Sparsest Cut

Computational and Metric Geometry

Instructor: Yury Makarychev

1 Overview

1.1 Sparsest Cut with Uniform Demands

Today we will talk about applications of metric embeddings in combinatorial optimization. We will first consider the Sparsest Cut problem and design an approximation algorithm for it. Then we will see how to design approximation algorithms for other problems using this algorithm as a basic building block.

Definition 1.1. *Consider a graph $G = (V, E)$. The sparsity of a cut $(S, \bar{S})$ equals*

$$\phi(S) = \frac{E(S, \bar{S})}{\min(|S|, |\bar{S}|)},$$

where $\bar{S} = V \setminus S$ and $E(S, \bar{S})$ is the number of cut edges, that is, the number of edges from S to $\bar{S}$. The Sparsest Cut problem asks to find a cut $(S, \bar{S})$ with smallest possible sparsity $\phi(S)$. We denote the value of the sparsest cut in G by $\phi^* = \phi^*_G$; ϕ^* is the Cheeger constant of G or the conductance of G.

The Sparsest Cut problem is NP-hard and therefore cannot be solved *exactly* in polynomial-time unless $P = NP$. We can hope only to find an *approximate solution* in polynomial-time.

1.2 Approximation Algorithms

An approximation algorithm is an efficient algorithm that finds an approximate solution with a provable approximation guarantee. The standard measure of the quality of an approximation algorithm is its *approximation factor*. The approximation factor of an algorithm is the worst case ratio between the value (cost) of the solution the algorithm finds and the value (cost) of the optimal solution. That is, an algorithm has approximation factor (at most) $\alpha = \alpha(n)$ if for every instance of the problem of size n the algorithm finds a feasible solution of value at most $\alpha\,\mathsf{OPT}$, if the problem is a minimization problem, and of value at least $\alpha\,\mathsf{OPT}$ if the problem is a maximization problem, where OPT is the value of the optimal solution. An algorithm is an α-approximation algorithm if it gives α-approximation.

Definition 1.2. *An algorithm is an $\alpha(n)$ approximation algorithm for Sparsest Cut if given a graph G on n vertices it finds a cut $(S, \bar{S})$ of sparsity at most $\alpha(n) \cdot \phi^*_G$.*

2 Metrics and Sparsest Cut

Let S be a cut in G. Consider the cut (semi-) metric δ_S on V. Observe that $E(S, \bar{S}) = \sum_{(u,v)\in E} \delta_S(u,v)$ and $|S||\bar{S}| = \frac{1}{2}\sum_{u,v\in V} \delta_S(u,v)$. We also have

$$\frac{|S||\bar{S}|}{n} \leq \min(|S|, |\bar{S}|) \leq \frac{2|S||\bar{S}|}{n}.$$

Therefore,

$$\frac{n\sum_{(u,v)\in E} \delta_S(u,v)}{\sum_{u,v\in V} \delta_S(u,v)} \leq \phi(S) \leq \frac{2n\sum_{(u,v)\in E} \delta_S(u,v)}{\sum_{u,v\in V} \delta_S(u,v)}.$$

We can reformulate the Sparsest Cut program as follows: find a cut metric d that minimizes the ratio

$$\phi'(d) = \frac{n\sum_{(u,v)\in E} d(u,v)}{\sum_{u,v\in V} d(u,v)}.$$

Note that the optimum value of this problem lies between $\phi^*/2$ and ϕ^*. Indeed let S^* be the optimal solution of Sparsest Cut and δ_S be the cut metric that minimizes $\phi'(d)$. Then

$$\frac{1}{2}\phi^* \leq \frac{1}{2}\phi(S) \leq \phi'(\delta_S) \leq \phi'(\delta_{S^*}) \leq \phi(S^*) = \phi^*.$$

If we *could* find the optimal cut metric δ_S, we would get a 2-approximation algorithm for Sparsest Cut (but we cannot). We consider the metric relaxation of the problem.

Metric Relaxation for Sparsest Cut. Find a (semi-) metric d that minimizes $\frac{n\sum_{(u,v)\in E}\delta(u,v)}{\sum_{u,v\in V}\delta(u,v)}$.

