

Network Design and Game Theory

Spring 2008

Lecture 6

Guest Lecturer: Aaron Archer

Instructor: Mohammad T. Hajiaghayi

Scribe: Fengming Wang

March 3, 2008

1 Overview

We study the Primal-dual, Lagrangian relaxation paradigm of approximation algorithms for the NP-hard problems: Steiner tree problem (both the standard and prize-collecting version) and k-Minimum spanning tree problem. Its corresponding reference is [1].

2 General Steiner Tree

2.1 Definition

Given the graph $G = (V, E)$, the cost function $C : E \rightarrow \mathbb{R}^+$ and a subset of terminal nodes $\mathcal{T} \subseteq V$, Find $T \subseteq E$ that spans \mathcal{T} of the total minimum cost. For example, let $\mathcal{T} = \{a, b, c\}$ in Figure 1. The extra node c is called a *steiner node*. It is easy to see that the corresponding steiner tree is the star centered at c of total weight 3 while any other spanning tree without using c costs 4 units.

Without loss of generality, we assume that G is always a complete graph, moreover, C obeys triangle inequalities, since the cost function could be transformed into a metric by taking the shortest path as the distance between any two nodes.

2.2 2-Approximation via MST

Compute the minimum spanning tree MST on the subgra

ph induced by \mathcal{T} , we claim that this 2-approximates the optimal steiner tree OPT. The proof consists of two steps. Since OPT is a tree, one could do the depth-first search on it from an arbitrary node, which constructs a Eulerian tour

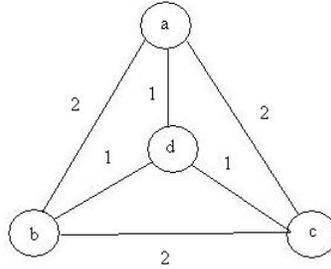


Figure 1: Steiner Tree

on the nodes spanned by OPT. Note the cost is doubled for the tour compared to OPT. Then revise the tree by skipping all of the steiner nodes and joining consecutive terminal nodes according to the tour, which returns a spanning tree on \mathcal{T} . By triangle inequalities on \mathcal{C} , the cost of this spanning tree is no larger than that of the tour, therefore, MST is a valid 2-approximation.

2.3 2-Approximation via LP

First we fix some notations.

- For any $S \subseteq V$, $\delta(S) = \{e_{u,v} \in E : u \in S, v \notin S\}$.
- $\mathcal{C} = \{S \subseteq V : S \cap \mathcal{T} \neq \emptyset, \neq \mathcal{T}\}$.
- For any $F \subseteq E$, $C(F) = \sum_{e \in F} C_e$.

The following observation is easy to verify and crucial for our algorithm.

Claim 1 $T \subseteq E$ spans \mathcal{T} if and only if $\forall S \in \mathcal{C}, T \cap \delta(S) \neq \emptyset$.

Proof: Apply max-flow/min-cut theorem. ■

We interpret this idea in terms of integer programming and define the following indicator variables for each edge: $\forall e \in E$, $\chi_e = 1$ if e is in our solution tree, 0 otherwise.

$$\begin{aligned} \min \quad & \sum_{e \in E} C_e \chi_e \\ \text{s.t.} \quad & \sum_{e \in \delta(S)} \chi_e \geq 1 \quad \forall S \subseteq \mathcal{C} \end{aligned}$$

Relax the condition that $\chi_e \in \{0, 1\}$ to make them nonnegative real variables. Now we obtain a linear program and then take a close look at its dual.

$$\begin{aligned} \max \quad & \sum_{S \in \mathcal{C}} y_S \\ \text{s.t.} \quad & \sum_{S: e \in \delta(S)} y_S \leq C_e \quad \forall e \in E \\ & y_S \geq 0 \end{aligned}$$

Apparently both the primal and the dual programs are feasible. So the complementary slackness constraints are satisfied for the pair of optimal solutions.

