CSE 521 – Algorithms – Winter 2010
Assignment 5 Suggested Solution

1 Multicommodity Flows

The value of the i-th commodity flow is $\sum_{u:(s_i,u)\in E} f_i(s_i,u)$. Then the LP is:

`Maximize` $\sum_{i=1}^{p} \sum_{u:(s_i,u)\in E} f_i(s_i,u)$
`s.t` $\sum_{u:(u,v)\in E} f_i(u,v) - \sum_{w:(v,w)\in E} f_i(v,w) = 0$ `for all` $1 \le i \le p, v \in V \setminus \{s_i,t_i\}$
$\sum_{i=1}^{p} f_i(u,v) \le c(u,v)$ `for all` $(u,v) \in E$
$f_i(u,v) \ge 0$ `for all` $i,(u,v)$

where the first constraint conserves flow at every vertex, and the second constraint upper bounds the total flow on every edge.

2 Scheduling

(a) Variables $x_{i,j}$ indicates whether job i is assigned to machine j.

`Minimize` T
`s.t.` $T - \sum_{i=1}^{n} x_{i,j}p_{i,j} \ge 0$ `for every machine` j
$\sum_{j=1}^{m} x_{i,j} \ge 1$ `for every job` i
$x_{i,j} \in \{0,1\}$

The second constraint says that every job must be assigned to some machine(s) (but optimality ensures that exactly one machine is assigned). As $p_{i,j} \ge 0$, it is not neccessary to put nonnegative constraint on T.

(b)

`Minimize` T
`s.t.` $T - \sum_{i=1}^{n} x_{i,j}p_{i,j} \ge 0$ `for every machine` j
$\sum_{j=1}^{m} x_{i,j} \ge 1$ `for every job` i
$x_{i,j} \ge 0$ `for every` i,j

(c) We have machine-variables b_j and job-variables a_i.

$$\begin{array}{ll} \texttt{Maximize} & \sum_{i=1}^{n} a_i \\ \texttt{s.t.} & \sum_{j=1}^{m} b_j = 1 \\ & -p_{i,j} b_j + a_i \leq \ 0 \quad \texttt{for every } i, j \\ & a_i, b_j \geq 0 \quad \texttt{for every } i, j \end{array}$$

(d) Since the primal LP is feasible and its optimality is bounded (must be nonnegative), strong duality holds. Therefore, $OPT_{IP} \geq OPT_{LP} = OPT_{dual-LP}$.

(e) Let's consider the following instance: there are 1 job and m machines, and the executing times of the job on all machines are the same, say $p = 1$. Obviously the optimal makespan is $T = 1$. However in the LP we are allowed to use fractional assignments. Assigning evenly to all machines, the optimal value of the LP is $1/m$ with all $x_{1,j} = 1/m$. The integrality gap of the LP is lower bounded by this example, which is m.

(f) Let $e_{i,j}$ be the random variable, where $e_{i,j} = 1$ if job i is assigned to machine j, and $e_{i,j} = 0$ otherwise. Thus $\mathbf{E}[e_{i,j}] = \Pr[e_{i,j} = 1] = x^*_{i,j}$. We also have $T_j = \sum_i e_{i,j} p_{i,j}$. Thus,

$$\mathbf{E}[T_j] = \mathbf{E}[\sum_i e_{i,j} p_{i,j}] = \sum_i \mathbf{E}[e_{i,j} p_{i,j}] = \sum_i x^*_{i,j} p_{i,j}.$$

3 Weighted Set-Cover

We will have $m_i^{(t)}, w_i^{(t)}, \phi^{(t)}, p_i^{(t)}$ as defined in the unweighted version. The only difference now is that at step t the adversary will pick the set S_{j_t} that maximizes

$$\frac{\sum_{i \in S_{j_t}} p_i^{(t)}}{c_{j_t}} = \frac{\sum_{i \in S_{j_t}} w_i^{(t)}}{c_{j_t} \phi^{(t)}}.$$

First let's bound the above quantity. Let $S_{OPT} = \{S_{OPT1}, S_{OPT2}, \ldots\}$ be the minimum weighted cover with weight $OPT = \sum_j c_{OPTj}$. Thus,

$$\frac{1}{OPT} = \frac{\sum_{i=1}^{n} p_i^{(t)}}{\sum_j c_{OPTj}} \leq \frac{\sum_j \sum_{i \in S_{OPTj}} p_i^{(t)}}{\sum_j c_{OPTj}} \leq \max_j \frac{\sum_{i \in S_{OPTj}} p_i^{(t)}}{c_{OPTj}},$$

where the first inequality follows since S_{OPT} is a cover, and the second inequality follows from the fact that $\frac{\sum_j x_j}{\sum_j y_j} \leq \max_j \frac{x_j}{y_j}$ for any positive numbers x_j, y_j.

From the way the adversary picks S_{j_t}, $1/OPT \leq \frac{\sum_{i \in S_{j_t}} w_i^{(t)}}{c_{j_t} \phi^{(t)}}$, or $\sum_{i \in S_{j_t}} w_i^{(t)} \geq c_{j_t} \phi^{(t)} / OPT$. Then,

$$\phi^{(t+1)} = \phi^{(t)} - \sum_{i \in S_{j_t}} w_i^{(t)} \leq \phi^{(t)} (1 - c_{j_t}/OPT) \leq \phi^{(t)} e^{-c_{j_t}/OPT},$$

where the last inequality follows from $1 - x \leq e^{-x}$ for $x \geq 0$.

Note that since $\sum_{i \in S_{j_t}} w_i^{(t)} \geq c_{j_t} \phi^{(t)} / OPT$, we also have $c_{j_t} \leq OPT$.

Since $\phi^{(1)} = n$, we have $\phi^{(t+1)} \le ne^{-(\sum_{k=1}^{t} c_{j_k})/OPT}$. Let t^* be the minimum integer such that $(\sum_{k=1}^{t^*} c_{j_k})/OPT > \ln n$. We then have $\phi^{(t^*+1)} < 1$ which means $\phi^{(t^*+1)} = 0$ and everyone is covered. The cost of this cover is

$$\sum_{k=1}^{t^*} c_{j_k} = \sum_{k=1}^{t^*-1} c_{j_k} + c_{j_{t^*}} \le (\ln n + 1)OPT.$$

This shows that the greedy algorithm (which chooses sets as the adversary does) has approximation ratio $\ln n + 1$.