in theory

"Marge, I agree with you - in theory. In theory, communism works. In theory." -- Homer Simpson

CS359G Lecture 5: Characters of Abelian Groups

January 28, 2011 in CS359G, math | Tags: characters, Fourier analysis

In which we introduce the theory of characters of finite abelian groups, which we will use to compute eigenvalues and eigenvectors of graphs such as the cycle and the hypercube

In the past lectures we have established the Cheeger inequalities

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

and the fact that the SpectralPartitioning algorithm, when given an eigenvector of λ_2, finds a cut $(S, V-S)$ such that $h(S) \leq 2\sqrt{h(G)}$. In the next lecture we will show that all such results are tight, up to constants, by proving that

- The dimension-d hypercube H_d has $\lambda_2 = 1 - \frac{2}{d}$ and $h(H_d) = \frac{1}{d}$, giving an infinite family of graphs for which $\frac{1-\lambda_2}{2} = h(G)$, showing that the first Cheeger inequality is exactly tight.
- The n-cycle C_n has $\lambda_2 = 1 - O(n^{-2})$, and $h(C_n) = \frac{2}{n}$, giving an infinite family of graphs for which $h(G) = \Omega(\sqrt{1-\lambda_2})$, showing that the second Cheeger inequality is tight up to a constant.
- There is an eigenvector of the 2nd eigenvalue of the hypercube H_d, such that the SpectralPartitioning algorithm, given such a vector, outputs a cut $(S, V-S)$ of expansion $h(S) = \Omega(1/\sqrt{d})$, showing that the analysis of the SpectralPartitioning algorithm is tight up to a constant.

In this lecture we will develop some theoretical machinery to find the eigenvalues and eigenvectors of *Cayley graphs of finite Abelian groups*, a class of graphs that includes the cycle and the hypercube, among several other interesting examples. This theory will also be useful later, as a starting point to talk about algebraic constructions of expanders.

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

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 1 (Character) *Let Γ be a group (we will also use Γ to refer to the set of group elements). A function $f : \Gamma \rightarrow \mathbb{C}$ 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 1 (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 $f : \Gamma \rightarrow \mathbb{C}$ 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 $f : \mathbb{Z} \rightarrow \mathbb{C}$ 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 2** *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 $a \rightarrow b + a$ 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 $\sum_{a\in\Gamma} \chi(a) = 0$. □

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

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

Lemma 3 *If $\chi_1, \chi_2 : \Gamma \rightarrow \mathbb{C}$ are two different characters of a finite group Γ, then*

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

We will prove Lemma 3 shortly, but before doing so we note that, for a finite group Γ, the set of functions $f : \Gamma \rightarrow \mathbb{C}$ is a $|\Gamma|$-dimensional vector space, and that Lemma 3 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 4 *If Γ is a finite group, then it has at most $|\Gamma|$ characters.*

It remains to prove Lemma 3, which follows from the next two statements, whose proof is immediate from the definitions.

Fact 5 *If χ_1, χ_2 are characters of a group Γ, then the mapping $x \rightarrow \chi_1(x)\cdot\chi_2(x)$ is also a character.*

Fact 6 *If χ is a character of a group Γ, then the mapping $x \rightarrow \chi^*(x)$ is also a character, and, for every x, we have $\chi(x)\cdot\chi^*(x) = 1$.*

To complete the proof of Lemma 3, 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 7 *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 8 *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 $(x, y) \rightarrow \chi_1(x) \cdot \chi_2(y)$ 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 $f : \Gamma \rightarrow \mathbb{C}$, 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 1 (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*
>
> $$\chi_{r1,\ldots,r_n}(x_1, \ldots, x_n) := (-1)^{r_1 x_1 + \cdots r_n x_n}$$
>
> *Every boolean function $f : \{0, 1\}^n \rightarrow \mathbb{C}$ 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 2 (The Cyclic Group) *To work out another example, consider the case* $\Gamma = \mathbb{Z}/N\mathbb{Z}$*. Then every function* $f : \{0, \ldots, N-1\} \rightarrow \mathbb{C}$ *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.

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 $f : \Gamma \rightarrow \mathbb{C}$.