1. $\forall e \in E, \chi_e > 0 \Rightarrow \sum_{S:e \in \delta(S)} y_S = C_e.$
2. $\forall S \in \mathcal{C}, y_S > 0 \Rightarrow \sum_{e \in \delta(S)} \chi_e = 1.$

These two implications have economic interpretations. The first postcondition says that for every edge, the primal buys exactly what the dual pays for it while the second one means that each cut pays exactly for one edge. Therefore, we get a sequence of equalities as follows:

$$\sum_{e \in E} C_e \chi_e = \sum_{e \in E} \chi_e \sum_{S:e \in \delta(S)} y_S = \sum_{S \in \mathcal{C}} y_S \underbrace{\sum_{e \in \delta(S)} \chi_e}_{= 1 \text{ whenever } y_S \neq 0} = \sum_{S \in \mathcal{C}} y_S$$

So any answer for the dual is always a lower bound for the primal. Next we would devise an algorithm which promises that the output never exceeds twice one specific dual solution.

Agrawal-Klein-Ravi (AKR) Algorithm For Steiner Tree

$\mathcal{F} = \phi.$
repeat
 repeat
 Grow at unit rate y_S for every component induced by \mathcal{F} that contains a terminal.
 until Some edge e goes tight (is paid for by the cut it crosses)[Break ties arbitrarily].
 $\mathcal{F} \leftarrow \mathcal{F} + e$ [This will never create a cycle because edge (u, v) stops being paid off once $u \& v$ are in the same component].
 Add edges without creating cycles and ties are broken arbitrarily.
until \mathcal{F} spans $\mathcal{T}.$
 Pruning: iteratively delete edges to steiner nodes that are leaves.

Consider the graph in Figure 2, where the terminal nodes are a, b, e and f . The first several steps of the algorithm are performed as follows:

1. Grow $y_{\{a\}}, y_{\{b\}}, y_{\{e\}}, y_{\{f\}}$ at rate 1. Merge ac, bd, eg at time 1.
2. Grow $y_{\{a,c\}}, y_{\{b,d\}}, y_{\{e,g\}}, y_{\{f\}}$. Merge $aceg, bdf$ at time 2.
3. ...

Remark: The y_S values produced by the algorithm always satisfies the constraints in the dual program. So they form a feasible solution.

Remark: The ratio between the optimal solution for the linear program and the output of the algorithm could be as large as 2. It is not hard to see that a simple cycle yields this gap.

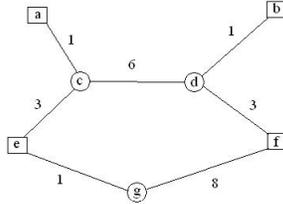


Figure 2: Example

Theorem 1 \mathcal{F} output by the algorithm satisfies $C(F) \leq 2(1 - \frac{1}{n})\sum y_S$, where the y_S values untouched by the algorithm are implicitly set to 0.

The proof does the backward analysis and utilizes the fact that in any tree, the total degree never exceeds twice the number of nodes. For more details, we refer the reader to chapter 22 of the book "Approximation Algorithms".

3 Prize-collecting Steiner Tree (PCST)

This is similar to the problem previously studied. In addition, there two more parameters, $\Pi : V \rightarrow \mathbb{R}^+$, the prize function which associates a nonnegative value to each node once reached, and $r \in V$, the root from which one tries to build the steiner tree. The goal is changed as well: we want to construct a tree T rooted at r such that the gain (the difference between the sum of prizes from the connected nodes and the total cost to realize the tree) is maximum. Formally the following quantity needs to be maximized:

$$\Pi(T) - C(T) \tag{1}$$

which is equivalent to maximize

$$\Pi(T) - C(T) - \Pi(V) \tag{2}$$

or

$$-C(T) - \Pi(V \setminus T) \tag{3}$$

Remark 1 This is the same as minimizing $C(T) + \Pi(V \setminus T)$. This seems like an unnatural objective function. The original motivation for studying it is that it is NP-hard to determine whether the optimum value for 1 is positive, and hence

no bounded approx ratio is possible unless $P = NP$. In contrast, we can get a 2-approx for \mathcal{Z} , as shown below. This algorithm has a "Lagrangian multiplier-preserving" property that turns out to have both theoretical and practical consequences that we will exploit.

