11 Obtain ℓ_1-embeddings from cutting schemes

In this lecture we show that the cutting schemes by Bartal and FRT that were used to get an embedding into a distribution over dominating trees can also be used to obtain a subset-embedding with low ℓ_1-distortion.

Definition 11.1 *A β-cutting scheme is a randomized procedure that gets as input a graph $G = (V, E)$ and a paramter δ, and outputs a partitioning $C_1, C_2, \ldots$ of V such that*

- *each component has weak diameter at most δ, and*
- *with constant probability a node is far away from the boundary of its component, i.e.,*

$$\mathbf{Pr}[B(x, \beta \cdot \delta) \text{ is cut}] \leq 1/2 \ ,$$

where $C(x)$ denotes the component that contains x.

Note that the cutting schemes presented in the previous class fulfill the above definition for $\beta = \Theta(\frac{1}{\log n})$. Next, we will show how to use the cutting scheme to generate good subset embeddings for metrics. In particular, one can prove an approximate version of Bourgain's theorem using the cutting schemes given in class.

We will work with the same assumptions as in the previous lectures, i.e., that all edge lengths are one, and the diameter of the graph G (which represents the metric (X, d)) is Δ. We will show how to use a partitioning scheme to prove the following theorem:

Theorem 11.2 *There exists a subset embedding ϕ_1 such that*

$$(X, d) \xhookrightarrow{O(\frac{\log \Delta}{\beta})} \ell_1^{O(\log n \log \Delta)} \tag{11.1}$$

Suppose, we wanted to maintain the distance between u and v which are at distance $\delta < d(u, v) \leq 2\delta$, we could do the following: perform a random partitioning with parameter δ and consider the sum of the distances of u and v from the boundary of the components they lie in. Since the partition ensures that the resulting components have diamater at most δ, the vertices u and v lie in different components in the partition; furthermore, the above observation implies that with constant probability, at least one of u and v lie at distance $\rho = \Omega(\beta \cdot \delta)$ from the boundary of the clusters.

Hence, if we look at the $\log \Delta$ partitions done with $\delta = 2^i$ for $i = 1, 2, \ldots, \log \Delta$, for the particular choice of $\delta \approx d(u, v)$, the sum of the distances of u and v from the boundaries of their clusters is a reasonable approximation ($\approx \beta \cdot d(u, v)$) to their real distance. We now have to figure out how to emulate this by embeddings into ℓ_1.

We define the mapping ϕ_1:

1. For each *level* $i = 0, \ldots, \log \Delta$,
2. For $j = 1, \ldots, O(\log n)$,

Use the above scheme to obtain a random partition Π_{ij} of the graph G with $\delta_i = 2^i$. For each connected component $C_1, C_2, \ldots, C_{k(i,j)}$ formed in the partition Π_{ij}, choose a random color $\mathrm{col}(C_l) \in \{0, 1\}$.

For $u \in V$, let $\mathrm{col}(u)$ be the color of the connected component in Π_{ij} that contains u. Let V_1 denote the set of all nodes x with $\mathrm{col}(x) = 1$.

$$f_{i,j}(u) := d(u, S_1.$$

3. Finally, set $\phi_1(u) := \bigoplus_{i,j} f_{i,j}(u)$.

What is the distortion of this embedding? First, let us observe that this subset embedding is a contraction.

Observation 11.3 *Since, we are dealing with a subset embedding every coordinate mapping f_{ij} is a contraction.*

This claim immediately implies that

$$\|\phi_1(u) - \phi_1(v)\|_1 = \sum_{i,j} |f_{i,j}(u) - f_{i,j}(v)| \leq O(\log n \log \Delta) \cdot d(u, v) \quad . \tag{11.2}$$

Next, we will argue that for all u and v, the distances in the embedding is at least a large fraction of $d(u, v)$.

Claim 11.4 $|\phi_1(u) - \phi_1(v)|_1 \geq \Omega(\beta \cdot d(u, v) \cdot \log n)$ *with probability* $1 - n^{-3}$.

