Spectral Graph Theory and Graph Partitioning

Luca Trevisan[1]

Contents

1 Overview **3**

2 Expander Graphs and Sparse Cuts **4**

3 Eigenvalues and Eigenvectors **7**

3.1 Basic definitions 7

3.2 More on Eigenvalues and Eigenvectors 9

3.3 Proof of Theorem 5 11

4 Expansion and The Second Eigenvalue **13**

4.1 The Easy Direction 14

4.2 Other Relaxations of $\phi(G)$ 15

4.3 Proof of the Difficult Direction 16

4.3.1 Proof of Lemma 15 17

5 Tightness of the Bounds * **20**

5.1 Characters 21

5.2 A Look Beyond 25

5.3 Cayley Graphs and Their Spectrum 27

5.4 The Cycle 29

5.5 The Hypercube 29

6 Computing Eigenvalues and Eigenvectors * **30**

6.1 The Power Method to Approximate the Largest Eigenvalue 31

6.2 Approximating the Second Eigenvalue 34

7 The Leighton-Rao Approximation of Sparsest Cut **36**

[1]Computer Science Department, Stanford University. `trevisan@stanford.edu`. This work was supported in part by the NSF under grant CCF-1161812.

7.1 Formulating the Leighton-Rao Relaxation as a Linear Program 37
7.2 An L1 Relaxation of Sparsest Cut 38
7.3 A Theorem of Bourgain 40
7.4 Proof of Bourgain's Theorem 40
7.4.1 Preliminary and Motivating Examples 41
7.4.2 The Proof of Bourgain's Theorem 46
7.5 Tightness of the Analysis of the Leighton-Rao Relaxation * 48
7.6 The Non-Uniform Sparsest Cut Problem 50

8 The Arora-Rao-Vazirani Relaxation 51
8.1 The Ellipsoid Algorithm and Semidefinite Programming 54
8.2 Rounding the Arora-Rao-Vazirani Relaxation 58

1 Overview

This series of lectures is about *spectral methods in graph theory* and *approximation algorithms for graph partitioning problems.* We will study approximation algorithms for the *sparsest cut* problem, in which one wants to find a cut (a partition into two sets) of the vertex set of a given graph so that a minimal number of edges cross the cut compared to the number of pairs of vertices that are disconnected by the removal of such edges.

This problem is related to estimating the edge expansion of a graph and to find *balanced separators*, that is, ways to disconnect a constant fraction of the pairs of vertices in a graph after removing a minimal number of edges.

Finding balanced separators and sparse cuts arises in *clustering* problems, in which the presence of an edge denotes a relation of similarity, and one wants to partition vertices into few clusters so that, for the most part, vertices in the same cluster are similar and vertices in different clusters are not. For example, sparse cut approximation algorithms are used for *image segmentation*, by reducing the image segmentation problem to a graph clustering problem in which the vertices are the pixels of the image and the (weights of the) edges represent similarities between nearby pixels.

Balanced separators are also useful in the design of divide-and-conquer algorithms for graph problems, in which one finds a small set of edges that disconnects the graph, recursively solves the problem on the connected components, and then patches the partial solutions and the edges of the cut, via either exact methods (usually dynamic programming) or approximate heuristic. The sparsity of the cut determines the running time of the exact algorithms and the quality of approximation of the heuristic ones.

We will study three approximation algorithms:

1. The *Spectral Partitioning Algorithm*, based on linear algebra;
2. The *Leighton-Rao Algorithm*, based on a linear programming relaxation;
3. The *Arora-Rao-Vazirani Algorithm*, based on a semidefinite programming relaxation.

The three approaches are related, because the continuous optimization problem that underlies the Spectral Partitioning algorithm is a relaxation of the ARV semidefinite programming relaxation, and so is the Leighton-Rao relaxation. Rounding the Leighton-Rao and the Arora-Rao-Vazirani relaxations raise interesting problems in metric geometry, some of which are still open.

2 Expander Graphs and Sparse Cuts

Before giving the definition of expander graph, it is helpful to consider examples of graphs that are not expanders, in order to gain intuition about the type of "bad examples" that the definition is designed to avoid.

Suppose that a communication network is shaped as a path, with the vertices representing the communicating devices and the edges representing the available links. The clearly undesirable feature of such a configuration is that the failure of a single edge can cause the network to be disconnected, and, in particular, the failure of the middle edge will disconnect half of the vertices from the other half.

This is a situation that can occur in reality. Most of Italian highway traffic is along the highway that connect Milan to Naples via Bologna, Florence and Rome. The section between Bologna and Florence goes through relatively high mountain passes, and snow and ice can cause road closure. When this happens, it is almost impossible to drive between Northern and Southern Italy. Closer to California, I was once driving from Banff, a mountain resort town in Alberta which hosts a mathematical institute, back to the US. Suddenly, traffic on Canada's highway 1 came to a stop. People from the other cars, after a while, got out of the cars and started hanging out and chatting on the side of the road. We asked if there was any other way to go in case whatever accident was ahead of us would cause a long road closure. They said no, this is the only highway here. Thankfully we started moving again in half an hour or so.

Now, consider a two-dimensional $\sqrt{n} \times \sqrt{n}$ grid. The removal of an edge cannot disconnect the graph, and the removal of a constant number of edges can only disconnected a constant number of vertices from the rest of the graph, but it is possible to remove just $\sqrt{n}$ edges, a $1/O(\sqrt{n})$ fraction of the total, and have half of the vertices be disconnected from the other half.

A k-dimensional hypercube with $n = 2^k$ is considerably better connected than a grid, although it is still possible to remove a vanishingly small fraction of edges (the edges of a dimension cut, which are a $1/k = 1/\log_2 n$ fraction of the total number of edges) and disconnect half of the vertices from the other half.

Clearly, the most reliable network layout is the clique; in a clique, if an adversary wants to disconnect a p fraction of vertices from the rest of the graph, he has to remove at least a $p \cdot (1-p)$ fraction of edges from the graph.

This property of the clique will be our "gold standard" for reliability. The expansion and the sparsest cut parameters of a graph measure how worse a graph is compared with a clique from this point.

Definition 1 (Sparsest Cut) *Let $G = (V, E)$ be a graph and let $(S, V - S)$ be a partition of the vertices (a cut). Then the sparsity of the cut is*

$$\phi(S) := \frac{E(S, V-S)}{|E|} \cdot \left(\frac{|S| \cdot |V-S|}{|V|^2/2} \right)^{-1}$$

where $E(S, V-S)$ is the number of edges in E that have one endpoint in S and one endpoint in $V-S$.

The sparsest cut is, given a graph, to find the set of minimal sparsity. The sparsity of a graph $G = (V, E)$ is

$$\phi(G) := \min_{S \subseteq V : S \neq \emptyset, S \neq V} \phi(S)$$

That is, we are looking at the ratio between the fraction of edges that need to be removed in order to disconnect S from $V-S$ and the fraction of pairs of vertices that would be so disconnected.

It is more common to define the sparsity as

$$\frac{E(S, V-S)}{|S| \cdot |V-S|}$$

without the normalizing factor $(V^2/2|E|)$; the normalized definition used above yields simpler formulas in some of the applications that we will discuss later.

Note that if G is a d-regular graph, then

$$\phi(S) := \frac{E(S, V-S)}{\frac{d}{|V|} \cdot |S| \cdot |V-S|}$$

In a d-regular graph, the *edge expansion* of a cut $(S, V-S)$ is the related quantity

$$h(S) := \frac{E(S, V-S)}{d \cdot \min\{|S|, |V-S|\}}$$

in which we look at the ratio between the number of edges between S and $V-S$ and the obvious upper bound given by the total number of edges incident on the smaller side of the cut.

The edge expansion $h(G)$ of a graph is the minimum of $h(S)$ over all non-trivial partitions $(S, V-S)$.

(It is common to define the edge expansion without the normalizing factor of d in the denominator.)

We note that for every regular graph G we have that, for every set S,

$$\phi(S) \leq h(S) \leq 2 \cdot \phi(S)$$

and hence

$$\phi(G) \leq h(G) \leq 2 \cdot \phi(G)$$

A family of constant degree expanders is a family of (multi-)graphs $\{G_n\}_{n \geq d}$ such that each graph G_n is a d-regular graph with n vertices and such that there is an absolute constant $h > 0$ such that $h(G_n) \geq h$ for every n.

Constant-degree graphs of constant expansion are sparse graphs with exceptionally good connectivity properties. For example, we have the following observation.

Lemma 2 *Let $G = (V, E)$ be a regular graph of expansion h. Then, after an $\epsilon < h$ fraction of the edges are adversarially removed, the graph has a connected component that spans at least $1 - \epsilon/2h$ fraction of the vertices.*

Proof: Let d be the degree of G, and let $E' \subseteq E$ be an arbitrary subset of $\leq \epsilon|E| = \epsilon \cdot d \cdot |V|/2$ edges. Let $C_1, \ldots, C_m$ be the connected components of the graph $(V, E - E')$, ordered so that $|C_1| \geq |C_2| \geq \cdots \geq |C_m|$. We want to prove that $|C_1| \geq |V| \cdot (1 - 2\epsilon/h)$. We have

$$|E'| \geq \frac{1}{2} \sum_{i \neq j} E(C_i, C_j) = \frac{1}{2} \sum_i E(C_i, V - C_i)$$

If $|C_1| \leq |V|/2$, then we have

$$|E'| \geq \frac{1}{2} \sum_i d \cdot h \cdot |C_i| = \frac{1}{2} \cdot d \cdot h \cdot |V|$$

but this is impossible if $h > \epsilon$.

If $|C_1| \geq |V|/2$, then define $S := C_2 \cup \cdots \cup C_m$. We have

$$|E'| \geq E(C_1, S) \geq d \cdot h \cdot |S|$$

which implies that $|S| \leq \frac{\epsilon}{2h} \cdot |V|$ and so $C_1 \geq \left(1 - \frac{\epsilon}{2h}\right) \cdot |V|$. $\square$

In words, this means that, in a d-regular expander, the removal of k edges can cause at most $O(k/d)$ vertices to be disconnected from the remaining "giant component." Clearly, it is always possible to disconnect k/d vertices after removing k edges, so the reliability of an expander is essentially best possible.

3 Eigenvalues and Eigenvectors

Spectral graph theory studies how the eigenvalues of the adjacency matrix of a graph, which are purely algebraic quantities, relate to combinatorial properties of the graph.

3.1 Basic definitions

We begin with a brief review of linear algebra.

If $x = a + ib$ is a complex number, then we let $x^* = a - ib$ denote its *conjugate*.

If $M \in \mathbb{C}^{n \times n}$ is a square matrix, $\lambda \in \mathbb{C}$ is a scalar, $\mathbf{v} \in \mathbb{C}^n - \{\mathbf{0}\}$ is a non-zero vector and we have

$$M\mathbf{v} = \lambda \mathbf{v} \tag{1}$$

then we say that λ is an *eigenvalue* of M and that $\mathbf{v}$ is *eigenvector* of M corresponding to the eigenvalue λ.

When (1) is satisfied, then we equivalently have

$$(M - \lambda I) \cdot \mathbf{v} = \mathbf{0}$$

for a non-zero vector $\mathbf{v}$, which is equivalent to

$$\det(M - \lambda I) = 0 \tag{2}$$

For a fixed matrix M, the function $\lambda \to \det(M - \lambda I)$ is a univariate polynomial of degree n in λ and so, over the complex numbers, the equation (2) has exactly n solutions, counting multiplicities.

If $G = (V, E)$ is a graph, then we will be interested in the adjacency matrix A of G, that is the matrix such that $A_{ij} = 1$ if $(i, j) \in E$ and $A_{ij} = 0$ otherwise. If G is a multigraph or a weighted graph, then A_{ij} is equal to the number of edges between (i, j), or the weight of the edge (i, j), respectively.

The adjacency matrix of an undirected graph is symmetric, and this implies that its eigenvalues are all real.

Definition 3 *A matrix $M \in \mathbb{C}^{n \times n}$ is Hermitian if $M_{ij} = M_{ji}^*$ for every i, j.*

Note that a real symmetric matrix is always Hermitian.

Lemma 4 *If M is Hermitian, then all the eigenvalues of M are real.*

PROOF: Let M be an Hermitian matrix and let λ be a scalar and $\mathbf{x}$ be a non-zero vector such that $M\mathbf{x} = \lambda\mathbf{x}$. We will show that $\lambda = \lambda^*$, which implies that λ is a real number. We define the following *inner product* operation over vectors in $\mathbb{C}^n$:

$$\langle \mathbf{v}, \mathbf{w} \rangle := \sum_i v_i^* \cdot w_i$$

Notice that, by definition, we have $\langle \mathbf{v}, \mathbf{w} \rangle = (\langle \mathbf{w}, \mathbf{v} \rangle)^*$ and $\langle \mathbf{v}, \mathbf{v} \rangle = ||\mathbf{v}||^2$. The lemma follows by observing that

$$\begin{aligned} &\langle M\mathbf{x}, \mathbf{x} \rangle \\ &= \sum_i \sum_j M_{ij}^* x_j^* x_i \\ &= \sum_i \sum_j M_{ji} x_i x_j^* \\ &= \langle \mathbf{x}, M\mathbf{x} \rangle \end{aligned}$$

where we use the fact that M is Hermitian, and that

$$\langle M\mathbf{x}, \mathbf{x} \rangle = \langle \lambda\mathbf{x}, \mathbf{x} \rangle = \lambda^* ||x||^2$$

and

$$\langle \mathbf{x}, M\mathbf{x} \rangle = \langle \mathbf{x}, \lambda\mathbf{x} \rangle = \lambda ||x||^2$$

so that $\lambda = \lambda^*$. $\square$

From the discussion so far, we have that if A is the adjacency matrix of an undirected graph then it has n real eigenvalues, counting multiplicities of the number of solutions to $\det(A - \lambda I) = 0$.

If G is a d-regular graph, then instead of working with the adjacency matrix of G it is somewhat more convenient to work with the normalized matrix $M := \frac{1}{d} \cdot A$.

In the rest of this section we shall prove the following relations between the eigenvalues of M and certain purely combinatorial properties of G.

Theorem 5 *Let G be a d-regular undirected graph, and $M = \frac{1}{d} \cdot A$ be its normalized adjacency matrix. Let $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$ be the real eigenvalues of M with multiplicities. Then*

1. *$\lambda_1 = 1$ and $\lambda_n \geq -1$.*

2. $\lambda_2 = 1$ *if and only if* G *is disconnected.*

3. $\lambda_n = -1$ *if and only if at least one of the connected components of* G *is bipartite.*

In the next lecture we will begin to explore an "approximate" version of the second claim, and to show that λ_2 is close to 1 if and only if G has a sparse cut.

3.2 More on Eigenvalues and Eigenvectors

In order to relate the eigenvalues of the adjacency matrix of a graph to combinatorial properties of the graph, we need to first express the eigenvalues and eigenvectors as solutions to *optimization problems*, rather than solutions to algebraic equations.

First, we observe that if M is a real symmetric matrix and λ is a real eigenvalue of M, then λ admits a real eigenvector. This is because if $M\mathbf{x} = \lambda\mathbf{x}$ for some $\mathbf{x} \in \mathbb{C}^n$, then we also have $M\mathbf{x}' = \lambda\mathbf{x}'$, where $\mathbf{x}' \in \mathbb{R}^n$ is the vector whose i-th coordinate is the real part of the i-th coordinate of $\mathbf{x}$. Now, if λ is a (real) eigenvalue of a symmetric real matrix M, then the set $\{\mathbf{x} \in \mathbb{R}^n : M\mathbf{x} = \lambda\mathbf{x}\}$ is a vector subspace of $\mathbb{R}^n$, called the *eigenspace* of λ.

Fact 6 *If* $\lambda \neq \lambda'$ *are two distinct eigenvalues of a symmetric real matrix* M*, then the eigenspaces of* λ *and* λ' *are orthogonal.*

Proof: Let $\mathbf{x}$ be an eigenvector of λ and $\mathbf{y}$ be an eigenvector of λ'. From the symmetry of M and the fact that M, $\mathbf{x}$ and $\mathbf{y}$ all have real entries we get

$$\langle M\mathbf{x}, \mathbf{y}\rangle = \langle \mathbf{x}, M\mathbf{y}\rangle$$

but

$$\langle M\mathbf{x}, \mathbf{y}\rangle = \lambda \cdot \langle \mathbf{x}, \mathbf{y}\rangle$$

and

$$\langle \mathbf{x}, M\mathbf{y}\rangle = \lambda' \cdot \langle \mathbf{x}, \mathbf{y}\rangle$$

so that

$$(\lambda - \lambda') \cdot \langle \mathbf{x}, \mathbf{y}\rangle = 0$$

which implies that $\langle \mathbf{x}, \mathbf{y}\rangle = 0$, that is, that $\mathbf{x}$ and $\mathbf{y}$ are orthogonal. $\square$

Definition 7 *The algebraic multiplicity of an eigenvalue* λ *of a matrix* M *is the multiplicity of* λ *as a root of the polynomial* $\det(M - \lambda I)$*. The geometric multiplicity of* λ *is the dimension of its eigenspace.*

The following is the only result of this section that we state without proof.

Fact 8 *If M is a symmetric real matrix and λ is an eigenvalue of M, then the geometric multiplicity and the algebraic multiplicity of λ are the same.*

This gives us the following "normal form" for the eigenvectors of a symmetric real matrix.

Fact 9 *If $M \in \mathbb{R}^{n\times n}$ is a symmetric real matrix, and $\lambda_1, \ldots, \lambda_n$ are its eigenvalues with multiplicities, and $\mathbf{v}_1$ is a length-1 eigenvector of λ_1, then there are vectors $\mathbf{v}_2, \ldots, \mathbf{v}_n$ such that $\mathbf{v}_i$ is an eigenvector of λ_i and $\mathbf{v}_1, \ldots, \mathbf{v}_n$ are orthonormal.*

Proof: For each eigenvalue, choose an orthonormal basis for its eigenspace. For λ_1, choose the basis so that it includes $\mathbf{v}_1$. $\square$

Finally, we get to our goal of seeing eigenvalue and eigenvectors as solutions to continuous optimization problems.

