

16.1 Introduction

The last few weeks we've done a lot of work on metric spaces, i.e., on problems where WLOG the input is a metric space. Then we focused on embedding these metrics, allowing us to solve problems on trees rather than on general graphs. We've also done a lot of work on LP-based approximation algorithms. Now we're going to consider some interesting settings in which metrics and LPs are combined in a different way: the *solution* to the LP is a metric, rather than the input being a metric. So then after solving the LP we'll have a metric space that we can manipulate algorithmically.

We're going to see this first in an exact algorithm (not an approximation algorithm) for the classical min $s - t$ cut problem. Then we'll see it used in an approximation algorithm for MULTIWAY CUT.

16.2 Min $s - t$ Cut

We recall the basic definition of the MIN CUT PROBLEM

$$\begin{aligned} \text{Input: } & \text{Graph } G = (V, E) \\ & \text{Costs } c : E \rightarrow \mathbb{R}^+ \\ & \text{Source } s \in V \text{ and sink } t \in V \\ \text{Feasible: } & A \subseteq E \text{ s.t. } G - A \text{ has no } s - t \text{ path} \\ \text{Objective: } & \min \sum_{e \in A} c(e) \end{aligned}$$

We note that this definition of MIN CUT PROBLEM can be written in an equivalent form:

$$\begin{aligned} \text{Input: } & \text{Graph } G = (V, E) \\ & \text{Costs } c : E \rightarrow \mathbb{R}^+ \\ & \text{Source } s \in V \text{ and sink } t \in T \\ \text{Feasible: } & S \subseteq V \text{ s.t. } s \in S \text{ and } t \notin S \\ \text{Objective: } & \min \sum_{e \in E(S, \bar{S})} c(e) \end{aligned}$$

Definition 16.2.1 $\mathcal{P}_{s,t} = \{\text{all } s - t \text{ paths}\}$

We define the following LP, whose integer solutions are solutions to the MIN CUT PROBLEM

$$\begin{aligned}
& \min && \sum_{e \in E} c(e)x_e \\
& \text{subject to} && \sum_{e \in P} x_e \geq 1 \quad \forall P \in \mathcal{P}_{s,t} \\
& && 0 \leq x_e \leq 1 \quad \forall e \in E
\end{aligned}$$

Intuitively, this states that along each path at least one edge must be in the cut.

Theorem 16.2.2 *This LP can be solved in polytime, even though it has an exponential number of constraints.*

Proof: Suppose \vec{x} is not a feasible solution to the LP. There must thus be a path $P \in \mathcal{P}_{s,t}$ s.t. $\sum_{e \in P} x_e < 1$. We now think of x_e as a length assigned to edge $e \in E$. We can easily find a shortest $s - t$ path (e.g., using Dijkstra's algorithm). Because it is the shortest path, and there is some $P \in \mathcal{P}_{s,t}$ with $\sum_{e \in P} x_e < 1$, the shortest $s - t$ path P^* also has $\sum_{e \in P^*} x_e < 1$, and thus P^* corresponds to a violated constraint. Thus we can use Dijkstra's algorithm as a separation oracle for the ellipsoid method, allowing us to solve the LP in polynomial time. ■

Now we can solve the LP, so let's define a rounding algorithm. We don't want to use *independent* randomized rounding, though – it's not hard to see that this will not do well. Let's instead use the interpretation of x as edge lengths, and round in a way that corresponds to the metric that this gives.

We begin by defining a few variables.

Definition 16.2.3 *Let $d(u)$ denote the shortest path distance from s to u under the edge lengths $x_e \in \vec{x}$.*

Definition 16.2.4 *Let $B(s, r) = \{v \in V : d(v) \leq r\}$.*

Definition 16.2.5 *If $S \subset V$, let $\delta(S) = E(S, \bar{S}) =$ set of edges with one endpoint in S and one endpoint in \bar{S} .*

Algorithm 1 LP rounding for Min Cut

Input: Graph $G = (V, E)$ and solution to the LP \vec{x}

Output: $S \subseteq E$

Choose r uniformly at random in $(0, 1)$

$S \leftarrow B(s, r)$

return $A \leftarrow \delta(S)$

Claim 16.2.6 *With probability 1, the algorithm returns a feasible solution.*

Proof: Since x is a feasible solution, every path from s to t has length at least 1. So $t \notin B(s, r)$ for any choice of r , and thus A cuts t from s . ■

Claim 16.2.7 $\Pr[e \in A] \leq x_e$ for all $e \in E$.