For every $U \subseteq V \setminus \{r\}$, z_U equals to 1 if U is exactly the set we don't span, 0 otherwise. Therefore, an easy observation leads to the following integer program.

$$\begin{aligned} \min \quad & \sum_{e \in E} C_e \chi_e + \sum_{U \subseteq V \setminus \{r\}} \Pi(U) z_U \\ \text{s.t.} \quad & \sum_{e \in \delta(S)} \chi_e + \sum_{U \supseteq S} z_U \geq 1 \quad \forall S \subseteq V \setminus \{r\} \\ & \chi_e, z_U \geq 0 \end{aligned}$$

The corresponding dual is

$$\begin{aligned} \max \quad & \sum y_S \\ \text{s.t.} \quad & \sum_{S: e \in \delta(S)} y_S \leq C_e \quad \forall e \in E \\ & \sum_{S \subseteq U} y_S \leq \Pi(U) \quad \forall U \subseteq V \\ & y_S \geq 0 \end{aligned} \tag{4}$$

We have one new implication added into the complementary slackness.

$$z_U > 0 \Rightarrow \sum_{S \subseteq U} y_S = \Pi(U)$$

Goemans-Williamson Algorithm For PCST

$\mathcal{F} = \phi$.
 $\mathcal{A} = \{\{v\} : \Pi_v > 0\}$
 $\mathcal{J} = \{\{v\} : \Pi_v = 0\}$ [This includes $\{r\}$ since we assume $\Pi_r = 0$]
 $\forall v \in V, P_{\{v\}} = \Pi_v$ [We call the P variables potentials.]
repeat
 repeat
 Grow y_S and shrink P_S for all $S \in \mathcal{A}$.
 until Some edge e goes tight or for some $S \in \mathcal{A}$, P_S hits 0. [Ties are broken arbitrarily.]
 if Edge e goes tight **then**
 $\mathcal{F} \leftarrow \mathcal{F} + e$.
 Merge S_1, S_2 that e joins and let $P_{S_1 \cup S_2} = P_{S_1} + P_{S_2}$
 Delete S_1, S_2 from \mathcal{A} (or \mathcal{J}) and add $S_1 \cup S_2$ to \mathcal{A} .
 if $r \in S_1 \cup S_2$ **then**
 Move $S_1 \cup S_2$ from \mathcal{A} to \mathcal{J} .
 end if
 else if For some $S \in \mathcal{A}$, P_S hits 0 **then**
 move S from \mathcal{A} to \mathcal{J}
 end if
until $\mathcal{A} = \phi$ or \mathcal{F} spans the whole graph.
 Pruning: throw away every component that was ever an inactive leaf.

Note 1 *GW should have to implement this using standard priority queue in $O(n^2 \log n)$ time.*

At the end of the algorithm, we have some tree T . Because the nodes we fail to span can be partitioned into components that ran out potential, and this was constructed to be equivalent to constraint 4 being tight for that component, we get

$$\sum_{S \subseteq V \setminus T} y_S = \Pi(V \setminus T)$$

Moreover, the analogous analysis in the last section gives

$$C(T) \leq 2 \sum_{S: T \cap \delta(S) \neq \emptyset} y_S$$

Combine them together and we get

Theorem 2 $C(T) + 2\Pi(V \setminus T) \leq 2\sum y_S$

To have a 2-approximation for PCST, it would have been sufficient to show $C(T) + 2\Pi(V \setminus T) \leq 2\sum y_S$.

The two sides being equal means that the Lagrangian-preserving Property is satisfied. If we multiply the prize function by a parameter λ to find some tree $T(\lambda)$ via GW PCST algorithm, this is very useful which amounts to the following claim.