Lemma 10 *If M is a symmetric matrix and λ_1 is its largest eigenvalue, then*

$$\lambda_1 = \sup_{\mathbf{x}\in\mathbb{R}^n : ||\mathbf{x}||=1} \mathbf{x}^T M \mathbf{x}$$

Furthermore, the sup *is achieved, and the vectors achieving it are precisely the eigenvectors of λ_1.*

Proof: That the sup is achieved follows from the fact that the set $\{\mathbf{x} \in \mathbb{R}^n : ||\mathbf{x}|| = 1\}$ is compact and that $\mathbf{x} \to \mathbf{x}^T M \mathbf{x}$ is a continuous function.

If $\mathbf{v}_1$ is a length-1 eigenvector of λ_1, then

$$\sup_{\mathbf{x}\in\mathbb{R}^n : ||\mathbf{x}||=1} \mathbf{x}^T M \mathbf{x} \geq \mathbf{v}_1^T M \mathbf{v}_1 = \lambda_1$$

If $\mathbf{y}$ is a length-1 vector that achieves the sup, then let $\mathbf{v}_1, \ldots, \mathbf{v}_n$ be as in Fact 9 and write

$$\mathbf{y} = \alpha_1 \mathbf{v}_1 + \ldots \alpha_n \mathbf{v}_n$$

Then

$$\sup_{\mathbf{x}\in\mathbb{R}^n : ||\mathbf{x}||=1} \mathbf{x}^T M \mathbf{x} = \mathbf{y}^T M \mathbf{y} = \sum_i \alpha_i^2 \lambda_i$$

Since $\sum_i \alpha_i^2 = ||\mathbf{y}||^2 = 1$, we have

$$\sup_{\mathbf{x}\in\mathbb{R}^n:||\mathbf{x}||=1} \mathbf{x}^T M\mathbf{x} = \sum_i \alpha_i^2 \lambda_i \leq \lambda_1 \cdot \sum_i \alpha_i^2 = \lambda_1$$

Finally, we see that we have $\mathbf{y}^T M\mathbf{y} = \lambda_1$ precisely when, for every i such that $\alpha_i \neq 0$ we have $\lambda_i = \lambda_1$, that is, precisely when $\mathbf{y}$ is in the eigenspace of λ_1. $\square$

Similarly we can prove

Lemma 11 *If M is a symmetric matrix, λ_1 is its largest eigenvalue, and $\mathbf{v}_1$ is an eigenvector of λ_1, then*

$$\lambda_2 = \sup_{\mathbf{x}\in\mathbb{R}^n:||\mathbf{x}||=1,\mathbf{x}\perp\mathbf{v}_1} \mathbf{x}^T M\mathbf{x}$$

Furthermore, the sup *is achieved, and the vectors achieving it are precisely the eigenvectors of λ_2. (If $\lambda_1 = \lambda_2$, then the vectors achieving the* sup *are the eigenvalues of $\lambda_1 = \lambda_2$ which are orthogonal to $\mathbf{v}_1$.)*

And

Lemma 12 *If M is a symmetric matrix and λ_n is its largest eigenvalue, then*

$$\lambda_n = \inf_{\mathbf{x}\in\mathbb{R}^n:||\mathbf{x}||=1} \mathbf{x}^T M\mathbf{x}$$

Furthermore, the inf *is achieved, and the vectors achieving it are precisely the eigenvectors of λ_n.*

3.3 Proof of Theorem 5

We will make repeated use of the following identity, whose proof is immediate: if M is the normalized adjacency matrix of a regular graph, and $\mathbf{x}$ is any vector, then

$$\sum_{i,j} M_{i,j}(x_i - x_j)^2 = 2\mathbf{x}^T\mathbf{x} - 2\mathbf{x}^T M\mathbf{x} \tag{3}$$

That is,

$$\mathbf{x}^T M\mathbf{x} = \mathbf{x}^T\mathbf{x} - \frac{1}{2}\sum_{i,j} M_{i,j}(x_i - x_j)^2 \leq \mathbf{x}^T\mathbf{x}$$

And so

$$\lambda_1 = \max_{\mathbf{x}\in\mathbb{R}^n:||\mathbf{x}||=1} \mathbf{x}^T M\mathbf{x} \leq 1$$

If we take $\mathbf{1} = (1, \ldots, 1)$ to be the all-one vector, we see that $\mathbf{1}^T M\mathbf{1} = 1$, and so 1 is the largest eigenvalue of M, with $\mathbf{1}$ being one of the vectors in the eigenspace of 1.

So we have the following formula for λ_2:

$$\lambda_2 = \sup_{\mathbf{x}\in\mathbb{R}^n:||\mathbf{x}||=1,\sum_i x_i=0} \mathbf{x}^T M\mathbf{x}$$

where we equivalently expressed the condition $\mathbf{x} \perp \mathbf{1}$ as $\sum_i x_i = 0$.

Using (3), we have

$$\lambda_2 = 1 - \inf_{\mathbf{x}\in\mathbb{R}^n:||\mathbf{x}||=1,\sum_i x_i=0} \frac{1}{2}\sum_{ij} M_{ij}(x_i - x_j)^2$$

So, if $\lambda_2 = 1$, there must exist a non-zero $\mathbf{v} \in \mathbb{R}^n$ such that $\sum_i v_i = 0$ and $\sum_{ij} M_{ij}(v_i - v_j)^2 = 0$, but this means that, for every edge $(i, j) \in E$ of positive weight we have $v_i = v_j$, and so $v_i = v_j$ for every i, j which are in the same connected component. The conditions $\sum_i v_i = 0$ and $\mathbf{v} \neq \mathbf{0}$ imply that $\mathbf{v}$ has both positive and negative coordinates, and so the sets $A := \{i : v_i > 0\}$ and $B := \{i : v_i < 0\}$ are non-empty and disconnected, and so G is not connected.

Conversely, if G is disconnected, and S and $V - S$ are non-empty sets such that $E(S, V - S) = 0$, then we can define $\mathbf{v}$ so that $v_i = |S|/(|V - S|)$ if $i \notin S$, and $v_i = -|V - S|/|S|$ if $\in S$, so that $\sum_i v_i = 0$. This gives us a non-zero vector such that $\sum_{ij} M_{ij}(v_i - v_j)^2 = 0$ and, after dividing every coordinate by $||\mathbf{v}||$, a length-1 vector proving that $\lambda_2 \geq 1$.

Finally, to study λ_n we observe that for every vector $\mathbf{x} \in \mathbb{R}^n$ we have

$$\sum_{i,j} M_{i,j}(x_i + x_j)^2 = 2\mathbf{x}^T\mathbf{x} + 2\mathbf{x}^T M\mathbf{x}$$

and so

$$\lambda_n = \min_{\mathbf{x}\in\mathbb{R}^n:||\mathbf{x}||=1} \mathbf{x}^T M\mathbf{x}$$

$$= \min_{\mathbf{x}\in\mathbb{R}^n:||\mathbf{x}||=1} -\mathbf{x}^T\mathbf{x} + \frac{1}{2}\sum_{i,j} M_{i,j}(x_i + x_j)^2$$

$$= -1 + \min_{\mathbf{x}\in\mathbb{R}^n : ||\mathbf{x}||=1} \frac{1}{2} \sum_{i,j} M_{i,j}(x_i + x_j)^2$$

From which we see that it is always $\lambda_n \geq -1$, and that if $\lambda_n = -1$ then there must be a non-zero vector $\mathbf{x}$ such that $x_i = -x_j$ for every edge $(i,j) \in E$. Let i be a vertex such that $x_i = a \neq 0$, and define the sets $A := \{j : x_j = a\}$, $B := \{j : x_j = -a\}$ and $R = \{j : x_j \neq \pm a\}$. The set $A \cup B$ is disconnected from the rest of the graph, because otherwise an edge with an endpoint in $A \cup B$ and an endpoint in R would give a positive contribution to $\sum_{i,j} M_{i,j}(x_i + x_j)^2$; furthermore, every edge incident on a vertex on A must have the other endpoint in B, and vice versa. Thus, $A \cup B$ is a connected component, or a collection of connected components, of G which is bipartite, with the bipartition A, B.

4 Expansion and The Second Eigenvalue

Let $G = (V, E)$ be an undirected d-regular graph, A its adjacency matrix, $M = \frac{1}{d} \cdot A$ its normalized adjacency matrix, and $1 = \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$ be the eigenvalues of M.

Recall that we defined the *edge expansion* of a cut $(S, V - S)$ of the vertices of G as

$$h(S) := \frac{E(S, V - S)}{d \cdot \min\{|S|, |V - S|\}}$$

and that the edge expansion of G is $h(G) := \min_{S \subseteq V} h(S)$.

We also defined the related notion of the *sparsity* of a cut $(S, V - S)$ as

$$\phi(S) := \frac{E(S, V - S)}{\frac{d}{n} \cdot |S| \cdot |V - S|}$$

and $\phi(G) := \min_S \phi(S)$; the *sparsest cut* problem is to find a cut of minimal sparsity.

Recall also that in the last lecture we proved that $\lambda_2 = 1$ if and only if G is disconnected. This is equivalent to saying that $1 - \lambda_2 = 0$ if and only if $h(G) = 0$. In this lecture and the next we will see that this statement admits an *approximate version* that, qualitatively, says that $1 - \lambda_2$ is small if and only if $h(G)$ is small. Quantitatively, we have

Theorem 13 (Cheeger's Inequalities)

$$\frac{1 - \lambda_2}{2} \leq h(G) \leq \sqrt{2 \cdot (1 - \lambda_2)} \tag{4}$$

4.1 The Easy Direction

In this section we prove

Lemma 14 $1 - \lambda_2 \leq \phi(G)$

From which we have one direction of Cheeger's inequality, after recalling that $\phi(G) \leq 2h(G)$.

Let us find an equivalent restatement of the sparsest cut problem. If represent a set $S \subseteq V$ as a bit-vector $x \in \{0,1\}^V$, then

$$E(S, V - S) = \frac{1}{2} \cdot \sum_{ij} A_{ij} \cdot |x_i - x_j|$$

and

$$|S| \cdot |V - S| = \frac{1}{2} \cdot \sum_{ij} |x_i - x_j|$$

so that, after some simplifications, we can write

$$\phi(G) = \min_{\mathbf{x} \in \{0,1\}^V - \{\mathbf{0},\mathbf{1}\}} \frac{\sum_{ij} M_{ij} |x_i - x_j|}{\frac{1}{n} \sum_{ij} |x_i - x_j|} \tag{5}$$

Note that, when x_i, x_j take boolean values, then so does $|x_i - x_j|$, so that we may also equivalently write

$$\phi(G) = \min_{\mathbf{x} \in \{0,1\}^V - \{\mathbf{0},\mathbf{1}\}} \frac{\sum_{ij} M_{ij} |x_i - x_j|^2}{\frac{1}{n} \sum_{ij} |x_i - x_j|^2} \tag{6}$$

In the last lecture, we gave the following characterization of $1 - \lambda_2$:

$$1 - \lambda_2 = \min_{\mathbf{x} \in \mathbb{R}^V - \{\mathbf{0}\}, \mathbf{x} \perp \mathbf{1}} \frac{\sum_{ij} M_{ij} |x_i - x_j|^2}{2 \cdot \sum_i x_i^2}$$

Now we claim that the following characterization is also true

$$1 - \lambda_2 = \min_{\mathbf{x} \in \mathbb{R}^V - \{\mathbf{0},\mathbf{1}\}} \frac{\sum_{ij} M_{ij} |x_i - x_j|^2}{\frac{1}{n} \sum_{ij} |x_i - x_j|^2} \tag{7}$$

This is because

$$\sum_{i,j} |x_i - x_j|^2$$

$$= \sum_{ij} x_i^2 + \sum_{ij} x_j^2 - 2\sum_{ij} x_i x_j$$
$$= 2n \sum_i x_i^2 - 2\left(\sum_i x_i\right)^2$$

so for every $\mathbf{x} \in \mathbb{R}^V - \{\mathbf{0}\}$ such that $\mathbf{x} \perp \mathbf{1}$ we have that $2 \cdot \sum_i x_i^2 = \frac{1}{n}\sum_{ij} |x_i - x_j|^2$, and so

$$\min_{\mathbf{x}\in\mathbb{R}^V-\{\mathbf{0}\},\mathbf{x}\perp\mathbf{1}} \frac{\sum_{ij} M_{ij}|x_i - x_j|^2}{2 \cdot \sum_i x_i^2} = \min_{\mathbf{x}\in\mathbb{R}^V-\{\mathbf{0}\},\mathbf{x}\perp\mathbf{1}} \frac{\sum_{ij} M_{ij}|x_i - x_j|^2}{\frac{1}{n}\sum_{ij} |x_i - x_j|^2}$$

To conclude the argument, take an $\mathbf{x}$ that maximized the right-hand side of (7), and observe that if we shift every coordinate by the same constant then we obtain another optimal solution, because the shift will cancel in all the expressions both in the numerator and the denominator. In particular, we can define $\mathbf{x}'$ such that $x'_i = x_i - \frac{1}{n}\sum_j x_j$ and note that the entries of $\mathbf{x}'$ sum to zero, and so $\mathbf{x}' \perp \mathbf{1}$. This proves that

$$\min_{\mathbf{x}\in\mathbb{R}^V-\{\mathbf{0}\},\mathbf{x}\perp\mathbf{1}} \frac{\sum_{ij} M_{ij}|x_i - x_j|^2}{\frac{1}{n}\sum_{ij} |x_i - x_j|^2} = \min_{\mathbf{x}\in\mathbb{R}^V-\{\mathbf{0},\mathbf{1}\}} \frac{\sum_{ij} M_{ij}|x_i - x_j|^2}{\frac{1}{n}\sum_{ij} |x_i - x_j|^2}$$

and so we have established (7).

Comparing (7) and (6), it is clear that the quantity $1 - \lambda_2$ is a *continuous relaxation* of $\phi(G)$, and hence $1 - \lambda_2 \leq \phi(G)$.

4.2 Other Relaxations of $\phi(G)$

Having established that we can view $1 - \lambda_2$ as a *relaxation* of $\phi(G)$, the proof that $h(G) \leq \sqrt{2 \cdot (1 - \lambda_2)}$ can be seen as a *rounding* algorithm, that given a real-valued solution to (7) finds a comparably good solution for (6).

Later in the course we will see two more approximation algorithms for sparsest cut and edge expansion. Both are based on continuous relaxations of ϕ starting from (5).

The algorithm of Leighton and Rao is based on a relaxation that is defined by observing that every bit-vector $\mathbf{x} \in \{0,1\}^V$ defines the semi-metric $d(i,j) := |x_i - x_j|$ over the vertices; the Leighton-Rao relaxation is obtained by allowing arbitrary semi-metrics:

$$LR(G) := \min_{\substack{d : V \times V \to \mathbb{R} \\ d \text{ semimetric}}} \frac{\sum_{ij} M_{ij} d(i,j)}{\frac{1}{n}\sum_{ij} d(i,j)}$$

It is not difficult to express $LR(G)$ as a linear programming problem.

The algorithm of Arora-Rao-Vazirani is obtained by noting that, for a bit-vector $x \in \{0,1\}^V$, the distances $d(i,j) := |x_i - x_j|$ define a metric which can also be seen as the Euclidean distance between the x_i, because $|x_i - x_j| = \sqrt{(x_i - x_j)^2}$, and such that $d^2(i,j)$ is also a semi-metric, trivially so because $d^2(i,j) = d(i,j)$. If a distance function $d(\cdot,\cdot)$ is a semi-metric such that $\sqrt{d(\cdot,\cdot)}$ is a Euclidean semi-metric, then $d(\cdot,\cdot)$ is called a *negative type* semi-metric. The Arora-Rao-Vazirani relaxation is

$$ARV(G) := \min_{\substack{d : V \times V \to \mathbb{R} \\ d \text{ negative type semimetric}}} \frac{\sum_{ij} M_{ij} d(i,j)}{\frac{1}{n}\sum_{ij} d(i,j)}$$

The Arora-Rao-Vazirani relaxation can be expressed as a semi-definite programming problem.

From this discussion it is clear that the Arora-Rao-Vazirani relaxation is a tightening of the Leigthon-Rao relaxation and that we have

$$\phi(G) \geq ARV(G) \geq LR(G)$$

It is less obvious in this treatment, and we will see it later, that the Arora-Rao-Vazirani is also a tightening of the relaxation of ϕ given by $1 - \lambda_2$, that is

$$\phi(G) \geq ARV(G) \geq 1 - \lambda_2$$

The relaxations $1 - \lambda_2$ and $LR(G)$ are incomparable.

4.3 Proof of the Difficult Direction

The proof can be seen as an analysis of the following algorithm.

Algorithm: *SpectralPartitioning*

- Input: graph $G = (V, E)$ and vector $\mathbf{x} \in \mathbb{R}^V$
- Sort the vertices of V in non-decreasing order of values of entries in $\mathbf{x}$, that is let $V = \{v_1, \ldots, v_n\}$ where $x_{v_1} \leq x_{v_2} \leq \ldots x_{v_n}$
- Let $i \in \{1, \ldots, n-1\}$ be such that $h(\{v_1, \ldots, v_i\})$ is minimal
- Output $S = \{v_1, \ldots, v_i\}$

We note that the algorithm can be implemented to run in time $O(|V|+|E|)$, assuming arithmetic operations and comparisons take constant time, because once we have computed $h(\{v_1, \ldots, v_i\})$ it only takes time $O(degree(v_{i+1}))$ to compute $h(\{v_1, \ldots, v_{i+1}\})$.

