Eigenvalue methods in extremal combinatorics: an overview

David Ellis

Eigenvalues and independent sets in graphs

Our aim in this lecture is to show how one may often gain combinatorial information about graphs by looking at the eigenvalues of certain linear operators (or matrices) associated with these graphs.

We start from the perspective of extremal set theory. Suppose we have a universe of objects, U, and we are interested in subsets of U with a certain property, P. Suppose we wish to determine the maximum possible size of a subset $S \subset U$ which has the property P.

If the universe U of objects is sufficiently simple, we can often determine the maximum possible size of a subset of U having property P using simple combinatorial techniques, such as induction, averaging and shifting.

However, when the universe of objects has a more complicated structure, these techniques can fail. Algebraic techniques which utilize this structure are sometimes more successful.

Often, our property P is a condition on all *pairs* of members of the subset S. (This is the case, for example, when P is the property that any two members of S are 't-intersecting'.) In other words, there is a binary relation R on elements of U such that

$$S \text{ has property } P \quad \Leftrightarrow \quad xRy \text{ for any two distinct } x, y \in S.$$

Let H be the graph on vertex-set U where we join x and y if xRy; then subsets of U with property P are precisely *cliques* in the graph H, so our problem is equivalent to finding the maximum possible size of a clique in the graph H. Or, considering the complement $\bar{H}$, our problem is to find the maximum possible size of an independent set in $\bar{H}$. (Recall that if G is a graph, an *independent set* in G is a set of vertices of G with no edges of G between them.)

It turns out that we can obtain upper bounds on the maximum possible size of an independent set in a graph in terms of the *eigenvalues* of certain linear operators (or matrices) associated with the graph. The simplest bound is Hoffman's theorem, which uses the *adjacency matrix* of the graph.

Recall that if $G = (V, E)$ is an n-vertex graph, the *adjacency matrix* of G is the 0-1 matrix with rows and columns indexed by V, and with

$$A_{x,y} = \begin{cases} 1 & \text{if } xy \in E(G); \\ 0 & \text{otherwise.} \end{cases}$$

Since A is a real symmetric matrix, all its eigenvalues are real, and for any inner product on $\mathbb{R}^V$, we can find an orthonormal basis consisting of eigenvectors of A. Let $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_n = \lambda_{\min}$ be the eigenvalues of A, repeated with their multiplicities. Note that if G is a d-regular graph, then the all 1's-vector $\mathbf{1}$ is an eigenvector of A with eigenvalue d. It is easily checked that the constant vectors are eigenvectors of A only if G is regular, and that if G is d-regular, then $\lambda_1 = d$.

Since for any graph G,

$$\sum_{i=1}^{n} \lambda_i = \text{Trace}(A) = 0,$$

we must have $\lambda_{\min} \leq 0$. Hoffman's theorem bounds the maximum possible size of an independent set in a d-regular graph in terms of the least eigenvalue of its adjacency matrix:

Theorem 1 (Hoffman)**.** *Let $G = (V, E)$ be a d-regular, n-vertex graph, and let A be the adjacency matrix of G. Let $\lambda_{\min}$ denote the least eigenvalue of A. If $S \subset V$ is an independent set in G, then*

$$\frac{|S|}{n} \leq \frac{-\lambda_{\min}}{d - \lambda_{\min}}.$$

If equality holds, then the characterstic function f_S of S satisfies:

$$f_S - \tfrac{|S|}{n}\mathbf{1} \in \mathit{Ker}(A - \lambda_{\min} I).$$

Proof. We equip $\mathbb{R}^V$, the vector space of real-valued functions on V, with the inner product induced by the uniform measure on V:

$$\langle f, g \rangle = \frac{1}{n} \sum_{x \in V} f(x)g(x).$$

We let $||.||_2$ denote the induced Euclidean norm:

$$||f||_2 = \sqrt{\langle f, f \rangle}.$$

Let $\{\mathbf{1} = v_1, v_2, \ldots, v_n\}$ be an orthonormal basis of real eigenvectors of A, with corresponding eigenvalues $d = \lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_n$.

Let $S \subset V$ be an independent set, let $f = f_S$ be its characteristic function, and let $\alpha = |S|/n$ be its measure. Write f as a linear combination of our eigenvectors:

$$f = \sum_{i=1}^{n} a_i v_i.$$

Then $a_1 = \langle f, v_1 \rangle = \langle f, \mathbf{1} \rangle = \alpha$. Moreover,

$$\sum_{i=1}^{n} a_i^2 = ||f||_2^2 = \alpha.$$

We have

$$\langle f, Af \rangle = \tfrac{1}{n} f^\top A f = \tfrac{1}{n} \sum_{x,y \in S} A_{x,y} = \tfrac{2}{n} e(G[S]) = 0,$$

since S is an independent set. Expanding the left-hand side in terms of our eigenvectors, we have

$$0 = \langle f, Af \rangle = \sum_{i=1}^{n} \lambda_i a_i^2 \geq \lambda_1 a_1^2 + \lambda_{\min} \sum_{i=2}^{n} a_i^2 = \lambda_1 \alpha^2 + (\alpha - \alpha^2)\lambda_{\min}.$$

Rearranging gives

$$\alpha \leq \frac{\lambda_{\min}}{\lambda_1 - \lambda_{\min}} = \frac{\lambda_{\min}}{d - \lambda_{\min}},$$

as required. If equality holds, then $a_i \neq 0$ implies that $i = 1$ or $\lambda_i = \lambda_{\min}$, completing the proof. □

