

Expander graphs based on GRH with an application to elliptic curve cryptography

David Jao*, Stephen D. Miller† and Ramarathnam Venkatesan

January 14, 2008

Abstract

We present a construction of expander graphs obtained from Cayley graphs of narrow ray class groups, whose eigenvalue bounds follow from the Generalized Riemann Hypothesis. Our result implies that the Cayley graph of $(\mathbb{Z}/q\mathbb{Z})^*$ with respect to small prime generators is an expander. As another application, we show that the graph of small prime degree isogenies between ordinary elliptic curves achieves non-negligible eigenvalue separation, and explain the relationship between the expansion properties of these graphs and the security of the elliptic curve discrete logarithm problem.

1 Introduction

Expander graphs are widely studied in many areas of mathematics and theoretical computer science, and such graphs are useful primarily because random walks along their edges quickly become uniformly distributed over their vertices. Several beautiful constructions of expanders have been based on deep tools from representation theory and arithmetic, for example Kazhdan's Property (T) [38] and the Ramanujan conjectures [35, 39].

The main contribution of this paper is a new, conditional construction of expanders based on the Generalized Riemann Hypothesis (GRH), which

*Partially supported by NSERC Discovery Grant #341769-07

†Partially supported by NSF grant DMS-0601009 and an Alfred P. Sloan Foundation Fellowship

arises naturally in the study of the elliptic curve discrete logarithm problem. This cryptographic connection is investigated in our parallel paper [25], where it is used to establish that the discrete logarithm problem has roughly uniform difficulty for equal sized curves. The present paper contains a generalization of the main theorem in that paper, along with explanations and applications of a more mathematical nature.

We briefly review some notions from graph theory, including that of *expander graph* from above. By an *undirected graph* $\Gamma = (\mathcal{V}, \mathcal{E})$ we mean a set of vertices \mathcal{V} and (unoriented) edges \mathcal{E} connecting specified pairs of vertices. Suppose that the graph is finite and is furthermore *k-regular*, meaning that there are exactly k edges incident to each vertex. The *adjacency operator* A acts on functions on \mathcal{V} by averaging them over neighbors:

$$(Af)(x) = \sum_{x \text{ and } y \text{ connected by an edge}} f(y). \quad (1.1)$$

Since the graph is regular, the constant function $\mathbb{1}(x) = 1$ is an eigenfunction of A with eigenvalue k , which is accordingly termed the *trivial eigenvalue* λ_{triv} of A . It is straightforward to see that the multiplicity of λ_{triv} is equal to the number of connected components of the graph, and that λ_{triv} is the largest eigenvalue of A in absolute value. An *expander* graph is a graph for which the nontrivial eigenvalues satisfy the bound

$$\lambda \leq \lambda_{\text{triv}} (1 - \delta) \quad \text{for some fixed constant } \delta > 0. \quad (1.2)$$

If the nontrivial eigenvalues further satisfy the stronger bound

$$|\lambda| \leq \lambda_{\text{triv}} (1 - \delta), \quad (1.3)$$

then a standard lemma (e.g. Lemma 2.1) shows that random walks of length $\frac{1}{\delta} \log 2|\mathcal{V}|$ are equidistributed in the sense that they land in arbitrary subsets of \mathcal{V} with probability at least proportional to their size. This rapid mixing of the random walk is at the heart of most, if not nearly all, applications of expanders.

A group G generated by a subset $S = S^{-1}$ can be made into the vertices of a *Cayley graph* $\text{Cay}(G, S)$ by defining edges from g to sg , for each $s \in S$ and $g \in G$.¹ For finite abelian groups, the eigenfunctions of A are precisely

¹Note that all graphs in this paper are undirected. We also allow for multiple edges by letting S be a multiset when necessary, such as in the statement of Theorem 1.1.

the characters $\chi : G \rightarrow \mathbb{C}^*$; indeed, the formula

$$(A\chi)(x) = \sum_{s \in S} \chi(sx) = \lambda_\chi \chi(x), \quad \text{where } \lambda_\chi = \sum_{s \in S} \chi(s), \quad (1.4)$$

shows that the spectrum consists of character sums ranging over the generating set. The trivial eigenvalue $\lambda_{\text{triv}} = |S|$ of course comes from the trivial character $\chi = \mathbb{1}$, and inequality (1.3) is satisfied if the character sums for λ_χ , $\chi \not\equiv \mathbb{1}$, have enough cancellation. Abelian Cayley graphs are a restricted yet important type of graph, and their expansion properties have been well studied (e.g. [2, 36]). To be expanders, they cannot have bounded degree but must have at least $\Omega(\log |G|)$ generators.