We have the following analysis of the quality of the solution:

Lemma 15 (Analysis of Spectral Partitioning) *Let $G = (V, E)$ be a d-regular graph, $\mathbf{x} \in \mathbb{R}^V$ be a vector such that $\mathbf{x} \perp \mathbf{1}$, let M be the normalized adjacency matrix of G, define*

$$\delta := \frac{\sum_{i,j} M_{i,j}|x_i - x_j|^2}{\frac{1}{n}\sum_{i,j}|x_i - x_j|^2}$$

and let S be the output of algorithm SpectralPartitioning on input G and x. Then

$$h(S) \leq \sqrt{2\delta}$$

Remark 16 *If we apply the lemma to the case in which $\mathbf{x}$ is an eigenvector of λ_2, then $\delta = 1 - \lambda_2$, and so we have*

$$h(G) \leq h(S) \leq \sqrt{2 \cdot (1 - \lambda_2)}$$

which is the difficult direction of Cheeger's inequalities.

Remark 17 *If we run the SpectralPartitioning algorithm with the eigenvector $\mathbf{x}$ of the second eigenvalue λ_2, we find a set S whose expansion is*

$$h(S) \leq \sqrt{2 \cdot (1 - \lambda_2)} \leq 2\sqrt{h(G)}$$

Even though this doesn't give a constant-factor approximation to the edge expansion, it gives a very efficient, and non-trivial, approximation.

As we will see in a later lecture, there is a nearly linear time algorithm that finds a vector x for which the expression δ in the lemma is very close to $1 - \lambda_2$, so, overall, for any graph G we can find a cut of expansion $O(\sqrt{h(G)})$ in nearly linear time.

4.3.1 Proof of Lemma 15

In the past lecture, we saw that $1 - \lambda_2$ can be seen as the optimum of a continuous relaxation of sparsest cut. Lemma 15 provides a *rounding algorithm* for the real vectors which are solutions of the relaxation. In this section we will think of it as a form of *randomized rounding*. Later, when we talk about the Leighton-Rao sparsest cut algorithm, we will revisit this proof and think of it in terms of *metric embeddings*.

To simplify notation, we will assume that $V = \{1, \ldots, n\}$ and that $x_1 \leq x_2 \leq \cdots x_n$. Thus our goal is to prove that there is an i such that $h(\{1, \ldots, i\}) \leq \sqrt{2\delta}$

We will derive Lemma 15 by showing that there is a distribution D over sets S of the form $\{1, \ldots, i\}$ such that

$$\frac{\mathbb{E}_{S\sim D} \frac{1}{d} Edges(S, V - S)}{\mathbb{E}_{S\sim D} \min\{|S|, |V - S|\}} \leq \sqrt{2\delta} \tag{8}$$

We need to be a bit careful in deriving the Lemma from (8). In general, it is not true that a ratio of averages is equal to the average of the ratios, so (8) does not imply that $\mathbb{E}\, h(S) \leq \sqrt{2\delta}$. We can, however, apply linearity of expectation and derive from (8) the inequality

$$\mathop{\mathbb{E}}_{S\sim D} \frac{1}{d} Edges(S, V - S) - \sqrt{2\delta} \min\{|S|, |V - S|\} \leq 0$$

So there must exist a set S in the sample space such that

$$\frac{1}{d} Edges(S, V - S) - \sqrt{2\delta} \min\{|S|, |V - S|\} \leq 0$$

meaning that, for that set S, we have $h(S) \leq \sqrt{2\delta}$. (Basically we are using the fact that, for random variables X, Y over the same sample space, although it might not be true that $\frac{\mathbb{E}X}{\mathbb{E}Y} = \mathbb{E}\frac{X}{Y}$, we always have $\mathbb{P}[\frac{X}{Y} \leq \frac{\mathbb{E}X}{\mathbb{E}Y}] > 0$, provided that $Y > 0$ over the entire sample space.)

From now on, we will assume that

1. $x_{\lfloor n/2 \rfloor} = 0$, that is, the median of the entries of $\mathbf{x}$ is zero
2. $x_1^2 + x_n^2 = 1$

which can be done without loss of generality because adding a fixed constant c to all entries of $\mathbf{x}$, or multiplying all the entries by a fixed constant does not change the value of δ nor does it change the property that $x_1 \leq \cdots \leq x_n$. The reason for these choices is that they allow us to define a distribution D over sets such that

$$\mathop{\mathbb{E}}_{S\sim D} \min\{|S|, |V - S|\} = \sum_i x_i^2 \tag{9}$$

We define the distribution D over sets of the form $\{1, \ldots, i\}$, $i \leq n-1$, as the outcome of the following probabilistic process:

- We pick a real value t in the range $[x_1, x_n]$ with probabily density function $f(t) = 2|t|$. That is, for $x_1 \leq a \leq b \leq x_n$, $\mathbb{P}[a \leq t \leq b] = \int_a^b 2|t|dt$.

  Doing the calculation, this means that $\mathbb{P}[a \leq t \leq b] = |a^2 - b^2|$ if a, b have the same sign, and $\mathbb{P}[a \leq t \leq b] = a^2 + b^2$ if they have different signs.

- We let $S := \{i : x_i \leq t\}$

According to this definition, the probability that an element $i \leq n/2$ belongs to the smallest of the sets $S, V - S$ is the same as the probability that it belongs to S, which is the probability that the threshold t is in the range $[x_i, 0]$, and that probability is x_i^2. Similarly, the probability that an element $i > n/2$ belongs to the smallest of $S, V - S$ is the same as the probability that it belongs to $V - S$, which is the probability that t is in the range $[0, x_i]$, which is again x_i^2. So we have established (9).

We will now estimate the expected number of edges between S and $V - S$.

$$\mathbb{E}\,\frac{1}{d}Edges(S, V - S) = \frac{1}{2}\sum_{i,j} M_{i,j}\,\mathbb{P}[(i,j) \text{ is cut by } (S, V - S)]$$

The event that the edge (i, j) is cut by the partition $(S, V - S)$ happens when the value t falls in the range between x_i and x_j. This means that

- If x_i, x_j have the same sign,

$$\mathbb{P}[(i,j) \text{ is cut by } (S, V - S)] = |x_i^2 - x_j^2|$$

- If x_i, x_j have different sign,

$$\mathbb{P}[(i,j) \text{ is cut by } (S, V - S)] = x_i^2 + x_j^2$$

Some attempts, show that a good expression to upper bound both cases is

$$\mathbb{P}[(i,j) \text{ is cut by } (S, V - S)] \leq |x_i - x_j| \cdot (|x_i| + |x_j|)$$

Plugging into our expression for the expected number of cut edges, and applying Cauchy-Schwarz

$$\begin{aligned}\mathbb{E}\,\frac{1}{d}Edges(S, V - S) &\leq \frac{1}{2}\sum_{i,j} M_{i,j}|x_i - x_j| \cdot (|x_i| + |x_j|) \\ &\leq \frac{1}{2}\sqrt{\sum_{ij} M_{ij}(x_i - x_j)^2} \cdot \sqrt{\sum_{ij} M_{ij}(|x_i| + |x_j|)^2}\end{aligned}$$

The assumption of the Lemma tell us that

$$\sum_{ij} M_{ij}(x_i - x_j)^2 = \delta \frac{1}{n} \sum_{ij} (x_i - x_j)^2$$

And we can rewrite

$$\sum_{ij} (x_i - x_j)^2 = 2n \sum_i x_i^2 - 2 \sum_{ij} x_i x_j = 2n \sum_i x_i^2 - 2 \left(\sum_i x_i \right)^2 \leq 2n \sum_i x_i^2$$

which gives us

$$\sum_{ij} M_{ij}(x_i - x_j)^2 \leq 2\delta \sum_i x_i^2$$

Finally, it remains to study the expression $\sum_{ij} M_{ij}(|x_i| + |x_j|)^2$. By applying the inequality $(a+b)^2 \leq 2a^2 + 2b^2$ (which follows by noting that $2a^2 + 2b^2 - (a+b)^2 = (a-b)^2 \geq 0$), we derive

$$\sum_{ij} M_{ij}(|x_i| + |x_j|)^2 \leq \sum_{ij} M_{ij}(2x_i^2 + 2x_j^2) = 4 \sum_i x_i^2$$

Putting all the pieces together we have

$$\mathbb{E} \frac{1}{d} Edges(S, V - S) \leq \sqrt{2\delta} \cdot \sum_i x_i^2 \tag{10}$$

which, together with (9) gives (8), which, as we already discussed, implies the Main Lemma 15.

5 Tightness of the Bounds *

In this section, which is not required material, we prove that the bounds of the Cheeger inequalities are tight for certain graphs. Toward this goal, we develop a general theory that allows us to compute the eigenvalues of certain families of graphs, the *Cayley graphs of abelian groups*, including cycles and hypercubes, which will be our tight examples.

For readers familiar with the Fourier analysis of Boolean functions, or the discrete Fourier analysis of functions $f : \mathbb{Z}/N\mathbb{Z} \to \mathbb{C}$, or the standard Fourier analysis of

periodic real functions, this theory will give a more general, and hopefully interesting, way to look at what they already know.

5.1 Characters

We will use additive notation for groups, so, if Γ is a group, its unit will be denoted by 0, its group operation by $+$, and the inverse of element a by $-a$. Unless, noted otherwise, however, the definitions and results apply to non-abelian groups as well.

Definition 18 (Character) *Let Γ be a group (we will also use Γ to refer to the set of group elements). A function $f : \Gamma \to \mathbb{C}$ is a character of Γ if*

- *f is a group homomorphism of Γ into the multiplicative group $\mathbb{C} - \{0\}$.*
- *for every $x \in \Gamma$, $|f(x)| = 1$*

Though this definition might seem to not bear the slightest connection to our goals, the reader should hang on because we will see next time that finding the eigenvectors and eigenvalues of the cycle C_n is immediate once we know the characters of the group $\mathbb{Z}/n\mathbb{Z}$, and finding the eigenvectors and eigenvalues of the hypercube H_d is immediate once we know the characters of the group $(\mathbb{Z}/2\mathbb{Z})^d$.

Remark 19 (About the Boundedness Condition) *If Γ is a finite group, and a is any element, then*

$$\underbrace{a + \cdots + a}_{|\Gamma| \text{ times}} = 0$$

and so if $f : \Gamma \to \mathbb{C}$ is a group homomorphism then

$$1 = f(0) = f(\underbrace{a + \cdots + a}_{|\Gamma| \text{ times}}) = f(a)^{|\Gamma|}$$

and so $f(a)$ is a root of unity and, in particular, $|f(a)| = 1$. This means that, for finite groups, the second condition in the definition of character is redundant.

In certain infinite groups, however, the second condition does not follow from the first, for example $f : \mathbb{Z} \to \mathbb{C}$ defined as $f(n) = e^n$ is a group homomorphism of $(\mathbb{Z}, +)$ into $(\mathbb{C} - \{0\}, \cdot)$ but it is not a character.

Just by looking at the definition, it might look like a finite group might have an infinite number of characters; the above remark, however, shows that a character of a finite group Γ must map into $|\Gamma|$-th roots of unity, of which there are only $|\Gamma|$, showing a finite $|\Gamma|^{|\Gamma|}$ upper bound to the number of characters. Indeed, a much stronger upper bound holds, as we will prove next, after some preliminaries.

Lemma 20 *If Γ is finite and χ is a character that is not identically equal to 1, then $\sum_{a\in\Gamma}\chi(a)=0$*

PROOF: Let b be such that $\chi(b)\neq 1$. Note that

$$\chi(b)\cdot\sum_{a\in\Gamma}\chi(a)=\sum_{a\in\Gamma}\chi(b+a)=\sum_{a\in\Gamma}\chi(a)$$

where we used the fact that the mapping $a\to b+a$ is a permutation. (We emphasize that even though we are using additive notation, the argument applies to non-abelian groups.) So we have

$$(\chi(b)-1)\cdot\sum_{a\in\Gamma}\chi(a)=0$$

and since we assumed $\chi(b)\neq 1$, it must be $\sum_{a\in\Gamma}\chi(a)=0$. $\square$

If Γ is finite, given two functions $f,g:\Gamma\to\mathbb{C}$, define the inner product

$$\langle f,g\rangle:=\sum_{a\in\Gamma}f(a)g^*(a)$$

Lemma 21 *If $\chi_1,\chi_2:\Gamma\to\mathbb{C}$ are two different characters of a finite group Γ, then*

$$\langle\chi_1,\chi_2\rangle=0$$

We will prove Lemma 21 shortly, but before doing so we note that, for a finite group Γ, the set of functions $f:\Gamma\to\mathbb{C}$ is a $|\Gamma|$-dimensional vector space, and that Lemma 21 implies that characters are orthogonal with respect to an inner product, and so they are linearly independent. In particular, we have established the following fact:

Corollary 22 *If Γ is a finite group, then it has at most $|\Gamma|$ characters.*

It remains to prove Lemma 21, which follows from the next two statements, whose proof is immediate from the definitions.

Fact 23 *If χ_1,χ_2 are characters of a group Γ, then the mapping $x\to\chi_1(x)\cdot\chi_2(x)$ is also a character.*

Fact 24 *If χ is a character of a group Γ, then the mapping $x\to\chi^*(x)$ is also a character, and, for every x, we have $\chi(x)\cdot\chi^*(x)=1$.*

To complete the proof of Lemma 21, observe that:

- the function $\chi(x) := \chi_1(x) \cdot \chi_2^*(x)$ is a character;
- the assumption of the lemma is that there is an a such that $\chi_1(a) \neq \chi_2(a)$, and so, for the same element a, $\chi(a) = \chi_1(a) \cdot \chi_2^*(a) \neq \chi_2(a) \cdot \chi_2^*(a) = 1$
- thus χ is a character that is not identically equal to 1, and so

$$0 = \sum_a \chi(a) = \langle \chi_1, \chi_2 \rangle$$

Notice that, along the way, we have also proved the following fact:

Fact 25 *If Γ is a group, then the set of characters of Γ is also a group, with respect to the group operation of pointwise multiplication. The unit of the group is the character mapping every element to 1, and the inverse of a character is the pointwise conjugate of the character.*

The group of characters is called the Pontryagin dual of Γ, and it is denoted by $\hat{\Gamma}$.

We now come to the punchline of this discussion.

Theorem 26 *If Γ is a finite abelian group, then it has exactly $|\Gamma|$ characters.*

Proof: We give a constructive proof. We know that every finite abelian group is isomorphic to a product of cyclic groups

$$(\mathbb{Z}/n_1\mathbb{Z}) \times (\mathbb{Z}/n_2\mathbb{Z}) \times \cdots \times (\mathbb{Z}/n_k\mathbb{Z})$$

so it will be enough to prove that

1. the cyclic group $\mathbb{Z}/n\mathbb{Z}$ has n characters;
2. if Γ_1 and Γ_2 are finite abelian groups with $|\Gamma_1|$ and $|\Gamma_2|$ characters, respectively, then their product has $|\Gamma_1| \cdot |\Gamma_2|$ characters.

For the first claim, consider, for every $r \in \{0, \ldots, n-1\}$, the function

$$\chi_r(x) := e^{2\pi i r x/n}$$

Each such function is clearly a character (0 maps to 1, $\chi_r(-x)$ is the multiplicative inverse of $\chi_r(x)$, and, recalling that $e^{2\pi i k} = 1$ for every integer k, we also have $\chi_r(a +$

$b \bmod n) = e^{2\pi i r a/n} \cdot e^{2\pi i r b/n})$, and the values of $\chi_r(1)$ are different for different values of r, so we get n distinct characters. This shows that $\mathbb{Z}/n\mathbb{Z}$ has at least n characters, and we already established that it can have at most n characters.

For the second claim, note that if χ_1 is a character of Γ_1 and χ_2 is a character of Γ_2, then it is easy to verify that the mapping $(x, y) \to \chi_1(x) \cdot \chi_2(y)$ is a character of $\Gamma_1 \times \Gamma_2$. Furthermore, if (χ_1, χ_2) and (χ_1', χ_2') are two distinct pairs of characters, then the mappings $\chi(x, y) := \chi_1(x) \cdot \chi_2(y)$ and $\chi'(x, y) := \chi_1'(x) \cdot \chi_2'(y)$ are two distinct characters of $\Gamma_1 \times \Gamma_2$, because we either have an a such that $\chi_1(a) \neq \chi_1'(a)$, in which case $\chi(a, 0) \neq \chi'(a, 0)$, or we have a b such that $\chi_2(b) \neq \chi_2'(b)$, in which case $\chi(0, b) \neq \chi'(0, b)$. This shows that $\Gamma_1 \times \Gamma_2$ has at least $|\Gamma_1| \cdot |\Gamma_2|$ characters, and we have already established that it can have at most that many $\square$

This means that the characters of a finite abelian group Γ form an orthogonal basis for the set of all functions $f : \Gamma \to \mathbb{C}$, so that any such function can be written as a linear combination

$$f(x) = \sum_{\chi} \hat{f}(\chi) \cdot \chi(x)$$

For every character χ, $\langle \chi, \chi \rangle = |\Gamma|$, and so the characters are actually a scaled-up orthonormal basis, and the coefficients can be computed as

$$\hat{f}(\chi) = \frac{1}{|\Gamma|} \sum_{x} f(x) \chi^*(x)$$

Example 27 (The Boolean Cube) *Consider the case $\Gamma = (\mathbb{Z}/2\mathbb{Z})^n$, that is the group elements are $\{0, 1\}^n$, and the operation is bitwise xor. Then there is a character for every bit-vector $(r_1, \ldots, r_n)$, which is the function*

$$\chi_{r_1, \ldots, r_n}(x_1, \ldots, x_n) := (-1)^{r_1 x_1 + \cdots r_n x_n}$$

Every boolean function $f : \{0, 1\}^n \to \mathbb{C}$ can thus be written as

$$f(x) = \sum_{r \in \{0,1\}^n} \hat{f}(r) \cdot (-1)^{\sum_i r_i x_i}$$

where

$$\hat{f}(r) = \frac{1}{2^n} \sum_{x \in \{0,1\}^n} f(x) \cdot (-1)^{\sum_i r_i x_i}$$

which is the boolean Fourier transform.