In fact, the proof above works just the same if we assume that G is an arbitrary (not necessarily regular) n-vertex graph, and A is *pseudo-adjacency matrix* for G, meaning a symmetric matrix with $A_{x,y} = 0$ whenever $xy \notin E(G)$, such that the all-1's vector is an eigenvector of A. Note that a symmetric matrix A has the all-1's vector as an eigenvector if and only if all the row and column sums of A are equal, so we have the following equivalent

Definition. *Let $G = (V, E)$ be an n-vertex graph. A matrix $A \in \mathbb{R}^{V \times V}$ is said to be a* pseudo-adjacency matrix for G *if it is symmetric, has $A_{x,y} = 0$ whenever $xy \notin E(G)$, and has all its row and column sums equal.*

In other words, a pseudo-adjacency matrix comes from choosing a function $w : E(G) \to \mathbb{R}$ which is 'vertex-regular', in the sense that

$$\sum_{y:\ xy \in E(G)} w(xy)$$

is the same for every vertex x.

Note that the adjacency matrix of any regular subgraph of G is a pseudo-adjacency matrix for G, and so is any real linear combination of such.

We have the following extension of Hoffman's theorem, due essentially to Delsarte:

Theorem 2 (Delsarte)**.** *Let $G = (V, E)$ be an n-vertex graph, and let A be a pseudo-adjacency matrix for G. Let $\lambda_{\min}$ denote the least eigenvalue of A, and let λ_{const} be the eigenvalue corresponding to the constant functions. If $S \subset V$ is an independent set in G, then*

$$\frac{|S|}{n} \leq \frac{-\lambda_{\min}}{\lambda_{\text{const}} - \lambda_{\min}}.$$

If equality holds, then the characteristic function f_S of S satisfies:

$$f_S - \tfrac{|S|}{n}\mathbf{1} \in \mathit{Ker}(A - \lambda_{\min} I).$$

This has been a crucial ingredient in the solution of several important problems in extremal combinatorics over the last thirty years. We now give a general description of how it is used.

Of course, when applying Theorem 2 to a specific problem, one needs to construct pseudo-adjacency matrices whose eigevalues can be feasibly calculated, or estimated. When doing this, it is useful to bear in mind that the set of pseudo-adjacency matrices for a graph G forms a linear space. Moreover, it often makes sense to restrict our attention to some subspace W of the space of all pseudo-adjacency matrices for G, such that the matrices in W are *simultaneously diagonalizable*, meaning that they have a simultaneous basis of eigenvectors. (This is true, for example, if W is in the span of the adjacency matrices of an *association scheme* — more on this later.) The problem of finding the matrix in W which gives the best possible bound in Theorem 2 is then a linear programming problem.

We may often wish to restrict our attention further, to pseudo-adjacency matrices whose eigenvalues are especially easy to calculate. Sometimes, the universe U in which we are working can be given a 'nice' group structure, and we can choose W to be the space spanned by the adjacency matrices of some Cayley graphs[1] on U. If U is an Abelian group, the *character group* of U is complete set of eigenfunctions for the adjacency matrix of *any* Cayley graph on U, so we can analyse eigenvalues using Fourier analysis on U (the Fourier transform of a function simply expresses the function as a linear combination of characters of U). Even if U is a non-Abelian group, the eigenspaces of any Cayley graph on U are linear representations of the group U, so the eigenvalues can be analysed using the representation theory of the group U. We will see examples of this later in the course.

We now give an overview of some applications of Theorems 1 and 2 in extremal combinatorics. One of the first was Lovász' eigenvalue proof of the Erdős-Ko-Rado theorem. Observe that an intersecting family of r-element subsets of $\{1, 2, \ldots, n\}$ is precisely an independent set in the *Kneser graph* $K(n, r)$, the graph with vertex-set $[n]^{(r)}$ where we join two r-sets iff they are disjoint. Lovász [5] showed that if $r \leq n/2$, then the least eigenvalue of the adjacency matrix of $K(n, r)$ is

$$-\binom{n-r-1}{r-1};$$

applying Hoffman's bound implies the Erdős-Ko-Rado theorem.

Erdős, Ko and Rado also showed that if n is sufficiently large depending on r and t, then a t-intersecting family of r-element subsets of $\{1, 2, \ldots, n\}$ has size at most $\binom{n-t}{r-t}$, the number of r-sets containing t fixed elements. The question now arises, for which n is this true? If $n < (r-t+1)(t+1)$, then the family

$$\{x \in [n]^{(r)} : \ |x \cap [t+2]| \geq t+1\}$$

is a t-intersecting family of size

$$(t+2)\binom{n-t-2}{r-t-1} + \binom{n-t-2}{r-t-2} > \binom{n-t}{r-t}.$$

The proof of Erdős, Ko and Rado only works for $n \geq 2tr^3$. Using Delsarte's theorem above, Wilson [6] was able to obtain the exact bound: a t-intersecting family in $[n]^{(r)}$ has size at most $\binom{n-t}{r-t}$ if and only if $n \geq (t+1)(r-t+1)$. A t-intersecting family is precisely an independent set in the graph $K(n, r, t)$,

[1] Recall that if Z is a group, and $Y \subset Z$ is inverse-closed ($y \in Y \Rightarrow y^{-1} \in Y$), the *Cayley graph on Z generated by Y* is the graph with vertex-set Z where we join z to zy for all $z \in Z$, $y \in Y$. It is denoted $\mathrm{Cay}(Y, Z)$.

where we join two r-sets if and only if their intersection has size less than t. For $t \geq 2$, applying Hoffman's bound to this graph does not give the desired bound, but Wilson was able to construct a suitable pseudo-adjacency matrix for $K(n, r, t)$ whose eigenvalues do yield the desired bound.

