

CSC 373: Assignment 3

Worth 10%. Due August 3 at the beginning of class (6pm).

The work you submit must be your own. You may discuss problems with each others; however, you should prepare written solutions alone. Copying assignments is a serious academic offence, and will be dealt with accordingly.

Question 1 [Based on Problem 16 (Chapter 7, pp 422–423 in the text) with explicit mention of flow network algorithm, correctness proof and running time.]

Back in the euphoric early days of the Web, people liked to claim that much of the enormous potential in a company like Yahoo! was in the “eyeball”—the simple fact that millions of people look at its pages every day. Further, by convincing people to register personal data with the site, a site like Yahoo! can show each user an extremely targeted advertisement whenever he or she visits the site, in a way that TV networks or magazines couldn’t hope to match. So if a user has told Yahoo! that he or she is a 20-year-old computer science major from U of T, the site can present a banner ad for apartments in downtown Toronto; on the other hand, if he or she is a 50-year-old investment banker from Greenwich, Connecticut, the site can display a banner ad pitching Lincoln Town Cars instead.

But deciding on which ads to show to which people involves some serious computation behind the scenes. Suppose that 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 example, G_1 can be all residents of Ontario, and G_2 can be all people with a degree in computer science.) The site has contracts with m different *advertisers* A_1, A_2, \dots, A_m , to show a certain number of copies of their ads to users of the site. Here’s what the contract with the i -th advertiser looks like:

- For a subset $X_i \subseteq \{G_1, \dots, G_k\}$ of the demographic groups, advertiser A_i wants its ads shown only to users who belong to at least one of the demographic groups in the set X_i .
- For a number r_i , the advertiser A_i wants its ads shown to at least r_i users each minute.

Now consider the problem of designing a good *advertising policy*—a way to show a single ad to each user of the site. Because we have registration information on each of these users, we know that user j (for $j = 1, 2, \dots, n$) belongs to a subset $U_j \subseteq \{G_1, G_2, \dots, G_k\}$ of the demographic groups. The problem is: Is there a way to show a single ad to each user so that the site’s contracts with each of the m advertisers is satisfied for this minute? (That is, for each $i = 1, 2, \dots, m$, can at least r_i of the n users, each belonging to at least one demographic group in X_i , be shown an ad provided by advertiser i ?)

Using flow network algorithm, give an algorithm to decide if this is possible, and if so, to actually choose an ad to show each user. Prove that your algorithm is correct by proving the relationship between the maximum flow and a way of showing ads to users that satisfies the advertisers contracts. What is the running time of your algorithm ?

Question 2 [Teaching Assignment]

Suppose that there are n professors p_1, \dots, p_n and m courses c_1, \dots, c_m . Each professor p_i has teaching (positive integer) load L_i (the number of sessions that he or she must teach) and a subset S_i of courses that he or she is capable of teaching. Each course c_j has k_j sessions ($k_j \geq 1$). The Teaching Assignment Problem is decide whether it is possible to assign sessions of the courses to professors that satisfies the following constraints, and if so, find one such assignment. The constraints are:

- Each professor is assigned only courses that he or she is capable of teaching;
- Each professor p_i is assigned exactly L_i sessions;
- Each session is assigned exactly one professor.

Using flow network, give an algorithm to solve the above problem. The algorithm should output an assignment if there exists one. Prove the correctness of your algorithm by proving the relationship between the maximum flow (output by the maximum flow algorithm) and a solution to the Teaching Assignment Problem.

Question 3

Give a linear program that, given a bipartite graph $G = (V, E)$, solves the maximum bipartite matching problem.

Question 4 [Minimum-Cost Multicommodity-Flow Problem]

Suppose that we are given a directed graph $G = (V, E)$ where each edge $e \in E$ has a nonnegative capacity $c(e) \geq 0$ and a weight $w(e)$. There are k commodities K_1, K_2, \dots, K_k , where commodity K_i has source s_i , sink t_i and demand d_i . For a commodity K_i we want a flow of value d_i which emanates from source s_i and goes to sink t_i obeying the Conservation Condition. Let $f_i(e)$ denote the flow of commodity K_i on edge e . The General Capacity Condition is that

$$f(e) \leq c(e)$$

for every edge e , where $f(e)$ the the *aggregate flow* on edge e :

$$f(e) = \sum_{i=1}^k f_i(e)$$

Thus, a feasible flow $f = (f_1, \dots, f_k)$ is such that

- the flow f_i (of commodity K_i) is generated at s_i and consumed at t_i , for $1 \leq i \leq k$;
- the value of the flow f_i (of commodity K_i) is d_i , for $1 \leq i \leq k$;
- the Conservation Condition holds for every flow;
- the General Capacity Condition holds.

The weight of a feasible flow f is

$$\sum_{e \in G} w(e)f(e)$$

The Minimum-Cost Multicommodity-Flow Problem is to find a feasible flow with minimum weight.

For simplicity, we may assume that there are no edge entering the sources s_i , and no edge leaving the sinks t_i .

Express this problem as a linear program. [You are *not* required to solve this problem.]