THE BOURGAIN-GAMBURD CONSTRUCTION OF EXPANDERS

NICK GILL

The aim of this lecture is to give an outline of the Bourgain-Gamburd construction of expanders using growth results [BG08]. We will not be able to give more than a basic outline of their argument - you should read the paper if you want to understand their ideas properly! If I have time in the future, I might dramatically expand this lecture; no time just now.

1. Introduction

Take A a finite set in $SL_2(\mathbb{Z})$. Define A_p a subset of $G_p = SL_2(\mathbb{Z}/p\mathbb{Z})$ to be the set we obtain by reducing all entries in elements of A modulo p; this yields an infinite family of graphs, $\mathcal{C}(G_p, A_p)_{p\,\mathrm{prime}}$.

Now we allow the possibility that we "throw away" the first few graphs - those corresponding to $p < C$ for some constant C. We write the resulting family of graphs as $\mathcal{C}(G_p, A_p)_{p\to\infty}$.

The question we are interested in is this:

$$\textbf{Is } \mathcal{C}(G_p, A_p)_{p\to\infty}\textbf{a family of expanders?} \tag{1.1}$$

We can give a complete answer to question (1.1):

Theorem 1. [BG08, Thm. 1] *The family of graphs $\mathcal{C}(G_p, A_p)_{p\to\infty}$ is a family of expanders if and only if the group $\langle A\rangle < SL_2(\mathbb{Z})$ does not contain a subgroup of finite index.*

Some notes about this result: Bourgain and Gamburd give a number of equivalent conditions to the one concerning a subgroup of finite index. They also give a result concerning *random* Cayley graphs of this form: i.e. pick k elements $g_1, \dots, g_k$ at random, and define the set A to be $\{g_1, g_1^{-1}, \dots, g_k, g_k^{-1}\}$.

We will not prove Thm. 1; nor will we prove the result about random graphs. Instead we focus on the following subsidiary result:

Theorem 2. *Fix $k \geq 2$, and suppose that*

(a) *we have a set $A \in SL_2(\mathbb{Z})$ such that $A = \{g_1, g_1^{-1}, \dots, g_k, g_k^{-1}\}$. Suppose further that $g_i^2 \neq 1$ for all $i = 1, \dots, k$.*

(b) *There is a constant $\tau > 0$ such that, for all primes p,*

$$\mathrm{girth}(\mathcal{C}(G_p, A_p) \geq \tau \log_{2k} p. \tag{1.2}$$

Then $\mathcal{C}(G_p, A_p)_{p\to\infty}$ is a family of expanders.

We will operate under the suppositions of Thm. 2 for the rest of this lecture. Note that supposition (a) implies that the Cayley graph $\mathcal{C}(G_p, A_p)$ is undirected, and $2k$-regular for p larger than some constant C. This supposition is very minor; we include it only for convenience.

Supposition (b), on the other hand, is significant. The proof of (the reverse implication in) Thm. 1 consists of two steps: take a set A such that $\langle A \rangle$ does not have a soluble subgroup of finite index; first prove that supposition (b) holds - we have a lower bound on the girth of $\mathcal{C}(G_p, A_p)$; then prove that, given this lower bound, the family $\mathcal{C}(G_p, A_p)_{p\to\infty}$ is a family of expanders. In this lecture, then, we are interested in the second step of this argument.

2. Connectedness

Remember that k, A, and τ are all fixed now. In this section we prove the following result:

Proposition 2.1. *There exists C such that for $p > C$ the graph $\mathcal{C}(G_p, A_p)$ is connected.*

The first step is a classical result of Dickson classifying the proper subgroups of $SL_2(q)$. We state an adapted version:

Theorem 3. *For $p \geq 5$, $H < G_p$, one of the following holds:*
(a) *H has a cyclic subgroup of index 2;*
(b) *H lies in a Borel subgroup of G_p;*
(c) $|H| \leq 120$.