However, in 1997, Ahlswede and Khachatrian [1] used purely combinatorial techniques to characterize the largest t-intersecting families in $[n]^{(r)}$ for *all* values of n, r and t, proving a conjecture of Frankl — namely, if

$$\mathcal{F}_i = \{x \in [n]^{(r)} : \ |x \cap [t + 2i]| \geq t + i\},$$

then for any n, r and t, a maximum-sized t-intersecting family in $[n]^{(r)}$ must be isomorphic to one of the $\mathcal{F}_i$'s. Their proof makes clever use of shifting, 'generating sets', double-counting, and passing between r- and $(n - r)$-uniform families by taking complements. It seems impossible to use eigenvalue techniques in the case $n < (t + 1)(r - t + 1)$. But the reader should not despair of eigenvalue methods just yet; we will soon encounter combinatorial problems where the only known solution relies on eigenvalue methods.

Let us digress for a moment to consider what is perhaps the 'simplest' t-intersection problem, where we remove the condition on the sizes of the sets:

Question. *What is the maximum possible size of a t-intersecting family of subsets of $\{1, 2, \ldots, n\}$?*

When $t = 1$, this is trivial: an intersecting family $\mathcal{A} \subset \mathcal{P}([n])$ cannot contain both x and x^c, so $|\mathcal{A}| \leq 2^{n-1}$. Equality holds if $\mathcal{A}$ is all sets containing 1, and also if $\mathcal{A}$ is all sets of size greater than $n/2$ (if n is odd), or if $\mathcal{A}$ is all sets containing at least 2 elements of $\{1, 2, 3\}$. For $t > 1$, Katona proved the following

Theorem 3 (Katona, 1964)**.** *Let $|X| = n$, and let $\mathcal{A} \subset \mathcal{P}([n])$ be a t-intersecting family. If $n + t$ is even, then*

$$|\mathcal{A}| \leq |[n]^{(\geq(n+t)/2)}| = \sum_{i=(n+t)/2}^{n} \binom{n}{i}.$$

If $n + t$ is odd, then

$$|\mathcal{A}| \leq |[n]^{(\geq(n+t+1)/2)} \cup [n-1]^{((n+t-1)/2)}| = \sum_{i=(n+t+1)/2}^{n} \binom{n}{i} + \binom{n-1}{(n+t-1)/2}.$$

This can be proved by induction on n, combined with ij-compressions. Alternatively, it can be proved using 'UV-compressions', with $|U| > |V|$.

In this problem, the ground set $[n]$ could be replaced by any n-element set; there is no 'structure' at all. What happens if we impose some structure on the ground set? Simonovits and Sós suggested taking the ground set to be the edge-set of the complete graph on n vertices; this leads us to consider *families of graphs.*

Definition. *Let F be a fixed, unlabelled graph. Let $\mathcal{G} \subset \mathcal{P}([n]^{(2)})$; so identifying graphs with their edge-sets, $\mathcal{G}$ is a family of (labelled) graphs on $\{1, 2, \ldots, n\}$. We say that $\mathcal{G}$ is* F-intersecting *if $G \cap H$ contains a copy of F for any $G, H \in \mathcal{G}$.*

This raises the question, what are the largest F-intersecting families of graphs on $\{1, 2, \ldots, n\}$? We write

$$m_n(F) = \max\{|\mathcal{G}| : \mathcal{G} \subset \mathcal{P}([n]^{(2)}),\ \mathcal{G} \text{ is } F\text{-intersecting}\}.$$

If F is a single edge, then we are simply looking at intersecting families of graphs, i.e. intersecting families of subsets of an $\binom{n}{2}$-element set. As above, the maximum possible size is $2^{\binom{n}{2}-1}$: we cannot have both a graph and its complement.

What happens when $F = S_d$, the star with d rays? By fixing a copy of S_d and taking all graphs containing that copy, we can obtain an S_d-intersecting family of graphs on $[n]$ with size $2^{\binom{n}{2}-d}$. But as observed by Simonovits and Sós, we can do much better: we can get almost half of all graphs, when n is large.

To see this, take $\mathcal{G}$ to be the family of all graphs on $[n]$ whose degree at the vertex 1 is at least $(n-1+d)/d$:

$$\mathcal{G} = \{G \in \mathcal{P}([n]^{(2)}) :\ \deg_G(1) \geq (n-1+d)/2\}.$$

Any two graphs in $\mathcal{G}$ share a d-ray star with centre 1. It is easy to see that

$$|\mathcal{G}| \geq (1-o(1))2^{\binom{n}{2}-1}.$$

Indeed, $|\mathcal{G}|/2^{\binom{n}{2}}$ is equal to the probability that a uniform random graph on $[n]$ has degree at least $(n-1+d)/2$ at the vertex 1. Note that a uniform random graph on $[n]$ is precisely the Erdős-Renyi random graph $G_{n,1/2}$, where each edge of $K_{[n]}$ is included independently with probability 1/2. If $G \sim G_{n,1/2}$, then $\deg_G(1) \sim \text{Bin}(n-1, 1/2)$. Since $\text{Bin}(n-1, 1/2)$ has standard deviation $\sqrt{n-1}/2$, we have

$$\text{Prob}\{\deg_G(1) \geq (n-1+d)/2\} = 1/2 - o(1).$$

Simonovits and Sós asked what happens when F is a triangle. They made the following

Conjecture 1 (Simonovits-Sós, 1976)**.** *Let $\mathcal{G}$ be a triangle-intersecting family of graphs on $\{1, 2, \ldots, n\}$. Then $|\mathcal{G}| \leq 2^{\binom{n}{2}-3}$. Equality holds if and only if $\mathcal{G}$ consists of all graphs containing a fixed triangle.*

Due to the structure on the ground set, purely combinatorial techniques do not seem to work for this problem. For example, there does not seem to be any suitable 'shifting' operation which preserves the property of being triangle-intersecting. Last year, the Simonovits-Sós conjecture was proved by Filmus, Friedgut and the author [2], using eigenvalue techniques (Delsarte's bound).