If Γ is an infinite abelian group, we can still define an inner product among functions $f : \Gamma \rightarrow \mathbb{C}$, 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 $f : \Gamma \rightarrow \mathbb{C}$, 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 $f : \Gamma \rightarrow \mathbb{C}$ are in $L^2(\Gamma)$, and a function $f : \mathbb{Z} \rightarrow \mathbb{C}$ 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 3 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 2, 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 2 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 8 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 $f:\Gamma\rightarrow\mathbb{C}$ 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 3 (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 $\chi_n:[0,1)\rightarrow\mathbb{C}$*
>
> $$\chi_n(x) := e^{2\pi ixn}$$
>
> *and it can be shown that there are no other characters. We thus have the Fourier series for continuous functions $f:[0,1)\rightarrow\mathbb{C}$,*
>
> $$f(x) = \sum_{n\in\mathbb{Z}} \hat f(n)e^{2\pi ixn}$$
>
> *where*

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

Share this: Twitter Digg Reddit Print

Like Be the first to like this post.

5 comments

Comments feed for this article

January 28, 2011 at 2:00 am

Tweets that mention CS359G Lecture 5: Characters of Abelian Groups « in theory -- Topsy.com

[...] This post was mentioned on Twitter by Aaron autofeed, TCS blog aggregator. TCS blog aggregator said: CS359G Lecture 5: Characters of Abelian Groups « in theory: http://bit.ly/fzPbDX [...]

January 28, 2011 at 6:15 am

Tyson Williams

At the end of the proof of Lemma 2, you are missing “= 0″.

January 29, 2011 at 1:17 pm

Marco

Thank you for the notes Prof. Trevisan! A couple of typo corrections:
- Corollary 2 “it as” -> “it has”
- In the proof of Lemma 3 the definition of X is missing an “(x)” at the end
- In property two 2 of the proof of Thm 8 there is a bar missing in |\Gamma_2|

January 29, 2011 at 1:38 pm

luca

thanks for the corrections!

April 17, 2011 at 6:13 am

Answering my own question on the Fourier Series | cartesian product

[...] CS359G Lecture 5: Characters of Abelian Groups (lucatrevisan.wordpress.com) [...]

LEAVE A REPLY

Your email address will not be published. Required fields are marked *

Name *

Email *

Website

Comment

You may use these HTML tags and attributes: `<a href="" title=""> <abbr title=""> <acronym title=""> <b> <blockquote cite=""> <cite> <code> <pre> <del datetime=""> <em> <i> <q cite=""> <strike> <strong>`

Post Comment

□ Notify me of follow-up comments via email.

□ Subscribe to this site by email

in theory

"Marge, I agree with you - in theory. In theory, communism works. In theory." -- Homer Simpson

CS359G Lecture 6: The Spectrum of the Cycle and of the Hypercube

January 29, 2011 in CS359G | Tags: Cayley graphs, eigenvalues, eigenvectors, hypercube

In which we talk about the spectrum of Cayley graphs of abelian groups, we compute the eigenvalues and eigenvectors of the cycle and of the hypercube, and we verify the tightness of the Cheeger inequalities and of the analysis of spectral partitioning

In this lecture we will prove the following results:

1. The dimension-d hypercube H_d has $\lambda_2 = 1 - \frac{2}{d}$ and $h(H_d) = \frac{1}{d}$, giving an infinite family of graphs for which $\frac{1-\lambda_2}{2} = h(G)$, showing that the first Cheeger inequality is exactly tight.
2. The n-cycle C_n has $\lambda_2 = 1 - O(n^{-2})$, and $h(C_n) \geq \frac{2}{n}$, giving an infinite family of graphs for which $h(G) = \Omega(\sqrt{1-\lambda_2})$, showing that the second Cheeger inequality is tight up to a constant.
3. There is an eigenvector of the second eigenvalue of the hypercube H_d, such that the SpectralPartitioning algorithm, given such a vector, outputs a cut $(S, V-S)$ of expansion $h(S) = \Omega(1/\sqrt{d})$, showing that the analysis of the SpectralPartitioning algorithm is tight up to a constant.

1. 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 1** *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 $w : \Gamma \rightarrow \mathbb{R}$ 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 1 (Cycle)** *The n-vertex cycle can be constructed as the Cayley graph $Cay(\mathbb{Z}/n\mathbb{Z}, \{-1, 1\})$.*

> **Example 2 (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 2 *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 $M\mathbf{x}$:

$$(M\mathbf{x})_a = \sum_b M_{a,b}x_b$$

$$= \frac{1}{|S|}\sum_{b:b-a\in S}\chi(b)$$

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

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

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.

