Lecture 9: Basic Metric Method in Approximation Algorithms

In this note, we show two $O(\log n)$ approximation algorithms for the undirected multicut problem based on two different (but closely related) techniques. We will use these techniques to prove a partitioning lemma about points in metric space. Such lemma can be used to round the feasible fractional solution given by a natural LP relaxation of undirected multicut (the solution can naturally be viewed as points in metric space).

Multicut problem: Suppose we are given a graph $G = (V, E)$, cost function $c : E \to \mathbb{R}$, and k terminal pairs $\{(s_i, t_i)\}_{i=1}^k$. Our goal is to remove subset of edges in order to separate all terminal pairs, while minimizing the cost of cut edges. Formally, we want to find a subset $E' \subseteq E$ such that there is no path from s_i to t_i in $(V, E \setminus E')$, while minimizing the cost $c(E') = \sum_{e \in E'} c_e$.

1 Partitioning Lemma

Let $G = (V, E)$ be a complete graph with weight function $w(e)$ and metric distance function $d(e)$ on edges. Let $\alpha, \Delta > 0$ be two parameters. We say that a partition $\pi = \{V_1, \dots, V_\ell\}$ of vertices is *α-cheap* Δ-small partition if

$$\sum_{i,j:i<j} \sum_{uv:u \in V_i, v \in V_j} w(uv) \leq \alpha \sum_{uv \in E} d(u,v) w(uv)$$

Moreover, the diameter of $G[V_i]$ is at most Δ. In other words, if we fix some parameter Δ, we are interested in partitioning the vertex set V into a number of low-diameter subsets such that the weight of edges across the cut is as small as possible.

Observe that, for a fixed value of Δ, the best α we can hope for is $\alpha = 1/\Delta$: Imagine a line metric with vertices corrsponding to points on the line, and the points are equally spaced. So it is reasonable to aim at computing Δ-small $(\frac{\beta}{\Delta})$-cheap partition with as small β as possible. For a general metric space, the best possible is $\beta = O(\log n)$, which will be shown below. More formally, we will prove the following.

Theorem 1.1. *Let $\Delta > 0$ be a parameter. There is a polynomial time algorithm that computes $O(\frac{\log n}{\Delta})$-cheap Δ-small partition.*

We first show how the partitioning lemma implies an $O(\log n)$ approximation algorithm for undirected multicut[1]. The lemma will be proved in the next two sections.

First we can write down the LP relaxation for minimum multicut problem as follows:

(LP)

$$\begin{aligned} \min \quad & \sum_{(u,v) \in E} c_e x_e \\ \text{s.t.} \quad & \sum_{e \in P} x_e \geq 1 \text{ for all } i \in [k] \text{, for all path } P \text{ from } s_i \text{ to } t_i \end{aligned}$$

[1]Applying the partitioning procedure directly gives $O(\log k)$ approximation

Separation Oracle: Even though this LP has exponential number of constraints, we can solve it in polynomial time by providing a separation oracle. Given a solution x, we consider a metric $d : V \times V \to \mathbb{R}$ where $d(u, v)$ is defined as the shortest path distance from u to v using x as a distance. We simply check, for all $i = 1, \ldots, k$, whether $d(s_i, t_i) \geq 1$. The constraints are satisfied if and only if for all $i = 1, \ldots, k$, we have $d(s_i, t_i) \geq 1$.

Algorithm: Now we describe an LP rounding algorithm. Given a feasible solution x to the above LP, we will be working with the metric (V, d) where d is a shortest path metric in x. Also, we define the weight $w(e) = c_e$ for all edges $e \in E$. Invoking the partition lemma with $\Delta = 1/2$, we obtain a partition $\pi = \{V_1, \ldots, V_\ell\}$ such that diameter of each V_j is at most $1/2$. Since $d(s_i, t_i) \geq 1$, the following claim is immediate.

Claim 1.1. *For each $i = 1, \ldots, k$, the pair s_i, t_i belong to different sets in the partition π.*

This means that, if we cut edges across the different sets in the partition π, we get a feasible solution for multicut. The partition lemma guarantees that the cost of edges across the cut is $O(\log n) \sum_{uv \in E} c_{uv} x_{uv} \leq O(\log n)\mathsf{OPT}_{LP}$. In the next sections, we will discuss two proofs of the lemma.

