Advanced Algorithms 16 October 2018

Lecture 5: Approximation Algorithms V

Lecturer: Mohsen Ghaffari *Scribe: Davin Choo*

1 Approximation algorithms via rounding ILPs

Linear programming (LP) and integer linear programming (ILP) are versatile models but with different solving complexities — LPs are solvable in polynomial time while ILPs are **NP**-hard.

Definition 1 (Linear program (LP)). *In canonical form, a LP is expressed as*

$$\begin{array}{ll} \textit{minimize} & \boldsymbol{c}^T\boldsymbol{x} \\ \textit{subject to} & A\boldsymbol{x} \leq \boldsymbol{b} \\ & \boldsymbol{x} \geq 0 \end{array}$$

where $\boldsymbol{x}$ is the vector of n variables (to be determined), $\boldsymbol{b}$ and $\boldsymbol{c}$ are vectors of (known) coefficients, and A is a (known) matrix of coefficients. $\boldsymbol{c}^T\boldsymbol{x}$ and $obj(\boldsymbol{x})$ are known as the objective function and objective value of the LP respectively. For an optimal variable assignment $\boldsymbol{x}^*$, $obj(\boldsymbol{x}^*)$ is the optimal objective.

ILPs are defined similarly with the additional constraint that variables take on integer values. As we will be relaxing ILPs into LPs, to avoid confusion, we use y for ILP variables to contrast against the x variables in LPs.

Definition 2 (Integer linear program (ILP)). *In canonical form, a ILP is expressed as*

$$\begin{array}{ll} \textit{minimize} & \boldsymbol{c}^T\boldsymbol{y} \\ \textit{subject to} & A\boldsymbol{y} \leq \boldsymbol{b} \\ & \boldsymbol{y} \geq 0 \\ & \boldsymbol{y} \in \mathbb{Z}^n \end{array}$$

where $\boldsymbol{y}$ is the vector of n variables (to be determined), $\boldsymbol{b}$ and $\boldsymbol{c}$ are vectors of (known) coefficients, and A is a (known) matrix of coefficients. $\boldsymbol{c}^T\boldsymbol{y}$ and $obj(\boldsymbol{y})$ are known as the objective function and objective value of the LP respectively. For an optimal variable assignment $\boldsymbol{y}^*$, $obj(\boldsymbol{y}^*)$ is the optimal objective.

In this lecture, we illustrate how we can model set cover and multi-commodity routing as ILPs and how to perform rounding to yield approximations for these problems. As per previous lectures, Chernoff bounds will be a useful inequality in our analysis toolbox.

Theorem 3 (Chernoff bound). *For independent Bernoulli variables $X_1, \ldots, X_n$, let $X = \sum_{i=1}^n X_i$. Then,*

$$\begin{array}{ll} \Pr[X \geq (1+\epsilon)\mathbb{E}(X)] \leq \exp(\frac{-\epsilon^2\mathbb{E}(X)}{3}) & \textit{for } 0 < \epsilon \\ \Pr[X \leq (1-\epsilon)\mathbb{E}(X)] \leq \exp(\frac{-\epsilon^2\mathbb{E}(X)}{2}) & \textit{for } 0 < \epsilon < 1 \end{array}$$

By union bound, for $0 < \epsilon < 1$, we get $\Pr[|X - \mathbb{E}(X)| \geq \epsilon\mathbb{E}(X)] \leq 2\exp(\frac{-\epsilon^2\mathbb{E}(X)}{3})$

Remark There is actually a tighter form of Chernoff bounds:

$$\forall \epsilon > 0, \Pr[X \geq (1+\epsilon)\mathbb{E}(X)] \leq (\frac{e^\epsilon}{(1+\epsilon)^{1+\epsilon}})^{\mathbb{E}(X)}$$

2 Set cover

Recall the minimum set cover problem.

