

DM582 Exercises - Sheet 4

Mads Anker Nielsen Tobias Samsøe Sørensen

February 25, 2025

This document contains exercises from the course DM582 (spring 2025). Most exercises are from the book *Introduction to Algorithms, 4th edition* by Cormen, Leiserson, Rivest, and Stein (CLRS), the book *Algorithm Design, 1st edition* by J. Kleinberg and E. Tardos (KT), and the book *Discrete Mathematics and its Applications, 8th edition* by K. Rosen.

References to CLRS refer to the book *Introduction to Algorithms, 4th edition* by Cormen, Leiserson, Rivest, and Stein.

References to KT refer to the book *Algorithm Design, 1st edition* by J. Kleinberg and E. Tardos.

References to Rosen refer to the book *Discrete Mathematics and its Applications, 8th edition* by K. Rosen.

This document will inevitably contain mistakes. If you find some, please report them to your TA so that we can correct them.

Sheet 4

KT, Exercise 1

Exercise

3-Coloring is a yes/no question, but we can phrase it as an optimization problem as follows. Suppose we are given a graph $G = (V, E)$, and we want to color each node with one of three colors, even if we aren't necessarily able to give different colors to every pair of adjacent nodes. Rather, we say that an edge (u, v) is satisfied if the colors assigned to u and v are different. Consider a 3-coloring that maximizes the number of satisfied edges, and let c^* denote this number. Give a polynomial-time algorithm that produces a 3-coloring that satisfies at least $\frac{2}{3}c^*$ edges. If you want, your algorithm can be randomized; in this case, the expected number of edges it satisfies should be at least $\frac{2}{3}c^*$.

KT, Exercise 2

Exercise

Consider a county in which 100,000 people vote in an election. There are only two candidates on the ballot: a Democratic candidate (denoted D) and a Republican candidate (denoted R). As it happens, this county is heavily Democratic, so 80,000 people go to the polls with the intention of voting for D , and 20,000 go to the polls with the intention of voting for R . However, the layout of the ballot is a little confusing, so each voter, independently and with probability $1/100$ votes for the wrong candidate – that is, the one that he or she didn't intend to vote for. (Remember that in this election, there are only two candidates on the ballot.) Let X denote the random variable equal to the number of votes received by the Democratic candidate D , when the voting is conducted with this process of error. Determine the expected value of X , and give an explanation of your derivation of this value.

KT, Exercise 3

Exercise

In Section 13.1, we saw a simple distributed protocol to solve a particular contention-resolution problem. Here is another setting in which ran-

domization can help with contention resolution, through the distributed construction of an independent set.

Suppose we have a system with n processes. Certain pairs of processes are in conflict, meaning that they both require access to a shared resource. In a given time interval, the goal is to schedule a large subset S of the processes to run – the rest will remain idle – so that no two conflicting processes are both in the scheduled set S . We'll call such a set S conflict-free.

One can picture this process in terms of a graph $G = (V, E)$ with a node representing each process and an edge joining pairs of processes that are in conflict. It is easy to check that a set of processes S is conflict-free if and only if it forms an independent set in G . This suggests that finding a maximum-size conflict-free set S , for an arbitrary conflict G , will be difficult (since the general Independent Set Problem is reducible to this problem). Nevertheless, we can still look for heuristics that find a reasonably large conflict-free set. Moreover, we'd like a simple method for achieving this without centralized control: Each process should communicate with only a small number of other processes and then decide whether or not it should belong to the set S .

We will suppose for purposes of this question that each node has exactly d neighbors in the graph G . (That is, each process is in conflict with exactly d other processes.)

- (a) Consider the following simple protocol.

Each process P_i independently picks a random value x_i ; it sets x_i to 1 with probability $\frac{1}{2}$ and sets x_i to 0 with probability $\frac{1}{2}$. It then decides to enter the set S if and only if it chooses the value 1, and each of the processes with which it is in conflict chooses the value 0.

Prove that the set S resulting from the execution of this protocol is conflict-free. Also, give a formula for the expected size of S in terms of n (the number of processes) and d (the number of conflicts per process).

- (b) The choice of the probability $\frac{1}{2}$ in the protocol above was fairly arbitrary, and it's not clear that it should give the best system performance. A more general specification of the protocol would replace the probability $\frac{1}{2}$ by a parameter p between 0 and 1, as follows.

Each process P_i independently picks a random value x_i ; it sets x_i to 1 with probability p and sets x_i to 0 with probability $1 - p$. It then decides to enter the set S if and only if it chooses the value 1, and each of the processes with which it is in conflict chooses the value 0.