2 First Proof: Random Cut

Denote the vertex set by $V = \{v_1, \ldots, v_n\}$. We define a notion of balls as normally used in geometry, i.e. $\mathsf{Ball}(v, r) = \{u : d(u, v) \leq r\}$. Notice that $\mathsf{Ball}(v, r)$ is simply the set of vertices that are within distance r to vertex v.

The algorithm proceeds as follows:

- Pick a random permutation $\sigma : \{1, \ldots, n\} \to \{1, \ldots, n\}$ and a random number $r \in (0, \Delta/2)$. Initially $X = V$ is the set of "unclaimed" vertices.
- We proceed in n iterations, where in iteration i, we construct a set of vertices B_i by $B_i = \mathsf{Ball}(v_{\sigma(i)}, r) \cap X$. These vertices are claimed by $v_{\sigma(i)}$, and so we update $X \leftarrow X \setminus B_i$. Observe that any vertex cannot be claimed twice and that B_i may not even include $v_{\sigma(i)}$ (if it was previously claimed by other vertices).

Now we analyze the probability that each edge $uv \in E$ is cut and show that this probability is at most $O(\log n/\Delta)d(u, v)$. This immediately implies the lemma (well, after writing down the expectation and applying Markov's inequality).

What is the probability that an edge $uv \in E$ is cut? We need to look at the probability that u and v are claimed by different sets B_i. Consider the process by which these vertices are claimed. Let $w_1, \ldots, w_n$ be the vertices, ordered such that $d(w_1, \{u, v\}) \leq d(w_2, \{u, v\}) \leq \ldots \leq d(w_n, \{u, v\})$. We say that uv is *settled* by the w_j, if the ball growing from w_j claims at least one vertex from $\{u, v\}$ for the first time. By definition, an edge can be settled only once.

Claim 2.1. *If uv is settled by w_i, then all other vertices in the set $\{w_1, \ldots, w_{i-1}\}$ cannot come before w_i in the permutation.*

Proof. Suppose not, and w_j comes before w_i in the permutation such that $j < i$. Then $\mathsf{Ball}(w_j, r)$ must include at least one of $\{u, v\}$, settling the edge uv prior to w_i, a contradiction. □

Lemma 2.1. *The probability that uv is cut is at most $O(\frac{\log n}{\Delta})d(u,v)$.*

Proof. We first write $\mathbf{Pr}\left[uv \text{ is cut }\right] = \sum_{i=1}^{n} \mathbf{Pr}\left[uv \text{ is cut by } w_i\right]$. In order for w_i to cut the edge, it must be the case that w_i settles uv and therefore comes before other vertices in $\{w_1, \ldots, w_{i-1}\}$ according to random permutation σ.

$$\begin{aligned}
\mathbf{Pr}_{\sigma,r}\left[uv \text{ is cut by } w_i\right] &= \mathbf{Pr}_{\sigma,r}\left[w_i \text{ is the first among } \{w_1,\ldots,w_i\} \text{ and } w_i \text{ cuts } uv\ \right] \\
&= \mathbf{Pr}_{\sigma}\left[w_i \text{ is the first}\right] \mathbf{Pr}_{r}\left[w_i \text{ cuts } uv \mid w_i \text{ is the first}\right] \\
&= \frac{1}{i}\mathbf{Pr}_{r}\left[r \in [d(w_i,u), d(w_i,v)]\right]
\end{aligned}$$

The last line follows because if r is out of such range, it would be impossible for w_i to cuts uv. More formally, let $\mathcal{E}$ is the event that w_i is the first among $\{w_1, \ldots, w_i\}$, so we can write the term

$$\mathbf{Pr}\left[w_i \text{ cuts } uv \mid \mathcal{E}\right] \leq \mathbf{Pr}\left[r \in [d(w_i,u), d(w_i,v)] \mid \mathcal{E}\right] + \mathbf{Pr}\left[w_i \text{ cuts } uv \wedge r \notin [d(w_i,u), d(w_i,v)] \mid \mathcal{E}\right]$$

The second term becomes zero (once r is outside of the range, it is impossible for w_i to cut uv), and the first term is at most $\frac{d(w_i,v)-d(w_i,u)}{\Delta/2} \leq 2d(u,v)/\Delta$. Summing over all i, we get the term $\sum_i \frac{1}{i} 2d(u,v)/\Delta \leq O(\log n)\frac{d(u,v)}{\Delta}$. ∎

3 Second Proof: Region Growing

The second proof will give a deterministic algorithm for computing the partitioning procedure. We define the notion of *volume*, which captures the weighted distance inside a particular subset of vertices. For technicality reasons, the volume of each ball consists of the *node volume* and the *edge volume*.

