Thin groups and superstrong approximation
MSRI Publications
Volume **61**, 2013

Superstrong approximation for monodromy groups

JORDAN S. ELLENBERG

Finitely generated subgroups of arithmetic lattices are frequently encountered in geometry, where they appear as *monodromy groups* attached to the variation of cohomology in families of manifolds or algebraic varieties. Here we survey briefly the construction of monodromy groups, discuss our (limited) knowledge about whether such groups are arithmetic, and summarize the results of [Ellenberg et al. 2012], which derive an application to arithmetic geometry from recent advances in superstrong approximation [Helfgott 2008; Pyber and Szabó 2011; Breuillard et al. 2011; Salehi Golsefidy and Varjú 2012]. We conclude by indulging ourselves in some speculations about more general contexts, asking: what are the interesting questions about "nonabelian superstrong approximation" and "superstrong approximation for Galois groups"?

1. What is monodromy?

The simplest example of monodromy is provided by the topological notion of a *covering space*. Suppose given a path-connected base space B, endowed with a choice of basepoint b, and let $f : X \to B$ be a covering space of degree n; that is, each point $x \in X$ has a neighborhood U such that $f^{-1}(U)$ is homeomorphic to n disjoint copies of U. In particular, $f^{-1}(b)$ consists of n points. A paradigmatic example is provided by $B = X = \mathbb{C}^\times$, where the map f is given by $z \mapsto z^n$, and $b = 1$.

Now let γ be a loop in B beginning and ending at b, and let $\tilde{b}$ be a point in $f^{-1}(b)$. It follows from the definition of covering space that there is a unique path $\tilde{\gamma}$ which starts at $\tilde{b}$ and projects to γ. Intuitively, one imagines the point in X "following its shadow" in B, whose motion is specified by γ. In the paradigmatic example, we can take $\gamma : [0, 1] \to B$ to be a counterclockwise traversal of the unit circle

$$\gamma(t) = e^{2\pi i t}$$

and $\tilde{b} = 1$. Then $\tilde{b}$ starts at 1, and at time t is required to satisfy

$$\tilde{\gamma}(t)^n = e^{2\pi i t}.$$

It follows that $\tilde{\gamma}(t) = e^{2\pi i t/n}$. The path $\tilde{\gamma}$ starts at 1, but it ends at $e^{2\pi i/n}$. In fact, starting at any point $\tilde{b}$ and following your shadow around γ puts you at $\tilde{b}e^{2\pi/n}$. The permutation of the n points in $f^{-1}(b)$ induced by this operation is called the *monodromy* of the cover around γ; in this case, it is a cyclic permutation of the n points. It is easy to check that the map sending a path γ based at b to its monodromy is a homomorphism

$$\mu : \pi_1(B, b) \to \mathrm{Sym}(f^{-1}(b)),$$

which we call the *monodromy map of the cover*. The image of μ is called the *monodromy group*.

Any identification of $f^{-1}(b)$ with the set $[1, \ldots, n]$ allows us to express the monodromy of the cover as a homomorphism from the fundamental group $\pi_1(B, b)$ to S_n; of course, changing the labeling on $f^{-1}(b)$ changes the monodromy by conjugation. Note that the monodromy map need not surject onto the symmetric group of $f^{-1}(b)$; for instance, in the paradigmatic example, the monodromy group is cyclic of order n. For that matter, the monodromy group can be trivial — this happens precisely when X is isomorphic to the trivial cover $B \times [1 \ldots, n]$. The monodromy group is transitive if and only if X is connected; in this case, there is a well-defined conjugacy class of point stabilizers in the monodromy group, and the preimage in $\pi_1(B, b)$ of such a stabilizer is an index-n subgroup of $\pi_1(B, b)$. In this setting, we are just rephrasing the usual Galois-theoretic identification between connected degree-n covers of B and index-n subgroups of $\pi_1(B, b)$.

We can generalize this story: suppose now that $f : X \to B$ is a *fibration* of manifolds. Denote the fiber over $b \in B$ by X_b. One can take hold of a homology class in $H_i(X_b, \mathbb{Z})$ and "move it along γ" just as we did in above with a single point of X, and this defines a monodromy map

$$\mu : \pi_1(B, b) \to \mathrm{GL}(H_i(X_b, \mathbb{Z})),$$

Again, the *monodromy group* of the fibration (or *the monodromy group on* H_i, if there is some ambiguity) is the image of $\pi_1(B, b)$ in the arithmetic group $\mathrm{GL}(H_i(X_b, \mathbb{Z}))$. Of course, we can define a monodromy group using the action on cohomology in just the same way.

Rather than define this process rigorously, we will sketch an illustrative case. Suppose that B is the moduli space $\mathcal{M}$ of lattices in $\mathbb{R}^2$ up to homothety, and let X be the universal torus over $\mathcal{M}$; in other words, for each point p of $\mathcal{M}$ parametrizing a lattice $\Lambda_p \subset \mathbb{R}^2$, the fiber X_p is the torus $\mathbb{R}^2/\Lambda_p$.

Now let the basepoint b be the point parametrizing the square lattice Λ_b

generated by $(1,0)$ and $(0,1)$, and define a loop $\gamma \subset \mathcal{M}$ by letting $\gamma(t)$ be the lattice generated by $(1,0)$ and $(t,1)$. Note that $\gamma(1) = \gamma(0) = b$.

The path from $(0,0)$ to $(0,1)$ represents a nontrivial homology class in $H_1(X_b, \mathbb{Z})$, which we call e_1. Write e_2 for the homology class of the path from $(0,0)$ to $(1,0)$; then e_1 and e_2 freely generate $H_1(X_b, \mathbb{Z})$.

What happens as we move along γ, carrying e_1 along with us as we go? At time t, our path moves along a straight line segment from $(0,0)$ to $(t,1)$. And when the loop is finished, our path goes from $(0,0)$ to $(1,1)$. The torus has returned to its starting point, but the homology class on X_b represented by our path is now $e_1 + e_2$.

What if we'd started with the path from $(0,0)$ to $(1,0)$? This path stays in place through our whole trip around γ.

In other words, the loop γ induces a *unipotent* monodromy transformation on $H_1(X_b, \mathbb{Z})$, sending e_1 to $e_1 + e_2$ and fixing e_2.

Below we present a few examples of monodromy maps attached to families of manifolds.

Example 1. If $X \to B$ is a finite covering space of degree n, then X_b is a set of n points, and the only nontrivial homology group of X_b is $H_0(X_b) \cong \mathbb{Z}^n$. In this case, after choosing an identification of X_b with $1, \ldots, n$, the monodromy map factors as

$$\mu : \pi_1(B, b) \to S_n \to \mathrm{GL}_n(\mathbb{Z}),$$

where the second map is the standard permutation representation.

Example 2. Let B be the moduli space of genus g Riemann surfaces, and let X be the universal curve over B; that is, for each point p of B, the fiber X_p is the Riemann surface parametrized by p. In this case, $\pi_1(B, b)$ is the *mapping class group* Γ_g of genus g. (See [Farb and Margalit 2011] for a thorough introduction to the mapping class group and its basic properties.) We can identify $H_1(X_b, \mathbb{Z})$ with $\mathbb{Z}^{2g}$; having done so, the monodromy map takes the form

$$\mu : \Gamma_g \to \mathrm{GL}_{2g}(\mathbb{Z}).$$

This map is definitely not surjective once $d > 2$, because $H_1(X_b, \mathbb{Z})$ carries a symplectic intersection form which is preserved when moving around in moduli space. So the image lies in the subgroup of matrices preserving this form, and we can write:

$$\mu : \Gamma_g \to \mathrm{Sp}_{2g}(\mathbb{Z}).$$

In fact, the image of μ is known to be all of $\mathrm{Sp}_{2g}(\mathbb{Z})$; see [Farb and Margalit 2011, §6.3.2], for example.

