Spectral Graph Theory Lecture 13

Cayley Graphs

Daniel A. Spielman October 14, 2009

13.1 Overview

In this lecture, I will explain how to make graphs from groups. I will begin by constructing a graph from a linear error-correcting code.

13.2 Graphs from Linear Codes

Consider a linear code over $\{0,1\}$ from m bits to n bits. We may assume that such a code is encoded by an m-by-n matrix M, and that its codewords are the vectors

$$\boldsymbol{b}M,$$

where $\boldsymbol{b} \in \{0,1\}^m$. Let d be the minimum distance of this code. We will use this code to construct an n-regular graph on 2^m vertices with $\lambda_2 = 2d$. The construction will be a generalization of the hypercube, and in fact the hypercube will be obtained if we take $M = I_m$.

We will take as the vertex set $V = \{0,1\}^m$. Thus, I will also write vertices as vectors, such as $\boldsymbol{x}$ and $\boldsymbol{y}$. Two vertices $\boldsymbol{x}$ and $\boldsymbol{y}$ will be connected by an edge if their sum modulo 2 is a column of M.

Let me say that again. Let $\boldsymbol{m}_1, \ldots, \boldsymbol{m}_n$ be the columns of M. Then, the graph has edge set

$$\{(\boldsymbol{x}, \boldsymbol{x} + \boldsymbol{m}_j) : \boldsymbol{x} \in V, 1 \leq j \leq n\}.$$

Of course, this addition is taken modulo 2. You should now verify that if M is the identity matrix, we get the hypercube. In the general case, it is like a hypercube with extra edges.

13.3 Analyzing the Eigenvectors and Eigenvalues

For each $\boldsymbol{b} \in \{0,1\}^m$, define the function $\boldsymbol{v_b}$ from V to the Reals given by

$$\boldsymbol{v_b}(\boldsymbol{x}) = (-1)^{\boldsymbol{b}^T \boldsymbol{x}}.$$

While it is natural to think of $\boldsymbol{b}$ as being a vertex, that is the wrong perspective. Instead, you should think of $\boldsymbol{b}$ as indexing a Fourier coefficient (if you don't know what a Fourier coefficient it, just don't think of it as a vertex).

The eigenvectors and eigenvalues of the graph are determined by the following theorem.

Theorem 13.3.1. *For each $\boldsymbol{b} \in \{0,1\}^m$ the vector $\boldsymbol{v_b}$ is an adjacency matrix eigenvector of with eigenvalue*

$$n - 2\left|\boldsymbol{b}M\right|.$$

Before proving the theorem, I will establish some elementary facts that could appear mysterious if they are presented in the middle of the proof. First, note that for two vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ in $\{0,1\}^m$,

$$\boldsymbol{b}^T(\boldsymbol{x}+\boldsymbol{y}) = \boldsymbol{b}^T\boldsymbol{x} + \boldsymbol{b}^T\boldsymbol{y}.$$

This is of course true, because $\{0,1\}^m$ is a vector space under addition modulo 2. This equality implies

$$(-1)^{\boldsymbol{b}^T(\boldsymbol{x}+\boldsymbol{y})} = (-1)^{\boldsymbol{b}^T\boldsymbol{x}}(-1)^{\boldsymbol{b}^T\boldsymbol{y}}.$$

For any $\boldsymbol{b}$, $|\boldsymbol{b}M|$ is the Hamming weight of the codeword $\boldsymbol{b}M$. We will also write it

$$\begin{aligned}\left|\boldsymbol{b}M\right| &= \sum_{i=1}^{m} \boldsymbol{b}^T\boldsymbol{m}_i \quad \text{taking the sum over the Reals, not modulo 2}\\ &= \sum_{i=1}^{m} \frac{1-(-1)^{\boldsymbol{b}^T\boldsymbol{m}_i}}{2}.\end{aligned}$$

We will use this in the form

$$\sum_{i=1}^{m}(-1)^{\boldsymbol{b}^T\boldsymbol{m}_i} = n - 2\left|\boldsymbol{b}M\right|.$$

Proof of Theorem 13.3.1. Let A be the adjacency matrix of the graph. For any vector $\boldsymbol{v_b}$ for $\boldsymbol{b} \in \{0,1\}^m$ and any vertex $\boldsymbol{x} \in V$, we compute

$$\begin{aligned}(A\boldsymbol{v_b})(\boldsymbol{x}) &= \sum_{i=1}^{m} \boldsymbol{v_b}(\boldsymbol{x}+\boldsymbol{m}_i)\\ &= \sum_{i=1}^{m}(-1)^{\boldsymbol{b}^T(\boldsymbol{x}+\boldsymbol{m}_i)}\\ &= \sum_{i=1}^{m}(-1)^{\boldsymbol{b}^T\boldsymbol{x}}(-1)^{\boldsymbol{b}^T\boldsymbol{m}_i}\\ &= (-1)^{\boldsymbol{b}^T\boldsymbol{x}}\sum_{i=1}^{m}(-1)^{\boldsymbol{b}^T\boldsymbol{m}_i}\\ &= \boldsymbol{v_b}(\boldsymbol{x})(n - 2\left|\boldsymbol{b}M\right|).\end{aligned}$$

So, $\boldsymbol{v_b}$ is an eigenvector of eigenvalue $n - 2\left|\boldsymbol{b}M\right|$. □

So, if d is the minimum weight of a non-zero codeword, then $\alpha_2 = n - 2d$, and $\lambda_2 = 2d$.

