College of Computer & Information Science — Fall 2007
Northeastern University — Handout 27 November 2007
CS G399: Algorithmic Power Tools I — Scribe: Shahzad Rajput

Lecture Outline:

- The Sparsest Cut Problem
- Previous Approximation Algorithms
- Arora-Rao-Vazirani Algorithm - $O(\sqrt{\log n})$ Approximation for the Sparsest Cut Problem
- Some Observations on High Dimension Geometry

In this lecture, we have introduced the $O(\sqrt{\log n})$ approximation algorithm for the sparsest cut problem. Good reference for the material covered in this lecture is an excellent text on graph partioning by Sanjeev Arora, Satish Rao and Umesh Vazirani [1].

1 The Sparsest Cut Problem

Recall the Sparsest Cut Problem. For a graph G = (V,E) and non-empty set S $\subset$ V of the vertices, (S,$\overline{S}$) denotes the set of edges of G with exactly one end vertex in S. An edge set of this form is called an *edge cut*, or cut. The *density* of an edge cut [S,$\overline{S}$] is the ratio between the number of edges that are present in the cut, and the maximum number of edges that are possible between S and $\overline{S}$: d(S,$\overline{S}$) = |(S,$\overline{S}$)|/(|S||$\overline{S}$|). An edge cut with minimum density is called a sparsest cut of the graph.

2 Previous Approximation Algorithms

- Eigenvalue approaches (Cheeger'70, Alon'85, Alon-Milman'85)
  Only yield factor n approximation.
- O(log n)-approximation via LP (multicommodity flows) - (Leighton-Rao'88)
- Embeddings of finite metric spaces into ℓ_1 - (Linial, London, Rabinovich'94, AR'94)
  Geometric approach; more general result (but still O(log n) approximation)

3 Arora-Rao-Vazirani algorithm - $O(\sqrt{\log n})$ approximation for the sparsest cut problem

The key idea underlying algorithms for graph partitioning is to spread out the vertices in some abstract space while not stretching the edges too much. Finding a good graph partition is then accomplished by partitioning this abstract space. [1]

The approach of ARV algorithm combines both geometry and the metric condition. In particular, this approach maps the vertices to points in an n dimensional space such that the average square distance between vertices is a fixed constant, but the average squared distance between the endpoints of edges is minimized. Furthermore, we insist that these squared distances form a metric. This embedding, which we refer to as an ℓ_2^2 representation, can be computed in polynomial time using semi-definite programming. It is shown in [1] that this ℓ_2^2 representation embeds into ℓ_1 with $O(\sqrt{\log n})$ distortion.

Definition 1. Given a graph G=(V,E), the goal in the UNIFORM SPARSEST CUT problem is to determine the cut $(S, \overline{S})$ (where |S|≤$|\overline{S}|$ without loss of generality) that minimizes

$$\frac{|E(S, \overline{S})|}{|S||\overline{S}|}$$

From previous lecture, we know that

$$\begin{aligned} & \min_S \frac{|E(S, \overline{S})|}{|S||\overline{S}|} \\ = & \min_{l_1} \frac{\sum_{i,j \epsilon E} d_{l_1}(i,j)}{\sum_{i,j} d_{l_1}(i,j)} \\ = & \min_{v_i \epsilon \{-1,+1\}} \frac{\sum_{i,j \epsilon E} |v_i - v_j|^2}{\sum_{i,j} |v_i - v_j|^2} \end{aligned}$$

Here

$$\sum_{i,j \epsilon E} |v_i - v_j|^2 = 4E(S, \overline{S})$$
$$\sum_{i,j} |v_i - v_j|^2 = 4|S||\overline{S}|$$

Relaxing v_i to arbitrary vector $(v_1, v_2, \cdots, v_n \epsilon \Re^n)$

$$\geq \min_{v_i \epsilon \Re^m} \frac{\sum_{i,j \epsilon E} \| v_i - v_j \|_2^2}{\sum_{i,j} \| v_i - v_j \|_2^2}$$

This can be written as following SDP:

$$\begin{aligned} & \min_{v_i \epsilon \Re^m} \sum_{i,j \epsilon E} \| v_i - v_j \|_2^2 \\ s.t. & \sum_{i,j} \| v_i - v_j \|_2^2 = n^2 \end{aligned}$$

n^2 is used for simplicity, as this would make the average distance equal to 1. We can use value other than n^2 and then scale the distance accordingly.