Example 3. Let B be the configuration space parametrizing unordered n-tuples of distinct elements of $\mathbb{C}$, so that $\pi_1(B, b)$ is the Artin braid group B_n on n strands. Let X be the family whose fiber over a configuration $\{p_1, \dots, p_n\}$ is the smooth proper algebraic curve with affine model

$$y^d = (x - p_1)(x - p_2) \dots (x - p_n).$$

Then the monodromy map

$$\mu : B_n \to \mathrm{GL}(H_1(X_b))$$

is a representation of the braid group, which is studied extensively in [McMullen 2013]. Each fiber X_b is an algebraic curve with a $\mathbb{Z}/d\mathbb{Z}$-action: namely, 1 sends y to $e^{2\pi i/d} y$ and leaves x unchanged. Thanks to the intersection form, the image of μ is again contained in the group of symplectomorphisms $H_1(X_b)$. By contrast with the previous example, the monodromy group is not the full symplectic group, because it commutes with the canonical action of $\mathbb{Z}/d\mathbb{Z}$ on the fiber. So the monodromy group is contained in the centralizer of $\mathbb{Z}/d\mathbb{Z}$ in the symplectic group; we denote this group by $\mathrm{Sp}_T(H_1(X_b))$.

Remark. Another classical route to monodromy groups is via solutions of differential equations. We won't pursue examples of this kind here, but see [Sarnak 2014; Fuchs 2014] in this volume for a fuller account.

2. Big monodromy and superstrong approximation

Given a fibration $X \to B$, we might aim to describe, as best we can, the corresponding monodromy group $\Gamma = \mu(\pi_1(B, b))$ in $\mathrm{GL}(H_i(X_b, \mathbb{Z}))$. A crucial invariant is the Zariski closure G of Γ, which is an algebraic group over $\mathbb{Q}$ contained in a general linear group GL_n, where $n = \dim_{\mathbb{Q}} H_i(X_b, \mathbb{Q})$. We refer to G as the *algebraic monodromy group*. By $G(\mathbb{Z})$ we shall mean the intersection of $G(\mathbb{C})$ with $\mathrm{GL}(H_i(X_b, \mathbb{Z}))$.[1]

As a first step, we would like to be able to say whether it is *big*. Of course, there are several senses in which the word "big" could be meant. For instance, we can ask whether G is all of GL_n. As we've seen, the answer is often negative, sometimes for readily perceived reasons — for instance, in the case where $X \to B$ is a family of genus g curves, the intersection form on $H_1(X_b, \mathbb{Z})$ constrains G to be contained in Sp_{2g}. So one might also ask, more vaguely, "is G as big as it could be, given all the reasons I can think of for it to be small"? For a typical family of curves, this amounts to asking whether G is the full symplectic group. In the context of Example 3, it would mean to ask whether G is the whole of Sp_T.

[1] In this article, we will always be in a context where $H_i(X_b, \mathbb{Z})$ is torsion-free, and where any subtleties about the "right" model for G over $\mathrm{Spec}\, \mathbb{Z}$ do not intervene.

It is not straightforward to compute the algebraic monodromy group G in general, but much progress has been made towards ensuring that G is big under various checkable conditions. See [Katz 2004] and [Hall 2008] for representative recent results.

2.1. ***Thin monodromy.*** Even if G is known, there are still important questions about the size of Γ itself. First of all, we can compare Γ with the full lattice $G(\mathbb{Z})$. The critical question is whether Γ has finite index in $G(\mathbb{Z})$.

Definition 4. A subgroup Γ of $G(\mathbb{Z})$ is *thin* if it is Zariski-dense and has infinite index.

There is, of course, a large literature on the geometry, group theory, and dynamics of arithmetic lattices — the main theme of this conference is to understand the extent to which some of these results extend to thin subgroups.

It can be difficult, in practice, to know whether the monodromy group of a family is thin, and both thin and nonthin cases do arise. For instance, the family of hyperelliptic curves of genus g has monodromy group which is finite-index in $\mathrm{Sp}_{2g}(\mathbb{Z})$, by a theorem of A'Campo [1979]. On the other hand, the family of curves described in Example 3 can have either thin or nonthin monodromy depending on the parameters; for instance, McMullen [2013, Section 11] shows (following Deligne and Mostow) that the monodromy is thin when $(n, d) = (4, 18)$. On the other hand, Venkataramana [2013] shows that the monodromy is an arithmetic lattice whenever $n \geq 2d$. Brav and Thomas [2012] have recently shown that a family of quintic three-folds constructed by Dwork has as its H^3-monodromy a thin subgroup of $\mathrm{Sp}_4(\mathbb{Z})$. Fuchs's article [2014] in this volume describes several senses in which thinness seems to be the "generic" situation, including many examples of differential equations whose monodromy groups appear on experimental grounds to be thin.

Example 5. Consider the family of complex algebraic curves given by

$$X_t : y^2 = f(x)(x - t),$$

where $f(x)$ is a fixed squarefree polynomial of degree $2g$, with $g > 1$, and t varies over the base B, which is the complement of the roots of f in $\mathbb{A}^1$. The monodromy group Γ of this family is Zariski dense in Sp_{2g}; see unpublished work of Yu [1997] and the more general theorem of Hall [2008, Theorem 4.1]. In fact, Yu proves that the monodromy group of this family is a specified finite-index subgroup of $\mathrm{Sp}_{2g}(\mathbb{Z})$, contained in the principal congruence subgroup of level 2 and containing the principal congruence subgroup of level 4. We take this opportunity to correct an error in Section 4.1 of [Ellenberg et al. 2012], where we make the incorrect assertion that the monodromy group is the whole of $\Gamma(2)$. (This error does not affect the argument in which it appears.)

Yu computes the monodromy group by direct computation with matrices. Ian Agol explained to us a very handsome pure-thought argument via the Margulis normal subgroup theorem; we include this argument here in the hope that the idea might be useful in other similar cases.

Let Conf_{2g+1} be the configuration space parametrizing squarefree complex polynomials of degree $2g+1$. The fundamental group of Conf_{2g+1} is the Artin braid group Br_{2g+1}, and there is a natural family of genus-g curves over Conf_{2g+1} whose fiber over a polynomial $g(x)$ is the hyperelliptic curve $y^2 = g(x)$. By the theorem of A'Campo already mentioned, the image of the corresponding monodromy map

$$\mathrm{Br}_{2g+1} \to \mathrm{Sp}_{2g}(\mathbb{Z})$$

is a finite-index subgroup of $\mathrm{Sp}_{2g}(\mathbb{Z})$ containing the full level-2 congruence subgroup $\Gamma(2)$.

The permutation action on strands induces a map $\mathrm{Br}_{2g+1} \to S_{2g+1}$ whose kernel P_{2g+1} is called the *pure braid group*. The image of P_{2g+1} under the monodromy image above is precisely $\Gamma(2)$.

