughly speaking, each piece contains a equilateral triangle with side length 1, in which the maximum distance is at most 1 (and the points in the remaining curved part also have distance at most 1 to the others). Therefore, two nodes in the same part including the rim have distance at most 1. So any node with distance > 1 from u must be outside the two parts in which u is located. But then there are only four parts left where one can place the remaining five points, so two of them have to be in the same part and therefore have distance at most 1, a contradiction.

Exercise 3: Coloring *(7 Points)*

Assume we have $C = \alpha(\Delta + 1) \in \mathbb{N}$ colors for some $\alpha \geq 1$. Consider the following algorithm in the synchronous message passing model to color the graph with C colors. Each node v repeats the following steps (corresponding to a phase) until it has a color:

- Let N_v be the set of yet uncolored neighbors of v and let C_v be the set of colors that v's neighbors already chose (initially N_v are all of v's neighbors and $C_v = \emptyset$).
- Node v picks a random number $r_c(v) \in [0, 1]$ for every remaining color $c \in \{1, \ldots, C\} \setminus C_v$ and informs its neighbors about those numbers.
- If $r_c(v) < r_c(u)$ for some $c \in \{1, \ldots, C\} \setminus C_v$ and every $u \in N_v$, then v colors itself with c, informs its neighbors and terminates (if this holds for several c, v chooses one of those arbitrarily).

(a) Show that the probability that a node obtains a color in a given phase is at least $1 - e^{-\alpha}$.

(b) Show that the algorithm terminates after $\mathcal{O}\big(1 + \frac{\log n}{\alpha}\big)$ rounds in expectation.

Sample Solution

(a) The algorithm uses a similar techniques as for computing an MIS seen in the lecture (Luby's algorithm). For each color we compute an independent set. A node joins one of those independent sets by having the smallest random number.

In some phase, let d be the number of uncolored neighbors of v (note that we have $d + |C_v| \leq \Delta$, i.e., $|C_v| \leq \Delta - d$). Then the number of colors that v can still compete for is

$$C - |C_v| \geq \alpha(\Delta + 1) - (\Delta - d) = (\alpha - 1)\Delta + \alpha + d \geq (\alpha - 1)d + \alpha + d = \alpha(d + 1)$$

The chance for v to join some remaining color c is at least $\frac{1}{d+1}$. The probability that a node does not get some color is at most

$$(1 - \tfrac{1}{d+1})^{\alpha(d+1)} = \big((1 - \tfrac{1}{d+1})^{(d+1)}\big)^{\alpha} \leq \big(\tfrac{1}{e}\big)^{\alpha}.$$

(b) In each phase i we expect that at most an $e^{-i\cdot\alpha}$ fraction of nodes remains uncolored.

$$\begin{aligned} n \cdot e^{-i\cdot\alpha} &\leq 1 \\ \iff \quad \ln n - i \cdot \alpha &\leq 0 \\ \iff \quad \tfrac{\ln n}{\alpha} &\geq i \end{aligned}$$

So we expect that there is at most one color left after $\frac{\ln n}{\alpha}$ phases. Coloring the remaining node takes at most one additional round, so the expected number of rounds is $1 + \frac{\ln n}{\alpha}$.