Thus, an asymptotically good error-correcting code gives us a good expander graph, although of logarithmic degree. Last week we learned that for every $\delta < 1/2$, there is an $r > 0$ such that, for

sufficiently large n, there exist codes of length n, rate r, and relative minimum distance δn. These provide graphs on 2^{rn} vertices of degree n with $\lambda_2 \geq 2\delta n$. So, these are expanders with degree logarithmic in the number of vertices. But, I should mention two caveats. In the proof of Theorem 10.2.1 and 10.3.1, we assumed an upper bound on λ_{max} as well. It turns out that an upper bound on λ_{max} is unnecessary for Theorem 10.3.1, and for one side of Theorem 10.2.1. But, if we wanted an upper bound on λ_{max}, it would be easy to modify the proof of Lemma 11.5.1 to provide one as well. The idea is to also show that it is unlikely that there are any codewords of very large Hamming weight.

13.4 Groups

This construction is a special case of a general construction of graphs from groups, called Cayley graphs.

In case you don't recall what a group is, I will remind you. A group consists of a set of elements Γ, together with a binary operations on these elements, often denoted $\circ$. For elements g and h of Γ, $g \circ h$ is always another element of Γ. The set Γ and operation $\circ$ form a group if

1. Γ contains a special element, called the identity and often written id, such that $g \circ id = g$ and $id \circ g = g$ for all $g \in \Gamma$.
2. For every element $g \in \Gamma$, there is another element $g^{-1} \in \Gamma$ such that $g \circ g^{-1} = id$ and $g^{-1} \circ g = id$.
3. For every f, g and h in Γ, $f \circ (g \circ h) = (f \circ g) \circ h$.

You won't need to know any group theory to follow this lecture, because the only groups we will use will be groups that you already know. Here are some examples of groups.

1. Let Γ be the integers, and $\circ$ be addition. The identity element is 0. This group is usually written $\mathbf{Z}$.
2. For any number $n > 0$, let Γ be the integers modulo n, and let $\circ$ be addition. Again, the identity element is 0. This group is usually written $\mathbf{Z}/n$.
3. For any prime p, let Γ be the integers modulo p, except for 0. Let $\circ$ be multiplication. This is a group, and the identity element is 1. If we included 0, it would not be a group because 0 does not have an inverse. Similarly, if we tried to use all the integers, instead of working modulo p, it would not be a group because 2 would not have an inverse under multiplication. If we tried the integers modulo a non-prime, say $n = ab$, it would not be a group because a would not have an inverse. To see this, assume that a^{-1} exists, and get a contradiction by

$$\begin{aligned} (a^{-1}a)b &= a^{-1}(ab) && \text{implies} \\ (id)b &= a^{-1}(n) && \text{implies} \\ b &= 0. \end{aligned}$$

If you haven't seen a proof of it before, I recommend that you figure out why you do get a group when n is a prime.

4. For some integer $k > 0$, let Γ be the set $\{0,1\}^k$, and let $\circ$ be component-wise addition, modulo 2. The identity is the all-zero element. This group is called $(\mathbf{Z}/2)^k$, and sometimes F_2^k or $GF(2)^k$.

   This is the example we just did in the previous section. In this group, every element is its own inverse!

5. For some integer $k > 0$, let Γ be the set of all k-by-k matrices over the integers, and let $\circ$ be addition. The identity is the all-zero matrix. We can also do this with matrices modulo an integer.

6. For some integer $k > 0$, let Γ be the set of all k-by-k non-singular matrices over the integers, and let $\circ$ be multiplication. The identity is the standard identity matrix. We can also do this with integers modulo a prime.

For this lecture, we will just consider the second and fourth of these groups. These groups have the advantage[1] of begin both finite and Abelian, where I recall that a group is abelian if $g \circ h = h \circ g$ for all g and h in Γ.

13.5 Cayley Graphs

A Cayley graph is defined by a group $(\Gamma, \circ)$ and a set of *generators*, a subset S of Γ that is closed under inverse. That is, for every $g \in S$, $g^{-1} \in S$. The vertex set of the Cayley graph is Γ, and the edges are the pairs

$$\{(g,h) : h = g \circ s \text{ for some } s \in S\} = \{(g, g+s) : s \in S\}.$$

For example, we get the ring graph on n vertices by taking $(\Gamma, \circ) = (\mathbf{Z}/n, +)$ and $S = \{1, -1\}$. By "-1", we of course mean modulo n; so this is the same thing as $n - 1$.