Now the base space B considered in the present example maps to Conf_{2g+1} by sending t to $f(x)(x-t)$, and the monodromy map $\pi_1(B,b) \to \mathrm{Sp}_{2g}(\mathbb{Z})$ thus factors as a composition

$$\mu : \pi_1(B,b) \to \pi_1(\mathrm{Conf}_{2g+1}, b) \cong \mathrm{Br}_{2g+1} \to \mathrm{Sp}_{2g}(\mathbb{Z}).$$

Topologically speaking, the image of $\pi_1(B,b)$ in Br_{2g+1} is the *point-pushing group* H consisting of all braids in which the first $2g$ strands are fixed in place, and the last is free to wind around the others. The point-pushing group is the kernel in the *Birman exact sequence*

$$H \to P_{2g+1} \to P_{2g} \to 1,$$

and in particular is a *normal* subgroup of the pure braid group. Thus, $\mu(H)$ is a normal subgroup of the lattice $\Gamma(2)$. It then follows form the Margulis normal subgroup theorem that $\mu(H)$ is either finite or finite-index; but it is easy to check that a loop carrying t once around a root of f has unipotent (and thus infinite-order) monodromy, so $\mu(H)$ must be finite index, as claimed.

To sum up: the monodromy group Γ can be "big" in the sense that its Zariski closure G is as large as possible, but it might well be the case that Γ is very "small" by virtue of being thin, and in a given explicit geometric situation we may not know which is the case. The question is then whether thin groups are "big enough" for applications.

2.2. ***Strong approximation.*** One may ask how well Γ approximates $G(\mathbb{Z})$ locally — that is, with respect to the p-adic Lie groups $G(\mathbb{Z}_p)$. A natural question in this vein is whether the inclusion $\Gamma \to G(\mathbb{Z}_p)$ is dense for almost all p; that is, does Γ satisfy *strong approximation?* Of course, this requires that $G(\mathbb{Z})$ itself satisfies strong approximation; this is the case whenever G is semisimple and simply connected and $G(\mathbb{R})$ is noncompact. It is a remarkable fact, due in its strongest form to Weisfeiler, that (under mild conditions on G) *any* finitely generated Zariski-dense subgroup Γ of $G(\mathbb{Z})$ satisfies strong approximation. For details and history concerning these results, see [Rapinchuk 2014] in this volume.

A central theme of the lectures in this workshop has been that, despite the name "strong approximation", density of Γ in almost all $G(\mathbb{Z}_p)$ is not the strongest local approximation condition we can place on Γ. Indeed, strong approximation is *not* strong enough to make sieving arguments work. For this we need a stronger condition, called *superstrong approximation*, which we explain below. See [Salehi Golsefidy 2014; Kowalski 2014] in this volume for a survey of the *affine sieve* and other arguments which use superstrong approximation to derive quantitative statements on thin subgroups of arithmetic groups.

2.3. ***Superstrong approximation.*** Let Γ be a finitely generated subgroup of $G(\mathbb{Z})$, and let $\underline{\gamma} = \gamma_1, \ldots, \gamma_r$ be a set of generators for Γ. If $\pi : G(\mathbb{Z}) \to Q$ is a finite quotient of $G(\mathbb{Z})$, then we can form a Cayley graph $X(Q, \underline{\gamma})$, whose vertices are labeled by the elements of Q, and which has q adjacent to q' precisely when $q' = q\pi(\gamma_i)^{\pm 1}$ for some generator γ_i.

In what follows, we will let Q range over the congruence quotients $G(\mathbb{Z}/p\mathbb{Z})$, and we will refer to the Cayley graph $X(G(\mathbb{Z}/p\mathbb{Z}), \underline{\gamma})$ simply as X_p.

Note first that Γ surjects on $G(Z/p\mathbb{Z})$ if and only if X_p is connected. The notion of connectedness, in turn, can be addressed spectrally: the (suitably normalized) Laplacian operator on X_p always has an eigenvalue of 0, attached to the space of constant eigenfunctions, and all other eigenvalues ≥ 0. In fact, the number of components of X_p is precisely the multiplicity of the eigenvalue 0. In other words, the second-largest eigenvalue λ_1 is 0 precisely when X_p is not connected. And strong approximation, from a spectral point of view, says that $\lambda_1(X_p) = 0$ for only finitely many p.

Replacing this condition with a quantitative bound yields the notion of *superstrong approximation*.

Definition 6. Let Γ be a finitely generated subgroup of $G(\mathbb{Z})$ and define a sequence of Cayley graphs X_p as above. We say Γ has *superstrong approximation* if there exists a constant $\epsilon > 0$ such that $\lambda_1 > \epsilon$ for all sufficiently large primes p.

In other words, superstrong approximation is the condition that the X_p form a family of *expander graphs*. By the usual theory of expander graphs, this

condition does not depend on the choice of the generating set $\underline{\gamma}$; in other words, it is a property of $\Gamma \subset G(\mathbb{Z})$, as our phrasing suggests. We note that superstrong approximation follows immediately if Γ has Kazhdan's property T, as it must if it is a lattice in a higher-rank arithmetic group. But when Γ is thin, superstrong approximation doesn't follow from any abstract group-theoretic property of Γ. Indeed, in practice Γ is often free, so has many actions on finite sets whose corresponding Cayley graphs have small λ_1; superstrong approximation demands that the $G(\mathbb{Z}/p\mathbb{Z})$ are not among those finite sets.

The remarkable progress of the last few years towards uniform bounds for expansion in Cayley graphs of linear groups, as described elsewhere in this volume [Pyber and Szabó 2014], has brought us to the point where we now have a superstrong analogue of Weisfeiler's theorem; under mild conditions on the ambient algebraic group G, thin groups have superstrong approximation merely by virtue of their Zariski density in G.

3. Application: gonality growth in families of covers

We now describe an arithmetic application of superstrong approximation for monodromy groups. All the material here is explained in substantially more detail in our joint paper with Hall and Kowalski [Ellenberg et al. 2012].

Let U be an smooth (but not necessarily proper) algebraic curve over a number field k, and let $A \to U$ be an abelian scheme over U (i.e., a family of abelian varieties parametrized by U.) If t is a point of $U(\bar{k})$, we denote by A_t the fiber of the family over t; it is an abelian variety defined over the finite extension $k(t)/k$. If n is an integer, $A_t[n]$ denotes the group of n-torsion points of A_t.

For instance, we could fix a squarefree polynomial f, let U be the complement in $\mathbb{A}^1$ of the roots of f, and take A to be the family whose fiber A_t is the Jacobian of the hyperelliptic curve $y^2 = f(x)(x-t)$. (Though be warned that to prove Theorem 7 in this case doesn't use new results on thin groups, since in this case the monodromy of the family is finite-index in $\mathrm{Sp}_{2g}(\mathbb{Z})$, as explained in Example 5.)

The next result follows from Theorem 7 of [Ellenberg et al. 2012].

Theorem 7. *For every $d \geq 1$ there exists a constant $l(d)$ such that, for all primes $l > l(d)$, there are only finitely many $t \in U(\bar{k})$ such that $[k(t) : k] \leq d$ and $A_t[l](k(t))$ contains a nontrivial point.*

In words, "there are only finitely many l-torsion points on fibers of A over points of bounded degree".