Example 28 (The Cyclic Group) *To work out another example, consider the case* $\Gamma = \mathbb{Z}/N\mathbb{Z}$. *Then every function* $f : \{0, \ldots, N-1\} \to \mathbb{C}$ *can be written as*

$$f(x) = \sum_{r \in \{0,\ldots,N-1\}} \hat{f}(r) e^{2\pi i r x/n}$$

where

$$\hat{f}(x) = \frac{1}{N} \sum_{x} f(x) e^{-2\pi i r x/n}$$

which is the discrete Fourier transform.

5.2 A Look Beyond

Why is the term "Fourier transform" used in this context? We will sketch an answer to this question, although what we say from this point on is not needed for our goal of finding the eigenvalues and eigenvectors of the cycle and the hypercube.

The point is that it is possible to set up a definitional framework that unifies both what we did in the previous section with finite Abelian groups, and the Fourier series and Fourier transforms of real and complex functions.

In the discussion of the previous section, we started to restrict ourselves to finite groups Γ when we defined an inner product among functions $f : \Gamma \to \mathbb{C}$.

If Γ is an infinite abelian group, we can still define an inner product among functions $f : \Gamma \to \mathbb{C}$, but we will need to define a measure over Γ and restrict ourselves in the choice of functions. A measure μ over (a sigma-algebra of subsets of) Γ is a Haar measure if, for every measurable subset A and element a we have $\mu(a + A) = \mu(A)$, where $a + A = \{a + b : b \in A\}$. For example, if Γ is finite, $\mu(A) = |A|$ is a Haar measure. If $\Gamma = (\mathbb{Z}, +)$, then $\mu(A) = |A|$ is also a Haar measure (it is ok for a measure to be infinite for some sets), and if $\Gamma = (\mathbb{R}, +)$ then the Lebesgue measure is a Haar measure. When a Haar measure exists, it is more or less unique up to multiplicative scaling. All *locally compact topological* abelian groups have a Haar measure, a very large class of abelian groups, that include all finite ones, $(\mathbb{Z}, +)$, $(\mathbb{R}, +)$, and so on.

Once we have a Haar measure μ over Γ, and we have defined an integral for functions $f : \Gamma \to \mathbb{C}$, we say that a function is an element of $L^2(\Gamma)$ if

$$\int_{\Gamma} |f(x)|^2 d\mu(x) < \infty$$

For example, if Γ is finite, then all functions $f : \Gamma \to \mathbb{C}$ are in $L^2(\Gamma)$, and a function $f : \mathbb{Z} \to \mathbb{C}$ is in $L^2(\mathbb{Z})$ if the series $\sum_{n \in \mathbb{Z}} |f(n)|^2$ converges.

If $f, g \in L^2(\Gamma)$, we can define their inner product

$$\langle f, g \rangle := \int_\Gamma f(x) g^*(x) d\mu(x)$$

and use Cauchy-Schwarz to see that $|\langle f, g \rangle| < \infty$.

Now we can repeat the proof of Lemma 21 that $\langle \chi_1, \chi_2 \rangle = 0$ for two different characters, and the only step of the proof that we need to verify for infinite groups is an analog of Lemma 20, that is we need to prove that if χ is a character that is not always equal to 1, then

$$\int_\Gamma \chi(x) d\mu(x) = 0$$

and the same proof as in Lemma 20 works, with the key step being that, for every group element a,

$$\int_\Gamma \chi(x+a) d\mu(x) = \int_\Gamma \chi(x) d\mu(x)$$

because of the property of μ being a Haar measure.

We don't have an analogous result to Theorem 26 showing that Γ and $\hat{\Gamma}$ are isomorphic, however it is possible to show that $\hat{\Gamma}$ itself has a Haar measure $\hat{\mu}$, that the dual of $\hat{\Gamma}$ is isomorphic to Γ, and that if $f : \Gamma \to \mathbb{C}$ is continuous, then it can be written as the "linear combination"

$$f(x) = \int_{\hat{\Gamma}} \hat{f}(\chi) \chi(x) d\hat{\mu}(x)$$

where

$$\hat{f}(\chi) = \int_\Gamma f(x) \chi^*(x) d\mu(x)$$

In the finite case, the examples that we developed before correspond to setting $\mu(A) := |A|/|\Gamma|$ and $\hat{\mu}(A) = |A|$.

Example 29 (Fourier Series) *The set of characters of the group $[0, 1)$ with the operation of addition modulo 1 is isomorphic to $\mathbb{Z}$, because for every integer n we can define the function $\chi_n : [0, 1) \to \mathbb{C}$*

$$\chi_n(x) := e^{2\pi i x n}$$

and it can be shown that there are no other characters. We thus have the Fourier series for continuous functions $f : [0, 1) \to \mathbb{C}$,

$$f(x) = \sum_{n \in \mathbb{Z}} \hat{f}(n) e^{2\pi i x n}$$

where

$$\hat{f}(n) = \int_0^1 f(x) e^{-2\pi i x n} dx$$

5.3 Cayley Graphs and Their Spectrum

Let Γ be a finite group. We will use additive notation, although the following definition applies to non-commutative groups as well. A subset $S \subseteq \Gamma$ is *symmetric* if $a \in S \Leftrightarrow -a \in S$.

Definition 30 *For a group Γ and a symmetric subset $S \subseteq \Gamma$, the Cayley graph $Cay(\Gamma, S)$ is the graph whose vertex set is Γ, and such that (a, b) is an edge if and only if $b - a \in S$. Note that the graph is undirected and $|S|$-regular.*

We can also define Cayley *weighted* graphs: if $w : \Gamma \to \mathbb{R}$ is a function such that $w(a) = w(-a)$ for every $a \in \Gamma$, then we can define the weighted graph $Cay(G, w)$ in which the edge (a, b) has weight $w(b - a)$. We will usually work with unweighted graphs.

Example 31 (Cycle) *The n-vertex cycle can be constructed as the Cayley graph $Cay(\mathbb{Z}/n\mathbb{Z}, \{-1, 1\})$.*

Example 32 (Hypercube) *The d-dimensional hypercube can be constructed as the Cayley graph*

$$Cay((\mathbb{Z}/2\mathbb{Z})^d, \{(1, 0, \ldots, 0), (0, 1, \ldots, 0), \ldots, (0, 0, \ldots, 1)\})$$

where the group is the set $\{0, 1\}^d$ with the operation of bit-wise xor, and the set S is the set of bit-vectors with exactly one 1.

If we construct a Cayley graph from a finite abelian group, then the eigenvectors are the characters of the groups, and the eigenvalues have a very simple description.

Lemma 33 *Let Γ be a finite abelian group, $\chi : \Gamma \to \mathbb{C}$ be a character of Γ, $S \subseteq \Gamma$ be a symmetric set. Let M be the normalized adjacency matrix of the Cayley graph $G = Cay(\Gamma, S)$. Consider the vector $\mathbf{x} \in \mathbb{C}^\Gamma$ such that $x_a = \chi(a)$.*

Then $\mathbf{x}$ is an eigenvector of G, with eigenvalue

$$\frac{1}{|S|} \sum_{s \in S} \chi(s)$$

PROOF: Consider the a-th entry of $M\mathbf{x}$:

$$\begin{aligned}
(M\mathbf{x})_a &= \sum_b M_{a,b} x_b \\
&= \frac{1}{|S|} \sum_{b : b-a \in S} \chi(b) \\
&= \frac{1}{|S|} \sum_{s \in S} \chi(a+s) \\
&= x_a \cdot \frac{1}{|S|} \cdot \sum_{s \in S} \chi(s)
\end{aligned}$$

And so

$$M\mathbf{x} = \left(\frac{1}{|S|} \sum_{s \in S} \chi(s) \right) \cdot \mathbf{x}$$

□

The eigenvalues of the form $\frac{1}{S} \sum_{s \in S} \chi(s)$, where χ is a character, enumerate all the eigenvalues of the graph, as can be deduced from the following observations:

1. Every character is an eigenvector;
2. The characters are linearly independent (as functions $\chi : \Gamma \to \mathbb{C}$ and, equivalently, as vectors in $\mathbb{C}^\Gamma$);
3. There are as many characters as group elements, and so as many characters as nodes in the corresponding Cayley graphs.

It is remarkable that, for a Cayley graph, a system of eigenvectors can be determined based solely on the underlying group, independently of the set S.

5.4 The Cycle

The n-cycle is the Cayley graph $Cay(\mathbb{Z}/n\mathbb{Z}, \{-1, +1\})$. Recall that, for every $n \in \{0, \ldots, n-1\}$, the group $\mathbb{Z}/n\mathbb{Z}$ has a character $\chi_r(x) = e^{2\pi i r x/n}$.

This means that for every $r \in \{0, \ldots, n-1\}$ we have the eigenvalue

$$\lambda_r = \frac{1}{2}e^{2\pi i r/n} + \frac{1}{2}e^{-2\pi i r/n} = \cos(2\pi r/n)$$

where we used the facts that $e^{ix} = \cos(x) + i\sin(x)$, that $\cos(x) = \cos(-x)$, and $\sin(x) = -\sin(-x)$.

For $r = 0$ we have the eigenvalue 1. For $r = 1$ we have the second largest eigenvalue $\cos(2\pi/n) = 1 - \Theta(1/n^2)$.

The expansion of the cycle is $h(C_n) \geq 2/n$, and so the cycle is an example in which the second Cheeger inequality is tight.

5.5 The Hypercube

The group $\{0,1\}^d$ with bitwise xor has 2^d characters; for every $r \in \{0,1\}^d$ there is a character $\chi_r : \{0,1\}^d \to \{-1, 1\}$ defined as

$$\chi_r(x) = (-1)^{\sum_i r_i x_i}$$

Let us denote the set S by $\{e^1, \ldots, e^d\}$, where we let $e^j \in \{0,1\}^d$ denote the bit-vector that has a 1 in the j-th position, and zeroes everywhere else. This means that, for every bit-vector $r \in \{0,1\}^d$, the hypercube has the eigenvalue

$$\frac{1}{d}\sum_j \chi_r(e^j) = \frac{1}{d}\sum_j (-1)^{r_j} = \frac{1}{d}(-|r| + d - |r|) = 1 - 2\frac{|r|}{d}$$

where we denote by $|r|$ the *weight* of r, that is, the number of ones in r.

Corresponding to $r = (0, \ldots, 0)$, we have the eigenvalue 1.

For each of the d vectors r with exactly one 1, we have the second largest eigenvalue$1 - 2/d$.

Let us compute the expansion of the hypercube. Consider "dimension cuts" of the form $S_i := \{x \in \{0,1\}^n : x_i = 0\}$. The set S_i contains half of the vertices, and the number of edges that cross the cut $(S_i, V - S_i)$ is also equal to half the number of vertices (because the edges form a perfect matching), so we have $h(S_i) = \frac{1}{d}$.

These calculations show that the first Cheeger inequality $(1 - \lambda_2)/2 \leq h(G)$ is tight for the hypercube.

Finally, we consider the tightness of the approximation analysis of the spectral partitioning algorithm.

We have seen that, in the d-dimensional hypercube, the second eigenvalue has multiplicity d, and that its eigenvectors are vectors $\mathbf{x}^j \in \mathbb{R}^{2^d}$ such that $x_a^j = (-1)^{a_j}$. Consider now the vector $\mathbf{x} := \sum_j \mathbf{x}^j$; this is still clearly an eigenvector of the second eigenvalue. The entries of the vector $\mathbf{x}$ are

$$x_a = \sum_j (-1)^{a_j} = d - 2|a|$$

Suppose now that we apply the spectral partitioning algorithm using $\mathbf{x}$ as our vector. This is equivalent to considering all the cuts $(S_t, V - S_t)$ in the hypercube in which we pick a threshold t and define $S_t := \{a \in \{0,1\}^n : |a| \geq t\}$.

Some calculations with binomial coefficients show that the best such "threshold cut" is the "majority cut" in which we pick $t = n/2$, and that the expansion of $S_{n/2}$ is

$$h(S_{n/2}) = \Omega\left(\frac{1}{\sqrt{d}}\right)$$

This gives an example of a graph and of a choice of eigenvector for the second eigenvalue that, given as input to the spectral partitioning algorithm, result in the output of a cut $(S, V-S)$ such that $h(S) \geq \Omega(\sqrt{h(G)})$. Recall that we proved $h(S) \leq 2\sqrt{h(G)}$, which is thus tight.

6 Computing Eigenvalues and Eigenvectors *

In this section, which is not required material, we describe a nearly-linear time algorithm that finds approximate eigenvalues and eigenvectors.

In past lectures, we showed that, if $G = (V, E)$ is a d-regular graph, and M is its normalized adjacency matrix with eigenvalues $1 = \lambda_1 \geq \lambda_2 \ldots \geq \lambda_n$, given an eigenvector of λ_2, the algorithm SpectralPartition finds, in nearly-linear time $O(|E| + |V| \log |V|)$, a cut $(S, V - S)$ such that $h(S) \leq 2\sqrt{h(G)}$.

More generally, if, instead of being given an eigenvector $\mathbf{x}$ such that $M\mathbf{x} = \lambda_2 \mathbf{x}$, we are given a vector $\mathbf{x} \perp \mathbf{1}$ such that $\mathbf{x}^T M \mathbf{x} \geq (\lambda_2 - \epsilon)\mathbf{x}^T\mathbf{x}$, then the algorithm finds a cut such that $h(S) \leq \sqrt{4h(G) + 2\epsilon}$. In this lecture we describe and analyze an algorithm that computes such a vector using $O((|V| + |E|) \cdot \frac{1}{\epsilon} \cdot \log \frac{|V|}{\epsilon})$ arithmetic operations.

A symmetric matrix is *positive semi-definite* (abbreviated PSD) if all its eigenvalues are nonnegative. We begin by describing an algorithm that approximates the largest

eigenvalue of a given symmetric PSD matrix. This might not seem to help very much because the adjacency matrix of a graph is not PSD, and because we want to compute the second largest, not the largest, eigenvalue. We will see, however, that the algorithm is easily modified to approximate the second eigenvalue of a PSD matrix (if an eigenvector of the first eigenvalue is known), and that the adjacency matrix of a graph can easily be modified to be PSD.

6.1 The Power Method to Approximate the Largest Eigenvalue

The algorithm works as follows

Algorithm *Power*

- Input: PSD symmetric matrix $M \in \mathbb{R}^{n \times n}$, positive integer t
- Pick uniformly at random $\mathbf{x}_0 \sim \{-1, 1\}^n$
- for $i := 1$ to t
  - $\mathbf{x}_i := M\mathbf{x}_{i-1}$
- return $\mathbf{x}_t$

That is, the algorithm simply picks uniformly at random a vector $\mathbf{x}$ with ± 1 coordinates, and outputs $M^t\mathbf{x}$.

Note that the algorithm performs $O(t \cdot (n + m))$ arithmetic operations, where m is the number of non-zero entries of the matrix M.

Theorem 34 *For every PSD matrix M, positive integer t and parameter $\epsilon > 0$, with probability $\geq 3/16$ over the choice of $\mathbf{x}_0$, the algorithm Power outputs a vector $\mathbf{x}_t$ such that*

$$\frac{\mathbf{x}_t^T M \mathbf{x}_t}{\mathbf{x}_t^T \mathbf{x}_t} \geq \lambda_1 \cdot (1 - \epsilon) \cdot \frac{1}{1 + 4n(1 - \epsilon)^{2t}}$$

where λ_1 is the largest eigenvalue of M.

Note that, in particular, we can have $t = O(\log n/\epsilon)$ and $\frac{\mathbf{x}_t^T M \mathbf{x}_t}{\mathbf{x}_t^T \mathbf{x}_t} \geq (1 - O(\epsilon)) \cdot \lambda_1$.

Let $\lambda_1 \geq \cdots \lambda_n$ be the eigenvalues of M, with multiplicities, and $\mathbf{v}_1, \ldots, \mathbf{v}_n$ be a system of orthonormal eigenvectors such that $M\mathbf{v}_i = \lambda_i \mathbf{v}_i$. Theorem 34 is implied by the following two lemmas

Lemma 35 *Let* $\mathbf{v} \in \mathbb{R}^n$ *be a vector such that* $||\mathbf{v}|| = 1$. *Sample uniformly* $\mathbf{x} \sim \{-1, 1\}^n$. *Then*

$$\mathbb{P}\left[|\langle \mathbf{x}, \mathbf{v}\rangle| \geq \frac{1}{2}\right] \geq \frac{3}{16}$$

Lemma 36 *Let* $\mathbf{x} \in \mathbb{R}^n$ *be a vector such that* $|\langle \mathbf{x}, \mathbf{v}_1\rangle| \geq \frac{1}{2}$. *Then, for every positive integer* t *and positive* $\epsilon > 0$, *if we define* $\mathbf{y} := M^t\mathbf{x}$, *we have*

$$\frac{\mathbf{y}^T M \mathbf{y}}{\mathbf{y}^T \mathbf{y}} \geq \lambda_1 \cdot (1-\epsilon) \cdot \frac{1}{1 + 4||\mathbf{x}||^2(1-\epsilon)^{2t}}$$

It remains to prove the two lemmas.