In terms of the parameters of the graph G , give a value of p so that the expected size of the resulting set S is as large as possible. Give a formula for the expected size of S when p is set to this optimal value.

KT, Exercise 4

Exercise

A number of peer-to-peer systems on the Internet are based on overlay networks. Rather than using the physical Internet topology as the network on which to perform computation, these systems run protocols by which nodes choose collections of virtual “neighbors” so as to define a higher-level graph whose structure may bear little or no relation to the underlying physical network. Such an overlay network is then used for sharing data and services, and it can be extremely flexible compared with a physical network, which is hard to modify in real time to adapt to changing conditions.

Peer-to-peer networks tend to grow through the arrival of new participants, who join by linking into the existing structure. This growth process has an intrinsic effect on the characteristics of the overall network. Recently, people have investigated simple abstract models for network growth that might provide insight into the way such processes behave, at a qualitative level, in real networks.

Here’s a simple example of such a model. The system begins with a single node v_1 . Nodes then join one at a time; as each node joins, it executes a protocol whereby it forms a directed link to a single other node chosen uniformly at random from those already in the system. More concretely, if the system already contains nodes v_1, v_2, \dots, v_{k-1} and node v_k wishes to join, it randomly selects one of v_1, v_2, \dots, v_{k-1} and links to this node.

Suppose we run this process until we have a system consisting of nodes v_1, v_2, \dots, v_n ; the random process described above will produce a directed network in which each node other than v_1 has exactly one outgoing edge. On the other hand, a node may have multiple incoming links, or none

at all. The incoming links to a node v_j reflect all the other nodes whose access into the system is via v_j ; so if v_j has many incoming links, this can place a large load on it. To keep the system load-balanced, then, we'd like all nodes to have a roughly comparable number of incoming links. That's unlikely to happen here, however, since nodes that join earlier in the process are likely to have more incoming links than nodes that join later. Let's try to quantify this imbalance as follows.

- (a) Given the random process described above, what is the expected number of incoming links to node v_j in the resulting network? Give an exact formula in terms of n and j , and also try to express this quantity asymptotically (via an expression without large summations) using $\Theta(\cdot)$ notation.
- (b) Part (a) makes precise a sense in which the nodes that arrive early carry an "unfair" share of the connections in the network. Another way to quantify the imbalance is to observe that, in a run of this random process, we expect many nodes to end up with no incoming links.

Give a formula for the expected number of nodes with no incoming links in a network grown randomly according to this model.

Exercise from course website

Exercise

In this problem, we consider an undirected graph $G = (V, E)$ to model the problem where there are small devices, each of which uses wireless communication to communicate with d of the other nearby devices. The vertices in V are the devices, and there are edges between vertices representing pairs of devices within communication range of each other. Thus, each vertex has exactly d neighbors. Some of the devices should be given an uplink transmitter so they can send data to the main station. We want to use as few uplink transmitters as possible while still getting all of the data from all of the devices. Thus, in our graph, we want a subset $S \subseteq V$ such that for all $v \in V$, either $v \in S$ or v has a neighbor u such that $u \in S$. Such a set S is called a *dominating set*. If the devices corresponding to vertices in S all get uplink transmitters, then all devices can either send data to the main station directly, or send data to a device that can send data directly.

Finding a minimum-sized dominating set is an NP-hard problem (this

means that no sub-exponential algorithm is known), so one does not expect to find a polynomial-time deterministic algorithm to solve it.

In this problem, consider a randomized algorithm that chooses a set S of $k = \frac{cn \ln(n)}{d+1}$ vertices from V uniformly at random. The constant c will be discussed later. You should show through the following that S is a dominating set with high probability.

In the following, assume that the vertices in S are chosen one at a time and choosing the same vertex more than once is allowed. A vertex v is said to dominate a vertex u if they are the same vertex or u is a neighbor of v . A set S' dominates u if some vertex in S' dominates u .

1. What can you say about the minimum size a dominating set in G must have?
2. Let $D[v, t]$ be the event that the t th random vertex chosen dominates v . Find the probability of $D[v, t]$.
3. Let D_v be the event that S dominates v . What is the probability of the complement of this event, $\overline{D_v}$?
4. Show that $\left(\frac{1}{e}\right)^{c \ln(n)}$ is an upper bound on the probability of $\overline{D_v}$ and that $\left(\frac{1}{e}\right)^{c \ln(n)} = \frac{1}{n^c}$.
5. Let A be the event that S is a dominating set. Use the Union Bound to give an upper bound on the probability of \overline{A} .
6. Discuss what value c should have.