Note that a triangle-intersecting family of graphs on $\{1, 2, \ldots, n\}$ is precisely an independent set in the graph Γ with vertex-set $\mathcal{P}([n]^{(2)})$, where we join G and H if and only if $G \cap H$ is triangle-free. We apply Delsarte's bound to a special class of pseudo-adjacency matrices for Γ, and use Fourier analysis on $\mathcal{P}([n]^{(2)})$ to analyse their eigenvalues. This will be described in upcoming lectures at the Newton Institute.

Another problem in which eigenvalue techniques have been crucial concerns *t-intersecting families of permutations.* We now work in S_n, the group of all permutations of $\{1, 2, \ldots, n\}$ under composition. (We can of course view S_n as the set of all orderings of the numbers $1, 2, \ldots, n$.)

Definition. *We say that a family of permutations $\mathcal{A} \subset S_n$ is* intersecting *if any two permutations in the family agree somewhere — for any $\sigma, \tau \in \mathcal{A}$, there exists some $i \in [n]$ such that $\sigma(i) = \tau(i)$.*

Our starting point is the Erdős-Ko-Rado question for permutations:

Question. *What is the maximum possible size of an intersecting family of permutations in S_n?*

This was answered by Deza and Frankl:

Theorem 4 (Deza, Frankl, 1977)**.** *Let $\mathcal{A} \subset S_n$ be an intersecting family; then $|\mathcal{A}| \leq (n-1)!$.*

Their proof is a simple partitioning argument, making crucial use of the group operation (composition):

Proof. Let $\mathcal{A} \subset S_n$ be an intersecting family. Let ρ be the n-cycle $(1\ 2\ \ldots n)$ (written in disjoint-cycle notation). Let H be the cyclic group of order n generated by ρ in S_n, i.e.

$$H = \{\mathrm{Id},\ (1\ 2\ \ldots n),\ (1\ 2\ \ldots n)^2, \ldots,\ (1\ 2\ \ldots n)^{n-1}\}.$$

Observe that if we write the permutations in H in sequence-notation,

$$\sigma(1), \sigma(2), \ldots, \sigma(n),$$

then they form the rows of an $n \times n$ latin square:

$$\begin{array}{ccccc} 1 & 2 & \ldots & n-1 & n \\ 2 & 3 & \ldots & n & 1 \\ & & \vdots & & \\ n & 1 & \ldots & n-2 & n-1 \end{array}$$

Hence, any two permutations in H disagree everywhere, and therefore any intersecting family $\mathcal{A} \subset S_n$ contains at most one permutation from H. The same is clearly true if H is replaced by any left coset σH of H: the permutations in σH also form the rows of an $n \times n$ latin square. Hence, $\mathcal{A}$ contains at most one permutation from each left coset of H. The left cosets of H partition S_n, by Lagrange's theorem; there are $(n-1)!$ of them, and therefore

$$|\mathcal{A}| \leq (n-1)!,$$

as required. □

Remark 1. *In the above proof, we could replace H by any set of n permutations which form the rows of an $n \times n$ latin square (when written in sequence-notation). The left-translates σH no longer necessarily partition S_n, but they cover S_n uniformly, so averaging gives $|\mathcal{A}| \leq (n-1)!$, as before.*

Deza and Frankl conjectured that equality holds only if $\mathcal{A}$ consists of all permutations mapping i to j, for some fixed i and j. This was eventually proved by Cameron and Ku in 2003, using the fact that an intersecting family of size $(n-1)!$ contains exactly one row of every latin square of order n, together with a highly non-trivial 'shifting' argument.

Perhaps the simplest proof of the Cameron-Ku theorem comes from applying Hoffman's theorem. Observe that an intersecting family in S_n is precisely an independent set in the graph on S_n where we join two permutations iff they disagree everywhere; this is known as the *derangement graph* on S_n.

Calculating the least eigenvalue of the derangement graph is non-trivial, but can be done using the representation theory of S_n. As conjectured by Ku, and first proved by Renteln [4],

$$\lambda_{\min} = -d_n/(n-1).$$

Plugging this into Hoffman's bound shows that an intersecting family $\mathcal{A} \subset S_n$ has size at most $(n-1)!$. Moreover, if equality holds, then the characteristic vector of $\mathcal{A}$ lies in the subspace spanned by the constant vectors and the $\lambda_{\min}$-eigenvectors of the derangement graph. This can be used to show that

$$\mathcal{A} = \{\sigma \in S_n : \ \sigma(i) = j\}$$

for some $i, j \in [n]$. We will sketch the argument at the end of the course.

