1 A 2-Approximation for Generalized Steiner Network Problem

Recall in the Generalized Steiner Network Problem (GSNP), we are given a graph $G = (V, E)$, a cost function defined over edges, $c : E \to (R)^+$, and requirements $r_{uv} \in \mathbb{Z}^+ \cup \{0\}$ for each pair $u, v \in V \times V$. We want to find a minimum cost set of edges $E' \subseteq E$ such that for every u, v, the graph $G[E']$ has r_{uv} edge disjoint paths between u and v. GSNP is a special case of the Abstract Network Design Problem where $f(A) := \max_{(u,v)\,:\,|A\cap\{u,v\}|=1} r_{uv}$. Recall our notation that $\delta(S)$ represents the set of edges crossing the cut $(S, \bar{S})$. And finally recall the LP for the Generalized Steiner Network Problem:

$$\begin{aligned} \min \sum_{e\in E} c_e x_e & \\ \text{such that} \sum_{e\in\delta(S)} x_e \geq f(S) & \qquad \forall\, S \subset V \\ 0 \leq x_e \leq 1 & \qquad \forall\, e \in E \end{aligned}$$

Our main result of this lecture will be the following *breakthrough* of Kamal Jain [2].

Theorem 1 (Jain [2]) *Let f be an integer valued skew-supermodular function. Then the integrality gap of the LP is 2. Moreover, there is a polynomial time 2-approximation for* GSNP *(for proper f).*

Note that for arbitrary skew-supermodular functions, we don't know a separation oracle for the above LP, but for the special case of Steiner Networks, we do. We will use the technique of Iterated Rounding, which relies on the following theorem:

Theorem 2 *Let f be a skew-supermodular function and x be a basic feasible solution to the above LP. Then $\exists e \in E$, such that $x_e \geq 1/2$.*

While we still will get a 2-approximation, note that in comparison with the LP for Vertex Cover, the vertex solution we get may not be half-integral. Instead, the idea will be to find some edge with $x_e \geq 1/2$, add that edge to our solution, remove it from the graph and recurse on the residual problem. We rely on the fact that the residual requirement function of a skew-supermodular function remains skew-supermodular.

The algorithm is as follows:

```
IteratedRoundingGSNP(G(V,E)):
F ← ∅.
While F is not feasible:
   g is residual requirement function w.r.t. F.
   Solve LP for g on G(V, E \ F) and get a basic feasible solution, x.
   If x_e ≥ 1/2, F ← F + e.
Output F.
```

To prove this algorithm works, we need to prove Theorem 2.

Let m be the number of variables in our LP. To get started we first note that we can find m tight sets, where a set $S \subset V$ is *tight* if $x(\delta(S)) = f(S)$. For, let x be a basic feasible solution to the above LP. We can assume, wlog that for each edge e, we have $0 < x_e < 1$, since if $x_e = 0$ we can throw out that edge, and if $x_e = 1$ then we're done. We then have that $|E| = m$, since each edge is in the vertex solution non-integrally. Since x is a basic feasible solution, we have that $\exists S_1, \ldots S_m$ such that x is the unique solution to the linear system $x(\delta(S_i)) = f(S_i)$, for all $i = 1 \ldots m$.

We now show that we can force the tight sets to belong to a special class:

Definition: Let $\mathcal{L}$ be a collection of subsets of V, $\mathcal{L}$ is called *laminar*, if $\forall A, B \in \mathcal{L}, A \cap B = \emptyset$ or $A \subset B$ or $B \subset A$.

Definition: For a graph G with m edges, and a set $S \subset V$, let $\chi_S \in \{0,1\}^m$ be the vector that has a 1 in position i iff edge $e_i \in \delta(S)$. χ_S is called the *incidence vector* of S.

