Doctoral mini-course on expander graphs
November 28, 2024 – Lecture 1

1 Preliminaries

Throughout the four lectures, d will be a fixed integer, and all graphs we consider will be d-regular (i.e., all vertices have degree d). Most of the theory can be developed for non-regular graphs but the regular case is much simpler, and this case is sufficient for most of the applications we will present.

All the graphs we consider will be finite. Intuitively, expander graphs are d-regular graphs of expansion at least α, for some $\alpha > 0$ and some notion of expansion which will made more precise below. Most of the times we will be interested in obtaining infinite families of d-regular graphs with expansion at least α (not just a single such graph).

The vertex set and edge set of a graph G are denoted by $V(G)$ and $E(G)$. We use in general n for the number $|V(G)|$ of vertices of G.

2 Definitions

We will consider two notions of expansion, one combinatorial and one spectral. These two notions will turn out to be equivalent for the purpose of defining expander graphs.

2.1 Cheeger constant

For two subsets S and T of vertices of a graph G, we denote by $e(S,T)$ the number of edges $uv \in E(G)$ with $u \in S$ and $v \in T$ (edges with both endpoints in $S \cap T$ are counted twice).

For a graph G on n vertices with define the *Cheeger constant* $h(G)$ as the minimum of $e(S,\bar{S})/|S|$, taken over all non-empty subsets S of vertices with $|S| \leq n/2$. Here $\bar{S} = V(G) - S$ denotes the complement of S.

So in a graph G, for any bipartition S, T of $V(G)$, the number of edges between S and T is at least $h(G)$ times the size of the smaller of the two sets.

Assume that G is a d-regular graph on n vertices with $h(G) \geq \alpha$, for some constants d and $\alpha > 0$. Observe that the diameter of G (the maximum distance between two vertices of G) is at most $O(\log n)$, where the implicit constant in the $O(\cdot)$ depends on d and α. This is because for any vertex v, the size of the ball $B(v,r)$ of radius r centered in v increases exponentially before reaching (at least) $n/2$. More precisely: if $S = B(v,r)$ and $|S| \leq n/2$, the number of edges between S and the set T of vertices at distance exactly $r+1$ from v is at least $\alpha|S|$. Since G has maximum degree d, it follows that $|T| \geq \frac{\alpha}{d}|S|$, and thus $|B(v,r+1)| \geq (1+\alpha/d)|B(v,r)|$. This implies that for $r = \Omega(\log n)$, $|B(v,r)| > n/2$. In particular, any two vertices u and v lie at distance at most $2r = O(\log n)$ apart.

2.2 Spectral expansion

Given a graph G with vertex set $v_1, \ldots, v_n$, the adjacency matrix $A_G = (a_{i,j})$ of G is an $n \times n$ $\{0,1\}$-matrix where $a_{i,j} = 1$ if and only if v_i and v_j are adjacent in G. Observe that A_G is a real symmetric matrix and therefore its n eigenvalues are real numbers. Let the eigenvalues be denoted by $\lambda_1 \geq \cdots \geq \lambda_n$.

When G is d-regular it can be checked that for each $1 \leq i \leq n$, $|\lambda_i| \leq d$. A quick way to see this is to view $x \mapsto A_G x$ (where $x \in \mathbb{R}^n$) as an operator that takes some real weights on the vertices of G (the entries of the vector x), and assigns to each vertex v the sum of the weighs of the neighbors of v. So if we take some eigenvector $x = (x_i)$ with eigenvalue λ and we look at the largest entry of the vector x, say on vertex v_i, then v_i is assigned the sum of the weights of its neighbors, which is at most $d\,|x_i|$ in absolute value. But this is also equal to $\lambda|x_i|$ by definition, and thus $|\lambda| \leq d$.

Note that $\lambda_1 = d$ (take for instance the all one vector), and a quick modification of the proof above shows that when G is connected, $\lambda_2 < d$ (so d has multiplicity 1, it is in essence the Perron-Frobenius theorem).

