COMPSCI 638: Graph Algorithms	November 15, 2019
Lecture 24	
Lecturer: Debmalya Panigrahi	*Scribe: Kevin Sun*

1 Overview

In this lecture, we revisit the sparsest cut problem. Recall that in Lecture 7, we gave an $O(\log n)$-approximation algorithm using a linear program. Today, we review that algorithm and prove that the integrality gap of this linear program is $\Omega(\log n)$.

2 Sparsest Cut

In the sparsest cut problem (see Lecture 6), we are given an undirected graph $G = (V, E)$ on n vertices, and each edge (i, j) has capacity $c(i, j) \geq 0$. The *sparsity* of a cut $S \subset V$ is defined as $\delta(S)/\min(|S|, |\overline{S}|)$, where $\delta(S)$ denotes the total capacity of edges crossing S and $\overline{S} = V \setminus S$. Our goal in this problem is to find a cut with minimum sparsity.

The *flux* of a cut $S \subset V$ is similar to the sparsity: it is defined as $\delta(S)/|S| \cdot |\overline{S}|$. Notice that $|S| \leq |\overline{S}|$ implies $n/2 \leq |\overline{S}| \leq n$, which means dividing the sparsity of S by its flux yields a 2-approximation of the sparsity. Thus, up to constant factors, the flux and sparsity of a cut are equivalent, so our goal is to find a cut with minimum sparsity.

2.1 An LP-based approach

Recall that in Lecture 7, we showed that finding the sparsest cut is equivalent to finding an elementary cut metric $d : V \times V \to \{0, 1\}$ that minimizes

$$\phi(d) = \frac{\sum_{i,j} c(i, j) d_{ij}}{\sum_{i,j} d_{ij}}.$$

Furthermore, we showed that the set of *cut* metrics on V is equivalent to the set of ℓ_1-metrics. Since our objective is a minimization, we can perform our search over ℓ_1-metrics to obtain an elementary cut metric that achieves the same objective value.

Finally, instead of optimizing of ℓ_1-metrics, we choose to optimize over general metrics. This results in a linear program that is equivalent to minimizing $\phi(d)$:

$$\text{(P): } \min \sum_{i,j} c(i, j) d_{ij}$$
$$\sum_{i,j} d_{i,j} \geq 1$$
$$d \text{ is a metric.}$$

As we saw, Bourgain's theorem allows us to approximate *any* metric by an ℓ_1-metric with $O(\log n)$ distortion, so the final result is an $O(\log n)$-approximation for the sparsest cut problem.

Constant degree expanders: To analyze the integrality gap of (P), consider the following procedure, which constructs a graph H from a graph G with the same vertex set. For each vertex v, select three neighbors of v in G uniformly at random and add the corresponding three edges in H. The resulting graph has $3n$ vertices, so it is very sparse. Intuitively, this suggests that there are vertices that are very far from each other, but as we will see, this is not true.

Consider the following "ideal" situation: the graph H generated above is a rooted tree, where each vertex has three children. Such a tree has $O(\log n)$ layers, so the distance between any two vertices in H is $O(\log n)$. This is not precisely true, but it can be shown that with high probability, the maximum pairwise distance in H is $\Theta(\log n)$.

Integrality gap: Notice that in the constant-degree expander H constructed above, the sparsity of any cut is roughly $3 = \Theta(1)$, so the flux of H (i.e., the flux of the minimum-flux cut in H) is $\Theta(1/n)$. We will now show that the optimum value of the linear program (P) is as low as $\Theta(1/n \log n)$, which implies that its integrality gap is $\Omega(\log n)$.

Instead of finding the optimal value of (P), we find the optimal value of its dual. Recall that the dual of sparsest cut is the maximum concurrent flow problem: we seek to find a set of flows whose sum is feasible (i.e., respects capacity constraints) and the minimum flow value within this set is maximized. Letting f_p denote the flow value on path p, the LP formulation is the following:

$$\begin{aligned} &\text{(D): } \max \lambda \\ &\sum_{p \in P_{ij}} f_p \geq \lambda \quad \forall i, j \\ &\sum_{p:(i,j) \in p} f_p \leq c(i,j) \quad \forall (i,j) \in E \\ &f_p \geq 0 \quad \forall p, \end{aligned}$$

where P_{ij} denotes the set of paths from vertex i to j. By strong duality, the optimal values of (P) and (D) are equal, so now, our goal bound the optimal value of (D) on a the graph H by $\Theta(\log n)$.

Recall that with high probability, the maximum distance between any pair of vertices in H is $\Theta(\log n)$. Furthermore, it can be shown that many pairs are this far apart: there are $\Theta(n^2)$ pairs of vertices that have a distance of $\Theta(\log n)$ from one another.

To support a flow of value λ between these pairs, the total capacity of all edges must be at least $\lambda n^2 \log n$. However, the if each edge has unit capacity, then the total capacity in H is $3n$. Therefore, we must have $\lambda = \theta(1/n \log n)$, so this is the optimal value of (D), as desired.

3 Summary

In this lecture, we showed that the integrality gap of the sparsest cut LP relaxation from Lecture 7 is $\Omega(\log n)$. Thus, any algorithm based on this LP cannot have a better approximation ratio than $O(\log n)$. In the next lecture, we will use an SDP to overcome this barrier.