The expander graphs produced by our construction are abelian Cayley graphs, and we give eigenvalue bounds for their character sums λ_χ using GRH. Before stating the construction, we briefly recall some terminology. For any integral ideal \mathfrak{m} in a number field K , let $I_{\mathfrak{m}}$ denote the group of fractional ideals relatively prime to \mathfrak{m} (i.e. those whose factorization into prime ideals contains no divisor of \mathfrak{m}). Let $P_{\mathfrak{m}}$ denote the principal ideals generated by an element $k \in K^*$ such that $k \equiv 1 \pmod{\mathfrak{m}}$, and let $P_{\mathfrak{m}}^+ \subset P_{\mathfrak{m}}$ denote those generated by such an element k which is furthermore totally positive (i.e. positive in all embeddings $K \hookrightarrow \mathbb{R}$). The quotients $I_{\mathfrak{m}}/P_{\mathfrak{m}}$ and $I_{\mathfrak{m}}/P_{\mathfrak{m}}^+$ are called, respectively, the ray and narrow ray class groups of K relative to \mathfrak{m} .

Theorem 1.1. (*“GRH Graphs”*). *Let K be a number field of degree n , \mathfrak{m} an integral ideal, and G the narrow ray class group of K relative to \mathfrak{m} . Let $q = D \cdot N\mathfrak{m}$, where D is the discriminant of K and $N\mathfrak{m}$ denotes the norm of \mathfrak{m} . Consider the set $\{\text{prime ideals } \mathfrak{p} \text{ coprime to } \mathfrak{m} \mid N\mathfrak{p} \leq x \text{ is prime}\}$, and let S_x denote the multiset consisting of its image and inverse in G (i.e., including multiplicities). Then assuming GRH for the characters of G , the graph $\Gamma_x = \text{Cay}(G, S_x)$ has*

$$\lambda_{\text{triv}} = 2 \operatorname{li}(x) + O(n x^{1/2} \log(xq)), \quad \operatorname{li}(x) = \int_2^x \frac{dt}{\log t}, \quad (1.5)$$

while the nontrivial eigenvalues λ obey the bound

$$|\lambda| = O(n x^{1/2} \log(xq)). \quad (1.6)$$

In particular, if $B > 2$ and $x \geq (\log q)^B$,

$$|\lambda| = O((\lambda_{\text{triv}} \log \lambda_{\text{triv}})^{1/2+1/B}). \quad (1.7)$$

The implied constants in (1.5) and (1.6) are absolute, while the one in (1.7) depends only on B and n .

Remark 1.2. a) The Theorem immediately applies to quotients of narrow ray class groups, such as ray class groups themselves. This is because the spectrum of the quotient Cayley graph consists of eigenvalues for those characters which factor through the quotient.

b) The parameter q should be thought of as large, in light of Minkowski's theorem that there are only a finite number of number fields with a given discriminant [32, p. 121].

c) The above bound on the spectral gap is worse than that for Ramanujan graphs [35], and thus abelian graphs are not optimal in this sense (see also [2, 36]). However, one gains explicit constructions that are simpler computationally; additionally, there are situations where these graphs occur naturally and the expansion bounds are helpful, as our following examples show.

From the abovementioned relationship between expander graphs and rapid mixing of random walks, we obtain the following application.

Corollary 1.3. *Fix $B > 2$ and $n \geq 1$, and assume the same hypotheses of the previous theorem, including the choice of $x \geq (\log q)^B$. Then there exists a positive constant C with the following property: for q sufficiently large, a random walk of length*

$$t \geq C \frac{\log |G|}{\log \log q}$$

from any starting vertex lands in any fixed subset $S \subset G$ with probability at least $\frac{1}{2} \frac{|S|}{|G|}$.

Let us now illustrate the theorem with a few examples. The first example is the field $K = \mathbb{Q}$, whose narrow ray class groups are of the form $(\mathbb{Z}/q\mathbb{Z})^*$, for $q > 1$. In this case the edges of the Cayley graph connect each vertex $v \in (\mathbb{Z}/q\mathbb{Z})^*$ to pv and $p^{-1}v \pmod{q}$, for all primes p such that $p \leq (\log q)^{2+\delta}$ and $p \nmid q$. Starting from any v and taking random steps of this form results in a uniformly distributed random element of $(\mathbb{Z}/q\mathbb{Z})^*$ in $O((\log q)/\log \log q)$ steps. The character sum (1.4) for λ_χ here amounts to the sum $2 \operatorname{Re} \sum_{p \leq (\log q)^{2+\delta}} \chi(p)$, so bounds on λ_χ yield statements about the distribution of small primes in residue classes modulo q . GRH, which is used in (1.7), is a natural tool for such problems. It seems difficult to obtain an unconditional result along these lines, because the special case when χ