2. The Cycle

The n-cycle is the Cayley graph $Cay(\mathbb{Z}/n\mathbb{Z}, \{-1, +1\})$. Recall that, for every $n \in \{0, \dots, n-1\}$, 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 $r \in \{0, \dots, n-1\}$ 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.

3. 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 $\{e^1, \dots, e^d\}$, 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 $r = (0, \dots, 0)$, 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 $x^j_a = (-1)^{a_j}$. 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.

Share this: Twitter Digg Reddit Print

Like Be the first to like this post.

No comments yet

Comments feed for this article

LEAVE A REPLY

Your email address will not be published. Required fields are marked *

Name *

Email *

Website

Comment

You may use these HTML tags and attributes: `<a href="" title=""> <abbr title=""> <acronym title=""> <b> <blockquote cite=""> <cite> <code> <pre> <del datetime=""> <em> <i> <q cite=""> <strike> <strong>`

Post Comment

□ Notify me of follow-up comments via email.

□ Notify me of new posts via email.

of B by the relative density of $|B \cap \ell|/|\mathbb{F}|$. This method requires $2\log(\mathbb{F}^m)$ bits and can be analyzed using the *pairwise independence* of points on a line. We present a randomness efficient line sampler (that is almost as good in terms of its error), based on ε-biased sets over $\mathbb{F}^m$ (Lemma 4.3). This efficient line sampler is crucial to our randomness efficient low degree testing (and the resulting PCPs and LTCs) and may be of further use in other contexts.

The Uniformity Lemma One of the first steps is PCP systems is a reduction of SAT to an algebraic constraint satisfaction problem (ACSP), where solutions are restricted to low degree polynomials (over some finite field $\mathbb{F}$), and variables and constraints are indexed by vectors in $\mathbb{F}^m$. The number of constraints resulting from this reduction bounds from below the overall length of the PCP. Previous reductions gave ACSPs that were quadratic in the input length [19]. We show that a small variant of the Harsha-Sudan reduction [19] results in ACSPs with a nearly linear number of constraints. The main observation (Lemma 5.9) is that a random input to a random constraint (in our small set of constraints) is (with high probability) a random point in $\mathbb{F}^m$.

1.7 Paper Organization

We start with an survey of ε-biased sets and their connections to expander graphs and error-correcting codes (Section 2). While many of the results there are folklore, we find it beneficial to state and prove them formally. Section 3 analyzes randomness efficient linearity tests over $\mathbb{Z}_p$. In the course of this analysis we also improve on previous results regarding linearity testing over these fields. In section 4 we analyze the randomness efficient low degree tests. Sections 5,6 give the applications of the low degree tests to short PCPs and LTCs, respectively.

2 Preliminaries: ε-Bias, Expanders and Codes

2.1 Characters and Fourier Representation

Let G be a finite Abelian group of order k, written additively. A *character* of G is a homomorphism $\chi : G \to \mathbb{C}^\times$ where $\mathbb{C}^\times$ is the multiplicative group of (nonzero) complex numbers, i.e. χ satisfies

$$\chi(a+b) = \chi(a) \cdot \chi(b)$$

for all $a, b \in G$. It follows that $\chi(a)$ is a k'th root of unity for all a. The *trivial* character of G is defined by $\chi_0(a) = 1$ for all $a \in G$. It can be shown that, for any character χ,

$$\sum_{a \in G} \chi(a) = \begin{cases} |G| & \chi \text{ is the trivial character} \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

The set of characters of G is denoted $\hat{G}$. Let $\mathbb{C}^G$ be the space of functions $f : G \to \mathbb{C}$. This is a $|G|$-dimensional linear space over $\mathbb{C}$. We use the standard inner product over this space:[4]

$$\langle f, g \rangle = \frac{1}{|G|} \sum_{a \in G} f(a)\overline{g(a)}$$

[4] Throughout the paper we normalize inner products by the dimension of the space. Namely, if u, v are n-dimensional vectors, then $\langle u, v \rangle = \frac{1}{n} \sum_{i=1}^{n} u_i \overline{v_i}$