We call the value $d - \lambda_2$ (the difference between the first two eigenvalues), the *spectral gap* of G, and this will be our spectral notion of expansion. It is related to the Cheeger constant by the following theorem.

Theorem 1. *For any d-regular graph G with second largest eigenvalue λ_2,*

$$\tfrac{1}{2}(d - \lambda_2) \leq h(G) \leq \sqrt{2d(d-\lambda_2)}.$$

It follows that $h(G)$ is bounded away from 0 if and only if $d - \lambda_2$ is bounded away from 0. So we can use either notion to define expansion in graphs.

We conclude this section on spectral expansion by mentioning that in a number of cases we will need a bit more that just $\lambda_2 < d$. We will instead be interested in $\lambda = \max(|\lambda_2|, |\lambda_n|)$ the second largest eigenvalue in absolute value, and we will require that $\lambda < d$. A reason for this is that we will sometimes want random walks to converge to the uniform distribution and this typically does not happen in bipartite graphs, graphs whose vertex set can be partitioned in two independent sets. We will come back to random walks in the next lecture. So what's the problem with bipartite graphs (from the spectral point of view)? Well the issue is that $\lambda_n = -d$ (as can be seen by assigning weight 1 to one part of the bipartition and -1 to the other part), and thus $\lambda = d$. This is not a huge deal and this is mostly a technical issue (there are ways to adapt some of the results we will see to the bipartite case).

We want to find graphs such that $d-\lambda$ (or just $d-\lambda_2$) is as large as possible, but how small can λ (and λ_2) be compared to d?

Theorem 2. *If G is d-regular, $\lambda \geq (1 - o(1))\sqrt{d}$.*

Proof. The trace of A_G^2 is equal to $\sum_{i=1}^n \lambda_i^2$. Observe that the entry (i, j) in A_G^2 counts the number of walks of length 2 between v_i and v_j, so in particular $a_{i,i} = d$ for any $1 \leq i \leq n$, and thus $dn = \sum_{i=1}^n \lambda_i^2 \leq d^2 + (n-1)\lambda^2$. It follows that $\lambda^2 \geq \frac{1}{n-1}(dn - d^2)$ and thus $\lambda \geq (1 - o(1))\sqrt{d}$. □

Note that a better bound $\lambda \geq \lambda_2 \geq 2\sqrt{d-1} - o(1)$ can be obtained (where the $o(1)$ again means some quantity that tends to 0 as $n \to \infty$), this was proved by Alon and Boppana. Graphs such that $\lambda \leq 2\sqrt{d-1}$ therefore have the best possible spectral expansion. These graphs are called *Ramanujan graphs*, and there are various highly non-trivial constructions (however constructing Ramanujan graphs for every degree d is a major open problem). We note that $2\sqrt{d-1}$ is precisely the spectral radius of the infinite d-regular tree (we are not defining this notion here, but the purpose is just to say that these objects are perfect expanders in some sense)

3 Examples

From now on, by a *family of expander graphs*, we mean an infinite family $\mathcal{F}$ such that for some integer d and some real $\alpha > 0$, all graphs from $\mathcal{F}$ are

d-regular with $h(G) \geq \alpha$. By Theorem 1, it can be equivalently defined as an infinite family $\mathcal{F}$ such that for some integer d and some real $\beta > 0$, all graphs from $\mathcal{F}$ are d-regular with spectral gap at least β.

Here are a few examples of families of expander graphs (beware that some of the graphs below might have multiple edges or sometimes even loops, which is something we would like to avoid in general).