Proof: Let $e = \{u, v\}$. Without loss of generality, let $d(u) \leq d(v)$. In order for e to be in A , we must have $r \in [d(u), d(v))$. So $\Pr[e \in A] = \Pr[r \in [d(u), d(v))] \leq d(v) - d(u) \leq d(u, v) \leq x_e$, where $d(u, v)$ denotes the length of the shortest $u - v$ path. ■

So now by linearity of expectations we know that $\mathbf{E}[c(A)] = \sum_{e \in E} c(e) \Pr[e \in A] = \sum_{e \in E} c(e) x_e = LP$. So this algorithm performs as well as the LP in expectation.

We can take this seemingly randomized algorithm and make it deterministic by simply trying all possibilities. Suppose there is no node w such that $d(u) \leq d(w) \leq d(v)$. Thus any $r \in [d(u), d(v)]$ will yield the same cut. Thus there are a linear number of possible cuts, which can each be tested in polynomial time. So our algorithm will be to solve the LP and try each of these cuts, taking whichever is best. Clearly this runs in polynomial time. Since the randomized algorithm does no worse than the LP in expectation, at least one of these cuts must do as well as the LP, so we will return a cut of cost at most $\sum_{e \in E} c(e) x_e$. Thus this algorithm returns the optimal solution, since $\sum_{e \in E} c(e) x_e \geq OPT$. So we have yet another polynomial-time algorithm for min $s - t$ cut.

16.3 Multiway Cut

Let's move to a problem that's actually NP-hard: MULTIWAY CUT.

$$\begin{aligned} \text{Input: } & \text{Graph } G = (V, E) \\ & \text{Costs } c : E \rightarrow \mathbb{R}^+ \\ & T = \{s_1, s_2, \dots, s_k\} \subseteq V \\ \text{Feasible: } & A \subseteq E \text{ s.t. } G - A \text{ has no } s_i - s_j \text{ path } \forall i, j \in \{1, 2, \dots, k\} \\ \text{Objective: } & \min \sum_{e \in A} c(e) \end{aligned}$$

In other words, rather than disconnecting s and t , we need to disconnect every terminal from every other terminal (nodes in T are usually called terminals). We define the following LP, whose integer solutions are solutions to MULTIWAY CUT.

$$\begin{aligned} \min & \sum_{e \in E} c(e) x_e \\ \text{subject to} & \sum_{e \in P} x_e \geq 1 \quad \forall i, j \in \{1, 2, \dots, k\}, \forall P \in \mathcal{P}_{s_i, s_j} \\ & 0 \leq x_e \leq 1 \quad \forall e \in E \end{aligned}$$

Note that the LP solution \vec{x} induces a metric d on the nodes through the shortest-path distances, and the constraints guarantee that $d(s_i, s_j) \geq 1$ for all $i, j \in [k]$. We use the same separation oracle as we did for MIN CUT example, but simply use it to check for each pair of $\{i, j\}$. Thus this LP can also be solved in polynomial time. As earlier, the LP solution \vec{x} , when interpreted as edge lengths, gives a metric d by shortest-paths.

Now consider the following rounding algorithm, which is the obvious generalization of the algorithm we designed for Min $s - t$ Cut.

Algorithm 2 LP rounding for Multiway Cut

Input: Graph $G = (V, E)$ and solution to the LP \vec{x}

Output: $S \subseteq E$

$A \leftarrow \emptyset$

Choose r uniformly at random from $[0, \frac{1}{2})$

for all $i \in \{1, 2, \dots, k\}$ **do**

$A_i \leftarrow \delta(B(s_i, r))$

$A \leftarrow A \cup A_i$

end for

return A

Claim 16.3.1 *This algorithm always returns a feasible solution.*

Proof: By the LP constraints on x , we know that $d(s_i, s_j) \geq 1$ for all terminals s_i, s_j . Since $r < 1/2$, this implies that A_i is enough to cut s_i from s_j (so is A_j , in fact). ■

Now let's analyze the cost.

Claim 16.3.2 $\Pr[e \in A] \leq 2x_e$ for all $e \in E$.

Proof: Let $w \in V$. By the triangle inequality, $d(s_i, w) + d(w, s_j) \geq d(s_i, s_j) \geq 1$ for all $s_i, s_j \in T$. This implies that for every $w \in V$, there is at most one $s_i \in T$ such that $d(s_i, w) < 1/2$. So we can define the following disjoint subsets:

Definition 16.3.3 Let $C_i = \{w \in V : d(s_i, w) < \frac{1}{2}\}$.

Let $e = \{u, v\} \in E$. To prove the claim, we break into two cases.