We obtain an immediate corollary:

Corollary 2.2. *For $g_1, g_2, g_3, g_4 \in H < G_p$, one of the following holds:*
(a) $[[g_1, g_2], [g_2, g_4]] = 1$*;*
(b) *for all i, there exists $n \leq 120$ such that $g_i^n = 1$.*

We can now prove Prop. 2.1:

Proof. Suppose that $\mathcal{C}(G_p, A_p)$ is not connected. Then $\langle A_p \rangle = H < G_p$ (by a basic property of Cayley graphs discussed in the last lecture).

Now take $g_1, g_2, g_3, g_4 \in A_p$ (not necessarily distinct). Now Cor. 2.2 implies that one of the following holds:
(a) $[[g_1, g_2], [g_3, g_4] = 1$. Note that $[[g_1, g_2], [g_3, g_4]]$ is a word of length 16. Thus we obtain that $\text{girth}(\mathcal{C}(G_p, A_p)) \leq 16$.
(b) $g_1^n = 1$ for some $n \leq 120$. Then we obtain that $\text{girth}(\mathcal{C}(G_p, A_p)) \leq 120$.

In all cases, then, $\text{girth}(\mathcal{C}(G_p, A_p)) \leq 120$. But now (1.2) implies that $\tau \log_{2k} p \leq 120$, and we obtain that $p \leq (2k)^{120\tau}$. Take $C = (2k)^{120\tau}$ and we are done. □

Prop. 2.1 allows us to add another supposition to those from Thm. 2. Specifically for the rest of this lecture we assume that

$$\text{There exists } C \text{ such that } p > C \Rightarrow \langle A_p \rangle = G_p. \tag{2.1}$$

Now we know a lot about the growth of generating sets in simple groups from earlier lectures. In this lecture we will make use of the following special case of the general theorem about growth in simple groups:

Theorem 4. *Fix $\epsilon > 0$. There exist constants $C, \delta > 0$ such that for any generating set B in G_p, either*

- $|B \cdot B \cdot B| \geq C|B|^{1+\delta}$*; or*
- $|B| \geq |SL_2(\mathbb{Z}/p\mathbb{Z})|^{1-\epsilon}$.

Note that the group G_p is not, strictly speaking, simple. Instead $G_p/Z(G_p)$ is simple and, for $p > 2$, $|Z(G_p)| = 2$. It is an easy matter to prove Thm. 4 from the general results about growth in simple groups already encountered.

3. Measures

We introduce some machinery that will be important in connecting growth in groups to expansion on graphs. For a finite group G, define $\mu : G \to \mathbb{R}^+$ to be a *probability measure* if $\sum_{x \in G} \mu(x) = 1$. We record the following definitions:

(a) We have a norm on measures: $||\mu||_2 = (\sum_{g \in G}(\mu(g))^2))^{\frac{1}{2}}$;

(b) We can *convolve* measures: let $\mu, \nu : G \to \mathbb{R}^+$ be two probability measures on a group G. Then

$$(\mu * \nu) : G \to \mathbb{R}^+, \ x \mapsto \sum_{g \in G} \mu(xg^{-1})\nu(g)$$

is a probability measure.

(c) Convolutions have a natural associativity property. Thus, for $l \in \mathbb{Z}^+$, we can write

$$\mu^{(l)} = \underbrace{\mu * \cdots * \mu}_{l}.$$

Now we will make use of a particular probability measure on G_p:

$$\mu_A : G \to \mathbb{R}^+, \ x \mapsto \begin{cases} \frac{1}{|A_p|}, & x \in A_p; \\ 0, & x \notin A_p. \end{cases}$$

Now the key point is that we can relate facts about walks on $\mathcal{C}(G_p, A_p)$ to properties of the measure μ_A.

Define W_{2l} to be the number of walks on $\mathcal{C}(G_p, A_p)$ from 1 to 1 of length $2l$. Write N for $|G_p| = \frac{1}{2}p(p-1)(p+1)$, write Adj for the adjacency matrix of $\mathcal{C}(G_p, A_p)$, and write $\lambda_0, \ldots, \lambda_{N-1}$ for the spectrum of Adj.