Definition 4 (Minimum set cover problem). *Given $\mathcal{U}$, $\mathcal{S}$, and $c : \mathcal{S} \to \mathbb{R}^+$, find a subset $S^* \subseteq \mathcal{S}$ such that:*

(i) (Set cover): $\bigcup_{S_i \in S^*} S_i = \mathcal{U}$

(ii) (Minimum cost): $c(S^*)$ *is minimized.*

Example

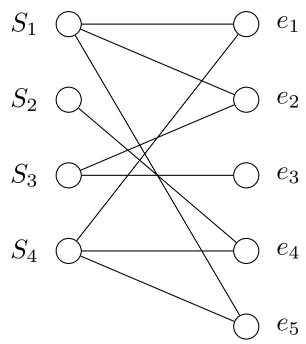

Suppose there are $n = 5$ vertices and $m = 4$ subsets $S = \{S_1, S_2, S_3, S_4\}$, where the cost function is defined as $c(S_i) = 2^i$. Then, the minimum set cover is $S = \{S_1, S_2, S_3\}$ with a cost of $c(S) = 14$.

In the first lecture, we saw a greedy selection of sets that maximize the number of remaining uncovered items gave a H_n-approximation for set cover. Furthermore, in the special cases where $\Delta = \max_{i \in \{1,...,m\}} deg(S_i)$ and $f = \max_{i \in \{1,...,n\}} deg(x_i)$ are small, one can obtain H_Δ-approximation and f-approximation respectively.

We now show how to formulate set cover as an ILP, reduce it into a LP, and how to round the solutions to yield an approximation to the original set cover instance. Consider the following ILP:

ILP$_{\textbf{Set cover}}$

$$\begin{array}{llll} \text{minimize} & \sum_{i=1}^{m} y_i \cdot c(S_i) & & \triangleleft \text{ Cost of chosen set cover} \\ \text{subject to} & \sum_{i=1, e_j \in S_i}^{m} y_i \geq 1 & \forall j \in \{1, \dots, n\} & \triangleleft \text{ Every item } e_j \text{ is covered by some set} \\ & y_i \in \{0, 1\} & \forall i \in \{1, \dots, m\} & \triangleleft \text{ Indicator variable for whether set } S_i \text{ is chosen} \end{array}$$

Upon solving ILP$_{\text{Set cover}}$, the set $\{S_i \in \{1, \dots, n\} : y_i^* = 1\}$ is the optimal solution for a given set cover instance. However, as solving ILPs are **NP**-hard, we consider relaxing the integral constraint by replacing binary $y_i \in \{0, 1\}$ variables by real-valued/fractional $x_i \in [0, 1]$. Such a relaxation will yield the corresponding LP:

LP$_{\textbf{Set cover}}$

$$\begin{array}{llll} \text{minimize} & \sum_{i=1}^{m} x_i \cdot c(S_i) & & \triangleleft \text{ Cost of chosen fractional set cover} \\ \text{subject to} & \sum_{i=1, e_j \in S_i}^{m} x_i \geq 1 & \forall j \in \{1, \dots, n\} & \triangleleft \text{ Every item } e_j \text{ is fractionally covered} \\ & 0 \leq x_i \leq 1 & \forall i \in \{1, \dots, m\} & \triangleleft \text{ Relaxed indicator variables} \end{array}$$

Since LPs can be solved in polynomial time, we can find the optimal fractional solution to LP$_{\text{Set cover}}$ in polynomial time.

Observation As the set of solutions of ILP$_{\text{Set cover}}$ is a subset of LP$_{\text{Set cover}}$, $obj(\boldsymbol{x}^*) \leq obj(\boldsymbol{y}^*)$.