The set of characters are an orthonormal basis for this space. Thus any $f \in \mathbb{C}^G$ can be represented as a linear combination of characters over $\mathbb{C}$. Such a representation is

$$f = \sum_{\chi \in \hat{G}} \hat{f}_\chi \cdot \chi$$

where $\hat{f}_\chi$ is called the *Fourier coefficient* corresponding to χ, and is given by

$$\hat{f}_\chi = \langle \chi, f \rangle$$

The linear transformation T sending f to $\hat{f}$ is unitary and hence length preserving. Thus we get Parseval's identity:

$$\langle f, f \rangle = |G| \cdot \langle \hat{f}, \hat{f} \rangle \tag{2}$$

The *convolution* of two functions f, g, denoted $f * g$ is a mapping of G to $\mathbb{C}$ defined by

$$(f * g)(x) = |G|^{-1} \sum_{u+v=x} f(u) \cdot g(v)$$

We need the following *convolution identity* which shows the relationship between the Fourier coefficients of two functions f, g and the Fourier coefficients of their convolution:

$$\forall \chi \in \hat{G} : \qquad \widehat{(f * g)}_\chi = \hat{f}_\chi \cdot \hat{g}_\chi \tag{3}$$

Of particular interest for our derandomization of the linearity test is the special case of the additive group $G = \mathbb{Z}_p^m$ (addition modulo a prime p), whose set of characters is $\hat{G} = \{\chi_\alpha : \alpha \in \mathbb{Z}_p^m\}$, where

$$\chi_\alpha(x) \overset{\text{def}}{=} \omega^{\sum_{i=1}^m \alpha_i \cdot x_i} \tag{4}$$

and ω is a primitive pth root of unity. When $G = \mathbb{Z}_p^m$ we use the standard notation $\hat{f}_\alpha$ to denote the Fourier coefficient $\hat{f}_{\chi_\alpha}$ corresponding to the character χ_α defined above.

2.2 ε-Biased Sets

For H a group, $S \subseteq H$ is called *symmetric* if it is closed under inverses (i.e. $s \in S \Rightarrow -s \in S$).

Definition 2.1 (ε-Bias) *Let H be a finite Abelian group, and let $S \subseteq H$ be a symmetric multiset. We say S is* ε-biased *if all of the nontrivial Fourier coefficients of the characteristic vector of S have absolute value at most $\varepsilon|S|/|H|$. That is, for every nontrivial character χ of H:*

$$\frac{1}{|S|} \left| \sum_{y \in S} \chi(y) \right| \leq \varepsilon$$

Requiring that S be symmetric guarantees that the character sums above are always real (because $\chi(-y) = \chi(y)^{-1} = \overline{\chi(y)}$). This is not a significant constraint, because making a set symmetric increases its size and the magnitudes of the character sum by at most a factor of 2. Note that by Equation (1), the entire space $S = H$ is 0-biased. But we will be interested in *small* ε-biased sets. A standard application of the Probabilistic Method yields such sets:

Proposition 2.2 *For every H, there exists an ε-biased set of size $O(\log |H|/\varepsilon^2)$. Indeed, most sets of this size are ε-biased.*

However, we need *explicit* constructions. We call S *explicit* if its elements can be enumerated in time $\mathrm{poly}(|S|, \log |H|)$.[5] The first definition and explicit construction of ε-biased sets was given by Naor and Naor [30]. Since then several other explicit constructions appeared for various groups. The groups we deal with in this paper are $\mathbb{Z}_p^m$ and $\mathbb{F}^m$ (where $\mathbb{F}$ is a finite field). For $\mathbb{Z}_2^m$ the state of the art is due to [1], and comes quite close to the probabilistic bounds.

Construction 2.3 *[1] For any integer n and any $\varepsilon > 0$ there exists an explicit ε-biased set over $\mathbb{Z}_2^m$ of size $\leq \min\left\{\left(\frac{n}{\varepsilon}\right)^2, O\left(\frac{n}{\varepsilon^3}\right)\right\}$.*

Notice that if $\mathbb{F} = \mathrm{GF}(2^k)$, then the additive group of $\mathbb{F}^m$ is isomorphic to $\mathbb{Z}_2^{mk}$. In particular, the characters of $\mathbb{F}^m$ coincide with those of $\mathbb{Z}_2^{mk}$. For this case the best construction is the one above. For $\mathbb{F} = \mathrm{GF}(p^k)$ where $p > 2$, a number of constructions appear in the literature, given by [24, 2, 34, 15, 4].