The following facts are easy (after you've thought for a while):

$$NW_{2l} = \mathrm{tr}(\mathrm{Adj}^{2l}) = \sum_{j=0}^{N-1} \lambda_j^{2l}; \tag{3.1}$$

$$\mu_A^{(2l)}(1) = \frac{W_{2l}}{(2k)^{2l}}; \tag{3.2}$$

$$\mu_A^{(l)}(g) = \mu_A^{(l)}(g^{-1}). \tag{3.3}$$

It is (3.2) that connects walks to measures. It will be central in what follows.

3.1. Walks on $\mathcal{C}(G_p, A_p)$.

Our discussion of measures to this point could be rewritten for any Cayley graph of a finite group. Now we focus in on the group G_p. We state the following proposition, which is proved in [BG08].

Proposition 3.1. *For all $\epsilon > 0$ there exists C such that, for $l \geq C \log_{2k}(p)$,*

$$||\mu_A^{(l)}||_2 < p^{-\frac{3}{2}+\epsilon}. \tag{3.4}$$

Prop. 3.1 is difficult - it takes about half of the work in [BG08]. We will return to it at the end, and see how it can be proved using Thm. 4 (it is in proving Prop. 3.1 that growth results enter the proof of Thm. 2). However, for now, we will assume Prop. 3.1 and see where it gets us.

Observe first that, for $l \geq C \log_{2k} p$,

$$\frac{W_{2L}}{(2k)^{2l}} = \mu_A^{(2l)}(1) = \sum_{g \in G} \mu^{(l)}(g)\mu^{(l)}(g^{-1}) = \sum_{g \in G} (\mu^{(l)}(g))^2 = ||\mu^{(l)}||_2^2 < p^{-3+2\epsilon}.$$

We summarize:

$$\text{For } l \geq C \log_{2k} p, \quad W_{2l} < \frac{(2k)^{2l}}{p^{3-2\epsilon}}. \tag{3.5}$$

4. Using representations

Next we use some representation theory. There is no time to give a proper introduction to this, so we will assume a basic working knowledge in what follows.

Recall that $A = \{g_1, g_1^{-1}, \dots, g_k, g_k^{-1}\}$. Once again write Adj for the adjacency matrix of $\mathcal{C}(G_p, A_p)$. We have

$$\text{Adj} = \pi_R(g_1) + \pi_R(g_1^{-1}) + \cdots \pi_R(g_k) + \pi_R(g_k^{-1})$$

where $\pi_R : G_p \to \mathbb{C}$ is the regular representation of G_p.

Now we can decompose π_R into irreducibles. Write $\hat{G}$ for the set of irreducible complex representations of G. Then

$$\pi_R = \bigoplus_{\rho \in \hat{G}} \rho.$$

Standard representation theory tells us that m_ρ, the multiplicity of ρ in π_R is equal to $\dim \rho$.

Now we appeal to a classical result of Frobenius:

Theorem 5. *The non-trivial irreducible complex representations of G_p have dimension at least $\frac{1}{2}(p-1)$.*

This yields an immediate corollary. Recall that $\lambda_0, \dots, \lambda_{N-1}$ is the spectrum for Adj.

Corollary 4.1. *All non-trivial eigenvalues of* Adj *occur in the spectrum with multiplicty at least $\frac{1}{2}(p-1)$.*

Note that Prop. 2.1 implies that, for p greater than some constant, the only trivial eigenvalue in the spectrum is $\lambda_0 = k$.

5. Proof of Thm. 2

Let us prove Thm. 2. Write m_{λ_i} for the multiplicity of λ_i in the spectrum of Adj. Then Cor. 4.1 implies, in particular, that $m_{\lambda_1} \geq \frac{1}{2}(p-1)$.