What about t-intersecting families of permutations in S_n? (We say that a family of permutations is t-intersecting if any two permutations in it agree on at least t points.) Deza and Frankl made the following

Conjecture 2 (Deza, Frankl, 1977)**.** *If n is sufficiently large depending on t, and $\mathcal{A} \subset S_n$ is a t-intersecting family, then $|\mathcal{A}| \leq (n-t)!$. If equality holds, then*

$$\mathcal{A} = \{\sigma \in S_n : \ \sigma(i_1) = j_1, \ \sigma(i_2) = j_2, \ldots, \sigma(i_t) = j_t\},$$

for some distinct $i_1, \ldots, i_t \in [n]$ and some distinct $j_1, \ldots, j_t \in [n]$ — in other words, $\mathcal{A}$ is a coset of the stabilizer of t points.

When there exists a sharply t-transitive subgroup of S_n (a subgroup of size $n(n-t)\ldots(n-t+1)$ in which any two distinct permutations agree on less than t points), one can prove this using the partitioning argument as in the proof of Theorem 4. When there exists a sharply t-transitive subset of S_n, one can use the averaging argument in Remark 1. However, sharply t-transitive subgroups of S_n exist only for a small number of values of n and t, and sharply t-transitive subsets of S_n are not known to exist for any additional values.

Again, no purely combinatorial proof of the full conjecture is known. It was proved by the author, and independently by Friedgut and Pilpel, in 2008, using a very similar method — Delsarte's bound, combined with the representation theory of the symmetric group. (We have written a joint paper, [3].)

Note that a t-intersecting family is precisely an independent set in the 't-derangement graph', the graph on S_n where we join two permutations if and only if they agree on less than t points. Calculating the eigenvalues of the adjacency matrix of this graph and applying Hoffman's theorem only shows that a t-intersecting family has size at most $O((n-1)!)$, which we know anyway from the Deza-Frankl theorem on 1-intersecting families. Instead, we construct a pseudo-adjacency matrix for the t-derangement graph which has the appropriate eigenvalues when n is sufficiently large. Again, the main work of the proof is in the construction of this matrix, whose eigenvalues are analyzed using the representation theory of S_n.

It would be interesting to determine the largest t-intersecting families in S_n for *all* values of n and t. The author made the following conjecture in 2008:

Conjecture 3. *For any $n \geq t$, a t-intersecting family $\mathcal{A} \subset S_n$ of maximum size must be a 'double translate' of one of the families*

$$\mathcal{F}_i = \{\sigma \in S_n : \ \sigma \text{ has at least } t+i \text{ fixed points in } [t+2i]\} \ (0 \leq i \leq (n-t)/2),$$

meaning that $\mathcal{A} = \pi \mathcal{F}_i \tau$ for some $\pi, \tau \in S_n$ and some i.

This remains open to date. As with the analogous question for t-intersecting families in $[n]^{(r)}$, it seems that eigenvalue techniques cannot be used to prove it.

Isoperimetric inequalities, expansion and eigenvalues

Isoperimetric problems are of ancient interest in mathematics. In general, they ask for the smallest possible 'boundary' of a set of a certain 'size'. For example, of all shapes in the plane with area 1, which has the smallest perimeter? The ancient Greeks 'knew' that the answer was a circle, but it was not until the 19th century that this was proved rigorously.

In extremal combinatorics, we are interested in *discrete* isoperimetric problems, which deal with discrete notions of boundary in finite spaces. Often, we have a graph $G = (V, E)$, and we are interested in the 'boundary' of subsets of V. For graphs, there are two different notions of 'boundary': the edge-boundary and the vertex-boundary.

Definition. *Let $G = (V, E)$ be a graph, and let $S \subset V$. The* edge-boundary *of S is the set of all edges of G between S and $V \setminus S = S^c$; it is denoted $\partial_G S$, or $E_G(S, S^c)$.*

Definition. *Let $G = (V, E)$ be a graph, and let $S \subset V$. The* vertex-boundary *$b_G(S)$ of S is the set of all vertices of S^c which are adjacent to a vertex in S, i.e.*

$$b_G(S) = \{x \in S^c : \ xy \in E(G) \text{ for some } y \in S\}.$$

The *edge-isoperimetric problem for G* asks for the determination of the minimum possible size of the edge-boundary of a set of k vertices in G, for each $k \in \mathbb{N}$. (We denote this minimum by $\Phi_G(k)$.)

Similarly, the *vertex-isoperimetric problem for G* asks for the determination of the minimum possible size of the vertex-boundary of a set of k vertices in G, for each $k \in \mathbb{N}$. (We denote this minimum by $\Psi_G(k)$.)

Isoperimetric problems for specific graphs are of great interest in combinatorics. If the graph G has a sufficiently nice structure, there may be an ordering of the vertex-set such that the boundary of a set of size k is minimized by the k-set consisting of the first k elements of this ordering, for any k. One can sometimes prove this using 'compressions': one defines a collection of compression operations that preserve the size of a set, cannot increase the size of the boundary of a set, and make sets 'more like' an initial segment of the ordering. Then take any set $S \subset V$, apply a suitable sequence of compression operations to S, and show that one ends up with the first $|S|$ elements of the ordering, or a set whose boundary is no smaller. It follows that S has boundary at least as large as that of the first $|S|$ elements of the ordering.

The classic examples of this phenomenon are the edge- and vertex- isoperimetric inequalities in the n-dimensional hypercube, Q_n (the graph with vertex-set $\mathcal{P}([n])$, where we join two sets if their symmetric difference has size 1). Let $B_{n,k}$ denote the first k elements of the *binary ordering* on $\mathcal{P}([n])$, defined by

$$x < y \quad \Leftrightarrow \quad \max(x\Delta y) \in y.$$

The edge-isoperimetric inequality of Harper, Bernstein, Lindsay and Hart states that among all subsets of $\mathcal{P}([n])$ of size k, $B_{n,k}$ has the smallest possible edge-boundary. This can be proved using suitable compressions, which are shown to have the desired property by induction on n. (Note that Hart proved it differently, using induction on n, combined with an inequality satisfied by the extremal function $\partial_{Q_n}(B_{n,k})$, which he proved by induction on k.)