We now sketch the proof. The first step is to construct finite covers of U as follows. For each integer n, let $f_n : U_n \to U$ be the map whose fiber is the finite set of nontrivial n-torsion points $A_t[n]$; this is a finite unramified covering

of degree $n^{2\dim A} - 1$. One might think of these covers as analogues of the classical modular curves; for if U is the moduli space of elliptic curves $Y(1)$[2], the cover U_n is precisely the moduli space $Y_1(n)$ of elliptic curves with an n-torsion point. In fact, the argument below is very much in the spirit of those of [Abramovich 1996] and [Zograf 1984] which gave a lower bound for the gonality of the classical modular curves; where we use new developments in superstrong approximation, they use Selberg's theorem that (in modern language) $\mathrm{SL}_2(\mathbb{Z})$ has property τ with respect to the family of congruence subgroups.

Theorem 7 is really a statement about algebraic points on the covers U_l: we are saying that, for all $l > l(d)$, the curve U_l has only finitely many points whose field of definition has degree at most d over k.

This statement might at first seem rather strong, since there are infinitely many degree-d field extensions of k. Even if a curve C has genus at least 2, so that $|C(k')| < \infty$ for each k'/k of degree d by Faltings' theorem, it is far from clear that the union of $C(k')$ over *all* such k' should be finite. Indeed, this is often not the case: for instance, the curve

$$y^2 = x^5 + x + 1$$

has infinitely many points over quadratic extensions of $\mathbb{Q}$, since one may produce such a point by specifying $x \in \mathbb{Q}$ arbitrarily. Of course, the story is the same for any hyperelliptic curve, which may have arbitrarily large genus. So the genus is not sufficient to distinguish curves with many degree-d algebraic points from curves with few such points. For that purpose, we need to study a different geometric invariant: the *gonality*.

Definition 8. The *gonality* of an algebraic curve C/k is the minimum integer γ such that there exists a morphism $C \to \mathbb{P}^1$ of degree γ.

(In this section, we take gonality to be a *geometric* invariant of the curve; that is, in the definition we allow the morphism $C \to \mathbb{P}^1$ to be defined over an extension of k.)

For example, a curve has gonality 1 precisely when it has genus 0, and gonality 2 precisely when it is hyperelliptic but not rational.

The gonality of a curve has a very strong effect on its points over fields of bounded degree over a fixed field, as we know from [Faltings 1983; Abramovich and Voloch 1996; Frey 1994]. Namely, if C/k has gonality greater than $2d$, then $C(\bar{k})$ has only finitely many points t such that $[k(t) : k] \leq d$.

[2]Warning: $Y(1)$ is not really an algebraic curve but a stack with nontrivial generic inertia. So it would be more precise to say that when $U = Y_1(p)$ for $p > 3$ and not dividing n, we have $U_n = Y_1(np)$.

To prove Theorem 7, it thus suffices to show that the gonality $\gamma(U_l)$ increases with l. This is accomplished using two main tools. First, we pass to the complex numbers and invoke an analytic theorem of Li and Yau [1982] which provides lower bounds for the gonality of a complex algebraic curve in terms of the spectrum of the Laplacian. In particular, it is enough to show that the family of Riemann surfaces U_l has a Laplacian spectral gap.

Why should this be so? This requires a combinatorial detour. The argument runs roughly as follows. Choose a basepoint $u \in U(\mathbb{C})$ and a basis $\gamma_1, \ldots, \gamma_r$ for $\pi_1(U(\mathbb{C}), u)$. (For simplicity we assume here that U is a punctured sphere, so that $\pi_1(U)$ is free.) We can represent each γ_i by a loop in $U(\mathbb{C})$ based at u, and we can arrange our loops to be mutually disjoint. We denote the union of these loops by X; it is a one-vertex graph embedded in the Riemann surface $U(\mathbb{C})$.

Now the preimage of X under the cover $U_l \to U$ is a graph $X(U_l)$ embedded in the Riemann surface $U(\mathbb{C})$. The vertices of $X(U_l)$ are the points of U_l lying over u; we recall that these correspond to the $l^{2g} - 1$ nontrivial l-torsion points on A_u.

On the other hand, we have a monodromy representation

$$\mu : \pi_1(U(\mathbb{C}), u) \to \mathrm{Sp}(H_1(A, \mathbb{Z})) \cong \mathrm{Sp}_{2g}(\mathbb{Z}).$$

The l-torsion in A_u is canonically identified with the mod-l homology group $H_1(A_u, \mathbb{Z}/l\mathbb{Z})$. So what happens if we start at a point $\tilde{u}$ in $U_l(\mathbb{C})$ lying over u, and make $\tilde{u}$ "follow its shadow" around a loop γ_i in $U(\mathbb{C})$? The description of monodromy from the first section tells us that the path ends at $\mu(\gamma_i)\tilde{u}$. In other words, the permutation on the preimages of u induced by γ_i is the same as that induced by the reduction mod l of the symplectomorphism $\mu(\gamma_i)$. This is precisely the same as saying that two vertices $\tilde{u}$ and $\tilde{u}'$ of $X(U_l)$ are adjacent precisely when $\tilde{u}' = \mu(\gamma_i)\tilde{u}$ for some i. In other words, $X(U_l)$ is a quotient of the Cayley graph X_l we described in Section 2.

This is the crucial point. For a given family of abelian variety we may have no easy way of checking whether the monodromy group Γ is thin or arithmetic. But it doesn't matter! The Helfgott, Pyber–Szabó, Breuillard–Green–Tao, and Golsefidy–Varjú theorems work equally well in either case to tell us that Γ has superstrong approximation. It is precisely this robustness that makes Theorem 7 possible.

Superstrong approximation for Γ tells us that the X_l have a spectral gap; that is, the nontrivial eigenvalues of the *discrete* Laplacians on these graphs is bounded away from 0 as l grows. (Hidden here is a certain amount of work necessary to pin down the possibilities for G, the Zariski closure of monodromy; this is necessary in order to ensure that superstrong approximation applies. This is covered in Section 5 of [Ellenberg et al. 2012].)

What is the relationship between the Laplacian of the graph and the geometry of the surface in which the graph is embedded? It can be illustrative to visualize the case where all the U_l have genus 0; that is, they are homeomorphic to the sphere. Now a graph embedded in a sphere is a planar graph, and it is known that a sequence of planar graphs with increasingly many vertices cannot have a spectral gap; intuitively, graphs with a spectral gap have very quickly dispersing random walks, while the random walk on a planar graph is constrained to linger for a while in the neighborhood of the plane where it starts out.

More generally, we should expect X_l to serve as a kind of combinatorial approximation to the geometry of U_l. Speaking roughly: the graph is like a wire frame along which U_l is stretched, and the graph separates U_l into many small components, each one of which is a copy of a patch on U. So the geometry of U_l is, in a sense, a hybrid of the geometry of X_l and the geometry of U. As l grows, and along with it the number of vertices of X_l, the graph plays a proportionately greater and greater role in the geometry of U_l and the shape of U a smaller and smaller one.

This intuition is beautifully confirmed and extended by results of Kelner, Brooks, and Burger, which together provide us with a "comparison principle" between the Laplacian spectrum of the Riemann surface $U_l(\mathbb{C})$ and the (discrete) Laplacian spectrum of the embedded graph X_l. In particular, the combination of their theorems shows that the sequence of Riemann surfaces $U_l(\mathbb{C})$ has a spectral gap whenever the sequence of graphs X_l does. This completes the proof.