Example The corresponding ILP for the above set cover instance is:

$$\begin{array}{lll} \text{minimize} & 2y_1 + 4y_2 + 8y_3 + 16y_4 & \\ \text{subject to} & y_1 + y_4 \geq 1 & \triangleleft \text{ Sets covering } e_1 \\ & y_1 + y_3 \geq 1 & \triangleleft \text{ Sets covering } e_2 \\ & y_3 \geq 1 & \triangleleft \text{ Sets covering } e_3 \\ & y_2 + y_4 \geq 1 & \triangleleft \text{ Sets covering } e_4 \\ & y_1 + y_4 \geq 1 & \triangleleft \text{ Sets covering } e_5 \\ & \forall i \in \{1, \dots, 4\}, y_i \in \{0, 1\} & \end{array}$$

After relaxing:

$$\begin{array}{lll} \text{minimize} & 2x_1 + 4x_2 + 8x_3 + 16x_4 & \\ \text{subject to} & x_1 + x_4 \geq 1 & \\ & x_1 + x_3 \geq 1 & \\ & x_3 \qquad \geq 1 & \\ & x_2 + x_4 \geq 1 & \\ & x_1 + x_4 \geq 1 & \\ & \forall i \in \{1, \dots, 4\}, 0 \leq x_i \leq 1 & \triangleleft \text{ Relaxed indicator variables} \end{array}$$

Solving it using a LP solver[1] yields: $x_1 = 1, x_2 = 1, x_3 = 1, x_4 = 0$. Since the solved $\boldsymbol{x}^*$ are integral, $\boldsymbol{x}^*$ is also the optimal solution for the original ILP. In general, the solved $\boldsymbol{x}^*$ may be fractional, which does not immediately yield a set selection.

In the following, we describe two ways to round the fractional assignments $\boldsymbol{x}^*$ into binary variables $\boldsymbol{y}$ so that we can interpret them as proper set selections.

2.1 (Deterministic) Rounding for small f

Recall that the vertex cover problem is a special case when $f = 2$. We round $\boldsymbol{x}^*$ as follows:

$$\forall i \in \{1, \dots, m\}, \text{set } y_i = \begin{cases} 1 & \text{if } x_i^* \geq \frac{1}{f} \\ 0 & \text{else} \end{cases}$$

Theorem 5. *The rounded $\boldsymbol{y}$ is a feasible solution to ILP$_{Set\ cover}$.*

Proof. Since $\boldsymbol{x}^*$ is a feasible (not to mention, optimal) solution for LP$_{\text{Set cover}}$, in each constraint, at least one $x_i^* \geq \frac{1}{f}$. Hence, every element is covered by some set y_i in the rounding. □

Theorem 6. *The rounded $\boldsymbol{y}$ is a f-approximation to ILP$_{Set\ cover}$. That is, $obj(\boldsymbol{y}) \leq f \cdot obj(\boldsymbol{y}^*)$.*

Proof. By the rounding, $y_i \leq f \cdot x_i^*, \forall i \in \{1, \dots, m\}$. Therefore, $obj(\boldsymbol{y}) \leq f \cdot obj(\boldsymbol{x}^*) \leq f \cdot obj(\boldsymbol{y}^*)$. □

2.2 (Randomized) Rounding for general f

If f is large, a f-approximation in the previous subsection is unsatisfactory. By introducing randomness in the rounding process, we show that one can obtain a $\ln(n)$-approximation (in expectation) with arbitrarily high probability through probability amplification.

Consider the following rounding procedure:

1. Interpret each fractional solution x_i^* as probability for picking S_i. That is, $\Pr(y_i = 1) = x_i^*$.
2. For each i, independently set y_i to 1 with probability x_i^*.