And recall the book's definition of submodularity (p. 214):

Definition: A function $f : 2^V \rightarrow \mathbb{Z}^+$ is called *submodular* if $f(V) = 0$, and for every two sets $A, B \subseteq V$, the following two conditions hold:

- $f(A) + f(B) \geq f(A \cap B) + f(A \cup B)$
- $f(A) + f(B) \geq f(A - B) + f(B - A)$

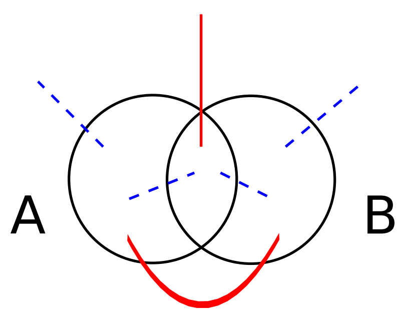

Figure 1: Cut function is submodular. $|\delta(A)| + |\delta(B)| \geq |\delta(A \cup B)| + |\delta(A \cap B)|$. The solid red edges contribute more to the left than the right, the others contribute equally to both sides of the inequality.

We also note that the size of a cut is a submodular function:

Lemma 3 *For any graph G on vertex set V, the function $|\delta(\cdot)|$ is submodular.*

We now need some lemmas about laminar families, building up to an uncrossing argument for the tight sets of a basic feasible solution of our LP.

Lemma 4 *There exist sets $S'_1, \ldots S'_m$ such that $\mathcal{L} = \{S'_1, \ldots S'_m\}$ is a laminar family and x is the unique solution to $x(\delta(S'_i)) = f(S'_i)$ for all $i = 1 \ldots m$.*

Proof: We skip this proof because the lemma follows from lemma 7, which we will prove. □

Lemma 5 *Let A, B be tight sets that properly intersect ($A \cap B \neq \emptyset$), then either $A - B, B - A$ are tight and $\chi_A + \chi_B = \chi_{A-B} + \chi_{B-A}$, or $A \cup B, B \cap A$ are tight and $\chi_A + \chi_B = \chi_{A \cup B} + \chi_{A \cap B}$.*

Proof: Because A, B are tight sets, we have $x(\delta(A)) = f(A)$ and $x(\delta(B)) = f(B)$. The definition of (f being) skew-supermodular implies that $f(A)+f(B) \leq f(A-B)+f(B-A)$ *or* $f(A)+f(B) \leq f(A \cup B) + f(A \cap B)$.

Suppose first that $f(A) + f(B) \leq f(A - B) + f(B - A)$, we have

$$x(\delta(A)) + x(\delta(B)) = f(A) + f(B) \leq f(A - B) + f(B - A) \leq x(\delta(A - B)) + x(\delta(B - A)) \quad .$$

Where the first equality follows because A, B are tight, the second inequality by assumption, and the third because x is feasible for f. But because the cut function is submodular, we have also that $x(\delta(A)) + x(\delta(B)) \geq x(\delta(A - B)) + x(\delta(B - A))$, therefore

$$f(A - B) + f(B - A) = x(\delta(A - B)) + x(\delta(B - A))$$

and $A - B, B - A$ are tight.

This implies that when tightness happens, edges between $A \cup B, A \cap B$ are not present, and therefore $\chi_A + \chi_B = \chi_{A-B} + \chi_{B-A}$. We can replace A, B by $A - B, B - A$ and preserve tightness.

The case for $A \cup B, A \cap B$ follows similarly. □

Lemma 6 *If $x \in (0, 1)^m$ and x is a basic feasible solution, then $\exists m$ sets, $S_1, \ldots S_m$ such that:*

1. *$x(\delta(S_i)) = f(S_i)$ for all $i = 1 \ldots m$.*
2. *The vectors χ_{S_i}, for $i = 1 \ldots m$, are linearly independent.*

Proof Sketch. Any basic feasible solution to the LP satisfies m linearly independent constraints with equality. Equality implies the sets corresponding to those constraints are tight, and linear independence implies the incidence vectors of the sets are linearly independent. □

Lemma 7 *If $x \in (0, 1)^m$ and x is a basic feasible solution, then $\exists$ a laminar family of sets, $\mathcal{L}$ such that:*

1. *$\forall A \in \mathcal{L}, x(\delta(A)) = f(A)$. That is, $\mathcal{L}$ is a collection of tight sets.*
2. *$|\mathcal{L}| = m$.*
3. *χ_A, $A \in \mathcal{L}$ are linearly independent.*

Proof:

We will iteratively build a laminar family $\mathcal{L}$ of tight sets, where the incidence vectors of sets in $\mathcal{L}$ are linearly independent. For any set $S \subseteq V$ we define the *crossing number* of S to be the number of sets in $\mathcal{L}$ that S crosses.