Construction 2.4 *[24, 2, 34, 15, 4] For any finite field $\mathbb{F}$, any integer $m \geq 1$ and any $\varepsilon > 0$, there exists an explicit ε-biased set over $\mathbb{F}^m$ of size $\mathrm{poly}(\log |\mathbb{F}|, m, \frac{1}{\varepsilon})$.*

2.3 ε-Biased Sets as Cayley Expanders

The connection between ε-biased sets and expanders will be pivotal in our derandomization of the low degree test (Section 4).

Definition 2.5 (Cayley Graph) *Let H be a finite Abelian group and let $S \subseteq H$ be symmetric. The Cayley graph generated by S over H is the (multi-)graph G_S with vertex set H and edge set $\{(a, a+s) : s \in S\}$.*

By symmetry of S, G_S is undirected. For G a d-regular graph, let A_G be its adjacency matrix. The largest eigenvalue of A_G is d, corresponding to the eigenvector $\vec{1}$. The *normalized second eigenvalue* of G is defined to be $|\lambda_2|/d$, where λ_2 is the second largest eigenvalue of A_G in absolute value. The following connection is folklore, and is implicit in [5, 30].

Lemma 2.6 *For any finite Abelian group H and any $S \subseteq H$, S is ε-biased iff the normalized second eigenvalue of G_S is at most ε.*

Proof: The key observation is that a character χ of H is an eigenvector of $A = A_{G_S}$, having eigenvalue $\lambda_\chi = \sum_{y \in S} \chi(y)$. To see this look at $\vec{v} = A \cdot \chi$. The i'th coordinate of $\vec{v}$ (where $i \in H$) is given by

$$\langle A_i, \chi \rangle = \sum_{y \in S} \chi(i+y) = \chi(i) \cdot \sum_{y \in S} \chi(y)$$

So $A \cdot \chi = \lambda_\chi \cdot \chi$. Recall that the characters of H form an orthogonal set of dimension $|H|$. Since A is an $|H| \times |H|$ matrix, we conclude that the eigenvalues of A are exactly $\{\lambda_\chi : \chi \text{ is a character of } H\}$.

[5] Clearly, this definition only makes sense when discussing infinite families of groups, e.g. $H_n = \{0,1\}^n$.

G_S is $|S|$-regular. Normalizing the eigenvalues by $|S|$ we get that the maximal bias of S equals the maximal normalized second eigenvalue of A. ■

One of the key properties of graphs with small second eigenvalue is given by the well-known Expander Mixing Lemma. For G a graph and A, B subsets of vertices of G, let $\mathbf{e}(A, B)$ be the number of edges connecting a vertex from A to a vertex in B.

Lemma 2.7 (Expander Mixing Lemma) *(cf., [7]) For $G = (V, E)$ a connected d-regular graph over n vertices with normalized second eigenvalue λ, and any two sets $A, B \subseteq V$ of densities $a = |A|/|V|$, $b = |B|/|V|$:* $\left|\frac{\mathbf{e}(A,B)}{|E|} - ab\right| \leq \lambda\sqrt{ab}$

That is, the fraction of edges between A and B is approximately the same as in the complete graph, namely the product of their densities.

2.4 ε-Biased Sets as Linear Codes

In our derandomization of the affineness test (section 3) we will use properties of an error-correcting code induced by an ε-biased set, first observed for the Boolean case by Naor and Naor [30]. Let $\mathcal{C}_S \subseteq \mathbb{C}^S$ be the truncation of the characters of $\mathbb{Z}_p^m$ to the index set S. I.e., if $(y_1, \ldots y_s)$ is an ordering of the elements of $S \subseteq \mathbb{Z}_p^m$, then $\mathcal{C}_S = \{(\chi(y_1), \ldots, \chi(y_s)) : \chi \in \hat{\mathbb{Z}_p^m}\}$.

An ε-biased set S induces a code $\mathcal{C}_S$ with good minimal distance. $\mathcal{C}$ is called *ε-balanced* if the distance between any two distinct codewords is $(1 - 1/p)(1 \pm \varepsilon)$. Formally, for any two distinct codewords $(c_1, \ldots, c_n) \neq (c'_1, \ldots, c'_n)$:

$$\left(1 - \frac{1}{p}\right)(1 - \varepsilon) \leq \Pr[c_i \neq c'_i] \leq \left(1 - \frac{1}{p}\right)(1 + \varepsilon)$$

The following lemma was originally proved for $\mathbb{Z}_2^m$ by Naor and Naor [30]. Although we do not use the Hamming distance, rather the related inner product measure (defined in Section 3.4) we include the following lemma for the sake of completeness.