Theorem 7. $\mathbb{E}(obj(\boldsymbol{y})) = obj(\boldsymbol{x}^*)$

Proof.

$$\begin{array}{rcll} \mathbb{E}(obj(\boldsymbol{y})) & = & \mathbb{E}(\sum_{i=1}^m y_i \cdot c(S_i)) & \\ & = & \sum_{i=1}^m \mathbb{E}(y_i) \cdot c(S_i) & \text{By linearity of expectation} \\ & = & \sum_{i=1}^m \Pr(y_i = 1) \cdot c(S_i) & \text{Since each } y_i \text{ is an indicator variable} \\ & = & \sum_{i=1}^m x_i^* \cdot c(S_i) & \text{Since } \Pr(y_i = 1) = x_i^* \\ & = & obj(\boldsymbol{x}^*) & \end{array}$$

□

Although we expect the rounded selection to yield an objective cost that is close to the optimum of the LP (which may be even better than the optimal of the ILP), we need to consider whether all constraints are satisfied.

Theorem 8. *For any $j \in \{1, \dots, n\}$, item e_j is not covered with probability $\leq e^{-1}$.*

[1]Using Microsoft Excel. See tutorial: http://faculty.sfasu.edu/fisherwarre/lp_solver.html
Or, use an online LP solver such as: http://online-optimizer.appspot.com/?model=builtin:default.mod

Proof. For any $j \in \{1, \ldots, n\}$,

$$\Pr[\text{Item } e_j \text{ not covered}] = \Pr[\sum_{i=1, e_j \in S_i}^{m} y_i = 0] = \Pi_{i=1, e_j \in S_i}^{m} (1 - x_i^*)$$
$$\leq \Pi_{i=1, e_j \in S_i}^{m} e^{-x_i^*} = e^{-\sum_{i=1, e_j \in S_i}^{m} -x_i^*} \leq e^{-1}$$

The first inequality is because $(1 - x) \leq e^x$ and the last inequality holds because the $\boldsymbol{x}^*$ satisfies the j^{th} constraint in the LP that $\sum_{i=1, e_j \in S_i}^{m} x_i^* \geq 1$. □

Since $e^{-1} \approx 0.37$, we would expect the rounded $\boldsymbol{y}$ not to cover several items. However, one can amplify the success probability by considering independent roundings and taking the union (See Algorithm 1).

Algorithm 1 ApproxSetCoverILP$(\mathcal{U}, \mathcal{S}, c)$

```
ILP_{Set cover} ← Construct ILP of problem instance
LP_{Set cover} ← Relax integral constraints on indicator variables x to y
x* ← Solve LP_{Set cover}
T ← ∅                                                   ▷ Selected subset of S
for k · ln(n) times (for any constant k > 1) do
    for i ∈ {1, ..., m} do
        y_i ← Set to 1 with probability x*_i
        if y_i = 1 then
            T ← T ∪ {S_i}                               ▷ Add to selected sets T
        end if
    end for
end for
return T
```

Similar to Theorem 6, we can see that $\mathbb{E}(obj(T)) \leq (k \cdot \ln(n)) \cdot obj(\boldsymbol{y}^*)$. Furthermore, Markov's inequality tells us that the probability of $obj(T)$ being z times larger than its expectation is at most $\frac{1}{z}$.

Theorem 9. ApproxSetCoverILP *gives a valid set cover with probability* $\geq 1 - n^{1-k}$.

Proof. $\forall j \in \{1, \ldots, n\}$,

$$\Pr[\text{Item } e_j \text{ not covered by } T] = \Pr[e_j \text{ not covered by } \textit{all } k \ln(n) \text{ roundings}] \leq (e^{-1})^{k \ln(n)} = n^{-k}$$

Taking union bound over all n items, $\Pr[T \text{ is not a valid set cover}] \leq \sum_{i=1}^{n} n^{-k} = n^{1-k}$. So, T covers all n items with probability $\geq 1 - n^{1-k}$. □

Note that the success probability of $1 - n^{1-k}$ can be further amplified by taking several (polynomial number of) *independent* samples ApproxSetCoverILP, then returning the lowest cost valid set cover sampled. With z samples, $\Pr[\textit{All} \text{ repetitions fail}] \leq n^{z(1-k)}$, so we succeed with probability $\geq 1 - n^{z(1-k)}$.