Now we put together a series of (in)equalities that have appeared in numbered equations throughout the lecture so far:

$$\frac{(2k)^{2l}}{p^{-2\epsilon}} > N\frac{(2k)^{2l}}{p^{3-2\epsilon}} > NW_{2l} = \sum_{j=1}^{N-1} \lambda_j^{2l} > m_{\lambda_1}\lambda_1^{2l} > \frac{1}{2}(p-1)\lambda_1^{2l}.$$

With some rearranging we get

$$\lambda_1^{2l} \leq 3\frac{(2k)^{2l}}{p^{1-2\epsilon}}.$$

Now fix $2l = C\log_{2k}(p)$ (where C is as given in Prop. 3.1). Then, with some more rearranging, we obtain that

$$\lambda_1 < (2k)^{1-\epsilon'} < 2k.$$

here ϵ' depends only on ϵ, τ, and k (in particular, it does not depend on p). Thus we have a spectral gap:

$$\lambda_0 - \lambda_1 > 2k - (2k)^{1-\epsilon'} > 0$$

for all graphs in the family $\mathcal{C}(G_p, A_p)_{p\to\infty}$. The result is proved.

6. The connection to growth

Recall that we made use of Prop. 3.1 without proving it. We do not have time to give a proper proof for Prop. 3.1 but we will spend a little time trying to justify how it might follow from results on growth, in particular from Thm. 4. We apologise in advance for the extreme sketchiness of this justification!

Let us recall the statement of Prop. 3.1.

Proposition 6.1. *For all $\epsilon > 0$ there exists C such that, for $l \geq C\log_{2k}(p)$,*

$$||\mu_A^{(l)}||_2 < p^{-\frac{3}{2}+\epsilon}. \tag{6.1}$$

Now (6.1) is an upper bound on the norm of the l-th convolution of the measure μ_A. It is reasonable to think that such a bound could be derived from a statement about a single convolution. A statement like this, for instance:

$$||\mu_A * \mu_A||_2 < p^{\epsilon}||\mu_A||_2. \tag{6.2}$$

So let us suppose that we can derive (6.1) from (6.2). How then can we prove (6.2)? (In fact (6.2) is not true without some extra conditions but we won't worry about this.)

We need to connect growth in groups to measures. For A, B subsets of a finite group G, we define the measure

$$\chi_A : G \to \mathbb{R}^+, \ x \mapsto \begin{cases} 1, & x \in A; \\ 0, & x \notin A. \end{cases}$$

We can define χ_B analagously. Note that χ_A bears a startling resemblance to μ_A; note too that χ_A is not a *probability* measure as it doesn't satisfy the property that $\sum_{x\in G}\chi_A(x) = 1$.

Now we define the *multiplicative energy* of A and B:

$$E(A,B) = |\{(x_1, x_2, y_1, y_2) \in A \times A \times B \times B \mid x_1y_1 = x_2y_2\}|.$$

The multiplicative energy counts the number of repeats when we multiply AB. Roughly speaking if the energy, $E(A,B)$ is large, then the set AB is small; in other words energy is in an inverse relationship to growth.

Now we connect the two in an easy lemma:

Lemma 6.2. $||\chi_A * \chi_B||_2 = E(A, B)$.

Now suppose that (6.2) does not hold. Then Lem. 6.2 immediately implies a lower bound on the additive energy:

$$E(A_p, A_p) \geq p^{\epsilon}||\mu_A||_2.$$

It turns out that, with a great deal of fiddling about, we can translate this statement into a statement about growth. Intuitively a lower bound on energy implies an upper bound on growth, and this is exactly what we get:

$$|A_p \cdot A_p \cdot A_p| < |A_p|^{1+\epsilon'}$$

where ϵ' depends only on ϵ, k, and τ. But now Thm. 4 implies that $|A_p| > |G_p|^{1-\delta}$; as p tends towards infinity this gives a contradiction (since $|A_p| = 2k$ is fixed). We conclude that (6.2) holds, as required.

The last paragraph should be taken as a *very rough* guide to the philosophy of Bourgain and Gamburd's argument. The key point is that we are able to connect growth to measures through the concept of multiplicative energy.

References

[BG08] J. Bourgain and A. Gamburd, *Uniform expansion bounds for Cayley graphs of* $sl_2(\mathbb{F}_p)$, Annals of Math. **167** (2008), 625–642.