PROOF: (Of Lemma 35) Let $\mathbf{v} = (v_1, \ldots, v_n)$. The inner product $\langle \mathbf{x}, \mathbf{v}\rangle$ is the random variable

$$S := \sum_i x_i v_i$$

Let us compute the first, second, and fourth moment of S.

$$\mathbb{E}\, S = 0$$

$$\mathbb{E}\, S^2 = \sum_i v_i^2 = 1$$

$$\mathbb{E}\, S^4 = 3\left(\sum_i v_i^2\right) - 2\sum_i v_i^4 \leq 3$$

Recall that the Paley-Zygmund inequality states that if Z is a non-negative random variable with finite variance, then, for every $0 \leq \delta \leq 1$, we have

$$\mathbb{P}[Z \geq \delta\, \mathbb{E}\, Z] \geq (1-\delta)^2 \cdot \frac{(\mathbb{E}\, Z)^2}{\mathbb{E}\, Z^2} \tag{11}$$

which follows by noting that

$$\mathbb{E}\, Z = \mathbb{E}[Z \cdot 1_{Z < \delta\, \mathbb{E}\, Z}] + \mathbb{E}[Z \cdot 1_{Z \geq \delta\, \mathbb{E}\, Z}]\ ,$$

that

$$\mathbb{E}[Z \cdot 1_{Z<\delta\,\mathbb{E}\,Z}] \leq \delta\,\mathbb{E}\,Z\ ,$$

and that

$$\begin{aligned}\mathbb{E}[Z \cdot 1_{Z\geq\delta\,\mathbb{E}\,Z}] &\leq \sqrt{\mathbb{E}\,Z^2} \cdot \sqrt{\mathbb{E}\,1_{Z\geq\delta\,\mathbb{E}\,Z}} \\ &= \sqrt{\mathbb{E}\,Z^2}\sqrt{\mathbb{P}[Z \geq \delta\,\mathbb{E}\,Z]}\end{aligned}$$

We apply the Paley-Zygmund inequality to the case $Z = S^2$ and $\delta = 1/4$, and we derive

$$\mathbb{P}\left[S^2 \geq \frac{1}{4}\right] \geq \left(\frac{3}{4}\right)^2 \cdot \frac{1}{3} = \frac{3}{16}$$

□

Remark 37 *The proof of Lemma 35 works even if* $\mathbf{x} \sim \{-1,1\}^n$ *is selected according to a 4-wise independent distribution. This means that the algorithm can be derandomized in polynomial time.*

Proof: (Of Lemma 36) Let us write $\mathbf{x}$ as a linear combination of the eigenvectors

$$\mathbf{x} = a_1\mathbf{v}_1 + \cdots + a_n\mathbf{v}_n$$

where the coefficients can be computed as $a_i = \langle \mathbf{x}, \mathbf{v}_i \rangle$. Note that, by assumption, $|a_1| \geq .5$, and that, by orthonormality of the eigenvectors, $||\mathbf{x}||^2 = \sum_i a_i^2$.

We have

$$\mathbf{y} = a_1\lambda_1^t\mathbf{v}_1 + \cdots + a_n\lambda_n^t\mathbf{v}_n$$

and so

$$\mathbf{y}^T M\mathbf{y} = \sum_i a_i^2\lambda_i^{2t+1}$$

and

$$\mathbf{y}^T\mathbf{y} = \sum_i a_i^2\lambda_i^{2t}$$

We need to prove a lower bound to the ratio of the above two quantities. We will compute a lower bound to the numerator and an upper bound to the denominator in terms of the same parameter.

Let k be the number of eigenvalues larger than $\lambda_1 \cdot (1-\epsilon)$. Then, recalling that the eigenvalues are sorted in non-increasing order, we have

$$\mathbf{y}^T M \mathbf{y} \geq \sum_{i=1}^{k} a_i^2 \lambda_i^{2t+1} \geq \lambda_1 (1-\epsilon) \sum_{i=1}^{k} a_i^2 \lambda_i^{2t}$$

We also see that

$$\begin{aligned}
&\sum_{i=k+1}^{n} a_i^2 \lambda_i^{2t} \\
&\leq \lambda_1^{2t} \cdot (1-\epsilon)^{2t} \sum_{i=k+1}^{n} a_i^2 \\
&\leq \lambda_1^{2t} \cdot (1-\epsilon)^{2t} \cdot ||\mathbf{x}||^2 \\
&\leq 4 a_1^2 \lambda_1^{2t} (1-\epsilon)^{2t} ||\mathbf{x}||^2 \\
&\leq 4 ||\mathbf{x}||^2 (1-\epsilon)^{2t} \sum_{i=1}^{k} a_i^2 \lambda_i^{2t}
\end{aligned}$$

So we have

$$\mathbf{y}^T \mathbf{y} \leq (1 + 4||\mathbf{x}||^2 (1-\epsilon)^{2t}) \cdot \sum_{i=1}^{k} a_i^2$$

giving

$$\frac{\mathbf{y}^T M \mathbf{y}}{\mathbf{y}^T \mathbf{y}} \geq \lambda_1 \cdot (1-\epsilon) \cdot \frac{1}{1 + 4||\mathbf{x}||^2 (1-\epsilon)^{2t}}$$

□

Remark 38 *Where did we use the assumption that M is positive semidefinite? What happens if we apply this algorithm to the adjacency matrix of a bipartite graph?*

6.2 Approximating the Second Eigenvalue

If M is a PSD matrix, and if we know a unit-length eigenvector $\mathbf{v}_1$ of the largest eigenvalue of M, we can approximately find the second eigenvalue with the following adaption of the algorithm from the previous section.

Algorithm *Power2*

- Input: PSD symmetric matrix $M \in \mathbb{R}^{n \times n}$, positive integer t, vector $\mathbf{v}_1$
- Pick uniformly at random $\mathbf{x} \sim \{-1, 1\}^n$
- $\mathbf{x}_0 := \mathbf{x} - \langle \mathbf{v}_1, \mathbf{x} \rangle \cdot \mathbf{v}_1$
- for $i := 1$ to t
  - $\mathbf{x}_i := M\mathbf{x}_{i-1}$
- return $\mathbf{x}_t$

If $\mathbf{v}_1, \ldots, \mathbf{v}_n$ is an orthonormal basis of eigenvectors for the eigenvalues $\lambda_1 \geq \cdots \geq \lambda_n$ of M, then, at the beginning, we pick a random vector

$$\mathbf{x} = a_1\mathbf{v}_1 + a_2\mathbf{v}_2 + \cdots a_n\mathbf{v}_n$$

that, with probability at least 3/16, satisfies $|a_2| \geq 1/2$. (Cf. Lemma 35.) Then we compute $\mathbf{x}_0$, which is the projection of $\mathbf{x}$ on the subspace orthogonal to $\mathbf{v}_1$, that is

$$\mathbf{x}_0 = a_2\mathbf{v}_2 + \cdots a_n\mathbf{v}_n$$

Note that $||\mathbf{x}||^2 = n$ and that $||\mathbf{x}_0||^2 \leq n$.

The output is the vector $\mathbf{x}_t$

$$\mathbf{x}_t = a_2\lambda_2^t\mathbf{v}^2 + \cdots a_n\lambda_n^t\mathbf{v}^n$$

If we apply Lemma 36 to subspace orthogonal to $\mathbf{v}_1$, we see that when $|a_2| \geq 1/2$ we have that, for every $0 < \epsilon < 1$,

$$\frac{\mathbf{x}_t^T M \mathbf{x}_t}{\mathbf{x}_t^T \mathbf{x}_t} \geq \lambda_2 \cdot (1-\epsilon) \cdot \frac{1}{4n(1-\epsilon)^{2t}}$$

We have thus established the following analysis.

Theorem 39 *For every PSD matrix M, positive integer t and parameter $\epsilon > 0$, if $\mathbf{v}_1$ is a length-1 eigenvector of the largest eigenvalue of M, then with probability $\geq 3/16$ over the choice of $\mathbf{x}_0$, the algorithm Power2 outputs a vector $\mathbf{x}_t \perp \mathbf{v}_1$ such that*

$$\frac{\mathbf{x}_t^T M \mathbf{x}_t}{\mathbf{x}_t^T \mathbf{x}_t} \geq \lambda_2 \cdot (1-\epsilon) \cdot \frac{1}{1 + 4n(1-\epsilon)^{2t}}$$

where λ_2 is the second largest eigenvalue of M, counting multiplicities.

Finally, we come to the case in which M is the normalized adjacency matrix of a regular graph.

We know that M has eigenvalues $1 = \lambda_1 \geq \cdots \lambda_n \geq -1$ and that $\frac{1}{\sqrt{n}} \cdot \mathbf{1}$ is an eigenvector of λ_1.

Consider now the matrix $M + I$. Every eigenvector of M with eigenvalue λ is clearly also an eigenvector of $M + I$ with eigenvalue $1 + \lambda$, and vice versa, thus $M + I$ has eigenvalues $2 = 1 + \lambda_1 \geq 1 + \lambda_2 \geq \cdots \geq 1 + \lambda_n \geq 0$ and it is PSD.

This means that we can run algorithm Power2 on the matrix $I + M$ using the vector $\mathbf{v}_1 = \frac{1}{\sqrt{n}}\mathbf{1}$ and a parameter $t = O(\epsilon^{-1} \log n/\epsilon))$. The algorithm finds, with probability $\geq 3/16$, a vector $\mathbf{x}_t \perp \mathbf{1}$ such that

$$\frac{\mathbf{x}_t^T (M + I)\mathbf{x}_t}{\mathbf{x}_t^T \mathbf{x}_t} \geq (1 + \lambda_2) \cdot (1 - 2\epsilon)$$

which is equivalent to

$$\frac{\mathbf{x}_t^T M \mathbf{x}_t}{\mathbf{x}_t^T \mathbf{x}_t} \geq \lambda_2 - 2\epsilon - 2\epsilon\lambda_2 \geq \lambda_2 - 4\epsilon$$

7 The Leighton-Rao Approximation of Sparsest Cut

Let $G = (V, E)$ be an undirected graph. Unlike past lectures, we will not need to assume that G is regular. We are interested in finding a *sparsest cut* in G, where the sparsity of a non-trivial bipartition $(S, V - S)$ of the vertices is

$$\phi_G(S) := \frac{\frac{1}{|E|} \cdot Edges(S, V - S)}{\frac{2}{V^2} \cdot |S| \cdot |V - S|}$$

which is the ratio between the fraction of edges that are cut by $(S, V - S)$ and the fraction of pairs of vertices that are disconnected by the removal of those edges.

Another way to write the sparsity of a cut is as

$$\phi_G(S) := \frac{|V|^2}{2|E|} \cdot \frac{\sum_{i,j} A_{i,j} |1_S(i) - 1_S(j)|}{\sum_{i,j} |1_S(i) - 1_S(j)|}$$

where A is the adjacency matrix of G and $1_S(\cdot)$ is the indicator function of the set S.

The observation that led us to see $1 - \lambda_2$ as the optimum of a continuous relaxation of ϕ was to observe that $|1_S(i) - 1_S(j)| = |1_S(i) - 1_S(j)|^2$, and then relax the problem by allowing arbitrary functions $x : V \to \mathbb{R}$ instead of indicator functions $1_S : V \to \{0, 1\}$.

The Leighton-Rao relaxation of sparsest cut is obtained using, instead, the following observation: if, for a set S, we define $d_S(i,j) := |1_S(i) - 1_S(j)|$, then $d_S(\cdot,\cdot)$ defines a semi-metric over the set V, because d_S is symmetric, $d_S(i,i) = 0$, and the triangle inequality holds. So we could think about allowing *arbitrary semi-metrics* in the expression for ϕ, and define

$$LR(G) := \min_{\substack{d : V \times V \to \mathbb{R} \\ d \text{ semi-metric}}} \frac{|V|^2}{2|E|} \cdot \frac{\sum_{i,j} A_{i,j} d(i,j)}{\sum_{i,j} d(i,j)} \tag{12}$$

This might seem like such a broad relaxation that there could be graphs on which $LR(G)$ bears no connection to $\phi(G)$. Instead, we will prove the fairly good estimate

$$LR(G) \leq \phi(G) \leq O(\log |V|) \cdot LR(G) \tag{13}$$

Furthermore, we will show that $LR(G)$, and an optimal solution $d(\cdot,\cdot)$ can be computed in polynomial time, and the second inequality above has a constructive proof, from which we derive a polynomial time $O(\log |V|)$-approximate algorithm for sparsest cut.

7.1 Formulating the Leighton-Rao Relaxation as a Linear Program

The value $LR(G)$ and an optimal $d(\cdot,\cdot)$ can be computed in polynomial time by solving the following linear program

$$\begin{array}{ll} \text{minimize} & \sum_{i,j} A_{i,j} d_{i,j} \\ \text{subject to} & \\ & \sum_{i,j} d_{i,j} = \frac{|V|^2}{2|E|} \\ & d_{i,k} \leq d_{i,j} + d_{j,k} \quad \forall i,j,k \in V \\ & d_{i,j} \geq 0 \qquad\qquad \forall i \in V \end{array} \tag{14}$$

that has a variable $d_{i,j}$ for every unordered pair of distinct vertices i, j. Clearly, every solution to the linear program (14) is also a solution to the right-hand side of the definition (12) of the Leighton-Rao parameter, with the same cost. Also every semi-metric can be normalized so that $\sum_{i,j} d(i,j) = |V|^2/2|E|$ by multiplying every distance by a fixed constant, and the normalization does not change the value of the right-hand side of (12); after the normalization, the semimetric is a feasible solution to the linear program (14), with the same cost.

In the rest of this lecture and the next, we will show how to round a solution to (14) into a cut, achieving the logarithmic approximation promised in (13).

7.2 An L1 Relaxation of Sparsest Cut

In the Leighton-Rao relaxation, we relax distance functions of the form $d(i,j) = |1_S(i) - 1_S(j)|$ to completely arbitrary distance functions. Let us consider an intermediate relaxation, in which we allow distance functions that can be realized by *an embedding of the vertices in an ℓ_1 space.*

Recall that, for a vector $\mathbf{x} \in \mathbb{R}^n$, its ℓ_1 norm is defined as $||\mathbf{x}||_1 := \sum_i |x_i|$, and that this norm makes $\mathbb{R}^n$ into a metric space with the ℓ_1 distance function

$$||\mathbf{x} - \mathbf{y}||_1 = \sum_i |x_i - y_i|$$

The distance function $d(i,j) = |1_S(i) - 1_S(j)|$ is an example of a distance function that can be realized by mapping each vertex to a real vector, and then defining the distance between two vertices as the ℓ_1 norm of the respective vectors. Of course it is an extremely restrictive special case, in which the dimension of the vectors is one, and in which every vertex is actually mapping to either zero or one. Let us consider the relaxation of sparsest cut to arbitrary ℓ_1 mappings, and define

$$\phi'(G) := \inf_{m,f:V\to\mathbb{R}^m} \frac{|V|^2}{2|E|} \cdot \frac{\sum_{i,j} A_{i,j}||f(i) - f(j)||_1}{\sum_{i,j} ||f(i) - f(j)||_1}$$

This may seem like another very broad relaxation of sparsest cut, whose optimum might bear no correlation with the sparsest cut optimum. The following theorem shows that this is not the case.

Theorem 40 *For every graph G, $\phi(G) = \phi'(G)$.*

Furthermore, there is a polynomial time algorithm that, given a mapping $f : V \to \mathbb{R}^m$, finds a cut S such that

$$\frac{\sum_{u,v} A_{u,v}|1_S(u) - 1_S(v)|}{\sum_{u,v} |1_S(u) - 1_S(v)|} \leq \frac{\sum_{u,v} A_{u,v}||f(u) - f(v)||_1}{\sum_{u,v} ||f(u) - f(v)||_1} \tag{15}$$

Proof: We use ideas that have already come up in the proof the difficult direction of Cheeger's inequality. First, we note that for every nonnegative reals $a_1, \ldots, a_m$ and positive reals $b_1, \ldots, b_m$ we have

$$\frac{a_1 + \cdots a_m}{b_1 + \cdots b_m} \geq \min_i \frac{a_i}{b_i} \tag{16}$$

as can be seen by noting that

$$\sum_j a_j = \sum_j b_j \cdot \frac{a_j}{b_j} \geq \left(\min_i \frac{a_i}{b_i}\right) \cdot \sum_j b_j$$

Let $f_i(v)$ be the i-th coordinate of the vector $f(v)$, thus $f(v) = (f_1(v), \ldots, f_m(v))$. Then we can decompose the right-hand side of (15) by coordinates, and write

$$\begin{aligned}
&\frac{\sum_{u,v} A_{u,v}||f(u) - f(v)||_1}{\sum_{u,v} ||f(u) - f(v)||_1} \\
&= \frac{\sum_i \sum_{u,v} A_{u,v}|f_i(u) - f_i(v)|}{\sum_i \sum_{u,v} |f_i(u) - f_i(v)|} \\
&\geq \min_i \frac{\sum_{u,v} A_{u,v}|f_i(u) - f_i(v)|}{\sum_{u,v} |f_i(u) - f_i(v)|}
\end{aligned}$$

This already shows that, in the definition of ϕ', we can map, with no loss of generality, to 1-dimensional ℓ_1 spaces.

Let i^* be the coordinate that achieves the minimum above. Because the cost function is invariant under the shifts and scalings (that is, the cost of a function $x \to f(x)$ is the same as the cost of $x \to af(x) + b$ for every two constants $a \neq 0$ and b) there is a function $g : V \to \mathbb{R}$ such that g has the same cost function as f_{i*} and it has a unit-length range $\max_v g(v) - \min_v g(v) = 1$.