Remarks. (1) For the arithmetic applications in [Ellenberg et al. 2012], the full strength of superstrong approximation was not needed, and we used a weaker condition instead. We call a sequence of graphs $\{X_l\}$ *esperantist* when $\lambda_1(X_l)$ decays no more quickly than a power of the logarithm of $|X_l|$. At the time [Ellenberg et al. 2012] was written, results guaranteeing superstrong approximation in the required generality were not available, but the results of [Pyber and Szabó 2011] established the esperantist condition in the contexts we needed. Though we now have superstrong approximation for thin groups in much greater generality, it still seems useful to keep in mind that in applications it is often sufficient to have at hand the easier condition on Γ that the Cayley graphs attached to Γ are esperantist ("slightly superstrong approximation…"?) On the other hand, the full strength of expansion allows one to conclude, not only that the gonality of U_l grows without bound as $l \to \infty$, but that the gonality is bounded below by a constant multiple of $\deg(U_l/U)$; in this case we say the gonality has *linear growth*.

(2) Many more applications of the basic method appear in [Ellenberg et al. 2012]. For instance, if the monodromy group Γ is Zariski dense in $\mathrm{Sp}_{2g}(\mathbb{Z})$, it follows

from Theorem 4 of [Ellenberg et al. 2012] that there are only finitely many t of degree at most d over k such that A_t is isogenous to a product of abelian varieties of smaller dimension.

(3) The theorems in [Ellenberg et al. 2012] require that the families of abelian varieties be defined over fields of characteristic 0, since everything rests on the ability to compare the geometry of Riemann surfaces with that of graphs drawn thereon. The *statements* of the theorems, however, largely make sense in characteristic p; it seems very interesting to wonder whether the results continue to be true in this case, and in particular whether any analogue of the argument by superstrong approximation can be made to work.

(4) Work of Cadoret and Tamagawa, carried out independently around the same time as [Ellenberg et al. 2012], gives lower bounds for the gonality of some of the covers discussed here, by very different methods [Cadoret and Tamagawa 2009; Cadoret 2012]. In particular, their results are purely algebraic, and can be applied in characteristic p as well as characteristic 0.

4. Speculations: nonabelian superstrong approximation

We can phrase the approximation criteria of the previous sections in a different way. Think of $\mathrm{GL}_n(\mathbb{Z})$ as the automorphism group of the free abelian group $\mathbb{Z}^n$. Then, for any finite abelian group A which can be generated by n elements, $\mathrm{GL}_n(\mathbb{Z})$ acts by permutations on $\mathrm{Epi}(\mathbb{Z}^n, A)$, the set of surjections from $\mathbb{Z}^n$ to A: we thus have a homomorphism

$$\mathrm{Aut}\ \mathbb{Z}^n \to \mathrm{Sym}(\mathrm{Epi}(\mathbb{Z}^n, A)).$$

Strong approximation for a subgroup Γ of $\mathrm{GL}_n(\mathbb{Z})$ is essentially the statement that the restriction of this action to Γ is *transitive* for all A; equivalently, the Cayley–Schreier graphs attached to this action are connected. Superstrong approximation, on the other hand, implies the stronger statement that those graphs form an expander family.

It is natural to ask about nonabelian analogues of this story; for instance, one can replace $\mathbb{Z}^n$ with a free group F_n on n generators, and replace the finite abelian group A with a general finite group G. It turns out that the analogues of strong and superstrong approximation have already appeared in various interesting contexts.

4.1. ***The product replacement algorithm.*** (Main references: [Lubotzky and Pak 2001; Lubotzky 2011].) If G is a finite group, then $\mathrm{Epi}(F^n, G)$ is precisely the set of n-element generating sets $(g_1, \ldots, g_n)$ of G. The *product replacement algorithm* of Leedham-Green and Soicher is a random walk on this set, which

proceeds as follows: choose distinct i, j in $1, \ldots, n$ at random, and then multiply g_i either on the left or right (chosen by coin flip) by g_j or g_j^{-1} (also chosen by coin flip.)

This random walk appears in practice to converge extremely quickly to the uniform distribution on generating n-tuples, making it an effective way to generate random elements in a group without having to enumerate the entire group. In fact, the product replacement algorithm is known to converge to uniform distribution in polynomial time when n is allowed to grow (slowly) with $|G|$, by a result of Pak [2000]. In the present paper we will concentrate mostly on the case where n is fixed and G changes.

The product replacement graph is in fact a Cayley–Schreier graph for the permutation action

$$A^+ \to \mathrm{Sym}(\mathrm{Epi}(F^n, G)),$$

where A^+ is a subgroup of Aut F_n of index 2, generated by automorphisms corresponding to the $4n(n-1)$ operations of the product replacement algorithm. (We will casually fail to distinguish between Aut F_2 and its finite-index subgroup A^+ in what follows.) The rapid convergence of the random walk naturally suggests the question: are these graphs expanders? In other words, if $G_1, G_2, \ldots$ is a sequence of finite groups, we may ask:

Question 9. Does Aut F^n have superstrong approximation with respect to its image in $\mathrm{Sym}(\mathrm{Epi}(F^n, G_i))$?

The above question would certainly have a positive answer if Aut F^n had property T (or even property τ for a suitable class of finite quotients.) Under such a hypothesis, quantitative bounds for the performance of the product replacement algorithm are given in [Lubotzky and Pak 2001]. Whether Aut F^n has property T in general is a well-known open question, about whose answer there is no clear consensus. We know that Aut F^n does *not* have property T when $n = 2, 3$ [Grunewald and Lubotzky 2009].

One very interesting partial result was proved in [Gamburd and Pak 2006]: Question 9 has a positive answer for the family of groups $\mathrm{PSL}_2(\mathbb{F}_p)$ under the hypothesis that $\mathrm{PSL}_2(\mathbb{F}_p)$ has *uniform expansion*; that is, there is a constant $\epsilon > 0$ such that, for all p, and all Cayley graphs of $\mathrm{PSL}_2(\mathbb{F}_p)$ with a specified number of generators, the spectral gap is at least ϵ. A very recent theorem of Breuillard and Gamburd shows that there is a family P of primes of density 1 such that $\mathrm{PSL}_2(\mathbb{F}_p)$ has uniform expansion as p ranges over P, but the general case remains mysterious.

It's important to keep in mind that, for all our talk about superstrong approximation, even the prior question of strong approximation is poorly understood in the nonabelian case. Recall that strong approximation has to do with the

transitivity of the action on the set of epimorphisms. There are many (G, n) for which we know that Aut F_n acts transitively on $\mathrm{Epi}(F^n, G)$. But there are many examples of (G, n) such that the action is nontransitive; see [Pak 2001; Lubotzky 2011] for a full discussion of what is known and what is conjectured about this problem.

So far, no thin groups have appeared in this discussion. It is not even precisely clear how the notion of "thin subgroup" of Aut F_n should be defined. Still, we can imagine what kind of questions might be asked in this setting. For instance, one could ask about the following strengthening of Question 9.

Question 10. Let Γ be a finitely generated subgroup of Aut F_n whose image in $\mathrm{Sym}\big(\mathrm{Epi}(F^n, \mathrm{PSL}_2(\mathbb{F}_p))\big)$ is the same as that of Aut F_n for all p. Do the Cayley–Schreier graphs of the action of Γ on $\mathrm{Epi}(F^n, \mathrm{PSL}_2(\mathbb{F}_p))$ form an expander family?