Lemma 2.8 *If $S \subset \mathbb{Z}_p^m$ is ε-biased then $\mathcal{C}_S$ is ε-balanced.*

Proof: Since $\mathcal{C}_S$ is linear, it suffices to show that for any fixed non-trivial character χ

$$\left(1 - \frac{1}{p}\right)(1 - \varepsilon) \leq \Pr_{y \in S}[\chi(y) \neq 1] \leq \left(1 - \frac{1}{p}\right)(1 + \varepsilon)$$

For $a \in \{1, \ldots, p-1\}$ let χ^a be the non-trivial character defined by $\chi^a(x) = (\chi(x))^a$. From the ε-bias of S we get

$$\varepsilon \geq \frac{1}{(p-1)} \left| \sum_{a \in \mathbb{Z}_p \setminus \{0\}} \left(\frac{1}{|S|} \sum_{y \in S} \chi^a(y) \right) \right| \tag{5}$$

Recall that for ζ a pth root of unity

$$\sum_{a=1}^{p-1} \zeta^a = \begin{cases} p-1 & \zeta = 1 \\ -1 & \text{otherwise} \end{cases}$$

Defining $S' = \{y \in S : \chi(y) \neq 1\}$ we change the order of summation in Eq. (5) and get

$$\begin{aligned} \varepsilon &\geq \frac{1}{(p-1)|S|}\left|\sum_{y\in S}\left(\sum_{a\in\mathbb{Z}_p\setminus\{0\}}\chi^a(y)\right)\right| && (6)\\ &= \left|\frac{|S'|-(|S|-|S'|)(p-1)}{(p-1)|S|}\right| = \left|1-\frac{p}{p-1}\cdot\frac{|S'|}{|S|}\right| && (7)\end{aligned}$$

Noticing $\Pr_{y\in S}[\chi(y) \neq 1] = |S'|/|S|$ completes the proof. ■

Remark 2.9 *If we demand S to be* closed under multiplication *(i.e. $y \in S$ iff $ay \in S$ for all $a \in \mathbb{Z}_p \setminus \{0\}$) then the converse of the previous lemma also holds. Namely, a set S closed under multiplication is ε-biased iff $\mathcal{C}_S$ is ε-balanced. In particular, since any $S \subset \mathbb{Z}_2^m$ is closed under multiplication, we conclude that ε-bias and ε-balance coincide for $\mathbb{Z}_2^m$.*

3 Affineness Testing

In this section we show how to test whether a function $f : \mathbb{Z}_p^m \to \mathbb{Z}_p$ is close to being affine, where the randomness needed is only $(1 + o(1))m \log p$, compared with $2m \log p$ in the previous tests.

3.1 The BLR Affineness Test

For Abelian groups G (additive) and H (multipilicative), an *affine function* from G to H is a function f such that

$\forall x, y \in G, f(x)f(y) = f(x+y)f(0)$. In the *affineness testing* problem we are given oracle access to a function $f : G \to H$, and wish to test whether it is close in Hamming distance to some affine function. For $f, g : G \to H$ two functions, we define the *agreement* of f and g as $\Pr_{x\in G}[f(x) = g(x)]$. We are interested in measuring the *maximal agreement* of f with some affine function. Blum, Luby, and Rubinfeld [14] suggested the following test:

BLR AffTest^f, on function $f : G \to H$

1. Select $x, y \in G$ uniformly at random.
2. Accept iff $f(x)f(y)f(-(x + y)) = f(0)$

The original test suggested by [14] was only for homomorphisms, for which $f(0) = 0$, but their techniques generalize to affine functions as well. Indeed, notice that f is affine iff the function $f' \overset{\text{def}}{=} f^{-1}(0) \cdot f$ is a homomorphism. Moreover the acceptance probability of the previous test on the two functions is the same. By definition f is affine if and only if the test accepts for *all pairs* $x, y \in G$. What is more surprising is that the acceptance probability of the test is a good estimate on the maximal agreement of f with an affine function. This has been shown in [14] by the following theorem.

Theorem 3.1 *[14] For any finite Abelian groups G, H and any function $f : G \to H$, if* **AffTest**^{f} *accpets with probablity $\geq 1 - \frac{2}{9}\delta$, then f has agreement $\geq 1 - \delta$ with some affine function.*