Let us now pick a threshold t uniformly at random from the interval $[\min_v g(v), \max_v g(v)]$, and define the random variables

$$S_t := \{v : g(v) \leq t$$

We observe that for every pairs of vertices u, v we have

$$\mathbb{E}\,|1_{S_t(}u) - 1_{S_t}(v)| = |g(u) - g(v)|$$

and so we get

$$\begin{aligned}
&\frac{\sum_{u,v} A_{u,v}||f(u) - f(v)||_1}{\sum_{u,v} ||f(u) - f(v)||_1} \\
&\geq \frac{\sum_{u,v} A_{u,v}|g(u) - g(v)|}{\sum_{u,v} |g(u) - g(v)|} \\
&= \frac{\mathbb{E} \sum_{u,v} A_{u,v}|1_{S_t}(u) - 1_{S_t}(v)|}{\mathbb{E} \sum_{u,v} |1_{S_t}(u) - 1_{S_t}(v)|}
\end{aligned}$$

Finally, by an application of (16), we see that there must be a set S among the possible values of S_t such that (15) holds.

Notice that the proof was completely constructive: we simply took the coordinate f_{i^*} of f with the lowest cost function, and then the "threshold cut" given by f_{i^*} with the smallest sparsity. $\square$

7.3 A Theorem of Bourgain

We will derive our main result (13) from the L1 "rounding" process of the previous section, and from the following theorem of Bourgain (the efficiency considerations are due to Linial, London and Rabinovich).

Theorem 41 (Bourgain) *Let $d : V \times V \to \mathbb{R}$ be a semimetric defined over a finite set V. Then there exists a mapping $f : V \to \mathbb{R}^m$ such that, for every two elements $u, v \in R$,*

$$||f(u) - f(v)||_1 \leq d(u, v) \leq ||f(u) - f(v)||_1 \cdot c \cdot \log |V|$$

where c is an absolute constant. Given d, the mapping f can be found with high probability in randomized polynomial time in $|V|$.

To see that the above theorem of Bourgain implies (13), consider a graph G, and let d be the optimal solution of the Leighton-Rao relaxation of the sparsest cut problem on G, and let $f : V \to \mathbb{R}$ be a mapping as in Bourgain's theorem applied to d. Then

$$\begin{aligned} LR(G) &= \frac{|V|^2}{|E|} \cdot \frac{\sum_{u,v} A_{u,v} d(u, v)}{\sum_{u,v} d(u, v)} \\ &\geq \frac{|V|^2}{|E|} \cdot \frac{\sum_{u,v} A_{u,v} ||f(u) - f(v)||_1}{c \cdot \log |V| \cdot \sum_{u,v} ||f(u) - f(v)||_1} \\ &\geq \frac{1}{c \cdot \log |U|} \cdot \phi(G) \end{aligned}$$

7.4 Proof of Bourgain's Theorem

The theorem has a rather short proof, but there is an element of "magic" to it. We will discuss several examples and we will see what approaches are suggested by the examples. At the end of the discussion, we will see the final proof as, hopefully, the "natural" outcome of the study of such examples and failed attempts.

7.4.1 Preliminary and Motivating Examples

A first observation is that embeddings of finite sets of points into L1 can be equivalently characterized as probabilistic embeddings into the real line.

Fact 42 *For every finite set V, dimension m, and mapping $F : V \to \mathbb{R}^m$, there is a finitely-supported probability distribution D over functions $f : V \to \mathbb{R}$ such that for every two points $u, v \in V$:*

$$\mathop{\mathbb{E}}_{f \sim D} |f(u) - f(v)| = ||F(u) - F(v)||_1$$

Conversely, for every finite set V and finitely supported distribution D over functions $f : V \to \mathbb{R}$, there is a dimension m and a mapping $F : V \to \mathbb{R}^m$ such that

$$\mathop{\mathbb{E}}_{f \sim D} |f(u) - f(v)| = ||F(u) - F(v)||_1$$

Proof: For the first claim, we write $F_i(v)$ for the i-th coordinate of $F(v)$, that is $F(v) = (F_1(v), \ldots, F_m(v))$, and we define D to be the uniform distribution over the m functions of the form $x \to m \cdot F_i(x)$.

For the second claim, if the support of D is the set of functions $\{f_1, \ldots, f_m\}$, where function f_i has probability p_i, then we define $F(v) := (p_1 f_1(v), \ldots, p_m f_m(v))$. $\square$

It will be easier to reason about probabilistic mappings into the line, so we will switch to the latter setting from now on.

Our task is to associate a number to every point v, and the information that we have about v is the list of distances $\{d(u, v)\}$. Probably the first idea that comes to mind is to pick a random reference vertex $r \in V$, and work with the mapping $v \to d(r, v)$, possibly scaled by a multiplicative constant. (Equivalently, we can think about the deterministic mapping $V \to \mathbb{R}^{|V|}$, in which the vertex v is mapped to the sequence $(d(u_1, v), \ldots, d(u_n, v)$, for some enumeration $u_1, \ldots, u_n$ of the elements of V.)

This works in certain simple cases.

Example 43 (Cycle) *Suppose that $d(\cdot, \cdot)$ is the shortest-path metric on a cycle, we can see that, for every two points on the cycle, $\mathbb{E}_{r \in V} |d(r, u) - d(r, v)|$ is within a constant factor of their distance $d(u, v)$. (Try proving it rigorously!)*

Example 44 (Simplex) *Suppose that $d(u, v) = 1$ for every $u \neq v$, and $d(u, u) = 0$. Then, for every $u \neq v$, we have $\mathbb{E}_{r \in V} |d(r, u) - d(r, v)| = \mathbb{P}[r = u \vee r = v] = 2/n$, so, up to scaling, the mapping incurs no error at all.*

But there are also simple examples in which this works very badly.

Example 45 (1-2 Metric) *Suppose that for every $u \neq v$ we have $d(u,v) \in \{1,2\}$ (any distance function satisfying this property is always a metric) and that, in particular, there is a special vertex z at distance 2 from all other vertices, while all other vertices are at distance 1 from each other. Then, for vertices u, v both different from z we have, as before*

$$\mathbb{E}[|d(r,u) - d(r,v)|] = \frac{2}{n}$$

but for every v different from z we have

$$\mathbb{E}[|d(r,z) - d(r,v)|] = \frac{n-2}{n} \cdot |2-1| + \frac{1}{n} \cdot |2-0| + \frac{1}{n} \cdot |0-2| = 1 + \frac{2}{n}$$

and so our error is going to be $\Omega(n)$ instead of the $O(\log n)$ that we are trying to establish.

Maybe the next simplest idea is that we should pick at random several reference points $r_1, \ldots, r_k$. But how do we combine the information $d(r_1,u), \ldots, d(r_k,u)$ into a single number to associate to u? If we just take the sum of the distances, we are back to the case of sampling a single reference point. (We are just scaling up the expectation by a factor of k.)

The next simplest way to combine the information is to take either the maximum or the minimum. If we take the minimum, we see that we have the very nice property that we immediately guarantee that our distances in the L1 embedding are no bigger than the original distances, so that it "only" remains to prove that the distances don't get compressed too much.

Fact 46 *Let $d : V \times V \to \mathbb{R}$ be a semimetric and $A \subseteq V$ be a non-empty subset of points. Define $f_A : V \to \mathbb{R}$ as*

$$f_A(v) := \min_{r \in A} d(r,v)$$

Then, for every two points u, v we have

$$|f_A(u) - f_A(v)| \leq d(u,v)$$

Proof: Let a be the point such that $d(a,u) = f_A(u)$ and b be the point such that $d(b,v) = f_A(v)$. (It's possible that $a = b$.) Then

$$f_A(u) = d(a,u) \geq d(v,a) - d(u,v) \geq d(v,b) - d(u,v) = f_A(v) - d(u,v)$$

and, similarly,

$$f_A(v) = d(b,v) \geq d(u,b) - d(u,v) \geq d(u,a) - d(u,v) = f_A(u) - d(u,v)$$

□

Is there a way to sample a set $A = \{r_1, \ldots, r_k\}$ such that, for every two points u, v, the expectation $\mathbb{E}\,|f_A(u) - f_A(v)|$ is not too much smaller than $d(u,v)$? How large should the set A be?

Example 47 (1-2 Metric Again) *Suppose that for every $u \neq v$ we have $d(u,v) \in \{1,2\}$, and that we pick a subset $A \subseteq V$ uniformly at random, that is, each event $r \in A$ has probability $1/2$ and the events are mutually independent.*

Then for every $u \neq v$:

$$\frac{1}{4} \cdot d(u,v) \leq |\,\mathbb{E}\,|f_A(u) - f_A(v)| \leq d(u,v)$$

because with probability $1/2$ the set A contains exactly one of the elements u, v, and conditioned on that event we have $|f_A(u) - f_A(v)| \geq 1$ (because one of $f_A(u), f_A(v)$ is zero and the other is at least one), which is at least $d(u,v)/2$.

If we pick A uniformly at random, however, we incur an $\Omega(n)$ distortion in the case of the shortest path metric on the cycle. In all the examples seen so far, we can achieve constant distortion if we "mix" the distribution in which A is a random set of size 1 and the one in which A is a chosen uniformly at random among all sets, say by sampling from the former probability with probability $1/2$ and from the latter with probability $1/2$.

Example 48 (Far-Away Clusters) *Suppose now that $d(\cdot,\cdot)$ has the following structure: V is partitioned into clusters $B_1, \ldots, B_k$, where $|B_i| = i$ (so $k \approx \sqrt{2n}$), and we have $d(u,v) = 1$ for vertices in the same cluster, and $d(u,v) = n$ for vertices in different clusters.*

If u, v are in the same cluster, then $d(u,v) = 1$ and

$$\mathbb{E}\,|f_A(u) - f_A(v)| = \mathbb{P}[A \textit{ contains exactly one of } u, v]$$

If u, v are in different clusters B_i, B_j, then $d(u,v) = n$ and

$$\mathbb{E}\,|f_A(u) - f_A(v)| \approx n\,\mathbb{P}[A \textit{ intersects exactly one of } B_i, B_j]$$

If we want to stick to this approach of picking a set A of reference elements according to a certain distribution, and then defining the map $f_A(v) := \min_{r \in A} d(r, v)$, then the set A must have the property that for every two sets S, T, there is at least a probability p that A intersects exactly one of S, T, and we would like p to be as large as possible, because the distortion caused by the mapping will be at least $1/p$.

This suggest the following distribution D:

1. Sample t uniformly at random in $\{0, \ldots, \log_2 n\}$
2. Sample $A \subseteq V$ by selecting each $v \in V$, independently, to be in A with probability 2^{-t} and to be in $V - A$ with probability $1 - 2^{-t}$.

This distribution guarantees the above property with $p = 1/O(\log n)$.

Indeed, the above distribution guarantees a distortion at most $O(\log n)$ in all the examples encountered so far, including the tricky example of the clusters of different size. In each example, in fact, we can prove the following claim: for every two vertices u, v, there is a scale t, such that conditioned on that scale being chosen, the expectation of $|f_A(u), f_A(v)|$ is at least a constant times $d(u, v)$. We could try to prove Bourgain's theorem by showing that this is true in every semimetric.

Let us call D_t the conditional distribution of D conditioned on the choice of a scale t. We would like to prove that for every semimetric $d(\cdot, \cdot)$ and every two points u, v there is a scale t such that

$$\mathop{\mathbb{E}}_{A \sim D_t} |f_A(u) - f_A(v)| \geq \Omega(d(u, v))$$

which, recalling that $|f_A(u) - f_A(v)| \leq d(u, v)$ for every set A, is equivalent to arguing that

$$\mathop{\mathbb{P}}_{A \sim D_t} [|f_A(u) - f_A(v)| \geq \Omega(d(u, v))] \geq \Omega(1)$$

For this to be true, there must be distances d_1, d_2 such that $d_1 - d_2 \geq \Omega(d(u, v))$ and such that, with constant probability according to D_t, we have $f_A(u) \geq d_1$ and $f_A(v) \leq d_2$ (or vice-versa). For this to happen, there must be a constant probability that A avoids the set $\{r : d(u, r) < d_1\}$ and intersects the set $\{r : d(v, r) \leq d_2\}$. For this to happen, both sets must have size $\approx 2^t$.

This means that if we want to make this "at least one good scale for every pair of points" argument work, we need to show that for every two vertices u, v there is a "large" distance d_1 and a "small" distance d_2 (whose difference is a constant times $d(u, v)$) such that a large-radius ball around one of the vertices and a small-radius ball around the other vertex contain roughly the same number of elements of V.

Consider, however, the following example.

Example 49 (Joined Trees) *Consider the graph obtained by taking two complete binary trees of the same size and identifying their leaves, as in the picture below.*

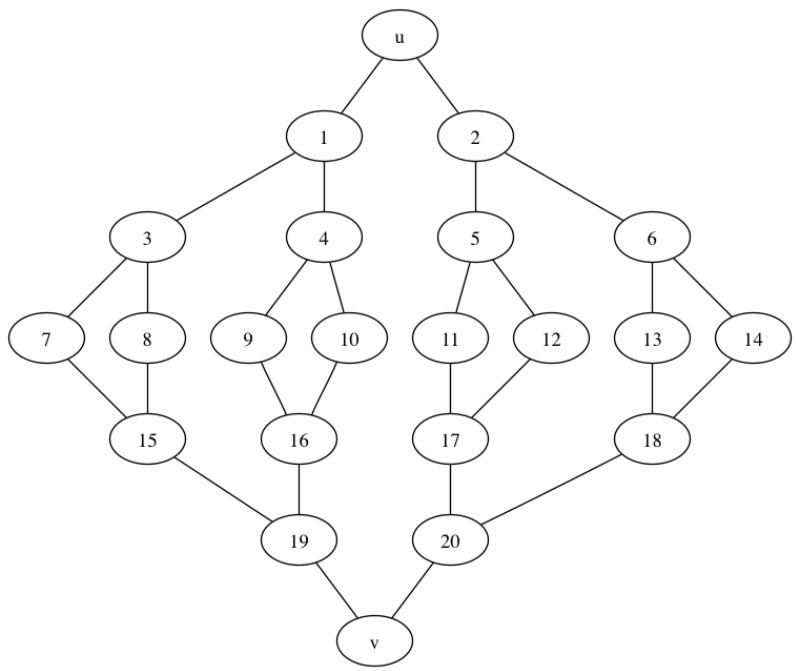

Consider the shortest-path metric $d(\cdot,\cdot)$ in the above graph. Consider the "root" vertices u and v. Their distance $d(u,v)$ is $\approx \log n$, but, at every scale t, both $f_A(u)$ and $f_A(v)$ are highly concentrated around t and, it can be calculated that, at every scale t, we have

$$\mathop{\mathbb{E}}_{A\sim D_t}[|f_A(u) - f_A(v)|] = \Theta(1)$$

This is still good, because averaging over all scales we still get

$$\mathop{\mathbb{E}}_{A\sim D}[|f_A(u) - f_A(v)|] \geq \Omega(1) = \frac{1}{O(\log n)} \cdot d(u,v)$$

but this example shows that the analysis cannot be restricted to one good scale but has, in some cases, to take into account the contribution to the expectation coming from all the scales.

In the above example, the only way to get a ball around u and a ball around v with approximately the same number of points is to get balls of roughly the same radius. No scale could then give a large contribution to the expectation $\mathbb{E}_{A\sim D}[|f_A(u)-f_A(v)|]$; every scale, however, gave a noticeable contribution, and adding them up we had a bounded distortion. The above example will be the template for the full proof, which

will do an "amortized analysis" of the contribution to the expectation coming from each scale t, by looking at the radii that define a ball around u and a ball around v with approximately 2^t elements.

7.4.2 The Proof of Bourgain's Theorem

Given Fact 42 and Fact 46, proving Bourgain's theorem reduces to proving the following theorem.

Theorem 50 *For a finite set of points V, consider the distribution D over subsets of V sampled by uniformly picking a scale $t \in \{0, \ldots, \log_2 |V|\}$ and then picking independently each $v \in V$ to be in A with probability 2^{-t}. Let $d : V \times V \to \mathbb{R}$ be a semimetric. Then for every $u, v \in V$,*

$$\mathop{\mathbb{E}}_{A \sim D}[|f_A(u) - f_A(v)|] \geq \frac{1}{c \log_2 |V|} \cdot d(u, v)$$

where c is an absolute constant.

Proof: For each t, let ru_t be the distance from u to the 2^t-th closest point to u (counting u). That is,

$$|\{w : d(u, w) < ru_t\}| < 2^t$$

and

$$|\{w : d(u, w) \leq ru_t\}| \geq 2^t$$

and define rv_t similarly. Let t^* be the scale such that both ru_{t^*} and rv_{t^*} are smaller than $d(u, v)/3$, but at least one of ru_{t^*+1} or rv_{t^*+1} are $\geq d(u, v)/3$.

Define

$$ru'_t := \min\{ru_t, d(u, v)/3\}$$

and similarly

$$rv'_t := \min\{rv_t, d(u, v)/3\}$$

We claim that there is an absolute constant c such that for every scale $t \in \{0, \ldots, t^*\}$, we have

$$\mathop{\mathbb{E}}_{A \sim D_t} |f_A(u) - f_A(v)| \geq c \cdot (ru'_{t+1} + rv'_{t+1} - ru'_t - rv'_t) \tag{17}$$

We prove the claim by showing that there are two disjoint events, each happening with probability $\geq c$, such that in one event $|f_A(u) - f_A(v)| \geq ru'_{t+1} - rv'_t$, and in the other event $|f_A(u) - f_A(v)| \geq rv'_{t+1} - ru'_t$.