Similarly, let $S_{n,k}$ denote the first k elements of the *simplicial ordering* on $\mathcal{P}([n])$, defined by

$$\begin{aligned} x < y \quad \Leftrightarrow \quad & |x| < |y|, \text{ or} \\ & |x| = |y| \text{ and } \min(x\Delta y) \in x. \end{aligned}$$

Harper's vertex-isoperimetric inequality states that among all subsets of $\mathcal{P}([n])$ of size k, $S_{n,k}$ has the smallest possible edge-boundary. Again, this can be proved using suitable compressions, which are shown to have the desired property by induction on n.

When the graph G we are interested in has a more complicated structure, there may not exist an ordering of the vertex-set with the property above, and purely combinatorial techniques may not give us any useful information. However, it turns out that there is a close link between the edge-isoperimetric problem for a graph and the eigenvalues of the *Laplacian* of the graph.

Let $G = (V, E)$ be an n-vertex graph. The *Laplacian operator* L_G is the linear operator defined by

$$(L_G f)(x) = \sum_{\substack{y \in V: \\ xy \in E(G)}} (f(x) - f(y)).$$

In other words, it takes a function $f : V \to \mathbb{R}$, and for each vertex, it returns the sum of the differences of the function along each edge incident to that vertex. If $d(x)$ denotes the degree of the vertex x, then we have

$$(L_G f)(x) = d(x)f(x) - \sum_{\substack{y \in V: \\ xy \in E(G)}} f(y) = d(x)f(x) - (Af)(x),$$

where A is the adjacency matrix of G. So we define the *Laplacian matrix* of G to be

$$L = D - A,$$

where $D \in \mathbb{R}^{V \times V}$ is the diagonal matrix with the degrees all along the diagonal, i.e.

$$D_{x,x} = d(x) \ \forall x \in V, \quad D_{x,y} = 0 \ \forall x \neq y \in V.$$

Note that L is a real symmetric matrix, so all its eigenvalues are real, and for any inner product on $\mathbb{R}^V$, there exists an orthonormal basis of eigenvectors of L. Clearly, the all-1's vector is always an eigenvector of L with eigenvalue

0. Moreover, L is positive semidefinite, meaning that all its eigenvalues are non-negative. To see this, observe that for any $f \in \mathbb{R}^V$, we have

$$\begin{aligned} f^\top L f &= \sum_{x \in V} f(x)(Lf)(x) \\ &= \sum_{x \in V} d(x) f(x)^2 - \sum_{\substack{x,y \in V: \\ xy \in E(G)}} f(x)f(y) \\ &= \sum_{xy \in E(G)} (f(x) - f(y))^2 \\ &\geq 0. \end{aligned}$$

Let $0 = \mu_1 \leq \mu_2 \leq \ldots \leq \mu_n$ be the eigenvalues of L, repeated with their multiplicities. It is easy to see that $\mu_2 > 0$ if and only if G is connected.

We call $\mu_2 = \mu_2 - \mu_1$ the *spectral gap*. The following result is very easy to prove, but is of great practical importance: it says that graphs with a large spectral gap have large 'edge expansion'.

Theorem 5. *Let $G = (V, E)$ be an n-vertex graph, let L be the Laplacian matrix of G, and let $0 = \mu_1 \leq \mu_2 \leq \ldots \leq \mu_n$ be the eigenvalues of L, repeated with their multiplicities. Then for any subset $S \subset V(G)$, we have*

$$e(S, S^c) \geq \mu_2 |S||S^c|/n.$$

If equality holds, then the characteristic function of S lies in the subspace of $\mathbb{R}^V$ spanned by the constant vectors and the μ_2-eigenvectors of L.

Proof. As before, we equip $\mathbb{R}^V$ with the inner product induced by the uniform measure on V,

$$\langle f, g \rangle = \frac{1}{n} \sum_{x \in V} f(x) g(x),$$

and we let $||.||_2$ denote the induced Euclidean norm:

$$||f||_2 = \sqrt{\langle f, f \rangle}.$$

Let $\{\mathbf{1} = v_1, v_2, \ldots, v_n\}$ be an orthonormal basis of real eigenvectors of A, with corresponding eigenvalues $0 = \mu_1 \leq \mu_2 \leq \ldots \leq \mu_n$.

Let $S \subset V$, let $\alpha = |S|/n$ denote the measure of S, and let $g = f_S - \alpha$ denote the shifted characteristic function of S. Then

$$\langle g, Lg \rangle = \sum_{xy \in E(G)} (g(x) - g(y))^2 = e(S, S^c).$$

But since g is orthogonal to the all 1's vector, we have $\langle g, Lg \rangle \geq \mu_2 ||g||_2^2$. To see this, simply expand g in terms of our eigenvectors. Write

$$g = \sum_{i=1}^n a_i v_i;$$

then we have

$$\langle g, Lg \rangle = \sum_{i=1}^n \mu_i a_i^2.$$

But $a_1 = \langle g, v_1 \rangle = \langle g, \mathbf{1} \rangle = 0$, and therefore

$$\langle g, Lg \rangle \geq \mu_2 \sum_{i=1}^{n} a_i^2 = \mu_2 ||g||_2^2.$$

Note that $||g||_2^2 = |S|(1-\alpha)^2 + (n-|S|)\alpha^2 = |S|(n-|S|)/n$, so

$$e(S, S^c) \geq \mu_2 |S|(n-|S|)/n,$$

as desired. If equality holds, then $a_i = 0$ whenever $\mu_i > \mu_2$. It follows that g (and therefore f_S) is in the direct sum of the μ_1 and μ_2-eigenspaces of L. □