It goes without saying that it would be of interest to ask all of the above questions with Aut F_n replaced by the braid group on n strands — or, for that matter, the mapping class group of any oriented surface whose fundamental group is isomorphic to F_n. This is more in the spirit of monodromy, since mapping class groups, unlike automorphism groups of free groups, are fundamental groups of natural moduli spaces — namely, moduli spaces of punctured curves.

4.2. ***Square-tiled surfaces and approximation for*** **Aut F_2.** The action of

$$\text{Aut } F_2 \quad \text{on } \mathrm{Epi}(F_2, G)$$

also appears in the theory of *square-tiled surfaces*. A square-tiled surface is just a connected degree-d cover of a torus branched at a single point. Such a cover can be thought of as a map from $\pi_{1,1}$ to S_d, where $\pi_{1,1}$ is the fundamental group of the once-punctured torus; it is a free group on 2 generators.

There is a natural obstruction to transitivity of the action of Aut F_2 on $\mathrm{Hom}(F_2, S_d)$; namely, the conjugacy class of the puncture on the torus (alternately: the commutator of a pair of free generators of F_2) is preserved up to conjugacy by Aut F_2. Since we have not demanded that our homomorphisms be surjective, we can break down $\mathrm{Hom}(F_2, S_d)$ further according to the image of the homomorphism.

Precisely: Let μ be a conjugacy class of S_d, that is, a partition of d, and let G be a transitive permutation group of rank d (i.e., a transitive subgroups of S_d defined up to conjugacy.) Then we denote by $Hom_\mu(F_2, S_d, G)$ the set of homomorphisms from F_2 to S_d which send the puncture class to μ and whose image is G. This set also carries an action of S_d, by conjugation on the right, and we denote the set of orbits of this action by $H^1_\mu(F_2, S_d, G)$.

Now Aut F_2 acts on $H^1_\mu(F_2, S_d, G)$, and in fact this action factors through its

quotient Out F_2. We denote the associated Cayley–Schreier graph by $X(d, \mu, G)$, and we can now ask questions of the same genre we have considered before: for which (d, μ) is $X(d, \mu, G)$ connected? For which sequences of (d, μ, G) do the $X(d, \mu, G)$ form an expander family?

We remark that Out F_2 is naturally identified with $\mathrm{GL}_2(\mathbb{Z})$ via its action on $F_2^{\mathrm{ab}} = \mathbb{Z}^2$, and its index 2 subgroup is in turn naturally identified with the mapping class group $\Gamma_{1,1}$ of the once-punctured torus, which is the fundamental group of $\mathcal{M}_{1,1}$, the moduli space of genus-1 curves with a single marked point. Under these identifications, the action of $\Gamma_{1,1}$ on the finite set $H^1_\mu(F_2, S_d, G)$ corresponds to a finite cover of $\mathcal{M}_{1,1}$; this is precisely a moduli space $\mathcal{H}_{d,\mu,G}$ of square-tiled surfaces, where the map $\mathcal{H}_{d,\mu,G} \to \mathcal{M}_{1,1}$ sends a square-tiled surface to the punctured torus which it covers.

The moduli space $\mathcal{H}_{d,\mu,G}$ is a *Hurwitz space* and is also an example of a *Teichmüller curve*. Their geometry and their relationship with moduli spaces of higher genus curves has been much studied — see, for example, [McMullen 2005; Chen 2011]. Any superstrong approximation result for a family of $X(d, \mu, G)$ could in principle be used to give lower bounds for the gonality of $\mathcal{H}_{d,\mu,G}$ along the lines explained in the previous section. And of course one could go further and ask about approximation results for infinite-index subgroups of $\Gamma_{1,1}$ (though it is not clear at first glance whether such results have interesting implications in this context.)

We note that it is not possible that only finitely many $\mathcal{H}_{d,\mu,G}$ have a component with low gonality, for the simple reason that every algebraic curve over a number field is birational to a connected component of $\mathcal{H}_{d,\mu,G}$ for some G [Ellenberg and McReynolds 2012]. So any general statement of growing gonality would have to be restricted to a special family of (d, μ, G).

5. Speculations: superstrong approximation for Galois groups and the Bogomolov property

In Section 3 we explained how results on superstrong approximation could provide lower bounds for gonality in a family of covers of algebraic curves over characteristic 0 fields. The group to which the approximation results were applied was the image of the topological monodromy group of the base (a free group) in the arithmetic lattice $\mathrm{Sp}_{2g}(\mathbb{Z})$.

If one tries to do the same thing over a more general base — for instance, a base in characteristic p — one immediately encounters the problem that there is no longer a *discrete* fundamental group; all that remains is the étale fundamental group, a topologically finitely generated profinite group.

What might superstrong approximation mean for a profinite, rather than a discrete group? Let Π be a profinite group which is topologically finitely generated. If $\pi_1, \dots, \pi_k$ are topological generators, the subgroup Γ they generate is dense in Π. We say Γ has property τ with respect to Π if the Cayley graphs induced by Γ on the finite quotients of Π form an expander family. We might wish to say that Π satisfies superstrong approximation if the dense subgroup Γ has property τ with respect to Π. But there's a problem; it follows from a result of Alon, Lubotzky and Wigderson [Alon et al. 2001] show that there are profinite groups Π and dense finitely generated subgroups $\Gamma, \Gamma' \subset \Pi$ such that Γ has property τ with respect to Π, but Γ' does not. Kassabov (personal communication) has shown that this can hold even when the profinite completions of Γ and Γ' are actually *isomorphic* to Π.

One can instead place on Π the stronger condition that *every* dense finitely generated subgroup Γ has property τ with respect to Π.[3] In this case we say that Π has property $\hat{\tau}$. This is very close to asking that the Cayley graphs attached to generating sets of size k of finite quotients Q of Π satisfy a spectral gap which is independent of Q. In other words, Π needs to have *uniform expansion* as discussed in the previous section. Uniform expansion is for the moment out of reach, but it is widely believed to hold, for example, when $\Pi = \mathrm{Sp}_{2g}(\hat{\mathbb{Z}})$, the relevant case for the current discussion.

Suppose $\mathrm{Sp}_{2g}(\hat{\mathbb{Z}})$ has property $\hat{\tau}$, let U be a complex algebraic curve, and let $\rho : \pi_1(U(\mathbb{C}), u) \to \mathrm{Sp}_{2g}(\hat{\mathbb{Z}})$ be a representation of the fundamental group of U. For each l, the reduction $\bar{\rho}_l : \pi_1(U(\mathbb{C}), u) \to \mathrm{Sp}_{2g}(\mathbb{Z}/l\mathbb{Z})$ gives rise to a cover U_l of U, and the arguments of the previous section show that the U_l grow in gonality.

One might ask the same question for a curve U over a finite field $\mathbb{F}_q$. We slightly modify the question in order to avoid introducing the machinery of the étale fundamental group. Suppose given, for an infinite sequence of primes l, an étale cover $U_l \to U$ such that $(U_l)_{\overline{\mathbb{F}}_q} \to U_{\overline{\mathbb{F}}_q}$ is a Galois cover with Galois group $\mathrm{Sp}_{2g}(\mathbb{Z}/l\mathbb{Z})$. For instance, you might make such a cover by adjoining the l-torsion points of an abelian g-fold A over U.