13.6 Eigenvectors of Cayley Graphs of Abelian Groups

The wonderful thing about Cayley graphs of Abelian groups is that we can construct an orthornormal basis of eigenvectors for these graphs without even knowing the set of generators S. That is, the eigenvectors only depend upon the group. Related results also hold for Cayley graphs of arbitrary groups, and are related to representations of the groups. See [Bab79] for details.

As Cayley graphs are regular, it won't matter which matrix we consider. For simplicity, we will consider adjacency matrices.

[1] This is an advantage in that they are easier to understand. It actually limits what one can do with them quite a bit.

Let n be an integer and let G be a Cayley graph on $\mathbf{Z}/n$ with generator set S. When $S = \{\pm 1\}$, we get the ring graphs. For general S, I think of these as generalized Ring graphs. Let's first see that they have the same eigenvectors as the Ring graphs.

Recall that we proved that the vectors $\boldsymbol{x}_k$ and $\boldsymbol{y}_k$ were eigenvectors of the ring graphs, where

$$\begin{aligned}\boldsymbol{x}_k(u) &= \sin(2\pi ku/n), \text{ and}\\ \boldsymbol{y}_k(u) &= \cos(2\pi ku/n),\end{aligned}$$

for $1 \leq k \leq n/2$.

Let's just do the computation for the $\boldsymbol{x}_k$, as the $\boldsymbol{y}_k$ are similar. For every u modulo n, we have

$$\begin{aligned}(A\boldsymbol{x}_k)(u) &= \sum_{g\in S} \boldsymbol{x}_k(u+g)\\ &= \frac{1}{2}\left(\sum_{g\in S} \boldsymbol{x}_k(u+g) + \boldsymbol{x}_k(u-g)\right)\\ &= \frac{1}{2}\left(\sum_{g\in S} \sin(2\pi k(u+g)/n) + \sin(2\pi k(u-g)/n)\right)\\ &= \frac{1}{2}\left(\sum_{g\in S} 2\sin(2\pi ku/n)\cos(2\pi kg/n)\right)\\ &= \sin(2\pi ku/n)\sum_{g\in S}\cos(2\pi kg/n)\\ &= \boldsymbol{x}_k(u)\sum_{g\in S}\cos(2\pi kg/n).\end{aligned}$$

So, the corresponding eigenvalue is

$$\sum_{g\in S}\cos(2\pi kg/n).$$

This makes it easy to bound the eigenvalues of a random generalized ring graph. If we choose a random generator set of size $\Theta(\log n)$, we will also obtain an expander, just as we did for the hypercubes. To see why, consider the value of

$$\sum_{g\in S}\cos(2\pi kg/n)$$

when we choose S at random, of size d. For $k > 0$, this eigenvalue is a sum of d random variables, each of absolute value at most 1 and symmetrically distributed around 0. So, their sum will be concentrated around 0, and the probability that is is greater than $\kappa\sqrt{d}$ will be exponentially small in κ^2. If we take κ larger enough that this probability is smaller than $1/n$, then we can get a bound on every eigenvalue. When $d = c\log n$, we will be able to prove such a bound with κ proportional to $\sqrt{\log n}$ as well, which will yield upper bounds on all eigenvalues other than the first. I won't go through exact settings of the parameters as they rely on Chernoff bounds. But, in a few weeks (maybe next week?) we will prove concentration bounds that will suffice for our analysis.

13.7 Non-Abelian Groups

In the homework, you will show that it is impossible to make constant-degree expander graphs from Cayley graphs of Abelian groups. The best expanders are constructed from Cayley graphs of 2-by-2 matrix groups. In particular, the Ramanujan expanders of Margulis [Mar88] and Lubotzky, Phillips and Sarnak [LPS88] are Cayley graphs over the Projective Special Linear Groups $\mathrm{PSL}(2, p)$, where p is a prime. These are the 2-by-2 matrices modulo p with determinant 1, in which we identify A with $-A$.

They provided a very concrete set of generators. For a prime q modulo to 1 modulo 4, it is known that there are $p + 1$ solutions to the equation

$$a_1^2 + a_2^2 + a_3^2 + a_4^2 = p,$$

where a_1 is odd and a_2, a_3 and a_4 are even. We obtain a generator for each such solution of the form:

$$\frac{1}{\sqrt{p}} \begin{bmatrix} a_0 + ia_1 & a_2 + ia_3 \\ -a_2 + ia_3 & a_0 - ia_1 \end{bmatrix},$$

where i is an integer that satisfies $i^2 = -1$ modulo p.

Even more explicit constructions, which do not require solving equations, may be found in [ABN+92].

References

[ABN+92] Noga Alon, Jehoshua Bruck, Joseph Naor, Moni Naor, and Ron M. Roth. Construction of asymptotically good low-rate error-correcting codes through pseudo-random graphs. *IEEE Transactions on Information Theory*, 38(2):509–516, March 1992.

[Bab79] László Babai. Spectra of cayley graphs. *Journal of Combinatorial Theory, Series B*, pages 180–189, 1979.

[LPS88] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.

[Mar88] G. A. Margulis. Explicit group theoretical constructions of combinatorial schemes and their application to the design of expanders and concentrators. *Problems of Information Transmission*, 24(1):39–46, July 1988.