Remark 2. *This is sharp for the n-dimensional hypercube Q_n, when S is a codimension-1 subcube. In this case, we have $\mu_2 = 2$, and therefore*

$$e(S, S^c) \geq 2^{-(n-1)}|S||S^c|.$$

Comparing this with the edge-isoperimetric inequality in Q_n, we see that it is sharp iff $|S| = 2^{n-1}$. (If $|S| = 2^{n-1}$, then $e(S, S^c) \geq 2^{n-1}$, with equality iff S is a codimension-1 subcube.)

Let $G = (V, E)$ be an n-vertex graph. We define the *edge-expansion ratio* of G to be

$$h(G) = \min \left\{ \frac{e(S, S^c)}{|S|} \ : \ S \subset V(G),\ 0 < |S| \leq n/2 \right\};$$

i.e. $h(G)$ is the minimum possible average out-degree of a set of at most $n/2$ vertices.

For a given n-vertex graph G, it is NP-hard to compute $h(G)$. But Theorem 5 implies that $h(G) \geq \mu_2/2$, so we have a lower bound for $h(G)$ in terms of the second eigenvalue of the Laplacian. Note that

$$\mu_2 = \max \left\{ \frac{\sum_{ij \in E(G)} (x_i - x_j)^2}{\sum_{i=1}^{n} x_i^2} \ : \ x \neq 0,\ \sum_{i=1}^{n} x_i = 0 \right\},$$

so μ_2 can be computed in polynomial time (to within an arbitrary degree of accuracy), as indeed can all the eigenvalues of any $n \times n$ matrix.

It turns out that for graphs of maximum degree at most d, we have $\mu_2 \geq h^2/(2d)$. So for bounded degree graphs, *having a large spectral gap is polynomially equivalent to having a large edge-expansion ratio*, and therefore if a bounded degree graph has large edge-expansion ratio, this can be verified efficiently by computing μ_2.

Theorem 6 (Alon)**.** *Let $G = (V, E)$ be an n-vertex graph of maximum degree at most d. Then*

$$\mu_2 \geq \frac{h(G)^2}{2d}.$$

Proof. Let $G = (V, E)$ be an n-vertex graph of maximum degree at most d, and write $h(G) = h$. Then the number of edges of any cut (S, S^c) in G satisfies

$$e(S, S^c) \geq h \min\{|S|, |S^c|\}.$$

Let w be the eigenvector of L associated with μ_2, as above. By definition,

$$\mu_2 = \langle w, Lw \rangle = \frac{w^\top L w}{w^\top w}.$$

We will show that $\mu_2 \geq h(G)^2/(2d)$ by applying the above condition to a collection of cuts (S, S^c) which are level sets of w.

Since $w \perp \mathbf{1}$, some of the entries of w are positive, and some are negative. Let

$$V^+ = \{i \in [n] : \; w_i > 0\}, \quad V^- = \{i \in [n] : \; w_i \leq 0\}.$$

By considering $-w$ if necessary, we may assume without loss of generality that $|V^+| \leq |V^-|$, so $|V^+| \leq n/2$. First, we define

$$u_i = \begin{cases} w_i, & \text{if } i \in V^+; \\ 0, & \text{if } i \in V^- \end{cases},$$

i.e. u is the positive part of w. We first observe that

$$\frac{u^\top L u}{u^\top u} \leq \mu_2.$$

Indeed, for every $i \in V^+$, we have

$$\begin{aligned} (Lu)_i &= d(i)u_i - \sum_{j \in V: \; ij \in E(G)} u_j \\ &= d(i)w_i - \sum_{j \in V^+: \; ij \in E(G)} w_j \\ &\leq d(i)w_i - \sum_{j \in V: \; ij \in E(G)} w_j \\ &= (Lw)_i \end{aligned}$$

and therefore

$$\begin{aligned} u^\top L u &= \sum_{i \in V^+} u_i (Lu)_i \\ &= \sum_{i \in V^+} w_i (Lu)_i \\ &\leq \sum_{i \in V^+} w_i (Lw)_i \\ &= \sum_{i \in V^+} w_i (\mu_2 w_i) \\ &= \mu_2 \sum_{i \in V} u_i^2 \\ &= \mu_2 \, u^\top u, \end{aligned}$$

as required.

Let $m = |V^+|$. Relabelleing the vertices of G if necessary, we may assume that $V^+ = [m]$ and $u_1 \geq u_2 \geq \ldots \geq u_m > u_{m+1} = 0$. We now obtain a lower bound on the quantity

$$B_u = \sum_{ij \in E(G): \; i<j} (u_i^2 - u_j^2),$$

by applying the edge-expansion property of G to a collection of cuts (S, S^c) where $S = [l]$, $l \leq m$, i.e. S ranges over level-sets of u. (We will then bound B_u from above in terms of $u^\top L u$ and $u^\top u$, using the Cauchy-Schwarz inequality.)

