



Principles of Distributed Computing

Exercise 11

1 How to ping a disrupted flow network

You are given a flow network with unit capacity edges, in which the maximum s - t flow is k . One morning the network administrator discovers that a hacker has attacked the network, and so now there is no path from s to t . That is, the maxflow has been reduced to zero. Suppose the hacker is a theorist, who managed to do this by deleting exactly k edges.

The network managers use a tool called *ping*, which has the property that the command *ping*(v) tells them whether the node v is reachable from s or not.

Your job is to design an algorithm that issues ping commands to various nodes in the network, and then reports the *full* set of nodes that are currently *not reachable* from s . Of course, this can be done by pinging every node in the graph, but the goal is to accomplish this by using as few pings as possible. Your algorithm can decide which node to ping next based on the past pings.

Show that this problem can be solved by using only $O(k \log n)$ pings.

2 Paid advertisements on search engines

Companies like Google and Yahoo derive most of their revenue from displaying ads from sponsors on web pages accessed by users. Deciding which ads to show to which users, however, requires sophisticated algorithms and computation. We consider a simplified version here.

The managers of a popular Web site have identified k distinct *demographic groups* G_1, G_2, \dots, G_k . These groups can overlap; for instance, G_1 can be all residents of California, while G_2 can be all lawyers. The web site has entered into a contract with m different '*advertisers*', to show their ads to users that visit this web site. The contract with an advertiser i has the following format:

- The advertiser specifies a subset of the demographic groups $X_i \subset \{G_1, G_2, \dots, G_k\}$, and wants its ads shown only to those users that belong to at least one group in X_i .
- The advertiser wants its ads shown to at least r_i users each minute.

Now the algorithmic problem facing the web site manager is to determine a good *advertising policy*: which ads to show to whom. In this simplified version, each user is shown exactly one ad.

Suppose in a given minute, there are n users visiting the web site. Because we have the *registration* information from each user, we know which group(s) each user belongs to.

Your task is to decide if it is feasible to meet the contracts with your advertisers for this minute? (That is, for each advertiser i , can at least r_i users, each belonging to demographic set X_i , be shown an ad from the advertiser i ?)

Give a polynomial time algorithm for this decision problem, and if the problem is feasible, also determine which ads are shown to which people.