Case 1: $u, v \in C_i$ for some $i \in 1, 2, \dots, k$. WLOG we assume $d(s_i, u) \leq d(s_i, v)$.

$$\Pr[e \in A] = \Pr[e \in A_i] = \Pr[r \in [d(s_i, u), d(s_i, v)]] = \frac{d(s_i, v) - d(s_i, u)}{\frac{1}{2}} \leq 2d(u, v) \leq 2x_e$$

Case 2: There is no i such that $u, v \in C_i$. If neither u nor v is in *any* of the C_i 's, then $\Pr[e \in A] = 0$ so we're finished. Otherwise, we know that $u \in C_i$, and either v is not in any C_j or $v \in C_j$ for $j \neq i$. WLOG, let's assume that $d(s_i, u) \leq d(s_j, v)$. Then e will be in A if and only if $r \geq d(s_i, u)$. So we get that

$$\Pr[e \in A] = \Pr[r \in [d(s_i, u), 1/2]] \leq \frac{\frac{1}{2} - d(s_i, u)}{\frac{1}{2}} = 2 \left(\frac{1}{2} - d(s_i, u) \right) \leq 2d(u, v) \leq 2x_e$$

So in all cases $\Pr[e \in A] \leq 2x_e$. ■

By linearity of expectations $E[c(A)] = \sum_{e \in E} c(e) \Pr[e \in A] \leq \sum_{e \in E} c(e) (2x_e) \leq 2 \sum_{e \in E} c(e) x_e = 2 \cdot LP$. So this is a 2-approximation in expectation, and we can derandomize as we did with min $s - t$ cut by trying out all relevant values of r .

16.3.1 Integrality Gap

Is our analysis tight? We consider the star with k nodes around the outside connected by a single node v . We choose as our k terminal nodes the outside nodes.

The optimal solution OPT is clearly given by cutting all but one edge. So the cost is $k - 1$.

The best LP solution is given by assigning all edges $\frac{1}{2}$. So the cost is $\frac{k}{2}$

Thus the gap is given by $\frac{OPT}{LP} = \frac{k-1}{\frac{k}{2}} = 2(1 - \frac{1}{k})$. So our analysis is tight.

16.3.2 A Better Algorithm

Despite this integrality gap, it turns out that we can do better by using a better LP relaxation! To construct a better relaxation, we first need to change our viewpoint a bit. Consider some optimal solution F . Then let $C_i = \{v \in V : v \text{ reachable from } s_i \text{ in } G \setminus F\}$. Clearly since F is a feasible multiway cut, $C_i \cap C_j = \emptyset$ for all $i \neq j$. But it turns out that they also form a partition: every node is in some C_i . To see this, suppose that it is false, and let S be the set of nodes that unreachable from any terminal in $G \setminus F$. Add S to C_1 to get C'_1 , which together with the other C_i 's now form a partition. Then any edge which is cut under the new sets was also cut by the old sets, so this new solution is just as good as the old solution.

Thus WLOG, we can assume that the optimal solution is actually a partition C_1, C_2, \dots, C_k where $s_i \in C_i$ for all $i \in [k]$, and the edges we cut are the edges between the parts of the partition.

This point of view suggests the following, different LP relaxation. We'll have the following variables.

$$x_u^i = \begin{cases} 1 & \text{if } u \in C_i \\ 0 & \text{else} \end{cases}$$

$$z_e^i = \begin{cases} 1 & \text{if } e \in \delta(C_i) \\ 0 & \text{else} \end{cases}$$

We now define an LP using these indicator variables

$$\begin{aligned} \min \quad & \frac{1}{2} \sum_{e \in E} \sum_{i=1}^k c(e) z_e^i \\ \text{subject to} \quad & \sum_{i=1}^k x_u^i = 1 && \forall u \in V \\ & z_e^i \geq x_u^i - x_v^i && \forall e = \{u, v\} \in E, \quad \forall i \in 1, 2, \dots, k \\ & z_e^i \geq x_v^i - x_u^i && \forall e = \{u, v\} \in E, \quad \forall i \in 1, 2, \dots, k \\ & x_{s_i}^i = 1 && \forall i \in 1, 2, \dots, k \\ & 0 \leq x_u^i \leq 1 && \forall u \in V, \quad \forall i \in [k] \\ & 0 \leq z_e^i \leq 1 && \forall e \in E, \quad \forall i \in [k] \end{aligned}$$

It is straightforward to verify that this is a valid relaxation of the multiway cut problem. We will give a more compact way of writing this LP which makes the connection to metrics clear.