1. The first event is that A avoids the set $\{z : d(u,z) < ru'_{t+1}\}$ and intersects the set $\{z : d(v,z) \leq rv'_t\}$. The former set has size $< 2^{t+1}$, and the latter set has size $\leq 2^t$; the sets are disjoint because we are looking at balls or radius $\leq d(u,v)/3$ around u and v; so the event happens with a probability that is at least an absolute constant. When the event happens,

$$|f_A(u) - f_A(v)| \geq f_A(u) - f_A(v) \geq ru'_{t+1} - rv'_t$$

2. The second event is that A avoids the set $\{z : d(v,z) < rv'_{t+1}\}$ and intersects the set $\{z : d(u,z) \leq ru'_t\}$. The former set has size $< 2^{t+1}$, and the latter set has size $\leq 2^t$; the sets are disjoint because we are looking at balls or radius $\leq d(u,v)/3$ around u and v; so the event happens with a probability that is at least an absolute constant. When the event happens,

$$|f_A(u) - f_A(v)| \geq f_A(v) - f_A(u) \geq rv'_{t+1} - ru'_t$$

So we have established (17). Averaging over all scales, we have

$$\begin{aligned} &\mathop{\mathbb{E}}_{A\sim D} |f_A(u) - f_A(v)| \\ \geq &\frac{c}{1+\log_2 n} \cdot (ru'_{t^*+1} + rv'_{t^*+1} - ru'_0 - rv'_0) \\ \geq &\frac{c}{1+\log_2 n} \cdot \frac{d(u,v)}{3} \end{aligned}$$

□

There is one remaining point to address. In Fact 42, we proved that a distribution over embeddings on the line can be turned into an L1 embeddings, in which the number of dimensions is equal to the size of the support of the distribution. In our proof, we have used a distribution that ranges over $2^{|V|}$ possible functions, so this would give rise to an embedding that uses a superpolynomial number of dimensions.

To fix this remaining problem, we sample $m = O(\log^3 |V|)$ sets $A_1, \ldots, A_m$ and we define the embedding $f(u) := (m^{-1} \cdot f_{A_1}(u), \ldots, m^{-1} \cdot f_{A_m}(u))$. It remains to prove that this randomized mapping has low distortion with high probability, which is an immediate consequence of the Chernoff bounds. Specifically, we use the following form of the Chernoff bound:

Lemma 51 *Let $Z_1, \ldots, Z_m$ be independent nonnegative random variables such that, with probability 1, $0 \leq Z_i \leq M$. Let $Z := \frac{1}{m}(Z_1 + \cdots + Z_m)$. Then*

$$\mathbb{P}[\mathbb{E}\, Z - Z \geq t] \leq e^{-2mt^2/M^2}$$

Let us look at any two vertices u, v. Clearly, for every choice of $A_1, \ldots, A_m$, we have $||f(u) - f(v)||_1 \leq d(u, v)$ so it remains to prove a lower bound to their L1 distance. Let us call Z the random variable denoting their L1 distance, that is

$$Z := ||f(u) - f(v)|| = \sum_{i=1}^{m} \frac{1}{m} |f_{A_i}(u) - f_{A_i}(v)|$$

We can write $Z = \frac{1}{m} \cdot (Z_1 + \cdots + Z_m)$ where $Z_i := |f_{A_i}(u) - f_{A_i}(v)|$, so that Z is the sum of identically distributed nonnegative random variables, such that

$$\begin{aligned} Z_i &\leq d(u, v) \\ \mathbb{E}\, Z_i &\geq \frac{c}{\log |V|} d(u, v) \end{aligned}$$

Applying the Chernoff bound with $M = d(u, v)$ and $t = \frac{c}{2 \log |V|} d(u, v)$, we have

$$\begin{aligned} &\mathbb{P}\left[Z \leq \frac{c}{2 \log |V|} d(u, v) \right] \\ &\leq \mathbb{P}\left[Z \leq \mathbb{E}\, Z - \frac{c}{2 \log |V|} d(u, v) \right] \\ &\leq 2^{-2mc^2/(2 \log |V|)^2} \end{aligned}$$

which is, say, $\leq 1/|V|^3$ if we choose $m = c' \log^3 |V|$ for an absolute constant c'.

By taking a union bound over all pairs of vertices,

$$\mathbb{P}\left[\forall u, v.\ ||f(u) - f(v)||_1 \geq \frac{c}{2 \log |V|} \cdot d(u, v) \right] \geq 1 - \frac{1}{|V|}$$

7.5 Tightness of the Analysis of the Leighton-Rao Relaxation *

If (X, d) and (X', d') are metric spaces, we say that a mapping $f : X \to X'$ is an *embedding of* (X, d) *into* (X', d) *with distortion at most* c if there are parameters c_1, c_2, with $c = c_1 c_2$ such that, for every $u, v \in X$, we have

$$\frac{1}{c_1} \cdot d'(u, v) \leq d(u, v) \leq c_2 \cdot d'(u, v)$$

The metric space $\mathbb{R}^m$ with distance $||u - v|| = \sqrt{\sum_i (u_i - v_i)^2}$ is denoted by ℓ_m^2, and the metric space $\mathbb{R}^m$ with distance $||u - v||_1 = \sum_i |u_i - v_i|$ is denoted by ℓ_m^1. In the past lecture we proved the following result.

Theorem 52 (Bourgain) *There is an absolute constant c such that every finite metric space (V, d) embeds into ℓ_m^1 with distortion at most $c \log |V|$, where $m = O(\log^3 |V|)$.*

If we solve the Leighton-Rao linear programming relaxation to approximate the sparsest cut of a graph $G = (V, E)$, and we let $d(\cdot, \cdot)$ be an optimal solution, we note that, if we weigh each edge $(u, v) \in E$ by $d(u, v)$, and then compute shortest paths in this weighted graph, then, for every two vertices x, y, the distance $d(x, y)$ is precisely the length of the shortest path from x to y. In particular, if we are using the Leighton-Rao relaxation in order to approximate the sparsest cut in a given *planar* graph, for example, then the solution $d(\cdot, \cdot)$ that we need to round is not an arbitrary metric space, but it is the shortest path metric of a weighted planar graph. It is conjectured that, in this case, the Leighton-Rao relaxation could deliver a constant-factor approximation.

Question 1 *Is there an absolute constant c such that every metric space (X, d) constructed as the shortest-path metric over the vertices of a planar graph can be embedded into ℓ_m^1 with distortion at most c, where $m = |V|^{O(1)}$?*

So far, it is known that k-*outerplanar* graphs, for constant k, embed in ℓ^1 with constant distortion.

This is just an example of a large family of questions that can be asked about the embeddability of various types of metric spaces into each other.

For general finite metric spaces, the logarithmic distortion of Bougain's theorem is best possible.

In order to prove the optimality of Bourgain's theorem, we will state without proof the existence of constant degree families of expanders.

Theorem 53 (Existence of Expanders) *There are absolute constants d and c such that, for infinitely many n, there is an n-vertex d-regular graph G_n such that $\phi(G_n) \geq c$.*

On such graphs, the Leighton-Rao relaxation is $LR(G_n) \leq O(1/\log n)$, showing that our proof that $LR(G) \geq \phi(G)O(logn)$ is tight.

For every two vertices u, v, define $d(u, v)$ as the length of (that is, the number of edges in) a shortest path from u to v in G_n.

Then

$$\sum_{u,v} A_{u,v} d(u, v) = 2|E|$$

Because each graph G_n is d-regular, it follows that for every vertex v there are $\leq 1 + d + \cdots + d^k < d^{k+1}$ vertices at distance $\leq k$ from v. In particular, at least half of the vertices have distance $\geq t$ from v, where $t = \lfloor \log_d n/2 \rfloor - 1$, which implies that

$$\sum_{u,v} d(u,v) \geq n \cdot \frac{n}{2} \cdot t = \Omega(n^2 \log n)$$

Recall that

$$LR(G) = \min_{d \text{ semimetric}} \frac{|V|^2}{2|E|} \frac{\sum_{u,v} A_{u,v} d(u,v)}{\sum_{u,v} d(u,v)}$$

and so

$$LR(G_n) \leq O\left(\frac{1}{\log n}\right)$$

even though

$$\phi(G_n) \geq \Omega(1)$$

Note that we have also shown that every embedding of the shortest-path metric $d(\cdot,\cdot)$ on G_n into ℓ^1 requires distortion $\Omega(\log n)$, and so we have proved the tightness of Bourgain's theorem.

7.6 The Non-Uniform Sparsest Cut Problem

In the *non-uniform sparsest cut* problem, we are given two graphs $G = (V, E_G)$ and $H = (V, E_H)$ over the same set of vertices, and we want to find a non-trivial cut $(S, V - S)$ of the vertex set that minimizes the *non-uniformity sparsity*

$$\phi(G, H, S) := \frac{|E_H|}{|E_G|} \cdot \frac{\sum_{(u,v)\in G} |1_S(u) - 1_S(v)|}{\sum_{(u,v)\in H} |1_S(u) - 1_S(v)|}$$

The non-uniformity sparsity of an optimal solution is denoted by $\phi(G, H)$. Note that the standard sparsest cut problem corresponds to taking H to be a clique.

The graph G is called the *capacity* graph and the graph H is called the *demand* graph. This terminology will become more clear when we discuss the dual of the Leighton-Rao relaxation of sparsest cut at a later point. The intuition for the problem is that H specifies which pairs of nodes we would like to see disconnected in G, and our goal is to find a set of edges to remove such that their number is small compared to the number of pairs of elements of H that are disconnected by their removal.

We can easily adapt the theory that we have developed for sparsest cut to the non-uniform version of the problem. In particular, we can define a Leighton-Rao relaxation

$$LR(G,H) := \min_{d \text{ semimetric}} \frac{|E_H|}{|E_G|} \frac{\sum_{u,v} A_{u,v} d(u,v)}{\sum_{u,v} B_{u,v} d(u,v)}$$

where A is the adjacency matrix of G and B is the adjacency matrix of H.

We can also define the special case of the Leighton-Rao relaxation in which the distance function is an ℓ^1 embedding:

$$\phi'(G,H) := \min_{m, f: V \to \mathbb{R}^m} \frac{|E_H|}{|E_G|} \frac{\sum_{u,v} A_{u,v} ||f(u) - f(v)||_1}{\sum_{u,v} B_{u,v} ||f(u) - f(v)||_1}$$

Exactly the same proof that established $\phi(G) = \phi'(G)$ can be used to prove

$$\phi'(G,H) = \phi'(G,H)$$

And, finally, Bourgain's embedding can be used to map a solution of the Leighton-Rao relaxation to an ℓ^1 solution, with $O(\log n)$ increase in the cost function, and then the ℓ^1 solution can be mapped to a cut with no loss, leading to the relation

$$LR(G,H) \leq \phi(G,H) \leq O(\log |V|) \cdot LR(G,H)$$

for every two graphs G, H over the same set of vertices.

8 The Arora-Rao-Vazirani Relaxation

Recall that the sparsest cut $\phi(G)$ of a graph $G = (V,E)$ with adjacency matrix A is defined as

$$\phi(G) = \min_{S \subseteq V} \frac{\frac{1}{2|E|} \sum_{u,v} A_{u,v} |1_S(u) - 1_S(v)|}{\frac{1}{|V|^2} \sum_{u,v} |1_S(u) - 1_S(v)|}$$

and the Leighton-Rao relaxation is obtained by noting that if we define $d(u,v) := |1_S(u) - 1_S(v)|$ then $d(\cdot,\cdot)$ is a semimetric over V, so that the following quantity is a relaxation of $\phi(G)$:

$$LR(G) = \min_{\substack{d: V \times V \to R \\ d \text{ semimetric}}} \frac{\frac{1}{2|E|} \sum_{u,v} A_{u,v} d(u,v)}{\frac{1}{|V|^2} \sum_{u,v} d(u,v)}$$

If G is d-regular, and we call $M := \frac{1}{d} \cdot A$ the normalized adjacency matrix of A, and we let $\lambda_1 = 1 \geq \lambda_2 \geq \cdots \geq \lambda_n$ be the eigenvalues of M with multiplicities, then we proved in a past lecture that

$$1 - \lambda_2 = \min_{x:V \to \mathbb{R}} \frac{\frac{1}{2|E|} \sum_{u,v} A_{u,v} |x(u) - x(v)|^2}{\frac{1}{|V|^2} \sum_{u,v} |x(u) - x(v)|^2} \tag{18}$$

which is also a relaxation of $\phi(G)$, because, for every S, every u and every v, $|1_S(u) - 1_S(v)| = |1_S(u) - 1_S(v)|^2$.

We note that if we further relax (18) by allowing V to be mapped into a higher dimension space $\mathbb{R}^m$ instead of $\mathbb{R}$, and we replace $|\cdot - \cdot|$ by $||\cdot - \cdot||^2$, the optimum remains the same.

Fact 54

$$1 - \lambda_2 = \min_{m, \mathbf{x}:V \to \mathbb{R}^m} \frac{\frac{1}{2|E|} \sum_{u,v} A_{u,v} ||\mathbf{x}(u) - \mathbf{x}(v)||^2}{\frac{1}{|V|^2} \sum_{u,v} ||\mathbf{x}(u) - \mathbf{x}(v)||^2}$$

Proof: For a mapping $\mathbf{x} : V \to \mathbb{R}^m$, define

$$\delta(\mathbf{x}) := \frac{\frac{1}{2|E|} \sum_{u,v} A_{u,v} ||\mathbf{x}(u) - \mathbf{x}(v)||^2}{\frac{1}{|V|^2} \sum_{u,v} ||\mathbf{x}(u) - \mathbf{x}(v)||^2}$$

It is enough to show that, for every $\mathbf{x}$, $1 - \lambda_2 \leq \delta(\mathbf{x})$. Let $x_i(v)$ be the i-th coordinate of $\mathbf{x}(v)$. Then

$$\begin{aligned}
\delta(\mathbf{x}) &= \frac{\frac{1}{2|E|} \sum_i \sum_{u,v} A_{u,v} |x_i(u) - x_i(v)|^2}{\frac{1}{|V|^2} \sum_i \sum_{u,v} |x_i(u) - x_i(v)|^2} \\
&\geq \min_i \frac{\frac{1}{2|E|} \sum_{u,v} A_{u,v} |x_i(u) - x_i(v)|^2}{\frac{1}{|V|^2} \sum_{u,v} |x_i(u) - x_i(v)|^2} \\
&\geq 1 - \lambda_2
\end{aligned}$$

where the second-to-last inequality follows from the fact, which we have already used before, that for nonnegative $a_1, \ldots, a_m$ and positive $b_1, \ldots, b_m$ we have

$$\frac{a_1 + \cdots a_m}{b_1 + \cdots + b_m} \geq \min_i \frac{a_i}{b_i}$$

□

The above observations give the following comparison between the Leighton-Rao relaxation and the spectral relaxation: both are obtained by replacing $|1_S(u) - 1_S(v)|$

with a "distance function" $d(u,v)$; in the Leighton-Rao relaxation, $d(u,v)$ is constrained to satisfy the triangle inequality; in the spectral relaxation, $d(u,v)$ is constrained to be the square of the Euclidean distance between $\mathbf{x}(u)$ and $\mathbf{x}(v)$ for some mapping $\mathbf{x} : V \to \mathbb{R}^m$.

The Arora-Rao-Vazirani relaxation is obtained by *enforcing both conditions*, that is, by considering distance functions $d(u,v)$ that satisfy the triangle inequality *and* can be realized of $||\mathbf{x}(u) - \mathbf{x}(v)||^2$ for some mapping $\mathbf{x} : V \to \mathbb{R}^m$.

Definition 55 *A semimetric $d : V \to V \to \mathbb{R}$ is called of negative type if there is a dimension m and a mapping $\mathbf{x} : V \to \mathbb{R}^m$ such that $d(u,v) = ||\mathbf{x}(u) - \mathbf{x}(v)||^2$ for every $u, v \in V$.*

With the above definition, we can formulate the Arora-Rao-Vazirani relaxation as

$$ARV(G) := \min_{\substack{d : V \times V \to R \\ d \text{ semimetric of negative type}}} \frac{\frac{1}{2|E|} \sum_{u,v} A_{u,v} d(u,v)}{\frac{1}{|V|^2} \sum_{u,v} d(u,v)} \tag{19}$$

Remark 56 *The relaxation (19) was first proposed by Goemans and Linial. Arora, Rao and Vazirani were the first to prove that it achieves an approximation guarantee which is better than the approximation guarantee of the Leighton-Rao relaxation.*

We have, by definition,

$$\phi(G) \leq ARV(G) \leq \min\{LR(G), 1 - \lambda_2(G)\}$$

and so the approximation results that we have proved for $1 - \lambda_2$ and LR apply to ARV. For every graph $G = (V, E)$

$$ARV(G) \leq O(\log |V|) \cdot \phi(G)$$

and for every regular graph

$$ARV(G) \leq \sqrt{8 \cdot \phi(G)}$$

Interestingly, the examples that we have given of graphs for which LR and $1 - \lambda_2$ give poor approximation are complementary. If G is a cycle, then $1 - \lambda_2$ is a poor approximation of $\phi(G)$, but $LR(G)$ is a good approximation of $\phi(G)$; if G is a constant-degree expander then $LR(G)$ is a poor approximation of $\phi(G)$, but $1 - \lambda_2$ is a good approximation.

When Goemans and Linial (separately) proposed to study the relaxation (19), they conjectured that it would always provide a constant-factor approximation of $\phi(G)$. Unfortunately, the conjecture turned out to be false, but Arora, Rao and Vazirani were able to prove that (19) does provide a strictly better approximation than the Leighton-Rao relaxation. In the next lectures, we will present parts of the proof of the following results.

Theorem 57 *There is a universal constant c such that, for every graph $G = (V, E)$,*

$$ARV(G) \leq c \cdot \sqrt{\log |V|} \cdot \phi(G)$$

Theorem 58 *There is an absolute constant c and an infinite family of graphs $G_n = (V_n, E_n)$ such that*

$$ARV(G) \geq c \cdot \log\log |V_n| \cdot \phi(G)$$

8.1 The Ellipsoid Algorithm and Semidefinite Programming

We now briefly discuss the polynomial time solvability of (19).

