CMPUT 675: Approximation Algorithms	**Winter 2005**

Lecture 14: Mar 4

Lecturer: Mohammad R. Salavatipour *Scribe: Frederick Vizeacoumar*

Last lecture, we defined the problem of Packing edge-disjoint Steiner trees. This lecture we will be discusing about the same problem but in directed graphs. As we will see the approximation factor we obtain is much worse than the one for undirected case (wehre we had an $O(1)$-approximation). The main reason behind it is a hardness result for this version of the problem that will be presented next lecture.

14.1 Packing Edge Capacitated Directed Steiner Trees (PED)

Packing edge capacitated directed Steiner tree probelm is defined as follows:

Input: Given a directed graph $G = (V, E)$ with a set $T \subseteq V$ of terminals which includes a vertex r called root. We also have capacity on the edge $c_e \in Z^*$on every edge $e \in E$.

Goal: Find the maximum number of directed Steiner trees covering all the terminals, such that at most c_e trees contain e, for every $e \in E$.

In general this problem is NP-hard. In fact, it can be shown that even finding two edge-disjoint directed Steiner trees when $|T| = 3$ is NP-hard. So the question is how good of an approximation algrorithm we can find for this problem. In this lecture, we will see an $O(m^{\frac{1}{2}+\epsilon})$-approximation where $m = |E|$.

14.1.1 The Primal-Dual formulation

We can formulate the problem as an IP in the same manner we did for the undirected version of the problem. Let $\mathcal{F}$ be the set of all directed Steiner Trees of G. We will have a 0/1 indicator variable x_F for every directed Steiner tree $F \in \mathcal{F}$. First, we obtain an integer programming formulation of the problem and then derive its LP-relaxations.

$$\begin{array}{rll} maximize & \sum_{F \in \mathcal{F}} & x_F \\ subject\ to\ \forall e \in E, & \sum_{\mathcal{F}: e \in \mathcal{F}} & x_F \leq C_e \\ & & x_F \in \{0, 1\} \end{array}$$

The LP-relaxation is obtained by replacing the constraint $x_F = \{0, 1\}$ with $x_F \geq 0$.

$$\begin{array}{rll} maximize & \sum & x_F \\ subject\ to\ \forall e \in E, & \sum_{\mathcal{F}: e \in \mathcal{F}} & x_F \leq C_e \\ & & x_F \geq 0 \end{array}$$

Next, we derive the dual LP:

$$\begin{array}{rcl} minimize & \sum & y_e \\ subject\ to\ \forall F \in \mathcal{F}, & \sum_{e \in \mathcal{F}} & y_e \geq 1 \\ & & y_e \geq 0 \end{array}$$

The separation oracle for this LP is the problem of finding minimum edge-weighted directed Steiner tree. We use Theorem 13.8 from the last lecture can be adpated to prove the following:

Theorem 14.1 *There is an α-approximation for fractional PED if and only if there is an α-approximation for minimum weighted directed Steiner tree problem.*

The problem is that minimum directed Steiner tree problem is a hard to approximate within a factor of $\Omega(\log^{2-\epsilon} n)$ for any $\epsilon > 0$. However, there is a polynomial time $O(n^\epsilon)$- approximation for this problem, for any fixed $\epsilon > 0$.

Theorem 14.2 (Charikar, Chekuri, Cheung, Dai, Goel, Guha, M. Li'99) *For any fixed $\epsilon > 0$, there is a polytime $O(n^\epsilon)$-approximation for minimum directed Steiner tree problem.*

Therefore, using this theorem and Theorem 14.1 we obtain:

Corollary 14.3 *There is an $O(n^\epsilon)$-approximation algorithm for fractional PED.*

A feature of algorithm of Theorem 14.1 is that it obtains an approximate solution to the LP in which only a polynomial x_F's are non-zero. We are going to do randomized rounding on this LP to find an integral solution with the desired approximation ratio. The following lemma is the captures the heart of the proof:

Lemma 14.4 (Main Lemma) *Let I be an instance of PED and I_f be the corresponding fractional problem (the primal LP) and and let φ^* be the (objective) value of a (not necessarily optimal) feasible solution $\{x_F^* : F \in \mathcal{F}\}$ to I_f such that the number of non-zero x_F^*'s is polynomially bounded and each $x_F^* < 1$. Then, we can find in polynomial time, a solution to I with value at least $O(\frac{\varphi^*}{\sqrt{m}})$.*

Proof: The idea is used randomized rounding and round every x_F^* to 1 (i.e. pick tree F) with probability $\frac{x_F^*}{\lambda}$ for some λ to be defined. This value of λ will help in adjusting so that there is no violation of constraints on edges.

We will use the following simple and well-known deviation bound.

Lemma 14.5 *(Chernoff-Hoeffding Bounds) Let $X_1, X_2, \ldots, X_q$ be a set of q independent random variables with $X_i \in \{0, 1\}$ and let $X = \sum_{i=1}^q X_i$. Then for $0 \leq \delta < 1$:*

$$\Pr[X < (1-\delta)\mathrm{E}[X]] \leq e^{-\delta^2 \mathrm{E}[X]/2}.$$

Also, the following simple lemma (whose proof is straightforward) will help us.