We first observe that if S is a set that crosses some $T \in \mathcal{L}$, then each of the sets $S - T, T - S, S \cup T, S \cap T$ have smaller crossing number than S. This is basically shown by picture, in figure 2. There are only three ways another set of $\mathcal{L}$ can cross one of $S - T, T - S, S \cup T, S \cap T$, and in all cases, it also crosses S. But since T doesn't cross any of $S - T, T - S, S \cup T, S \cap T$, the crossing number of each of those four sets is at least one less than S.

Next we show that if S is a tight set such that $\chi_S \notin \mathrm{span}(\mathcal{L})$ and S crosses some set $T \in \mathcal{L}$, then there is some tight set S' with smaller crossing number than S and $\chi_{S'} \notin \mathrm{span}(\mathcal{L})$.

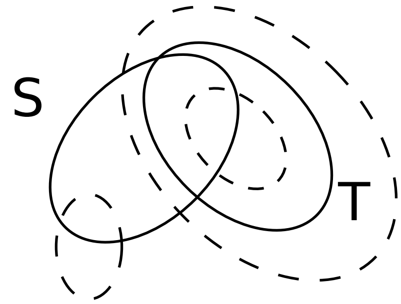

Figure 2: S crosses a set $T \in \mathcal{L}$. Dotted ovals are other sets of $\mathcal{L}$.

Lemma 5 implies that either $S - T$ and $T - S$ are tight or $S \cup T$ and $S \cap T$ are tight. Suppose the first possibility holds, the case for the second is similar. We then also have that $\chi_S + \chi_T = \chi_{S-T} + \chi_{T-S}$. From this and the fact that $\chi_S \notin \text{span}(\mathcal{L})$, we cannot have that both $\chi_{S-T}, \chi_{T-S} \in \text{span}(\mathcal{L})$. Also both $S - T$ and $T - S$ have smaller crossing number than S, so one of the two satisfy the conditions for S'.

Finally we note that if $\text{span}(\mathcal{L}) \neq \mathbb{R}^m$, then we can find some tight set S where $\chi_S \notin \text{span}(\mathcal{L})$ and $\mathcal{L} \cup \{S\}$ remains laminar. Therefore if $\mathcal{L}$ is a maximal laminar family constructed by finding such sets S and including them in $\mathcal{L}$, then $\mathcal{L}$ consists only of tight sets with linearly independent incidence vectors and $|\mathcal{L}| = m$.

□

1.1 A counting argument

From the above lemmas we will now find an edge e with $x_e \geq 1/2$. Suppose that $\forall e, x_e < 1/2$.

Observation: $\forall A \in \mathcal{L}, |\delta(A)| \geq 3$. Otherwise, if $|\delta(A)| < 3$, one of the edges must have $x_e \geq 1/2$ to maintain feasibility, since $f(A) \geq 1$.

We will distribute tokens, from edges to sets of $\mathcal{L}$ with these properties:

1. Each edge e gives out 1 unit of tokens to sets in $\mathcal{L}$ according to rules given below.
2. Each set $A \in \mathcal{L}$ receives a positive, non-zero, integral amount of tokens (that is, at least 1 unit of tokens).
3. Some non-zero amount of tokens are left over.

Rule 1: If $e = (u, v), e$ gives x_e tokens to the smallest set containing u and x_e to the smallest set containing v.

Rule 2: e gives $1 - 2x_e$ tokens to the smallest set containing both u, v. If no set contains both u, v tokens are unused.

Our contradiction will come from the fact that every set gets at least 1 unit of tokens, and some tokens must be left over, but we have m sets and only m edges (each of which give out at most 1 unit of tokens), so we must have used all our tokens.

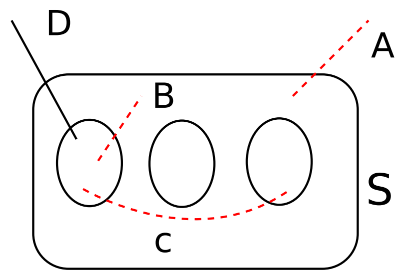

Figure 3: Set S containing sets $R_1, \ldots R_k$. Edges of type D do not give any tokens to S, but edges of type A, B, C do.