3 Minimizing congestion in multi-commodity routing

A multi-commodity routing (MCR) problem involves routing multiple (s_i, t_i) flows across a network with the goal of minimizing congestion, where congestion is defined as the largest ratio of flow over capacity of any edge in the network. In this section, we discuss two variants of the multi-commodity routing problem. In the first variant (special case), we are given the set of possible paths $\mathcal{P}_i$ for each (s_i, t_i) source-target pairs. In the second variant (general case), we are given only the network. In both cases, [RT87] showed that one can obtain an approximation of $\mathcal{O}(\frac{\log(n)}{\log\log(n)})$ with high probability.

Definition 10 (Multi-commodity routing problem)**.** *Consider a directed graph $G = (V, E)$ where each edge $e = (u, v) \in E$ has a capacity $c(u, v)$. The in-set/out-set of a vertex v is denoted as $in(v) = \{(u, v) \in E : u \in V\}$ and $out(v) = \{(v, u) \in E : u \in V\}$ respectively. Given k triplets (s_i, t_i, d_i), where $s_i \in V$ is the source, $t_i \in V$ is the target, and $d_i \geq 0$ is the demand for the i^{th} commodity respectively, denote $f(e, i)$ as the fraction of d_i that is flowing through edge e. The task is to minimize the congestion parameter λ by finding paths p_i for each $i \in [k]$, such that:*

(i) *(Valid sources):* $\sum_{e \in out(s_i)} f(e,i) - \sum_{e \in in(s_i)} f(e,i) = 1, \forall i \in [k]$

(ii) *(Valid sinks):* $\sum_{e \in in(t_i)} f(e,i) - \sum_{e \in out(t_i)} f(e,i) = 1, \forall i \in [k]$

(iii) *(Flow conservation):* $\sum_{e \in out(v)} f(e,i) - \sum_{e \in in(v)} f(e,i) = 0, \forall e \in E, \forall i \in [k], \forall v \in V \setminus \cup_{i=1}^{k} \{s_i \cup t_i\}$

(iv) *(Single path): All demand for commodity i passes through a single path p_i (no repeated vertices).*

(v) *(Congestion factor):* $\forall e \in E, \sum_{i=1}^{k} d_i \mathbb{1}_{e \in p_i} \leq \lambda \cdot c(e)$*, where indicator* $\mathbb{1}_{e \in p_i} = 1 \iff e \in p_i$.

(vi) *(Minimum congestion): λ is minimized.*

Example Consider the following flow network with $k = 3$ commodities with edge capacities as labelled:

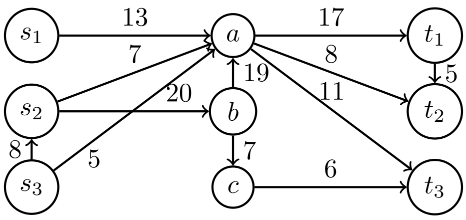

For demands $d_1 = d_2 = d_3 = 10$, there exists a flow assignment such that the total demands flowing on each edge is below its capacity:

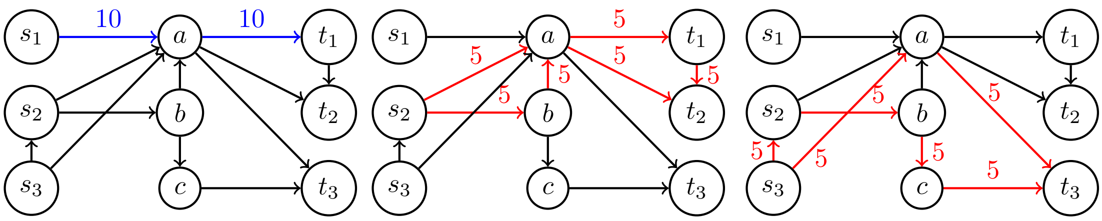

Although the assignment attains congestion $\lambda = 1$ (due to edge (s_3, a)), the path assignments for commodities 2 and 3 violate the property of "single path". Forcing all demand of each commodity to flow through a single path, we have a minimum congestion of $\lambda = 1.25$ (due to edges (s_3, s_2) and (a, t_2)):