Lemma 14.6 *Assume that $A = \{a_1, \ldots, a_n\}$ is a set of n non-negative reals and let $\mathcal{A}_k$ be the set of all subsets of size k of A. If $\sum_{i=1}^{n} a_i \le Q$, then $\sum_{\{a_{i_1},\ldots,a_{i_k}\}\in\mathcal{A}_k} a_{i_1}a_{i_2}\ldots a_{i_k} \le \binom{n}{k}(Q/n)^k$.*

Let X_F be the random variable that is 1 if we pick tree F and 0 otherwise. Then for $X = \sum_{F\in\mathcal{F}} X_F$ (i.e. the total number of trees picked by the algorithm), we have:

$$\mathrm{E}[X] = \sum_{F\in\mathcal{F}} \Pr[X_F = 1] = \sum_{F\in\mathcal{F}} \frac{x_F^*}{\lambda} = \frac{\varphi^*}{\lambda}.$$

For every edge $e \in E$, define the bad event A_e to be the event that the capacity constraint of e is violated, i.e. more than c_e trees containing e are picked. Our goal is to show that with some positive probability, none of these bad events happen (i.e. all $\overline{A}_e$'s hold) *and* that the total number of trees picked is not too small. We want to find a good upper bound for $\Pr[A_e]$. For every edge e, denote the number of trees F with $x_F^* > 0$ that contain e by ψ_e. By this definition:

$$\Pr[A_e] \le \sum \prod_{i=1}^{c_e+1} x_{T_{a_i}}^*/\lambda,$$

where the summation is over all subsets $\{F_{a_1}, \ldots, F_{a_{c_e+1}}\}$ of size $c_e + 1$ of trees with $x_{F_{a_i}}^* > 0$ that contain edge e. Therefore, using Lemma 14.6:

$$\Pr[A_e] \le \binom{\psi_e}{c_e+1}\left(\frac{c_e}{\lambda\psi_e}\right)^{c_e+1} \le \left(\frac{e\psi_e}{c_e+1}\right)^{c_e+1}\left(\frac{c_e}{\lambda\psi_e}\right)^{c_e+1} \le \frac{e^2}{\lambda^2},$$

where we have used the fact $\binom{n}{k} \le (\frac{en}{k})^k$ for the second inequality. It is intuitively clear that if $\overline{A_e}$ holds then it does not increase the probability of any other $A_{e'}$. In other words, events $\overline{A_e}$ are "positively correlated". Therefore:

$$\Pr[\bigwedge_{e\in E} \overline{A}_e] \ge \prod_{e\in E} \Pr[\overline{A}_e] \ge (1 - \frac{e^2}{\lambda^2})^m.$$

So, the probability that at least one event A_e happens is at most $(1 - \frac{e^2}{\lambda^2})^m$. Also, by Lemma 14.5, for $0 \le \delta < 1$: $\Pr[X < (1-\delta)\mathrm{E}[X]] \le e^{-\delta^2\varphi^*/2\lambda}$. Thus:

$$\Pr[(X < (1-\delta)\mathrm{E}[X]) \vee (\exists e \in E : A_e)] \le e^{-\delta^2\varphi^*/2\lambda} + 1 - (1 - e^2/\lambda^2)^m.$$

If we can show that for suitable δ and λ: $(1 - e^2/\lambda^2)^m > e^{-\delta^2\varphi^*/2\lambda}$ then we can efficiently find a selection of trees such that $X \ge (1-\delta)\varphi^*/\lambda$ and that no edge constraint is violated, using the method of conditional probability. So we need to show that $(1 - \frac{e^2}{\lambda^2})^m > e^{-\frac{\delta^2\varphi}{2\lambda}}$.

We can assume that $\phi^* \ge 20.e.\sqrt{m}$. Otherwise, we can find just one Steiner tree and return it. This is clearly within a factor $O(\sqrt{m})$ of the solution φ^*.

We choose $\delta = \frac{1}{2}$ and $\lambda = e\sqrt{m}$. Then

$$(1 - \frac{e^2}{\lambda^2})^m = (1 - \frac{e^2}{e^2.m})^m = (1 - \frac{1}{m})^m \ge \frac{1}{4}.$$

On the other hand

$$e^{-\frac{\delta^2\varphi}{2\lambda}} = e^{-\frac{\varphi^*}{8.e.\sqrt{m}}} \le e^{-2.5}.$$

Therefore $(1 - \frac{e^2}{\lambda^2})^m > e^{-\frac{\delta^2 \varphi}{2\lambda}}$. Thus with positive probability, X is at least $\frac{1}{2}$ of its expected value which is $\frac{\varphi^*}{20e\sqrt{m}}$ and no bad event A_e happens. ■.

The overall algorithm for the problem of packing edge-capacitated directed Steiner trees will be as follows. First we use Corollary 14.3 to obtain an approximate fractional solution I_f with objective value φ^* such that $\varphi^* \geq c\varphi_f / m^{\frac{\epsilon}{2}}$ for some constant c and the given $\epsilon > 0$. Then we apply a preprocessing step to the fractional solution. For every Steiner tree F with $x_F \geq 1$ we "take out" $\lfloor x_F \rfloor$ copies of that tree and put it in the final integral solution, we decrease x_F by $\lfloor x_F \rfloor$, and also we update the capacities of the edges accordingly. This decomposes x into a (multi)set of Steiner trees $\mathcal{F}_1$ and a fractional part (with each entry $x_F < 1$). We will "round" the fractional part x to an integer solution (using Lemma 14.4). For the rest of the proof we may assume that the fractional solution x has each entry < 1, since the other case reduces to this one.

Note that the approximate fractional solution x contains only a polynomial number of Steiner trees with non-zero fractional values. If we substitute φ^* in Lemma 14.4 we obtain an approximation algorithm that finds a set $\mathcal{F}'$ of directed Steiner trees such that $\mathcal{F}'$ has the required size. The total approximation factor of this algorithm is $O(m^{\frac{1}{2}+\epsilon})$.