is a quadratic character modulo q is related to the problem of estimating the smallest prime quadratic nonresidue modulo q . Finding such a prime is equivalent to obtaining any cancellation at all in the sum $\sum_{p \leq x} (\frac{p}{q})$, and even this problem seems to require a strong hypothesis such as GRH. However, it is possible to use the Large Sieve to prove unconditional results for typical values of q , such as [17, Theorem 3], which shows that $\frac{\lambda_x}{\lambda_{\text{triv}}}$ goes to zero outside of a sparse subset of moduli q .

The next example, when K is an imaginary quadratic number field, is related to elliptic curves over finite fields. Using the correspondence between ordinary elliptic curves and ideal classes in orders of imaginary quadratic number fields, we prove the following theorem.

Definition 1.4. We say that two ordinary elliptic curves E_1, E_2 defined over \mathbb{F}_q have the *same level* if their rings of endomorphisms $\text{End}(E_i)$ are isomorphic. (In this paper, we follow the standard convention that $\text{End}(E)$ refers to $\bar{\mathbb{F}}_q$ -endomorphisms.)

Theorem 1.5. *Consider the set $S_{N,q}$ of $\bar{\mathbb{F}}_q$ -isomorphism classes of ordinary elliptic curves defined over \mathbb{F}_q having N points. Fix² an $E \in S_{N,q}$ and let \mathcal{V} be the set of all curves in $S_{N,q}$ having the same level as E . Form a graph on the set of vertices \mathcal{V} by connecting curves E_1 and E_2 with an edge if there exists an isogeny of prime degree less than $(\log 4q)^B$ between them, for some fixed $B > 2$. Then, assuming GRH, this graph is an expander graph in the sense that its nontrivial eigenvalues satisfy the bound (1.7).*

Theorem 1.5 has implications for the security of the elliptic curve discrete logarithm problem. Recall that the discrete logarithm problem (DLOG) asks to recover the exponent a of a power g^a of a known element g . Its presumed difficulty serves as the basis of several cryptosystems, for example the Diffie-Hellman key exchange. Though many difficult problems in computer science are only hard in rare instances, good cryptosystems typically must be based on problems which are almost always hard. We recall that the DLOG problem on a given group has *random self-reducibility*: that means given an algorithm $\mathcal{A}(g^a) = a$ which solves DLOG on, say, half of all input values y , we may easily find a random value of r such that \mathcal{A} works on $y' = g^r y$, and deduce that $\mathcal{A}(y) = \mathcal{A}(y') - r$. Therefore, if DLOG is hard for some values of y , it must

²We will frequently treat the elements of $S_{N,q}$ as curves, though strictly speaking they are isomorphism classes of curves. This distinction does not affect Theorem 1.6 because isomorphisms between curves in $S_{N,q}$ can be computed in time polynomial in $\log q$.

be hard for almost all values. Though this result says nothing about the absolute difficulty of the problem, it is a comforting assurance regarding the relative difficulty of multiple instances of the problem.

Elliptic curve cryptography [7, 28, 41] is based on the conjectured difficulty of DLOG problems within the group of points of an elliptic curve over a finite field. At present, cryptographers typically select elliptic curves in the following way: a large finite field \mathbb{F}_q is selected, and an elliptic curve E/\mathbb{F}_q is generated at random. Its order $\#E(\mathbb{F}_q)$ is quickly computed [13, 45], and the curve is discarded unless the order has a large prime factor (because otherwise DLOG is much easier). It is also checked from the point count whether or not E is supersingular or has other weaknesses, and if it is then the curve is discarded.³ The above practice efficiently yields elliptic curves thought to be suitable for cryptographic purposes. An obvious question is whether or not other considerations are important, i.e. whether the point count is the only factor influencing the difficulty of DLOG on an elliptic curve over a fixed finite field.