- sequences of random d-regular graphs with $n \to \infty$ (a classical result of Friedman shows that for these graphs $\lambda \leq 2\sqrt{d-1} + o(1)$ with high probability as $n \to \infty$);

- sequences of random d-regular bipartite graphs with $n \to \infty$ (somewhat equivalently, take two sets U and V of size $n/2$ and add d random perfect matchings between U and V);

- (Bilu-Linial) Take any good d-regular expander, and take a random 2-lift of it (i.e. double every vertex, and replace every edge by a random perfect matching between the two corresponding pairs of vertices). With non-zero probability all new eigenvalues are at most $O(\sqrt{d \log^3 d})$ in absolute value, so the resulting graph is still a good expander (a major conjecture is that this can the bound can be replaced by $2\sqrt{d-1}$, that is the construction produces a Ramanujan graph with non-zero probability, if it starts with a Ramanujan graph).

- for $p \to \infty$, p prime, look at the graph with vertex set $\mathbb{Z}_p$, in which each x is adjacent to $x+1$, $x-1$, and x^{-1}.

- for $p \to \infty$, p prime, look at the graph with vertex set the 2×2 matrices with coefficients in $\mathbb{Z}_p$, with determinant 1. For any such matrix A, and any C in the set of matrices

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$$

  and their inverses, add an edge between A and AC.

The two final items have a very useful property that we will usually require in most applications: these expander graphs are *strongly explicit*, in the sense that there is an algorithm that given a vertex v in input, produces the list of the neighbors of v in time polynomial in the description of v. Typically if G has n vertices we can describe its vertices with $O(\log n)$-bit strings and the goal is to describe the neighbors of any vertex v in time polylogarithmic in n (instead of just polynomial in n). The main reason is that in most applications, we will consider huge expander graphs (of order 2^n, say), and we will then want all computations to take time polynomial in n.

A d-regular graph with n vertices and second largest eigenvalue λ in absolute value is said to be an (n, d, λ)-*graph*. The main result that we will use is the following.

Theorem 3. *For any n, d there exist strongly explicit constructions of (n, d, λ)-graph with $\lambda = O(\sqrt{d})$.*

4 The Expander Mixing Lemma

We will need the following simple yet extremely useful result, called the *Expander Mixing Lemma*.

Lemma 4. *Let G be an (n, d, λ)-graph and let S, T be two subsets of vertices of G. Then*

$$\left|e(S,T) - \tfrac{d}{n}|S||T|\right| \leq \lambda\sqrt{|S||T|}.$$

Dividing both sides by dn we obtain $|e(S,T)/dn - |S||T|/n^2| \leq \lambda/d$. Note that

- $e(S,T)/dn$ is the probability that if we take a random (directed) edge (u, v) in G, we have $u \in S$ and $v \in T$.
- $|S||T|/n^2$ is the probability that if we take a random pair (u, v) of vertices of G, we have $u \in S$ and $v \in T$.
- the error term is at most λ/d which can be as low as $O(1/\sqrt{d})$ if we apply Theorem 3.

Another way of seeing Lemma 4 is that the number of edges $e(S,T)$ is close to the expectation of the number of edges between S and T in the random graph with edge probability d/n, or in the random d-regular graph (with the same vertex set), and this holds for any choice of S and T.

We mention that Lemma 4 does not apply directly to bipartite graphs, but there is a natural version of the lemma where S is a subset of one part of the bipartition, and T a subset of the other part.

We now describe a simple application of Lemma 4.

Let S be a subset of vertices of size $s = 2\lambda n/d$. The number of edges with both ends in S is $e(S,S)/2 \geq (ds^2/n - \lambda s)/2 = \Omega(n)$. It follows that S contains at least an edge whenever n is sufficiently large. In particular G has no independent set (i.e., no set of pairwise non-adjacent vertices) of size s, and thus G has chromatic number at least $n/s = d/2\lambda$. The existence of large girth expander graphs thus implies the existence of large girth graphs of large chromatic number. Note that when $\lambda = O(\sqrt{d})$, as is the case for non-bipartite Ramanujan graphs, we obtain graphs of chromatic number $\Omega(\sqrt{d})$. This is typically the order of magnitude of the chromatic number of classical explicit constructions of Ramanujan graphs (Lubotzky-Philips-Sarnak), while the chromatic number of random d-regular graphs is closer to $d/\log d$.