Proof. The proof uses Chernoff bounds. Let $d(u, v) \in (2^i, 2^{i+1}]$, and consider the partition Π_{ij}. Let X_j be the indicator variable for the event

$$\mathcal{E} = \mathsf{Ball}(u, \beta \cdot 2^i) \text{ is not cut by } \Pi_{ij} \quad .$$

Note that the properties of the cutting scheme imply that $\mathbf{Pr}[X_j = 1] \geq 1/2$.

Now if event $\mathcal{E}$ were to happen, then it holds that $|f_{i,j}(u) - f_{i,j}(v)| \geq beta \cdot 2^i$ with probability $1/2$; indeed, with probability $1/2$, the node u is assigned color "1" and with probability $1/2$ the node v is assigned color "0". In this case $|f_{i,j}(u) - f_{i,j}(v)| \geq \beta \cdot 2^i$. Since, all events are indepent we conclude that with probability $1/8$ the coordinate ij gives a contribution of $\beta \cdot 2^i$ to the distance between u and v.

Let Y_j denote the indicator variable for the event that the coordinate ij gives contribution $\beta \cdot 2^i$. We can use the Chernoff-Hoeffding bounds and argue that, since each $Y_j = 1$ with constant probability, it holds that for some constant $c > 0$

$$\mathbf{Pr}\left[\sum_{j=1}^{O(\log n)} Y_j \geq c \log n\right] \geq 1 - \frac{1}{n^3} \quad .$$

This immediately implies that

$$\mathbf{Pr}\left[\sum_{j=1}^{O(\log n)} |f_{i,j}(u) - f_{i,j}(v)| \geq \Omega(\beta \cdot d(u,v) \cdot \log n)\right] \geq 1 - \frac{1}{n^3} \ ,$$

and hence $|\phi_1(u) - \phi_1(v)|_1 \geq \Omega(d(u,v))$ with probability $1 - 1/n^3$, thus proving the claim. ■

By taking a union bound we get that with probability at least $1 - 1/n$ all pairs have an ℓ_1-distance at least $\Theta(\beta \cdot d(u,v) \cdot \log n)$. To get the distortion of the embedding we have to divide the upper bound (Equation 11.2) on the length by the lower bound which gives $\Theta(\frac{\log \Delta}{\beta})$.

We can improve on this result by interpreting ϕ as an embedding into ℓ_2. Then the length is bounded by

$$\|\phi(u) - \phi(v)\|_2^2 \leq \sqrt{\log n \log \Delta} \cdot d(u,v) \ ,$$

since each coordinate is contracting. For the lower bound we again observe that at least a logarithmic number of coordinates have $|f_{ij}(u) - f_{ij}(v)| \geq \Theta(\beta \cdot d(u,v))$, which results in an ℓ_2 distance of at least $\sqrt{O(\log n)} \cdot \beta \cdot d(u,v)$. Dividing the upper bound by the lower bound gives an embedding with distortion $O(\sqrt{\log \Delta}/\beta) = O(\sqrt{\log n}/\beta)$. Since ℓ_2 embeds into ℓ_1 isometrically we also get an ℓ_1 embedding with this distortion.

In the coming lectures we will first see how to improve this result to a distortion of $O(\sqrt{\log n/\beta})$, then we will show how to construct cutting schemes with a $\beta = \Theta(1)$ for planar graphs, and finally we show how to obtain cutting schemes with $\beta = \frac{1}{\sqrt{\log n}}$ for squared ℓ_2-metrics.

References

[CGR05] Shuchi Chawla, Anupam Gupta, and Harald Räcke. An improved approximation to sparsest cut. In *Proceedings of the 16th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 102–111, 2005.

[CKR01] Joseph Cheriyan, Howard Karloff, and Yuval Rabani. Approximating directed multicuts. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 320–328, 2001.

[FRT03] Jittat Fakcharoenphol, Satish B. Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. In *Proceedings of the 35th ACM Symposium on Theory of Computing (STOC)*, pages 448–455, 2003.

[KLMN04] Robert Krauthgamer, James Lee, Manor Mendel, and Assaf Naor. Measured descent: A new embedding method for finite metrics. In *Proceedings of the 45th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 434–443, 2004.