In studying this question, the random self-reducibility fact from above does not apply, because it pertains only to a single curve, and says nothing about the comparative difficulty of DLOG between two different curves. However, we can instead use the fact that an efficiently computable isogeny provides a reduction of the DLOG problems between two curves. Furthermore, a theorem of Tate [49] states that all curves of cardinality N defined over \mathbb{F}_q are isogenous, but unfortunately not all isogenies are efficiently computable, so the theorem does not immediately imply that all curves in $S_{N,q}$ have equivalent DLOG problems. On the other hand, isogenies of low degree are efficiently computable, and the rapid mixing in Theorem 1.5 says that their random compositions become uniformly distributed over curves within each level in $S_{N,q}$. This property allows us to establish that the difficulty of the elliptic curve DLOG problem is in a sense uniform over any given level. More precisely:

Theorem 1.6. *With the hypotheses of Theorem 1.5, assume there is an algorithm \mathcal{A} which solves the discrete logarithm problem on a positive fraction*

³Supersingular curves are thought to be cryptographically weaker, because of the existence of subexponential attacks on their DLOG problems [40]. This is not to say that no subexponential attacks exist for ordinary curves; in fact, some are known to succeed on a very modest proportion of them [16, 47], and of course other unknown ones may yet be discovered. The supersingular analog of Theorems 1.5 and 1.6 are given in [25, Appendix].

μ of the elliptic curves in a given level. There exists an absolute polynomial $p(x)$ such that one can probabilistically solve the discrete logarithm problem on any curve in the same level with expected runtime $\frac{1}{\mu}p(\log q)$ times the maximal runtime of \mathcal{A} .

In practice, the level restriction in Theorem 1.6 is actually irrelevant. Indeed, if two curves in $S_{N,q}$ are not of the same level, then their levels must differ at either a small prime or a large prime.⁴ In the small prime case, we can still obtain DLOG reductions using low degree isogenies (cf. Section 5), and in the large prime case, no constructible examples of such pairs of curves are known. Several interesting theoretical questions remain concerning the large prime case and the true value of the isogeny degrees needed to achieve expansion. We describe some open problems in Section 7.

2 Expander Graphs

In this section we recall a standard bound for the mixing time of a random walk on an expander graph, discuss the lack of nontrivial short cycles on the GRH graphs, and prove Theorem 1.1 and Corollary 1.3. We keep the notation and definitions of the introduction.

Lemma 2.1. *Let Γ be a finite k -regular graph for which the nontrivial eigenvalues λ of the adjacency matrix A are bounded by $|\lambda| \leq c$, for some $c < k$. Let S be any subset of the vertices of Γ , and v be any vertex in Γ . Then a random walk of any length at least $\frac{\log 2|\Gamma|/|S|^{1/2}}{\log k/c}$ starting from v will end in S with probability between $\frac{1}{2} \frac{|S|}{|\Gamma|}$ and $\frac{3}{2} \frac{|S|}{|\Gamma|}$.*

Of course the probability range can be significantly narrowed by lengthening the walk, as it turns out even by a slight amount.

Proof. Letting χ_S and $\chi_{\{v\}}$ denote the characteristic functions of the sets S and $\{v\}$, respectively, the number of paths of length t which start at v and end in S is given by the L^2 -inner product $\langle \chi_S, A^t \chi_{\{v\}} \rangle$. Let P denote the projection from $L^2(\Gamma)$ onto the orthogonal complement of the constant

⁴There is possibly an intermediate range, though its existence is somewhat fluid depending on hardware and software developments (see Section 7.1).

functions; the operator A preserves this space and its operator norm on it is bounded by c because of our eigenvalue assumption. Then

$$\langle \chi_S, A^t \chi_{\{v\}} \rangle = \frac{|S|}{|\Gamma|} k^t + \langle P \chi_S, A^t P \chi_{\{v\}} \rangle. \quad (2.1)$$

The latter term is bounded by

$$\begin{aligned} |\langle P \chi_S, A^t P \chi_{\{v\}} \rangle| &\leq \|P \chi_S\| \|A^t P \chi_{\{v\}}\| \leq c^t \|P \chi_S\| \|P \chi_{\{v\}}\| \\ &\leq c^t \|\chi_S\| \|\chi_{\{v\}}\| = c^t |S|^{1/2}. \end{aligned} \quad (2.2)$$

For $t \geq \frac{\log 2|\Gamma|/|S|^{1/2}}{\log k/c}$ this is at most half the size of the main term $k^t |S|/|\Gamma|$ from (2.1), as was to be shown. \square

Next we come to the topic of *girth*, the length of the shortest closed cycle on the graph. Graphs with large girth are important in many applications, for example to the design of collision resistant hash functions and stream ciphers (see, for example, [18, 22, 27]). The girth of a k -regular graph cannot be larger than $2 \log_{k-1} |\Gamma|$. This inequality comes from counting the number of points $b(r)$ in a ball of radius r in a k -regular tree; a graph with girth γ satisfies the inequality $b(\gamma) \leq |\Gamma|$, which gives an upper bound on γ . Random graphs tend to have small girth, but one can use probabilistic methods to show the existence of graphs having girth at least $(1 + o(1)) \log_{k-1} |\Gamma|$, i.e. roughly half the optimal size. The LPS Ramanujan graphs have the largest known girths: $(4/3 + o(1)) \log_{k-1} |\Gamma|$ [5, 35]. It is an open question as to how large the girth can be.