Definition 16.3.4 Let $x, y \in \mathbb{R}^k$ then their ℓ_1 -distance is $\|x - y\|_1 = \sum_{i=1}^k |x^i - y^i|$ where x^i is the i^{th} coordinate of the vector x .

Definition 16.3.5 The k -simplex is $\Delta_k = \{x \in \mathbb{R}^k : \sum_{i=1}^k x^i = 1 \wedge x^i \geq 0 \forall i\}$ where x^i is the i^{th} coordinate of the vector x .

Definition 16.3.6 Let e_i be a vector with a 1 in the i^{th} coordinate and zeros elsewhere

Let $x_u = (x_u^1, x_u^2, \dots, x_u^k)$ and $x_v = (x_v^1, x_v^2, \dots, x_v^k)$. Note that $z_e^i = |x_u^i - x_v^i|$ in any optimal LP solution because of the constraints on z_e^i and because we are minimizing the objective function. So, $\|x_u - x_v\|_1 = \sum_{i=1}^k z_e^i$. Using all these definitions we can rewrite the new LP as follows

$$\begin{aligned} \min \quad & \frac{1}{2} \sum_{e=\{u,v\} \in E} c(e) \|x_u - x_v\|_1 \\ \text{subject to} \quad & x_{s_i} = e_i \quad \forall i \in 1, 2, \dots, k \\ & x_u \in \Delta_k \end{aligned}$$

Note that $\|x_{s_i} - x_{s_j}\|_1 = \|e_i - e_j\|_1 = 2$ for all $i \neq j \in [k]$.

Now we need a way of rounding this LP. We're going to use a technique due to Calinescu, Karloff, and Rabani [CKR00], which might look familiar: it's a lot like one level of FRT! Formally, we'll use the following algorithm. We'll define balls as one would expect, i.e., we'll let $B(s_i, r) = \{v \in V : \|e_i - x_v\|_1 \leq r\}$.

1. Initially we set $C_i = \emptyset$ for all $i \in [k]$, and set another set $X = \emptyset$.
2. Pick $r \in (0, 2)$ uniformly at random, and pick a permutation π of $[k]$ uniformly at random.
3. For $i = 1$ to $k - 1$
 - (a) Set $C_{\pi(i)} \leftarrow B(s_{\pi(i)}, r) \setminus X$
 - (b) $X \leftarrow X \cup C_{\pi(i)}$
4. Set $C_{\pi(k)} = V \setminus X$
5. Return $A = \cup_{i=1}^k \delta(C_i)$ (the edges cut by the partition that we built)

We're going to prove the following theorem.

Theorem 16.3.7 *This is a 3/2-approximation.*

To prove this, for every edge $\{u, v\}$ let $Z_{u,v}$ be an indicator random variable which will be 1 if this rounding algorithm separates u and v (i.e., includes e in the cut) and is 0 otherwise. We'll first state the following lemma.

Lemma 16.3.8 $\Pr[Z_{u,v} = 1] \leq \frac{3}{4}\|x_u - x_v\|_1$ for all $\{u, v\} \in E$.

It's straightforward to see that this lemma implies the theorem:

$$\begin{aligned} \mathbf{E}[ALG] &= \mathbf{E}\left[\sum_{e=\{u,v\} \in E} c(e)Z_{u,v}\right] = \sum_{e=\{u,v\} \in E} c(e)\mathbf{E}[Z_{u,v}] \\ &\leq \sum_{e=\{u,v\} \in E} c(e)\frac{3}{4}\|x_u - x_v\|_1 = \frac{3}{2} \cdot \frac{1}{2} \cdot \sum_{e=\{u,v\} \in E} c(e)\|x_u - x_v\|_1 \\ &\leq \frac{3}{2} \cdot OPT. \end{aligned}$$

So we just need to prove Lemma 16.3.8. This is going to look an awful lot like a simplified version of FRT. Fix some $\{u, v\} \in E$. Define S_i to be 1 if i is the first index in π such that $|B(s_i, r) \cap \{u, v\}| \geq 1$ (i.e., s_i settles $\{u, v\}$) and 0 otherwise. Define X_i to be 1 if $|B(s_i, r) \cap \{u, v\}| = 1$ (i.e., s_i cuts $\{u, v\}$) and 0 otherwise.