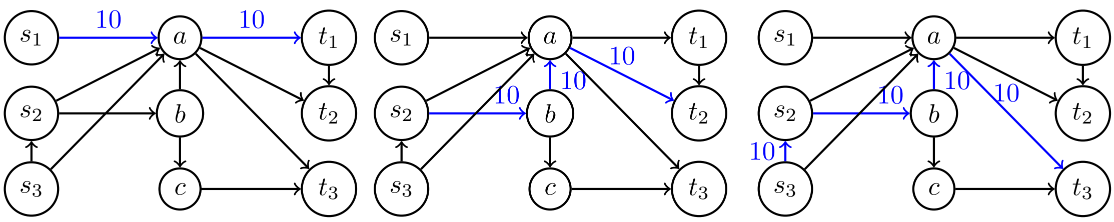

3.1 Special case: Given sets of $s_i - t_i$ paths $\mathcal{P}_i$

For each commodity $i \in [k]$, we are to select a path p_i from a given set of valid paths $\mathcal{P}_i$, where each edge in all paths in $\mathcal{P}_i$ have capacities $\geq d_i$. Because we intend to pick a single path for each commodity to send *all* demands through, constraints (i)-(iii) of MCR are fulfilled trivially. Using $y_{i,p}$ as indicator variables whether path $p \in \mathcal{P}_i$ is chosen, we can model the following ILP:

ILP$_{\textbf{MCR-Given-Paths}}$

minimize	λ		$\triangleleft$ Congestion parameter λ
subject to	$\sum_{i=1}^{k} d_i \sum_{p \in \mathcal{P}_i, e \in p} y_{i,p} \leq \lambda \cdot c(e)$	$\forall e \in E$	$\triangleleft$ Congestion factor relative to selected paths
	$\sum_{p \in \mathcal{P}_i} y_{i,p} = 1$	$\forall i \in [k]$	$\triangleleft$ Exactly one path chosen from each $\mathcal{P}_i$
	$y_{i,p} \in \{0, 1\}$	$\forall i \in [k], p \in \mathcal{P}_i$	$\triangleleft$ Indicator variable for path $p \in \mathcal{P}_i$

Relax the integral constraint on $y_{i,p}$ to $x_{i,p} \in [0,1]$ and solve the corresponding LP. Let us denote $\lambda^* = obj(\text{LP}_{\text{MCR-Given-Paths}})$ and $\boldsymbol{x}^*$ as a fractional path selections that achieves λ^*. Then, to obtain a valid path selection, for each $i \in [k]$, pick path $p \in \mathcal{P}_i$ with weighted probability $\frac{x^*_{i,p}}{\sum_{p \in \mathcal{P}_i} x^*_{i,p}}$.

Remark 1 For a fixed i, a path is selected *exclusively* (only one!) (cf. set cover's roundings where we may pick multiple sets for an item).

Remark 2 The weighted sampling is independent across different commodities. That is, the choice of path amongst $\mathcal{P}_i$ does not influence the choice of path amongst $\mathcal{P}_j$ for $i \neq j$.

Theorem 11. $\mathbb{E}(obj(\boldsymbol{y})) \leq obj(\boldsymbol{x}^*)$

Proof. Fix an arbitrary edge $e \in E$. For each commodity i, define an indicator variable $Y_{e,i}$ whether edge e is part of the chosen path for commodity i. By randomized rounding, $\Pr[Y_{e,i} = 1] = \sum_{p \in \mathcal{P}_i, e \in p} x_{i,p}$. Denoting $Y_e = \sum_{i=1}^{k} d_i \cdot Y_{e,i}$ as the total demand on edge e in the *all* k chosen paths:

$$\begin{array}{rcll} \mathbb{E}(Y_e) & = & \mathbb{E}(\sum_{i=1}^{k} d_i \cdot Y_{e,i}) & \\ & = & \sum_{i=1}^{k} d_i \cdot \mathbb{E}(Y_{e,i}) & \text{By linearity of expectation} \\ & = & \sum_{i=1}^{k} d_i \sum_{p \in \mathcal{P}_i, e \in p} x_{i,p} & \text{Since } \Pr[Y_{e,i} = 1] = \sum_{p \in \mathcal{P}_i, e \in p} x_{i,p} \\ & \leq & \lambda^* \cdot c(e) & \text{By MCR constraint and optimality of the solved LP} \end{array}$$

Since the above analysis holds for all edges, $\frac{\mathbb{E}(Y_e)}{c(e)} \leq \lambda^*$ for any edge e, so $\mathbb{E}(obj(\boldsymbol{y})) \leq \lambda^* = obj(\boldsymbol{x}^*)$. □

Theorem 12. $\Pr[obj(\boldsymbol{y}) \geq \frac{2c \log n}{\log \log n} \max\{1, \lambda^*\}] \leq \frac{1}{n^c}$

Proof. Since each edge in all paths in $\mathcal{P}_i$ have capacities $\geq d_i$, $\frac{Y_e}{c(e)} \in [0,1]$. Recall $\mathbb{E}(obj(\boldsymbol{y})) \leq \lambda^*$. Then, apply[2] the tight form of Chernoff bounds on $\{\frac{Y_e}{c(e)}\}_{e \in E}$ with $(1+e) = \frac{2 \log n}{\log \log n}$. □

3.2 General: Given only a network

In the general case, we may not be given path sets $\mathcal{P}_i$ and there may be exponentially many $s_i - t_i$ paths in the network. However, we show that one can still formulate an ILP and round it (slightly differently) to yield the same approximation factor. Consider the following:

ILP$_{\text{MCR-Given-Network}}$ t

$$\begin{array}{llll} \text{minimize} & \lambda & & \triangleleft \text{ Congestion parameter } \lambda \\ \text{subject to} & \sum_{e \in out(s_i)} f(e,i) - \sum_{e \in in(s_i)} f(e,i) = 1 & \forall i \in [k] & \triangleleft \text{ Valid sources} \\ & \sum_{e \in in(t_i)} f(e,i) - \sum_{e \in out(t_i)} f(e,i) = 1 & \forall i \in [k] & \triangleleft \text{ Valid sinks} \\ & \sum_{e \in out(v)} f(e,i) - \sum_{e \in in(v)} f(e,i) = 0 & \forall e \in E, & \triangleleft \text{ Flow conservation} \\ & & \forall i \in [k], & \\ & & \forall v \in V \setminus \cup_{i=1}^{k} \{s_i \cup t_i\} & \end{array}$$

Relax the integral constraint on $y_{i,p}$ to $x_{i,p} \in [0,1]$ and solve the corresponding LP. To extract the path candidates $\mathcal{P}_i$ for each commodity, perform flow decomposition[3]. For each extracted path p_i for commodity i, treat the minimum $\min_{e \in p_i} f(e,i)$ on the path as the selection probability (as per $x_{e,i}$ in the previous section). By selecting the path p_i with probability $\min_{e \in p_i} f(e,i)$, one can show by similar arguments as before that $\mathbb{E}(obj(\boldsymbol{y})) \leq obj(\boldsymbol{x}^*) \leq obj(\boldsymbol{y}^*)$.

References

[RT87] Prabhakar Raghavan and Clark D Tompson. Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7(4):365–374, 1987.

[2]See Corollary 2 of https://courses.engr.illinois.edu/cs598csc/sp2011/Lectures/lecture_9.pdf for details.

[3]See https://www.youtube.com/watch?v=zgutyzA9JM4&t=1020s (17:00 to 29:50) for a recap on flow decomposition.