Let $W = \sum_{uv \in E} d(u,v)w(uv)$. The node volume of each vertex is W/n. The edge volume inside the ball $B = \mathsf{Ball}(s,r)$ is $\sum_{uv \in B} w(u,v)d(u,v) + \sum_{uv \in \delta(B)} w(u,v)(r - d(s,\{u,v\}))$. In other words, the edge volume sums over all weighted distance of the edges inside the ball, and "partial" cost of the edges that cross from inside to outside of the ball.

Even if we do not talk about the ball $\mathsf{Ball}(s_i, r)$, the notion of volume is still well-defined: $vol(B) = \sum_{uv \in G[B] \cup \delta(B)} w(u,v)d(u,v) + |B|W/n$. The following lemma gives the upper bound on the total volume of disjoint balls.

Claim 3.1. *Let $B_1, \ldots, B_l$ be a collection of disjoint subsets of vertices. Then*

$$\sum_j vol(B_j) \leq O(W)$$

Proof. Notice that for each edge e inside some ball B_j, we have the term $c_e x_e$ that appears once, and for each edge crossing B_j, the term appears at most twice. So the contribution from edge volume is at most $2W$, and the contribution from node volume is at most W as well. The total volume is at most $3W$. ∎

Our goal is to partition the graph G into $V(G) = \bigcup_{i=1}^{\ell} B_i$ such that each set B_i has radius less than $\Delta/2$, i.e. it is a ball of radius at most $\Delta/2$. The cost of edges leaving each set B_i in the partition will be charged to the volume $vol(B_i)$.

Lemma 3.1. *Let $V' \subseteq V$ be any subset of vertices, and vertex $v \in V'$ is a designated vertex in the graph. Then there is a ball $B = \mathsf{Ball}_{G[V']}(v,r)$ for some $r < \Delta/2$ such that $w(\delta(B)) = \sum_{ab \in \delta(B)} w(a,b) \leq O(\frac{\log n}{\Delta})vol(B)$. Moreover, such ball can be computed deterministically in polynomial time.*

Proof. Imagine the process of growing the ball of radius r continuously from $r = 0$ at vertex v, and consider how the total volume of the ball changes in terms of r. Notice that, when $r = 0$, we have the volume of $vol(B(v, 0)) = W/n$, and this volume is at most $2W$ for any value of r.

For convenience, we abbreviate the term $Ball(v, r)$ by B_r. Notice that the rate of change of volume is at least $dV \geq w(Ball(v, r))dr$. Dividing both sides of the inequality by $V = vol(B_r)$, we get $\int \frac{dV}{V} \geq \int \frac{w(B_r)}{V} dr$. The LHS is equal to $\ln V$. By taking the limiting value of $r \in (0, \Delta/2)$, we get

$$|\ln(3W) - \ln(W/n)| \geq \int_0^{\Delta/2} \frac{w(B_r)}{V} dr$$

By averaging, there is a value $r \in (0, \Delta/2)$ such that $\frac{w(B_r)}{vol(B_r)} \leq O(\frac{\log n}{\Delta})$. □

We show how this lemma implies an algorithm for computing low-diameter low-cost partitioning. Starting from $X = V$, we proceed in iterations. In iteration i, as long as $diam(G[X]) > \Delta$, we pick arbitrary $v_i \in X$ and compute the ball B_i by invoking the lemma. Then we update $X \leftarrow X \setminus B_i$. In the end, we have obtained $B_1, \ldots, B_{n'}$ for some $n' \leq n$ such that the diameter of each component $G[B_j]$ is at most Δ. The total cost of edges across the cut can be bounded easily: Each edge uv in the cut must cross the ball at some iteration, so the total cost is at most $O(\frac{\log n}{\Delta}) \sum_{i=1}^{n'} vol(B_i)$ which is at most $O(\frac{\log n}{\Delta})W$, as desired.