Steiner Tree Approximation via Iterative Randomized Rounding

Nikolidaki Aikaterini

Graduate Program in Logic, Algorithms and Computation

$\mu\Pi\lambda\forall$

Network Algorithms and Complexity

June 18, 2013

Overview

Scope

- The Steiner tree problem is one of the most fundamental NP-hard problems
  (weighted undirected graph, subset of terminal nodes $\Rightarrow$ minimum-cost tree spanning the terminals).
- Applications in:
  * VLSI
  * Optical and Wireless Communication Systems
  * Transportation and Distribution Networks
- A long-standing open problem is whether there is an LP relaxation of Steiner tree with integrality gap smaller than 2.

Scope

- An LP-based approximation algorithm for Steiner tree with an improved approximation factor.
  - * Iterative randomized rounding technique.
  - * An LP relaxation of the problem, which is based on the notion of directed components.
- A solution of cost at most $\ln(4) + \epsilon < 1.39$ times the cost of an optimal Steiner tree.
- The algorithm can be derandomized using the method of limited independence.
- The integrality gap of LP is showed at most 1.55.

Scope

- Based on the directed-component cut relaxation for the Steiner tree problem.

DCR

$$
\begin{aligned}
&min \sum_{C \in C_n} c(C) x_C \\
&s.t. \sum_{C \in \delta_{C_n}{}^{+}(U)} x_C \geq 1, \quad \forall U \subseteq R \backslash \{r\}, U \neq \emptyset \\
&x_c \geq 0 \qquad\qquad\qquad\qquad \forall C \in C_n
\end{aligned}
$$

- The algorithm combines features randomized rounding and iterative rounding.
- Used the Bridge Lemma for the reduction of the cost of the optimal Steiner tree in each iteration

Related work

- A minimum-cost terminal spanning tree is
  - a 2-approximation for the Steiner tree problem [Gilbert and Pollak 1968; Vazirani 2001].
  - the famous $1 + \ln(3)/2 + \epsilon < 1.55$ approximation algorithm [Robins and Zelikovsky 2005].
- All these improvements are based on the notion of k-restricted Steiner tree.
  - A collection of components, with at most k terminals each (k-components), whose union induces a Steiner tree.

Theorem 1 [Borchers and Du 1997]

Let r and s be the nonnegative integers satisfying $k = 2^r + s$ and $s < 2^r$. Then $\rho_k = \frac{(r+1)2^r + s}{r2^r + s} \leq 1 + \frac{1}{\lfloor \log_2 k \rfloor}$

Related work

- LP relaxation (2 intgrality gap) of undirected cut formulation [Goemans and Williamson 1995 and Vazirani 2001]
  - 2-approximation using primal-dual schemes or
  - 2-approximation using iterative rounding
- Another is the bidirected cut relaxation [Chakrabarty et al. 2008; Edmonds 1967; Rajagopalan and Vazirani 1999].

Bidirected Cut Relaxation (BCR)

$$\begin{array}{ll} min \sum_{e \in E} c(e) z_e & \\ s.t. \sum_{e \in \delta_{Ck}{}^{+}(U)} z_e \geq 1, & \forall U \subseteq V \setminus \{r\}, U \cap R \neq \emptyset \\ z_e \geq 0 & \forall e \in E \end{array}$$

Related work

Theorem 2 [Edmonds 1967]

For R = V, the polyhedron of BCR is integral.

- The best-known lower bound on the integrality gap of BCR is 8/7 [Knemann et al.2011; Vazirani 2001].
- The best-known upper bound is 2, though BCR is believed to have a smaller integrality gap than the undirected cut relaxation [Rajagopalan and Vazirani 1999].

Approximation DCR

Computed a $(1+\epsilon)$-approximate solution to DCR, for any given constant $\epsilon > 0$, in polynomial time.

Introduced a relaxation k-DCR of the k-restricted Steiner tree problem and can be solved exactly in polynomial time for any constant value of the parameter k.

The optimal solutions to k-DCR and DCR are close for large-enough k.