Within our laminar family, consider a set S which contains sets $R_1, \ldots R_k$, such that for no set R do we have $R_i \subset R \subset S$, that is the smallest set containing any R_i is S. We will count the tokens given to S. Four types of edges may exist, they are illustrated in figure 3. Note that edges of type D do not give any tokens to S but edges of type A, B, C do. We have:

$$tokens(S) = \sum_{e \in A} x_e + \sum_{e \in B}(1 - x_e) + \sum_{e \in C}(1 - 2x_e) = x(A) + |B| - x(B) + |C| - 2x(C) \quad (1)$$

Claim 8 $A \cup B \cup C \neq \emptyset$.

Proof Sketch. If $A \cup B \cup C = \emptyset$, then all edges of $\delta(S), \delta(R_i)$, for all i, are of type D. But then $\chi_S = \sum \chi_{R_i}$, and we would contradict the fact that the incidence vectors of our laminar family are linearly independent. □

Claim 9 *$tokens(S) > 0$ and $tokens(S)$ is a (positive) integer.*

Proof: Clearly $tokens(S)$ cannot be negative, since we never take tokens from a set. We have that

$$x(\delta(S)) - \sum_{i=1}^{k} x(\delta(R_i)) = f(S) - \sum_i f(R_i) = x(A) - x(B) - 2x(C)$$

is an integer, since all the $f(\cdot)$ are integral. The first equality follows from tightness, and the second from the way we assign tokens. Of course $|B|, |C|$ are integral, so we have that equation 1 above, is integral. □

Claim 10 *Some non-zero amount of tokens are left over.*

Proof: We observed earlier that $\forall S \in \mathcal{L}, |\delta(S)| \geq 3$. The only edges that contribute to $\delta(S)$ are of type A or D. Consider in particular some set $S \in \mathcal{L}$ such that no for no $S' \in \mathcal{L}$ do we have $S \subseteq S'$. Since no set of $\mathcal{L}$ contains both endpoints of the ≥ 3 edges of type A or D leaving S (if there was, that set would have to contain S), each edge e among these leave at least $1 - 2x_e$ tokens unused, which is positive as $x_e < 1/2$. □

If you look back, you've noticed we've now contradicted our assumption that for each edge $e, x_e < 1/2$. We're done!

2 Extra Theorems

Theorem 11 [1] *There is an $O(r_{max}^2 \log n)$ approximation for Node Connectivity Steiner Network, where r_{max} is the maximum requirement.*

Note that problems about vertex connectivity are harder than problems about edge connectivity.

Theorem 12 *Unless $\mathcal{P} = \mathcal{NP}$, no $r_{max}^{1-\epsilon}$ approximation exists.*

Theorem 13 [3] *Given degree bounds B_v on each v, there is an algorithm to output a spanning tree of cost OPT where $deg(v) \leq B_v + 1$ for each v, where OPT is the cost of the minimum cost spanning tree that respects degree bounds B_v.*

Theorem 14 [4] *For Steiner Network with degree requirements B_v, there is an algorithm that outputs a solution of cost $\leq 2OPT$ and $deg(v) \leq 2B_v + 3$ or $deg(v) \leq 6r_{max} + 3$, where OPT is the cost of the minimal cost Steiner Network satisfying degree constraints.*

References

[1] J. Chuzoy, S. Khanna. An $O(k^3 logn)$-Approximation Algorithm for Vertex-Connectivity Survivable Network Design. `http://arxiv.org/abs/0812.4442`

[2] K. Jain. A Factor 2 Approximation Algorithm for the Generalized Steiner Network Problem. *Combinatorica* 21(1): 39–60, 2001; preliminary version in FOCS 1998.

[3] L. Lau, M. Singh. Approximating Minimum Bounded Degree Spanning Trees to within One of Optimal. STOC 2007.

[4] L. Lau, M. Singh Additive Approximation for Bounded Degree Survivable Network Design. STOC 2008.