Definition 59 *If $C \subseteq \mathbb{R}^m$ is a set, then a separation oracle for C is a procedure that, on input $\mathbf{x} \in R^m$,*

- *If $\mathbf{x} \in C$, outputs "yes"*
- *If $\mathbf{x} \notin C$, outputs coefficients $a_1, \ldots, a_m, b$ such that*

$$\sum_i x_i a_i < b$$

  but, for every $\mathbf{z} \in C$,

$$\sum_i z_i a_i \geq b$$

Note that a set can have a separation oracle only if it is convex. Under certain additional mild conditions, if C has a polynomial time computable separation oracle, then the optimization problem

$$\begin{array}{ll} \text{minimize} & \sum_i \mathbf{c}^T \mathbf{x} \\ \text{subject to} & \\ & A\mathbf{x} \geq b \\ & \mathbf{x} \in C \end{array}$$

is solvable in polynomial time using the Ellipsoid Algorithm.

It remains to see how to put the Arora-Rao-Vazirani relaxation into the above form.

Recall that a matrix $X \in \mathbb{R}^{n \times n}$ is *positive semidefinite* if all its eigenvalues are nonnegative. We will use the set of all $n \times n$ positive semidefinite matrices as our set C (thinking of them as n^2-dimensional vectors). If we think of two matrices $M, M' \in \mathbb{R}^{n \times n}$ as n^2-dimensional vectors, then their "inner product" is

$$M \bullet M' := \sum_{i,j} M_{i,j} \cdot M'_{i,j}$$

Lemma 60 *The set of $n \times n$ positive semidefinite matrices has a separation oracle computable in time polynomial in n.*

PROOF: Given a symmetric matrix X, its smallest eigenvalue is

$$\min_{\mathbf{z} \in \mathbb{R}^n,\ ||\mathbf{z}||=1} \mathbf{z}^T X \mathbf{z}$$

the vector achieving the minimum is a corresponding eigenvector, and both the smallest eigenvalue and the corresponding eigenvector can be computed in polynomial time.

If we find that the smallest eigenvalue of X is non-negative, then we answer "yes." Otherwise, if $\mathbf{z}$ is an eigenvector of the smallest eigenvalue we output the matrix $A = \mathbf{z}^T \mathbf{z}$. We see that we have

$$A \bullet X = \mathbf{z}^T X \mathbf{z} < 0$$

but that, for every positive semidefinite matrix M, we have

$$A \bullet M = \mathbf{z}^T M \mathbf{z} \geq 0$$

□

This implies that any optimization problem of the following form can be solved in polynomial time

$$\begin{array}{ll} \text{minimize} & C \bullet X \\ \text{subject to} & \\ & A^1 \bullet X \geq b_1 \\ & \cdots \\ & A^m \bullet X \geq b_m \\ & X \succeq 0 \end{array} \tag{20}$$

where $C, A^1, \ldots, A^m$ are square matrices of coefficients, $b_1, \ldots, b_m$ are scalars, and X is a square matrix of variables. An optimization problem like the one above is called a *semidefinite program.*

It remains to see how to cast the Arora-Rao-Vazirani relaxation as a semidefinite program.

Lemma 61 *For a symmetric matrix $M \in \mathbb{R}^{n \times n}$, the following properties are equivalent:*

1. *M is positive semidefinite;*
2. *there are vectors $\mathbf{x}_1, \ldots, \mathbf{x}_n \in \mathbb{R}^d$ such that, for all i, j, $M_{i,j} = \langle \mathbf{x}_i, \mathbf{x}_j \rangle$;*
3. *for every vector $\mathbf{z} \in \mathbb{R}^n$, $\mathbf{z}^T M \mathbf{z} \geq 0$*

Proof: That (1) and (3) are equivalent follows from the characterization of the smallest eigenvalue of M as the minimum of $\mathbf{z}^T M \mathbf{z}$ over all unit vectors $\mathbf{z}$.

To see that (2) $\Rightarrow$ (3), suppose that vectors $\mathbf{x}_1, \ldots, \mathbf{x}_n$ exist as asserted in (2), take any vector $\mathbf{z}$, and see that

$$\mathbf{z}^T M \mathbf{z} = \sum_{i,j} z(i) M_{i,j} z(j)$$

$$= \sum_{i,j,k} z(i) x_i(k) x_j(k) z(j) = \sum_k \left(\sum_i z(i) x_i(k) \right)^2 \geq 0$$

Finally, to see that (1) $\Rightarrow$ (2), let $\lambda_1, \ldots, \lambda_n$ be the eigenvalues of M with multiplicities, and let $\mathbf{v}_1, \ldots, \mathbf{v}_n$ be a corresponding orthonormal set of eigenvectors. Then

$$M = \sum_i \lambda_k \mathbf{v}_k \mathbf{v}_k^T$$

that is,

$$M_{i,j} = \sum_k \lambda_k v_k(i) v_k(j) = \langle \mathbf{x}_i, \mathbf{x}_j \rangle$$

if we define $\mathbf{x}_1, \ldots, \mathbf{x}_n$ as the vectors such that $x_i(k) := \sqrt{\lambda_k} v_k(i)$. $\square$

This means that the generic semidefinite program (20) can be rewritten as an optimization problem in which the variables are the vectors $\mathbf{x}_1, \ldots, \mathbf{x}_n$ as in part (2) of the above lemma.

$$\begin{array}{ll} \text{minimize} & \sum_{i,j} C_{i,j} \langle \mathbf{x}_i, \mathbf{x}_j \rangle \\ \text{subject to} & \\ & \sum_{i,j} A^1_{i,j} \langle \mathbf{x}_i, \mathbf{x}_j \rangle \geq b_1 \\ & \cdots \\ & \sum_{i,j} A^m_{i,j} \langle \mathbf{x}_i, \mathbf{x}_j \rangle \geq b_m \\ & \mathbf{x}_i \in \mathbb{R}^d \qquad\qquad \forall i \in \{1, \ldots, n\} \end{array} \tag{21}$$

where the dimension d is itself a variable (although one could fix it, without loss of generality, to be equal to n). In this view, a semidefinite program is an optimization problem in which we wish to select n vectors such that their pairwise inner products satisfy certain linear inequalities, while optimizing a cost function that is linear in their pairwise inner product.

The square of the Euclidean distance between two vectors is a linear function of inner products

$$||\mathbf{x} - \mathbf{y}||^2 = \langle \mathbf{x} - \mathbf{y}, \mathbf{x} - \mathbf{y} \rangle = \langle \mathbf{x}, \mathbf{x} \rangle - 2\langle \mathbf{x}, \mathbf{y} \rangle + \langle \mathbf{y}, \mathbf{y} \rangle$$

and so, in a semidefinite program, we can include expressions that are linear in the pairwise squared distances (or squared norms) of the vectors. The ARV relaxation can be written as follows

$$\begin{array}{ll} \text{minimize} & \sum_{u,v} A_{u,v} ||\mathbf{x}_u - \mathbf{x}_v||^2 \\ \text{subject to} & \\ & \sum_{u,v} ||\mathbf{x}_u - \mathbf{x}_v||^2 = \frac{|V|^2}{2|E|} \\ & ||\mathbf{x}_u - \mathbf{x}_v||^2 \leq ||\mathbf{x}_u - \mathbf{x}_w||^2 + ||\mathbf{x}_w - \mathbf{x}_v||^2 \quad \forall u, v, w \in V \\ & \mathbf{x}_u \in \mathbb{R}^d \qquad\qquad\qquad\qquad\qquad\qquad \forall u \in V \end{array}$$

and so it is a semidefinite program, and it can be solved in polynomial time.

Remark 62 *Our discussion of polynomial time solvability glossed over important issues about numerical precision. To run the Ellipsoid Algorithm one needs, besides the separation oracle, to be given a ball that is entirely contained in the set of feasible solutions and a ball that entirely contains the set of feasible solutions, and the running time of the algorithm is polynomial in the size of the input, polylogarithmic in the ratio of the volumes of the two balls, and polylogarithmic in the desired amount of precision. At the end, one doesn't get an optimal solution, which might not have a finite-precision exact representation, but an approximation within the desired precision. The algorithm is able to tolerate a bounded amount of imprecision in the*

separation oracle, which is an important feature because we do not have exact algorithms to compute eigenvalues and eigenvectors (the entries in the eigenvector might not have a finite-precision representation).

The Ellipsoid algorithm is typically not a practical algorithm. Algorithms based on the *interior point* method have been adapted to semidefinite programming, and run both in worst-case polynomial time and in reasonable time in practice.

Arora and Kale have developed an $\tilde{O}((|V|+|E|)^2/\epsilon^{O(1)})$ time algorithm to solve the ARV relaxation within a multiplicative error $(1+\epsilon)$. The dependency on the error is worse than that of generic algorithms, which achieve polylogarithmic dependency, but this is not a problem in this application, because we are going to lose an $O(\sqrt{\log |V|})$ factor in the rounding, so an extra constant factor coming from an approximate solution of the relaxation is a low-order consideration.

8.2 Rounding the Arora-Rao-Vazirani Relaxation

Given the equivalence between the sparsest cut problem and the "ℓ_1 relaxation" of sparsest cut, it will be enough to prove the following result.

Theorem 63 (Rounding of ARV) *Let G be a graph, A its adjacency matrix, and $\{\mathbf{x}_v\}_{v\in V}$ be a feasible solution to the ARV relaxation.*

Then there is a mapping $f : V \to \mathbb{R}$ such that

$$\frac{\sum_{u,v} A_{u,v}|f(u)-f(v)|}{\sum_{u,v}|f(u)-f(v)|} \leq O(\sqrt{\log |V|}) \cdot \frac{\sum_{u,v} A_{u,v}||\mathbf{x}_u-\mathbf{x}_v||^2}{\sum_{u,v}||\mathbf{x}_u-\mathbf{x}_v||^2}$$

As in the rounding of the Leighton-Rao relaxation via Bourgain's theorem, we will identify a set $S \subseteq V$, and define

$$f_S(v) := \min_{s\in S} ||\mathbf{x}_s - \mathbf{x}_v||^2 \tag{22}$$

Recall that, as we saw in the proof of Bourgain's embedding theorem, no matter how we choose the set S we have

$$|f_S(u) - f_S(v)| \leq ||\mathbf{x}_u - \mathbf{x}_v||^2 \tag{23}$$

where we are not using any facts about $||\cdot - \cdot||^2$ other than the fact that, for solutions of the ARV relaxation, it is a distance function that obeys the triangle inequality.

This means that, in order to prove the theorem, we just have to find a set $S \subseteq V$ such that

$$\sum_{u,v} |f_S(u) - f_S(v)| \geq \frac{1}{O(\sqrt{\log |V|})} \cdot \sum_{u,v} ||\mathbf{x}_u - \mathbf{x}_v||^2 \tag{24}$$

and this is a considerable simplification because the above expression is completely independent of the graph! The remaining problem is purely one about geometry.

Recall that if we have a set of vectors $\{\mathbf{x}_v\}_{v \in V}$ such that the distance function $d(u,v) := ||\mathbf{x}_u - \mathbf{x}_v||^2$ satisfies the triangle inequality, then we say that $d(\cdot,\cdot)$ is a (semi-)metric of *negative type*.

After these preliminaries observations, our goal is to prove the following theorem.

Theorem 64 (Rounding of ARV – Revisited) *If $d(\cdot,\cdot)$ is a semimetric of negative type over a set V, then there is a set S such that if we define*

$$f_S(v) := \min_{s \in S} \{d(s,v)\}$$

we have

$$\sum_{u,v} |f_S(u) - f_S(v)| \geq \frac{1}{O(\sqrt{\log |V|})} \cdot \sum_{u,v} d(u,v)$$

Furthermore, the set S can be found in randomized polynomial time with high probability given a set of vector $\{\mathbf{x}_v\}_{v \in V}$ such that $d(u,v) = ||\mathbf{x}_u - \mathbf{x}_v||^2$.

Since the statement is scale-invariant, we can restrict ourselves, with no loss of generality, to the case $\sum_{u,v} d(u,v) = |V|^2$.

Remark 65 Let us discuss some intuition before continuing with the proof.

As our experience proving Bourgain's embedding theorem shows us, it is rather difficult to pick sets such that $|f_S(u) - f_S(v)|$ is not much smaller than $d(u,v)$. Here we have a somewhat simpler case to solve because we are not trying to preserve all distances, but only the average pairwise distance. A simple observation is that if we find a set S which contains $\Omega(|V|)$ elements and such that $\Omega(|V|)$ elements of V are at distance $\Omega(\delta)$ from S, then we immediately get $\sum_{u,v} |f_S(u) - f_S(v)| \geq \Omega(\delta |V|^2)$, because there will be $\Omega(|V|^2)$ pairs u, v such that $f_S(u) = 0$ and $f_S(v) \geq \delta$. In particular, if we could find such a set with $\delta = 1/O(\sqrt{\log |V|})$ then we would be done. Unfortunately this is too much to ask for in general, because we always have $|f_S(u) - f_S(v)| \leq d(u,v)$, which means that if we want $\sum_{u,v} |f_S(u) - f_S(v)|$ to have $\Omega(V^2)$ noticeably large terms we must also have that $d(u,v)$ is noticeably large for $\Omega(|V|^2)$ pairs of points, which is not always true.

There is, however, the following argument, which goes back to Leighton and Rao: either there are $\Omega(|V|)$ points concentrated in a ball whose radius is a quarter (say) of the

average pairwise distance, and then we can use that ball to get an ℓ_1 mapping with only *constant* error; or there are $\Omega(|V|)$ points in a ball of radius twice the average pairwise distance, such that the pairwise distances of the points in the ball account for a constant fraction of all pairwise distances. In particular, the sum of pairwise distances includes $\Omega(|V|^2)$ terms which are $\Omega(1)$.

After we do this reduction and some scaling, we are left with the task of proving the following theorem: suppose we are given an n-point negative type metric in which the points are contained in a ball of radius 1 and are such that the sum of pairwise distances is $\Omega(n^2)$; then there is a subset S of size $\Omega(n)$ such that there are $\Omega(n)$ points whose distance from the set is $1/O(\sqrt{\log n})$. This theorem is the main result of the Arora-Rao-Vazirani paper. (Strictly speaking, this form of the theorem was proved later by Lee – Arora, Rao and Vazirani had a slightly weaker formulation.)

We begin by considering the case in which a constant fraction of the points are concentrated in a small ball.

Definition 66 (Ball) *For a point $z \in V$ and a radius $r > 0$, the ball of radius r and center z is the set*

$$B(z, r) := \{v : d(z, v) \leq r\}$$

Lemma 67 *For every vertex z, if we define $S := B(z, 1/4)$, then*

$$\sum_{u,v} |f_S(u) - f_S(v)| \geq \frac{|S|}{2|V|} \sum_{u,v} d(u, v)$$

Proof: Our first calculation is to show that the typical value of $f_S(u)$ is rather large. We note that for every two vertices u and v, if we call a a closest vertex in S to u, and b a closest vertex to v in S, we have

$$\begin{aligned} d(u, v) &\leq d(u, a) + d(a, z) + d(z, b) + d(b, v) \\ &\leq f_S(u) + f_S(v) + \frac{1}{2} \end{aligned}$$

and so

$$|V|^2 = \sum_{u,v} d(u, v) \leq 2|V| \cdot \sum_{v} f_S(v) + \frac{|V|^2}{2}$$

that is,

$$\sum_v f_S(v) \geq \frac{|V|}{2}$$

Now we can get a lower bound to the sum of ℓ_1 distances given by the embedding $f_S(\cdot)$.

$$\begin{aligned}
&\sum_{u,v} |f_S(u) - f_S(v)| \\
&\geq \sum_{u \in S, v \in V} |f_S(v)| \\
&= |S| \sum_v f_S(v) \\
&\geq \frac{1}{2}|S| \cdot |V|
\end{aligned}$$

□

This means that if there is a vertex z such that $|B(z, 1/4)| = \Omega(|V|)$, or even $|B(z, 1/4)| = \Omega(|V|/\sqrt{\log |V|})$, then we are done.

Otherwise, we will find a set of $\Omega(|V|)$ vertices such that their average pairwise distances are within a constant factor of their maximum pairwise distances, and then we will work on finding an embedding for such a set of points. (The condition that the average distance is a constant fraction of the maximal distance will be very helpful in subsequent calculations.)

Lemma 68 *Suppose that for every vertex z we have $|B(z, 1/4)| \leq |V|/4$. Then there is a vertex w such that, if we set $S = B(w, 2)$, we have*

- $|S| \geq \frac{1}{2} \cdot |V|$
- $\sum_{u,v \in S} d(u, v) \geq \frac{1}{8}|S|^2$

Proof: Let w be a vertex that maximizes $|B(w, 2)|$; then $|B(w, 2)| \geq |V|/2$, because if we had $|B(u, 2)| < |V|/2$ for every vertex u, then we would have

$$\sum_{u,v} d(u, v) > \sum_u 2 \cdot (|V - B(u, 2)|) > |V|^2$$

Regarding the sum of pairwise distances of elements of S, we have

$$\sum_{u,v \in S} d(u, v) > \sum_{u \in S} \frac{1}{4}(|S - B(u, 1/4)|) \geq |S| \cdot \frac{1}{4} \cdot \frac{|S|}{2}$$

□

The proof of the main theorem now reduces to proving the following geometric fact, which we state without its (difficult) proof.

Theorem 69 *Let d be a negative-type metric over a set V such that the points are contained in a unit ball and have constant average distance, that is,*

- *there is a vertex z such that $d(v, z) \leq 1$ for every $v \in V$*
- $\sum_{u,v \in V} d(u, v) \geq c \cdot |V|^2$

Then there are sets $S, T \subseteq V$ such that

- $|S|, |T| \geq \Omega(|V|)$*;*
- *for every $u \in S$ and every $v \in S$, $d(u, v) \geq 1/O(\sqrt{\log |V|})$*

where the multiplicative factors hidden in the $O(\cdot)$ and $\Omega(\cdot)$ notations depend only on c.