Note that every cut metric is a feasible solution of the metric relaxation for Sparsest Cut (but not the other way around). Therefore, the value of the relaxation is at most the value of the sparsest cut. We find the optimal solution of the relaxation using linear programming (LP). We introduce an LP variable d_{uv} for every pair of vertices (u,v) and write the following program.

$$\text{minimize} \sum_{(u,v)\in E} d_{uv}$$

subject to

$$\sum_{u,v\in V} d_{uv} = n \qquad \text{spreading constraint}$$

$$d_{uv} + d_{vw} \geq d_{uw} \qquad \text{metric constraints}$$

$$d_{uv} = d_{vu}$$

$$d_{uv} \geq 0$$

$$d_{uu} = 0.$$

The optimal solution of this LP defines a metric $d(u,v) = d_{uv}$ on V. We denote the value of this LP by LP.

Claim 2.1. *The value of the optimal solution for the metric relaxation for Sparsest Cut equals* LP. *Thus,* $\mathsf{LP} \leq \phi^*$.

3 Approximation Algorithm for Sparsest Cut

We solve the LP relaxation for Sparsest Cut and get the optimal solution $d(u,v)$ of value LP. Now our goal is to *round* the LP solution to a combinatorial solution. To this end, we embed (V,d) into ℓ_1^N with distortion $O(\log n)$ (where $N = O(\log^2 n)$).

Assume for a moment that our embedding embeds (V,d) into a space of random variables with norm $\|\cdot\|_1$ and furthermore it maps every vertex u to a random variable ξ_u that only takes values 0 and 1. Since the distortion of the embedding is $O(\log n)$, we have for $D = O(\log n)$, some $\alpha > 0$ and every vertices u and v,

$$\alpha \cdot d(u,v) \leq \|\xi_u - \xi_v\|_1 \leq \alpha \cdot D \cdot d(u,v).$$

Assume additionally that for every sample $\{\xi_u : u \in V\}$ at least one r.v. ξ_u is equal to 0, and one r.v. ξ_v is equal to 1.

Let $S = \{u : \xi_u = 0\}$ and $\bar{S} = V \setminus S = \{v : \xi_v = 1\}$ (note that S is a proper subset of V and $(S,\bar{S})$ is a cut in G). We have, $|\xi_u - \xi_v| = \delta_S(u,v)$. Thus

$$\|\xi_u - \xi_v\|_1 = \mathbb{E}\left[\delta_S(u,v)\right]. \tag{1}$$

By the linearity of expectation,

$$\mathbb{E}\left[E(S,\bar{S})\right] = \mathbb{E}\left[\sum_{(u,v)\in E} \delta_S(u,v)\right] = \sum_{(u,v)\in E} \|\xi_u - \xi_v\|_1 \leq \alpha D \sum_{(u,v)\in E} d(u,v) = \alpha D \cdot \mathsf{LP},$$

$$\begin{aligned}\mathbb{E}\left[\min(|S|,|\bar{S}|)\right] &\geq \mathbb{E}\left[\frac{|S||\bar{S}|}{n}\right] = \mathbb{E}\left[\frac{1}{2n}\sum_{u,v\in V} \delta_S(u,v)\right] = \frac{1}{2n}\sum_{u,v\in V} \|\xi_u - \xi_v\|_1 \\ &\geq \frac{1}{2n}\sum_{u,v\in V} \alpha d(u,v) = \frac{\alpha}{2}.\end{aligned}$$

Then, $\mathbb{E}\left[E(S,\bar{S}) - 2D \cdot \mathsf{LP} \cdot \min(|S|,|\bar{S}|)\right] \leq 0$. Therefore with positive probability, we have $E(S,\bar{S}) - 2D \cdot \mathsf{LP} \cdot \min(|S|,|\bar{S}|) \leq 0$, and

$$\phi(S) = \frac{E(S,\bar{S})}{\min(|S|,|\bar{S}|)} \leq 2D \cdot \mathsf{LP} \leq 2D \cdot \phi^*.$$

That is, the cut $(S,\bar{S})$ is a $2D = O(\log n)$ approximate solution for Sparsest Cut.