Question 11. Is the gonality of U_l bounded below by a positive constant multiple of $[U_l : U]$? What if we impose the extra condition that the U_l are obtained by adjoining the l-torsion points on a fixed abelian g-fold A/U?

We have focused on the case of curves over finite fields because of an intriguing connection with a different topic in number theory, which we now explain. Write K_l for the function field of U_l. To say U_l has gonality greater than n is to say

[3] There is a natural temptation to call this criterion *superduperstrong approximation*, but this terminology is perhaps better kept quarantined in a footnote.

that there is no rational function f on U_l with at most n poles and n zeroes. This is equivalent to the assertion that there is no nonconstant element f of K_l^* whose (logarithmic) Weil height $h_{K_l}(f)$ is at most n. The *absolute height* of an element f of an extension of K_l is just $h_{K_l}(f)/[U_l : U]$. In other words, Question 11 asks whether the *absolute* height of an element of K_l^* is bounded below by a constant if it is nonzero.

This condition already has a name — it is called the *Bogomolov property*.

Definition 12. Let K be a global field. A set S of algebraic numbers over K has the *Bogomolov property* if the set of positive absolute heights of elements of S is bounded away from 0.

Note that this definition makes sense not only when K is a function field (as in the case described above) but also when K is a number field; in fact, it is the case $K = \mathbb{Q}$ that has been most energetically studied. Some sets known to have the Bogomolov property over $\mathbb{Q}$ are the field of totally real algebraic numbers (Schinzel), the maximal abelian extension of $\mathbb{Q}$ (Amoroso–Zannier) and the maximal extension of $\mathbb{Q}$ totally split at a finite place p (Bombieri–Zannier). The recent paper of Amoroso, David, and Zannier [Amoroso et al. 2011] contains a good summary of what is known along with several striking new results.

Speculation piled atop speculation: If we believe that $\mathrm{Sp}_{2g}(\mathbb{F}_l)$ has uniform expansion, and if we believe by analogy with the complex case that uniform expansion of Galois groups implies the Bogomolov property for the corresponding extensions of function fields, might this be the case for number fields as well? There are many ways one could analogize Question 11 to the number field case. Here is a minimalist version:

Question 13. Let $A/\mathbb{Q}$ be an abelian variety and let K_l be the number field $\mathbb{Q}(A[l])$. Does the compositum of all the K_l have the Bogomolov property?

And here is a maximalist version:

Question 14. Let F/Q be an infinite Galois extension such that $\mathrm{Gal}(F\mathbb{Q}^{\mathrm{ab}}/\mathbb{Q}^{\mathrm{ab}})$ has property $\hat{\tau}$. Does F have the Bogomolov property?

A striking result of Habegger [2013] shows that the answer to Question 13 is affirmative when A is an elliptic curve. The theorems from [Ellenberg et al. 2012] described in Section 3 can be used to show that the analogue of Question 13 with $\mathbb{Q}$ replaced by $\mathbb{C}(t)$ has an affirmative answer as well; this requires lower bounds for the gonality, not only of the covers U_l, but of fiber products of multiple U_{l_i} over U. In the case of $\mathbb{C}(t)$, the relevant expansion theorem (which was not available when [Ellenberg et al. 2012] was written) came with [Salehi Golsefidy and Varjú 2012]. It shows that expansion holds not only for the Cayley graphs

associated to reduction modulo primes, but also modulo arbitrary squarefree integers.

As for the maximalist question, its interest is somewhat diminished by our lack of examples of groups known to have property $\hat{\tau}$. This need not keep us from speculating. For instance, S be a finite set of primes, and let K_S^t be the maximal extension of $\mathbb{Q}$ which is unramified outside S and *tamely* ramified at S, and let G_S^t be its Galois group. Then G_S^t has property *FAb*, which is to say that every finite-index subgroup of G_S^t has finite abelianization. This makes G_S^t a reasonable candidate for having some form of spectral expansion property. Caution is required, since G_S^t is only conjectured to be topologically finitely generated. One can circumvent this problem as follows. For each integer k, we can consider the family of all Cayley graphs on finite quotients G of G_S^t endowed with subsets $(g_1, \ldots, g_k)$ generating G. We say that G_S^t has $\hat{\tau}_k$ if there is a uniform spectral gap for all such Cayley graphs.

Question 15. Does G_S^t have property $\hat{\tau}_k$ for all k? Does K_S^t have the Bogomolov property?

Note that for G_S^t to have property $\hat{\tau}_2$ would require that there are only finitely many n such that G_S^t has a quotient isomorphic to S_n. It is not even known whether there are arbitrarily large S_n-extensions of $\mathbb{Q}$ unramified away from S; a recent preprint of Roberts and Venkatesh [$\geq$ 2012] argues that such extensions might indeed exist and proposes a construction, but at the same time suggests that these extensions are *not* likely to be tamely ramified at the primes in S. Another hint that the first question is not completely unreasonable comes from Ershov's work in [2011] on Golod–Shafarevich groups with infinite quotients with property T. But we should emphasize that we have no strong opinion as to the answer to either question.

Acknowledgements

This document is an expanded version of a lecture presented by the author at a conference on "Thin Groups and Superstrong Approximation" held at the Mathematical Sciences Research Institute in February 2012. The work described here was partially supported by NSF Grant DMS-1101267 and a Romnes Faculty Fellowship. The author is grateful to Ian Agol, Frank Calegari, and Peter Sarnak for useful conversations about the mathematics discussed here, and to MathOverflow for useful input concerning the image of the point-pushing subgroup in the symplectic group and the meaning of superstrong approximation for Galois groups.

References

[Abramovich 1996] D. Abramovich, "A linear lower bound on the gonality of modular curves", *Int. Math. Res. Not.* **1996**:20 (1996), 1005–1011.

[Abramovich and Voloch 1996] D. Abramovich and J. F. Voloch, "Lang's conjectures, fibered powers, and uniformity", *New York J. Math.* **2** (1996), 20–34.

[A'Campo 1979] N. A'Campo, "Tresses, monodromie et le groupe symplectique", *Comment. Math. Helv.* **54**:2 (1979), 318–327.

[Alon et al. 2001] N. Alon, A. Lubotzky, and A. Wigderson, "Semi-direct product in groups and zig-zag product in graphs: Connections and applications", pp. 630–637 in *42nd IEEE Symposium on Foundations of Computer Science* (Las Vegas, NV, 2001), IEEE Computer Soc., Los Alamitos, CA, 2001.

[Amoroso et al. 2011] F. Amoroso, S. David, and U. Zannier, "On fields with the property (b)", preprint, 2011, http://hal.archives-ouvertes.fr/docs/00/64/99/54/PDF/FieldsB.pdf. To appear in *Proc. Amer. Math. Soc.*

[Brav and Thomas 2012] C. Brav and H. Thomas, "Thin monodromy in Sp(4)", preprint, 2012. arXiv 1210.0523

[Breuillard et al. 2011] E. Breuillard, B. Green, and T. Tao, "Approximate subgroups of linear groups", *Geom. Funct. Anal.* **21**:4 (2011), 774–819.

[Cadoret 2012] A. Cadoret, "Note on the gonality of abstract modular curves", pp. 89–106 in *The arithmetic of fundamental groups*, edited by J. Stix, Contributions in Mathematical and Computational Sciences **2**, 2012.