The relaxation of the k-restricted Steiner tree problem:

$$\begin{aligned} min \sum_{C \in Ck} c(C) x_C & & \text{(k-DCR)} \\ s.t. \sum_{C \in \delta_{Ck}(U)} x_C \geq 1, & \quad \forall U \subseteq R \backslash \{r\}, U \neq \emptyset \\ x_C \geq 0 & \quad \forall C \in Ck \end{aligned}$$

Approximation DCR

Lemma 1

$opt_{f,k} \leq \rho_k \cdot opt_f$

Proof

Let (x,**C**) be an optimal fractional solution for DCR. We show how to construct a solution (x',$\mathbf{C}'$) to k-DCR with the claimed property. For any component C $\in$ **C**, we can apply Theorem 1 to obtain a list of undirected components $C_1, ..., C_l$ such that: (a) $\bigcup_{i=1 \to l} \mathbf{C}_i$ connects the terminals in **C**, (b) any $\mathbf{C}_i$ contains at most k terminals and (c) $\sum_{i=1}^{l} \leq \rho_k \cdot c(C)$. Directed the edges of all C_is consistently towards sink(C) and increase the value of x' Ci by x_C for each Ci. The resulting solution (x', C') satisfies the claim.

Approximation DCR

Lemma 2

The optimal solution to **k-DCR** can be computed in polynomial time for any constant k.

Lemma 3

For any fixed $\epsilon \geq 0$, a $(1 + \epsilon)$-approximate solution (x,**C**) to **DCR** can be computed in polynomial time.

The Bridge Lemma

- The Bridge Lemma relates the cost of any terminal spanning tree to the cost of any fractional solution to DCR via the notion of bridges.
- Constructed of a proper weighted terminal spanning tree Y. Consider Steiner tree S on terminals R.
- Defined a bridge weight function w : RR $\rightarrow$ Q+ as follows: For any terminal pair u, v $\in$ R, the quantity w(u, v) is the maximum cost of any edge in the unique u-v path in S.
- $Br_S(R')$ is the set of bridges of S with respect to terminals R', and $br_S(R')$ denotes its cost.

The Bridge Lemma

Lemma 4

Let S be any Steiner tree on terminals R, and w : RR $\rightarrow$ Q+ be the associated bridge weight function. For any subset $R' \subseteq$ R of terminals, there is a tree Y $\subseteq$ $R'R'$ such that
(a) Y spans R'
(b) $w(Y) = br_S(R')$
(c) For any {u,v}$\in$ Y, the u-v path in S contains exactly one edge from $br_S(R')$

The Bridge Lemma

The Bridge Lemma

Let T be a terminal spanning tree and (x,C) be a feasible solution to **DCR**. Then

$$c(T) \leq \sum_{C \in \mathbf{C}} x_C br_T(C)$$

Proof

Constructed a spanning tree Y_C with weight $w(Y_C) = br_T(C)$. We direct the edges of Y_C towards sink(C). We obtain a directed capacity reservation y: RR$\rightarrow$ Q$+$ with y(u,v)$= \sum_{(u,v) \in Y_C} x_C$.

$$c(T) \leq w(F) \leq \sum_{e \in R \times R} w(e)y(e) = \sum_{C \in \mathbf{C}} x_C w(Y_C) =$$
$$= \sum_{C \in \mathbf{C}} x_C br_T(C)$$

A First Bound

The Bridge Lemma shows that given a terminal spanning tree and contracting a random component from any feasible fractional solution, one can remove a 1/M fraction of the edges and still obtain a terminal spanning tree. Showed that the cost of the minimum terminal spanning tree decreases by a factor (1-1/M) per iteration in expectation

Lemma 5

$E[{opt_f}^t] \leq (1 - \frac{1}{M})^{t-1} \cdot 2opt_f$

Theorem

For any fixed $\epsilon > 0$, there is a randomized polynomial-time algorithm that computes a solution to the Steiner tree problem of expected cost at most $(1+\ln(2)+\epsilon)\cdot opt_f$

A Second Bound

Showed that in each iteration, the cost opt^t of the optimal (integral) Steiner tree of the current instance decreases by a factor (1-1/2M) in expectation.