The cut integrality gap is very large, therefore they added *triangle inequality.*

Definition 2. (ℓ_2^2 REPRESENTATION) *An ℓ_2^2 representation of a graph is an assignment of a point (vector) to each node, say v_i assigned to node i, such that for all i, j, k:*

$$|v_i - v_j|^2 + |v_j - v_k|^2 \geq |v_i - v_k|^2 \quad \textit{(triangle inequality)}$$

An ℓ_2^2 representation of a graph is called a unit-ℓ_2^2 representation if all points lie on the unit sphere (or equivalently, all vectors have unit length.)

We will further strengthen it by adding another constraint that all the angles (in the triangle formed by these vectors) are acute angles.

This SDP is then solved and the collection of points returned can then be embedded in ℓ_1 with average distortion $O(\sqrt{\log n})$. For embedding ℓ_2^2 into ℓ_1, we rely on the Master Structure Theorem.

The Master Structure Theorem: Given $v_1, v_2, \cdots, v_n \epsilon \Re^m$

$$s.t. \frac{1}{n^2} \sum i, j \|v_i - v_j\|^2 = \Omega(1)$$
$$Squared\ distances\ form\ a\ metric\ (i.e.\ all\ three\ points\ subtend\ acute\ angle)$$
$$\forall\ i\ \|v_i\|^2 \leq 1$$

then $\exists$ S,T s.t. $|S| = \Omega(n)$, $|T| = \Omega(n)$ and d(S,T) = $\Omega(\frac{1}{\sqrt{\log n}})$

A tight example of $\Omega(\frac{1}{\sqrt{\log n}})$ is a hypercube with n points in log n dimensions.

4 Some Observations on High Dimension Geometry

- What is the maximum number of points in n-dimensions such that no three points subtend obtuse angle?
  - 2-dimensions 4 points (square)
  - 3-dimensions 8 points (cube)
  - ... ...
  - n-dimensions 2^n points
- High-Dimensional Spheres in Cubes

  In high dimensions, 99% of the mass of the sphere is concentrated at its center.

  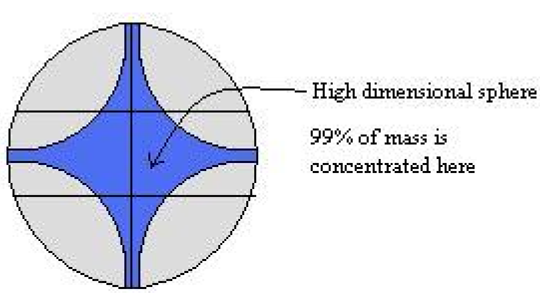

  We'll make the numbers easy by using a cube (in dimension N) with side length 2. Then when we cut the cube into orthants, each sub-cube has side length 1. Let's just look at one of those sub-cubes. The sphere S inscribed in that sub-cube has diameter 1. When we draw the central sphere, its center is on a corner of that subcube. Draw the diagonal from that corner to the opposite corner of the sub-cube. That diagonal has length $\sqrt{N}$. The part of the diagonal going through the sphere S has length 1 because that is the diameter of S. Of the part that is left over, half of it is in the central sphere, and in fact forms the radius of that central sphere. So the central sphere has radius $\frac{\sqrt{N}-1}{2}$. If N=9, the radius of the central sphere is 1, so it is just tangent to the cube. If N>9, then part of the central sphere bulges outside the cube! And, eventually the volume of the central sphere is actually larger than the cube. [2]

References

[1] Sanjeev Arora, Satish Rao and Umesh Vazirani. *Expander Flows, Geometric Embeddings and Graph Partitioning.* April 2007.

[2] http://www.math.hmc.edu/funfacts/ffiles/20007.2.shtml