[Cadoret and Tamagawa 2009] A. Cadoret and A. Tamagawa, "A uniform open image theorem for l-adic representations, II", preprint, 2009, http://www.math.polytechnique.fr/∼cadoret/uoi2_Final.pdf.

[Chen 2011] D. Chen, "Square-tiled surfaces and rigid curves on moduli spaces", *Adv. Math.* **228**:2 (2011), 1135–1162.

[Ellenberg and McReynolds 2012] J. S. Ellenberg and D. B. McReynolds, "Arithmetic Veech sublattices of SL(2, **Z**)", *Duke Math. J.* **161**:3 (2012), 415–429.

[Ellenberg et al. 2012] J. S. Ellenberg, C. Hall, and E. Kowalski, "Expander graphs, gonality, and variation of Galois representations", *Duke Math. J.* **161**:7 (2012), 1233–1275.

[Ershov 2011] M. Ershov, "Kazhdan quotients of Golod–Shafarevich groups", *Proc. Lond. Math. Soc.* (3) **102**:4 (2011), 599–636.

[Faltings 1983] G. Faltings, "Endlichkeitssätze für abelsche Varietäten über Zahlkörpern", *Invent. Math.* **73**:3 (1983), 349–366.

[Farb and Margalit 2011] B. Farb and D. Margalit, *A primer on mapping class groups*, Princeton Mathematical Series **49**, Princeton University Press, 2011.

[Frey 1994] G. Frey, "Curves with infinitely many points of fixed degree", *Israel J. Math.* **85**:1-3 (1994), 79–83.

[Fuchs 2014] E. Fuchs, "The ubiquity of thin groups", pp. 73–92 in *Thin groups and superstrong approximation*, edited by H. Oh and E. Breuillard, Math. Sci. Res. Inst. Publ. **61**, Cambridge, New York, 2014.

[Gamburd and Pak 2006] A. Gamburd and I. Pak, "Expansion of product replacement graphs", *Combinatorica* **26**:4 (2006), 411–429.

[Grunewald and Lubotzky 2009] F. Grunewald and A. Lubotzky, "Linear representations of the automorphism group of a free group", *Geom. Funct. Anal.* **18**:5 (2009), 1564–1608.

[Habegger 2013] P. Habegger, "Small height and infinite non-abelian extensions", preprint, 2013. arXiv 1109.5859

[Hall 2008] C. Hall, "Big symplectic or orthogonal monodromy modulo l", *Duke Math. J.* **141**:1 (2008), 179–203.

[Helfgott 2008] H. A. Helfgott, "Growth and generation in $\mathrm{SL}_2(\mathbf{Z}p\mathbf{Z})$", *Ann. of Math.* (2) **167**:2 (2008), 601–623.

[Katz 2004] N. M. Katz, "Larsen's alternative, moments, and the monodromy of Lefschetz pencils", pp. 521–560 in *Contributions to automorphic forms, geometry, and number theory* (Baltimore, MD, 2002), edited by H. Hida et al., Johns Hopkins Univ. Press, Baltimore, MD, 2004.

[Kowalski 2014] E. Kowalski, "Sieves in discrete groups, especially sparse", pp. 107–140 in *Thin groups and superstrong approximation*, edited by H. Oh and E. Breuillard, Math. Sci. Res. Inst. Publ. **61**, Cambridge, New York, 2014.

[Li and Yau 1982] P. Li and S. T. Yau, "A new conformal invariant and its applications to the Willmore conjecture and the first eigenvalue of compact surfaces", *Invent. Math.* **69**:2 (1982), 269–291.

[Lubotzky 2011] A. Lubotzky, "Dynamics of $\mathrm{Aut}(F_N)$ actions on group presentations and representations", pp. 609–643 in *Geometry, rigidity, and group actions*, edited by B. Farb and D. Fisher, Univ. Chicago Press, Chicago, IL, 2011.

[Lubotzky and Pak 2001] A. Lubotzky and I. Pak, "The product replacement algorithm and Kazhdan's property (T)", *J. Amer. Math. Soc.* **14**:2 (2001), 347–363.

[McMullen 2005] C. T. McMullen, "Teichmüller curves in genus two: Discriminant and spin", *Math. Ann.* **333**:1 (2005), 87–130.

[McMullen 2013] C. T. McMullen, "Braid groups and Hodge theory", *Math. Ann.* **355**:3 (2013), 893–946.

[Pak 2000] I. Pak, "The product replacement algorithm is polynomial", pp. 476–485 in *41st Annual Symposium on Foundations of Computer Science* (Redondo Beach, CA, 2000), IEEE Comput. Soc. Press, Los Alamitos, CA, 2000.

[Pak 2001] I. Pak, "What do we know about the product replacement algorithm?", pp. 301–347 in *Groups and computation, III* (Columbus, OH, 1999), edited by W. M. Kantor and Á. Seress, Ohio State Univ. Math. Res. Inst. Publ. **8**, de Gruyter, Berlin, 2001.

[Pyber and Szabó 2011] L. Pyber and E. Szabó, "Growth in finite simple groups of lie type of bounded rank", preprint, 2011. arXiv 1005.1858

[Pyber and Szabó 2014] L. Pyber and E. Szabó, "Growth in linear groups", pp. 253–268 in *Thin groups and superstrong approximation*, edited by H. Oh and E. Breuillard, Math. Sci. Res. Inst. Publ. **61**, Cambridge, New York, 2014.

[Rapinchuk 2014] A. Rapinchuk, "Strong approximation for algebraic groups", pp. 269–298 in *Thin groups and superstrong approximation*, edited by H. Oh and E. Breuillard, Math. Sci. Res. Inst. Publ. **61**, Cambridge, New York, 2014.

[Roberts and Venkatesh $\geq$ 2012] D. P. Roberts and A. Venkatesh, "Hurwitz number fields", in preparation; see http://cda.morris.umn.edu/∼roberts/research/HNF.pdf for slides of a related lecture.

[Salehi Golsefidy 2014] A. Salehi Golsefidy, "Affine sieve and expanders", pp. 325–342 in *Thin groups and superstrong approximation*, edited by H. Oh and E. Breuillard, Math. Sci. Res. Inst. Publ. **61**, Cambridge, New York, 2014.

[Salehi Golsefidy and Varjú 2012] A. Salehi Golsefidy and P. P. Varjú, "Expansion in perfect groups", *Geom. Funct. Anal.* **22**:6 (2012), 1832–1891.

[Sarnak 2014] P. Sarnak, "Notes on thin matrix groups", pp. 343–362 in *Thin groups and superstrong approximation*, edited by H. Oh and E. Breuillard, Math. Sci. Res. Inst. Publ. **61**, Cambridge, New York, 2014.

[Venkataramana 2013] T. N. Venkataramana, "Monodromy of cyclic coverings of the projective line", preprint, 2013. To appear in *Invent. Math.* arXiv 1204.4778

[Yu 1997] J.-K. Yu, "Toward a proof of the Cohen–Lenstra conjecture in the function field case", preprint, 1997.

[Zograf 1984] P. G. Zograf, "Small eigenvalues of automorphic Laplacians in spaces of cusp forms", *Zap. Nauchn. Sem. Leningrad. Otdel. Mat. Inst. Steklov.* (*LOMI*) **134** (1984), 157–168. In Russian.

ellenber@math.wisc.edu *Department of Mathematics, University of Wisconsin, 480 Lincoln Drive, Madison, WI 53706, United States*