We conclude that if we sample random cut S sufficiently many times, we will find a $O(\log n)$ approximate solution. Let us now consider the general case. Suppose that we are given an embedding of (V,d) in space ℓ_1^N. We say that a metric $\rho(u,v)$ is a convex combination of cut metrics if there exists a distribution of cuts $\mathcal{D}$ such that

$$\rho(u,v) = \mathbb{E}_{S\sim\mathcal{D}}[\delta_S(u,v)].$$

Lemma 3.1. *Let (X, ρ) is a bounded metric space. Suppose that X embeds isometrically into ℓ_1^N. Then for some $\beta > 0$, the metric $\beta\rho$ is a convex combination of cuts.*

Proof. Let M be the diameter of X. Let f be embedding of X in ℓ_1^N. Without loss of generality all coordinates of $f(u)$ lie between 0 and M for every u. Define the distribution $\mathcal{D}$ of cuts $(S, \bar{S})$ as follows

- Choose k uniformly at random from $\{1, \dots, N\}$.
- Choose t uniformly at random from $[0, M]$.
- Let $S = \{u : f(u)_k \le t\}$.

We have for $\beta = 1/(MN)$,

$$\mathbb{E}_{S\sim\mathcal{D}}[\delta_S(u,v)] = \frac{1}{N}\sum_{k=1}^{N}\frac{|f(u)_k - f(v)_k|}{M} = \frac{1}{MN}\|f(u) - f(v)\|_1 = \frac{\rho(u,v)}{MN} = \beta\rho(u,v).$$

There is a small flaw in this proof. The set S may be equal to $\varnothing$ or V. In this case, $(S, \bar{S})$ is not a cut! We condition $\mathcal{D}$ on the event that $S \ne \varnothing$ and $S \ne V$ and get a distribution $\mathcal{D}'$ s.t.

$$\mathbb{E}_{S\sim\mathcal{D}'}[\delta_S(u,v)] = \beta'\rho(u,v),$$

where $\beta' = \beta / \Pr_{S\sim\mathcal{D}}(S \ne \varnothing \text{ and } S \ne V)$. □

The lemma shows that there exists a distribution of cuts $(S, \bar{S})$ s.t. $\mathbb{E}_{S\sim\mathcal{D}'}[\delta_S(u,v)] = \beta'\rho(u,v)$. Up to scaling, this is the same identity as (1). The argument for $\{0,1\}$–valued r.v. shows that for some cut $(S, \bar{S})$ in the support of the distribution $\mathcal{D}$

$$\phi(S) \le 2D \cdot \mathsf{LP} \le 2D \cdot \phi^*.$$

3.1 Efficient Algorithm

We showed how to sample random cuts $(S, \bar{S})$ so that with positive probability $\phi(S) \le 2D \cdot \phi^*$. However, we do not have a lower bound on the success probability. Hypothetically, we may need to perform exponentially many trials before we find a cut $(S, \bar{S})$ with sparsity at most $2D \cdot \phi^*$.

Let us say that a cut $(S, \bar{S})$ is a *threshold cut* if $S = \{u : f(u)_k < t\}$ and $\bar{S} = \{u : f(u)_k \ge t\}$ for some $k \in \{1, \dots, n\}$ and $t \in \mathbb{R}$. Note that our distribution of cuts is supported on the set of threshold cuts. Therefore, our proof shows that there is a threshold cut $(S, \bar{S})$ with $\phi(S) \le 2D \cdot \phi^*$. We formally present our algorithm.

Approximation Algorithm for Sparsest Cut
Input: graph $G = (V, E)$.
Output: a cut $(S, \bar{S})$ with $\phi(S) \leq O(\log n)\phi^*_G$.

1. Solve the LP relaxation for Sparsest Cut. Define $d(x, y) = d_{xy}$.
2. Find an embedding $f : (V, d) \hookrightarrow \ell_1^N$ with distortion $D = O(\log n)$ and $N = O(\log^2 n)$.
3. Go over all threshold cuts $(S, \bar{S})$ and compute $\phi(S)$ for each of them.
4. Output the sparsest cut among all threshold cuts.

The algorithm examines at most $O(n \log^2 n)$ different threshold cuts and thus clearly runs in polynomial-time.

3.2 Notes

Leighton and Rao designed an $O(\log n)$ approximation algorithm for Sparsest Cut in 1988. Later, Aumann and Rabani, and Linial, London, and Rabinovich showed the connection between the Sparsest Cut problem and low-distortion metric embedding into ℓ_1.