

CME 305: Discrete Mathematics and Algorithms

Instructor: Reza Zadeh (rezab@stanford.edu)

HW#4 – Due at the beginning of class Thursday 03/13/14

1. We are given n jobs that each take one unit of processing time. All jobs are available at time 0, and job j has a profit of c_j and a deadline d_j . The profit for job j will only be earned if the job completes by time d_j . The problem is to find an ordering of the jobs that maximizes the total profit. First, prove that if a subset of the jobs can be completed on time, then they can also be completed on time if they are scheduled in the order of their deadlines. Now, let $E = \{1, \dots, n\}$ and let

$$I = \{J \subseteq E : J \text{ can be completed on time}\}$$

Prove that $M = (E, I)$ is a matroid and describe how to find an optimal ordering for the jobs.

2. Let $G = (V, E)$ be a c -edge connected graph. In other words, assume that the size of minimum cut in G is at least c . Construct a graph $G'(V, E')$ by sampling each edge of G with probability p independently at random and reweighing each edge with weight $1/p$. Suppose $c > \log n$, and ϵ is such that $\frac{\log(n)}{c\epsilon^2} \leq 1$. Show that as long as $p \geq \frac{\log(n)}{c\epsilon^2}$, with high probability the size of every cut in G' is within $(1 \pm \epsilon)$ of the cut in the original graph G .
3. Consider a random bipartite graph on two vertex sets L, R of size n each, in which every vertex in L independently picks d uniformly independent random neighbors in R . Given $\epsilon > 1/d$, prove that with high probability, every set S in L of size $\alpha n/d$ has at least $|S|(1 - \epsilon)d$ neighbors in R , where α depends only on ϵ and d .
4. Consider scheduling n jobs to m identical machines to minimize the time taken by the machine with the heaviest load (i.e. to minimize the makespan). One algorithm is to order the jobs by decreasing processing times $t_1 \geq t_2 \geq \dots \geq t_n$, then greedily assign jobs to the machine whose load is the smallest so far (starting with the heaviest job).
 - (a) Show that this algorithm is a $4/3$ approximation to the optimal makespan.
 - (b) Show that the $4/3$ approximation is tight for this algorithm.
5. Suppose you are given a directed graph $G = (V, E)$ in which each edge has a positive cost. Also suppose that G has a node r such that there is a path from r to every other node in G . Give a polynomial-time algorithm that constructs a minimum cost arborescence rooted at r and prove it is correct.
6. A square integer matrix A is **unimodular** if and only if its determinant is -1 or 1 . A matrix (not necessarily square) M is **totally unimodular** iff every square submatrix has determinant $1, -1$, or 0 , i.e. every non-singular square submatrix is unimodular.
Show that for a linear program with totally unimodular constraint matrix M and integral right-hand side c , all extreme points must be integral.

7. Given a list of personnel (n persons) and of list of k vacation periods, each period spanning several contiguous vacation days. Let D_j be the set of days included in the j th vacation period. You need to produce a schedule satisfying:
- For a given parameter c , each tech support person should be assigned to work at most c vacation days total.
 - For each vacation period j , each person should be assigned to work at most one of the days during the period.
 - Each vacation day should be assigned a single tech support person.
 - For each person, only certain vacation periods are viable.

Describe a polynomial time algorithm to generate an assignment or output that no assignment exists. Prove correctness.

8. Online social networks carry a huge potential for online advertising. After a recent controversy, a popular social networking platform does not allow advertisers to target the users individually. However, it is allowed to run ads on user communities.

Formally, let X be the set of all users on a social network, and S_1, S_2, \dots, S_m be subsets of X , where each S_i represents a user community. Notice that a user can belong to several communities. Suppose the advertiser can afford placing ads on at most k communities. The goal is to show the ads to as many users as possible, i.e. to find $S_{i_1}, S_{i_2}, \dots, S_{i_k}$ such that $|\cup_{j=1}^k S_{i_j}|$ is maximized.

Unfortunately, this problem is NP-hard and therefore we are interested in designing efficient approximation algorithms to solve it. Consider the following greedy approach: pick the k communities one at a time, and in each iteration pick the community that contains the largest number of users that have not been covered yet. In other words, choose the community that maximizes the current coverage. Show that this greedy approach yields at least $1 - (1 - 1/k)^k > 1 - 1/e$ fraction of the optimal solution.

Hint: Let x_i denote the number of new elements covered by the algorithm in the i -th set that it picks. Also, let $y_i = \sum_{j=1}^i x_j$, and $z_i = OPT - y_i$. Show $x_{i+1} \geq z_i/k$ and prove by induction that $z_i \leq (1 - 1/k)^i OPT$.