Claim 2 *Among all trees that capture at least $\lambda\Pi(T(\lambda))$ prizes, the cost of $T(\lambda)$ is no worse than twice that of the optimal.*

Proof: As the previous linear program

$$\begin{aligned} \min \quad & \sum C_e \chi_e \\ \text{s.t.} \quad & \sum_{e \in \delta(S)} \chi_e + \sum_{U \supseteq S} z_U \geq 1 \quad \forall S \subseteq V \setminus r \\ & \sum_{U \subseteq V \setminus r} \Pi(U) z_U \leq G \\ & \chi_e, z_U \geq 0 \end{aligned}$$

Its dual is

$$\begin{aligned} \max \quad & \sum_{S \subseteq V \setminus r} y_S - \lambda G \\ \text{s.t.} \quad & \sum_{S: e \in \delta(S)} y_S \leq C_e \quad \forall e \in E \\ & \sum_{S \subseteq U} y_S \leq \lambda \Pi(U) \quad \forall U \subseteq V \setminus r \\ & \lambda, y_S \geq 0 \end{aligned}$$

Think of G as $\Pi(V \setminus T(\lambda))$. Then similar to the previous analysis, $C(T(\lambda)) \leq 2(\sum y_S - \lambda \Pi(V \setminus T(\lambda)))$. ■

The theorem above gives us a base for the algorithmic paradigm which "turns the λ knob" to search for a better approximation.

4 k-MST problem

Given $G = (V, E)$, the cost function $C : E \rightarrow \mathbb{R}^+$, the root $r \in V$ and $k \in \mathbb{Z}^+$, find the tree T rooted at r spanning at least k nodes of the smallest total edge-cost.

Notation:

- $T(\lambda)$ = tree given by GW PCST run with edge profit λ .
- OPT_k = cost of the optimal k-MST.
- LB_k = a valid lower bound on OPT_k

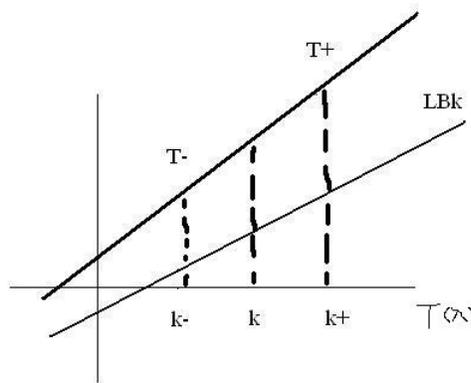


Figure 3: Search

The sketch of the idea is that find a value of λ (using binary search, Meggido parametric search, etc.) such that $T(\lambda^-) \leq k$ and $T(\lambda^+) > k$, where $T(\lambda^-)$ and $T(\lambda^+)$ are 2-approximate k' -MSTs for $k' = |T(\lambda^-)|$ and $k' = T(\lambda^+)$. If $k = T(\lambda^-)$, we are done and have a 2-approximate solution. Otherwise, we must combine the trees $T(\lambda^-)$ and $T(\lambda^+)$ to get our final solution T . If $c(T)$ are exactly the linear interpretation of $c(T(\lambda^-))$ and $c(T(\lambda^+))$, then we would have a 2-approximation. This is essentially how a 2-approximation is obtained, although it is possibly to get a 5 or a 3 much more simply.

5 Real world Application

At AT&T a team that includes Aaron has built a tool that uses the Goemans-Williamson PCST algorithm to help decide where to lay new cables. The tree produced by the GW algorithm is a first cut that is fed into some heuristics to augment the tree into a 2-connected network. We have found this approach to be practical and effective. The multiplier λ can be set to reflect the target payback period for this capital investment.

References

- [1] David S. Johnson, Maria Minkoff, and Steven Phillips. The prize collecting steiner tree problem: theory and practice. In *SODA*, pages 760–769, 2000.