Lemma 6

Let **S** be any Steiner tree and (x,**C**) be a feasible solution to **DCR**. Sample a component $C \in DCR$ such that C $= C'$ with probability $x_{C'}$/M. Then, there is a subgraph $S' \subseteq S$ such that $S' \cup C$ spans R and
$E[c(S')] \leq (1 - \frac{1}{2M})^{t-1} \cdot c(S)$

Corollary

For every $t \geq 1$
$E[opt^t] \leq (1 - \frac{1}{2M})^{t-1} \cdot opt$

A Second Bound

Theorem

For any $\epsilon > 0$, there is a polynomial-time randomized approximation algorithm for Steiner tree with expected approximation ratio $3/2 + \epsilon$.

$$E\left[\frac{\sum_{t \geq 1} c(C^t)}{opt}\right] \leq 3/2 + \epsilon$$

Upper Bound

An LP relaxation of the problem does not imply a ln(4) (nor even a 1.5) upper bound on the integrality gap of the studied LP. The LP changes during the iterations of the algorithm, and its solution is only bounded with respect to the initial optimal integral solution. Then the LP has integrality gap at most $1 + \ln(3)/2 < 1.55$.

Algorithm RZ (Robins and Zelikovsky)

- Constructed a sequence $T^0, T^1, ..., T^\mu$ of terminal spanning trees. T^0 is a minimum cost terminal spanning tree in graph.
- At iteration t we are given a tree T^t and a cost function c^t on the edges of the tree.
- Any component C with at least 2 and at most k terminals.

Upper Bound

- $T^t[C]$ denote the minimum spanning tree of the graph $T^t \cup C$, where the edges e $\in$ C have weight 0 and the edges f $\in$ T^t weight c^t(f).
- The subset of edges in T^t but not in T^t[C] are denoted by Br_{T^t}(C).
- Loss(C) the minimum-cost subforest of C with the property that there is a path between each Steiner node in C and some terminal in R(C). So, Loss(C) is the complement of the set of bridges of the subtree C after contracting R(C)
- loss(C) = c(Loss(C)).
- Selected the component C^{t+1} that maximizes $gain^t(C)/loss(C)$
- It halts if the quantity is nonpositive

Upper Bound

- Otherwise, it considers the graph $T^t \cup C^{t+1}$, and contracts Loss(C^{t+1}).The tree T^{t+1} is a minimum-cost terminal spanning tree in the resulting graph.
- In case that parallel edges are created this way, the algorithm only keeps the cheapest of such edges. This way we obtain the cost function c^{t+1} on the edges of T_{t+1}.

$$gain^t(C) = br_{T^t}(C) - c(C) \text{ and } sgain^t(C) = gain^t(C) + loss(C)$$

Lemma [Robins and Zelikovsky 2005]

For t=1,2,...,μ, $c^t(T^l) = c^{t-1}(T^{l-1}) - sgain^{t-1}(C^t)$

Lemma [Robins and Zelikovsky 2005]

For any $l \leq \mu$, $apx_k \leq \sum_{t=1}^{l} loss(C^t) + c^l(T^l)$

Lower Bound

Theorem

The integrality gap of DCR is at least $8/7 > 1.142$.

Proof

Using Skutella's graph. Consider a Set Cover instance with elements U=1,...7 and sets S1,...,S7 and a vector b(i), where $Sj = i \in U | b(i) \cdot b(j) \equiv 1$. It needs 3 sets to cover all elements, but choosing each set to an extent of $1/4$ gives a fractional Set Cover solution of cost $7/4$. In graph, each element forms a terminal and each set is a nonterminal node connected to the root and to the contained elements by unit cost edges.

Lower Bound

continue...

The graph has 7 edge disjoint components, each one containing one non-terminal node and the 5 edges incident into it. On one hand, installing 1/4 on each of these components gives a fractional solution of cost 35/4, while on the other hand, at least 3 Steiner nodes must be included for an integer solution. Consequently opt $= 10$ and we obtain the promised gap of $\frac{10}{35/4} = 87$.

Comparison with BCR

DCR is a relaxation strictly stronger than BCR.

Lemma

Let opt_{DCR} and opt_{BCR} be the optimal fractional solutions to **DCR** and **BCR**, respectively, for a given input instance. Then, $opt_{DCR} \geq opt_{BCR}$.

Theorem

The integrality gap of BCR is at least $36/31 > 1.161$.

References

Byrka, Jaros-law and Grandoni, Fabrizio and Rothvoss, Thomas and Sanità, Laura
Steiner Tree Approximation via Iterative Randomized Rounding
J. ACM February 2013, 6:1–6:33.

Questions?