We have

$$\begin{aligned} B_u &= \sum_{\substack{i<j:\\ ij \in E(G)}} (u_i^2 - u_j^2) \\ &= \sum_{\substack{i<j:\\ ij \in E(G)}} \sum_{l=i}^{j-1} (u_l^2 - u_{l+1}^2) \\ &= \sum_{l=1}^{n-1} (u_l^2 - u_{l+1}^2) e([l],[l]^c) \\ &= \sum_{l=1}^{m} (u_l^2 - u_{l+1}^2) e([l],[l]^c) \\ &\geq h \sum_{l=1}^{m} l(u_l^2 - u_{l+1}^2) \\ &= h \sum_{i=1}^{m} u_i^2 \\ &= h\ u^\top u. \end{aligned}$$

In the other direction, using the Cauchy-Schwarz Inequality, we have

$$\begin{aligned} B_u &= \sum_{ij \in E(G):\ i<j} (u_i - u_j)(u_i + u_j) \\ &\leq \sqrt{\sum_{ij \in E(G)} (u_i - u_j)^2} \sqrt{\sum_{ij \in E(G)} (u_i + u_j)^2} \\ &= \sqrt{u^\top L u} \sqrt{\sum_{ij \in E(G)} (u_i + u_j)^2} \\ &\leq \sqrt{u^\top L u} \sqrt{2 \sum_{ij \in E(G)} (u_i^2 + u_j^2)} \\ &= \sqrt{u^\top L u} \sqrt{2 \sum_{i \in V} d(i) u_i^2} \\ &\leq \sqrt{u^\top L u} \sqrt{2d \sum_{i \in V} u_i^2} \\ &= \sqrt{u^\top L u} \sqrt{2d\ u^\top u}. \end{aligned}$$

Combining these inequalities, we have

$$\mu_2 \geq \frac{u^\top L u}{u^\top u} \geq \frac{h^2}{2d},$$

as required. □

Remark 3. *Considering the cycle C_{2n} shows that Theorem 6 is tight up to a constant factor; this has $d = 2$, $h = 2/n$, and $\mu_2 = (1 + O(1/n^2))\pi^2/n^2$.*

Expander graphs

We have seen the close link between edge-expansion and the spectral gap. We are often interested in graphs with good vertex-expansion. If G is an n-vertex graph, we define the *vertex-expansion ratio* of G to be

$$\beta(G) = \min\left\{\frac{|b(S)|}{|S|} : \ S \subset V(G),\ 0 < |S| \leq n/2\right\}.$$

Clearly, for any graph G, we have

$$|\partial S| \geq |b(S)|$$

for all subsets $S \subset V(G)$, and therefore

$$h(G) \geq \beta(G).$$

If G has maximum degree d, then for any set S,

$$|b(S)| \geq |\partial S|/d,$$

since any vertex in $V \setminus S$ meets at most d edges in ∂S. Hence,

$$c(G) \geq h(G)/d.$$

So for bounded-degree graphs, having large edge-expansion ratio is linearly equivalent to having large vertex-expansion ratio, and these are both polynomially equivalent to having large spectral gap.

Bounded-degree graphs with large vertex-expansion are of great importance in theoretical computer science. For example,

Exercise. *If G is an n-vertex graph with vertex-expansion ratio β, then*

$$\operatorname{diam}(G) \leq 2\lceil(\log n)/\log(1+\beta)\rceil.$$

On the other hand, it is easy to show that the diameter of any bounded-degree graph is at least logarithmic in the number of vertices:

Exercise. *If G is an n-vertex graph with maximum degree at most d, then*

$$\operatorname{diam}(G) \geq \log_{d-1}(n-1) + \log_{d-1}(1 - 2/d).$$

We now come to the crucial question:

Question. *Do there exist $d \in \mathbb{N}$ and $c > 0$ such that there exist n-vertex graphs with maximum degree at most d and edge-expansion ratio at least c, for infinitely many n?*

We call such graphs (n, d, c)-*edge-expanders*. They turn out to be of great importance in theoretical computer science. For example, as mentioned in Lecture 1, they can be used to construct sparse *superconcentrators*. They were first considered by Valiant, who at first believed that the answer to the above question was no. However, as we will see next lecture, if $d \geq 3$, *almost all* d-regular graphs on n vertices (for nd even) have edge-expansion ratio at least $d/16$. Formally, if $G(n, d)$ denotes a *uniform random d-regular graph on* $[n]$, meaning a d-regular graph on $[n]$ chosen uniformly at random from the set of all d-regular graphs on $[n]$, then

$$\text{Prob}\{h(G_{n,d}) \geq d/16\} \to 1 \quad \text{as } n \to \infty.$$

References

[1] Ahlswede, R., Khachatrian, L. H., The complete intersection theorem for systems of finite sets, *European Journal of Combinatorics* Volume 18 (1997), 125-136.

[2] Ellis, D., Filmus, Y., Friedgut, E., Triangle-intersecting families of graphs, submitted. Available at `http://arxiv.org/abs/1010.4909`.

[3] Ellis, D., Friedgut, E., Pilpel, H., Intersecting families of permutations, *J. Amer. Math. Soc.*, posted on 31st January 2011, PII: S 0894-0347 (2011) 00690-5 (to appear in print).

[4] Renteln, P., On the spectrum of the derangement graph, *Electronic Journal of Combinatorics* 14 (2007) #R82.

[5] Lovász, L., On the Shannon capacity of a graph, *IEEE Tram. Inform. Theory*, Volume IT-25, pp. 1-7.

[6] Wilson, R.M., The exact bound in the Erdős-Ko-Rado Theorem, *Combinatorica* Volume 4 (1984) 247-257.