Similar to FRT (but simpler since there's only one "level"), we can note that $Z_{u,v} = 1$ if and only if there is some $i \in [k]$ such that $S_i = X_i = 1$. And by the definition of S_i there is at most one i which settles $\{u, v\}$, and thus some $i \in [k]$ has $S_i = X_i = 1$ if and only if $\sum_{i=1}^k S_i X_i = 1$. Thus $Z_{u,v} = \sum_{i=1}^k S_i X_i$, so by linearity of expectations we get that $\mathbf{E}[Z_{u,v}] = \sum_{i=1}^k \mathbf{E}[S_i X_i]$, and thus $\Pr[Z_{u,v} = 1] = \sum_{i=1}^k \Pr[S_i = X_i = 1]$. Less formally, the probability that $\{u, v\}$ is cut by the algorithm is equal to the sum over all terminals of the probability that the algorithm cut the edge because of that terminal.

Let's start by analyzing the cut probabilities. To do this, let's prove another easy lemma.

Lemma 16.3.9 For all $w \in V$ and $i \in [k]$, $w \in B(s_i, r)$ if and only if $1 - x_w^i \leq r/2$.

Proof: We have the following series of if and only if statements:

$$\begin{aligned} w \in B(s_i, r) &\Leftrightarrow \|e_i - x_w\|_1 \leq r \\ &\Leftrightarrow \sum_{j=1}^k |e_i^j - x_w^j| \leq r \\ &\Leftrightarrow 1 - x_w^i + \sum_{j \neq i} x_w^j \leq r \\ &\Leftrightarrow 2(1 - x_w^i) \leq r \\ &\Leftrightarrow 1 - x_w^i \leq r/2, \end{aligned}$$

as claimed. ■

To interpret this lemma, it says that we can think of the distance from w to s_i as being $2(1 - x_w^i)$. This lemma implies that

$$\begin{aligned} \Pr[X_i = 1] &= \Pr[\min(1 - x_v^i, 1 - x_u^i) \leq r/2 < \max(1 - x_u^i, 1 - x_v^i)] \\ &= |x_u^i - x_v^i|. \end{aligned}$$

Let $\ell \in [k]$ be the index minimizing $\min_{i=1}^k \min(1 - x_u^i, 1 - x_v^i)$, i.e., s_ℓ is the terminal that is closest to $\{u, v\}$ in the ℓ_1 -metric given by the LP. For this special index, we can bound the probability that the algorithm cuts $\{u, v\}$ in because of s_ℓ by

$$\Pr[S_\ell = X_\ell = 1] \leq \Pr[X_\ell = 1] \leq |x_u^\ell - x_v^\ell|$$

Now consider some $i \neq \ell$. If i cuts $\{u, v\}$, then $|B(s_\ell, r) \cap \{u, v\}| \geq 1$ since ℓ is closer to $\{u, v\}$ than i is. This if i cuts $\{u, v\}$, a necessary condition for it to settle $\{u, v\}$ is for it to be before ℓ in π . Since π is chosen uniformly at random, this happens with probability $1/2$. Thus we get that

$$\Pr[S_i = X_i = 1] = \Pr[S_i = 1 | X_i = 1] \Pr[X_i = 1] \leq \frac{1}{2} |x_u^i - x_v^i|.$$

Putting this all together we have

$$\begin{aligned} \Pr[Z_{u,v} = 1] &= \sum_{i=1}^k \Pr[X_i = S_i = 1] \leq |x_u^\ell - x_v^\ell| + \frac{1}{2} \sum_{i \neq \ell} |x_u^i - x_v^i| \\ &= \frac{1}{2} |x_u^\ell - x_v^\ell| + \frac{1}{2} \|x_u - x_v\|_1 \\ &\leq \frac{3}{4} \|x_u - x_v\|_1. \end{aligned}$$

which completes the proof of Lemma 16.3.8 and thus the proof of Theorem 16.3.7. However, to prove the last inequality, we actually need one final lemma.

Lemma 16.3.10 *For any index j and any $u, v \in V$, $|x_u^j - x_v^j| \leq \frac{1}{2} \|x_u - x_v\|_1$.*

Proof: Without loss of generality, let's assume that $x_u^j \geq x_v^j$. Then

$$|x_u^j - x_v^j| = x_u^j - x_v^j = \left(1 - \sum_{i \neq j} x_u^i\right) - \left(1 - \sum_{i \neq j} x_v^i\right) = \sum_{i \neq j} (x_v^i - x_u^i) \leq \sum_{i \neq j} |x_v^i - x_u^i|.$$

Adding $|x_u^j - x_v^j|$ to both sides and then dividing by two gives the lemma. ■

References

- [CKR00] Grigori Călinescu, Howard Karloff, and Yuval Rabani. An improved approximation algorithm for multiway cut. *Journal of Computer and System Sciences*, 60(3):564 – 574, 2000.