Abelian Cayley graphs cannot have large girth because they have many short cycles of the form $xyx^{-1}y^{-1}$. To rule out these, one can speak of the *nonabelian girth*, which is the shortest cycle not having steps both of the form x^a and x^{-b} for $a, b > 0$ and $x \in S$. We remark that the graphs on $(\mathbb{Z}/q\mathbb{Z})^*$ described just after Theorem 1.1 have

$$\text{nonabelian girth of } \Gamma \geq (1 + o(1)) \log_{k-1} |\Gamma|. \quad (2.3)$$

Indeed, a cycle amounts to two products of small primes which are equal modulo q ; by unique factorization, at least one of these products must be larger than q , which gives a lower bound on the number of factors. This argument also gives the same lower bound for the *odd girth* of Γ (i.e. the shortest closed cycle of odd length), which again is relatively large. It should

be noted, however, that this is not optimal; in fact there are code-based constructions [1, 2] which have nonabelian girth at least $(2 + o(1)) \log_{k-1} |\Gamma|$. The reason for mentioning this, though, is that explicit examples of graphs with large nonabelian girth are important for cryptographic applications.

We conclude this section with the proofs of Theorem 1.1 and Corollary 1.3.

Proof of Theorem 1.1. We explained in (1.4) and in the remarks following it that (1.5) and (1.6) follow from the following estimates for sums of characters χ of G :

$$\begin{aligned} \sum_{N\mathfrak{p} \leq x} (\chi(\mathfrak{p}) + \chi(\mathfrak{p})^{-1}) &= 2 \operatorname{Re} \sum_{N\mathfrak{p} \leq x} \chi(\mathfrak{p}) \\ &= 2r \operatorname{li}(x) + O(n x^{1/2} \log(xq)) , \end{aligned} \tag{2.4}$$

with an absolute implied constant. Here $r = 1$ if χ is the trivial character, and 0 otherwise. Of course, χ may be viewed as a Hecke Grossencharacter on $I_{\mathfrak{m}}$ which is trivial on $P_{\mathfrak{m}}^+$. Hecke proved that its L -function

$$L(s, \chi) = \sum_{\mathfrak{a} \text{ integral ideal}} \chi(\mathfrak{a}) (N\mathfrak{a})^{-s} = \prod_{\mathfrak{p} \text{ prime ideal}} (1 - \chi(\mathfrak{p}) (N\mathfrak{p})^{-s})^{-1} \tag{2.5}$$

analytically continues to a holomorphic function on $\mathbb{C} - \{1\}$ of order 1, with at most a simple pole at $s = 1$ which occurs only when χ is the trivial character. Furthermore, he also established a standard functional equation for its completed L -function, which is a product of $L(s, \chi)$, Γ -factors of the form $\Gamma(\frac{s}{2})$, $\Gamma(\frac{s+1}{2})$, and $\Gamma(s)$, and a power $Q^{s/2}$ of some integer $Q > 0$ [23, p. 211]. The value of Q varies with different characters, but is always bounded above by $q = D \cdot N\mathfrak{m}$. The Dirichlet series coefficients of $L(s, \chi)$, like those of any Artin L -function, satisfy the Ramanujan-Petersson conjecture.

Using these analytic properties, along with the assumption of GRH, one can derive the following standard estimate (which is found in [24, p. 114]):

$$\sum_{N\mathfrak{p} \leq x} \chi(\mathfrak{p}) \log(N\mathfrak{p}) = rx + O(n x^{1/2} \log(x) \log(xq)) \tag{2.6}$$

for primitive characters χ , again with an absolute implied constant. If χ is imprimitive, one must also include terms for prime ideals \mathfrak{p} dividing \mathfrak{m} . There are at most $O(\log N\mathfrak{m}) = O(\log q)$ of these, so both their contribution and the existing error term in (2.6) can be safely absorbed into the enlarged error term $O(nx^{1/2} \log(x) \log(xq))$. This variant of (2.6) in turn implies (2.4) by a simple application of partial summation. \square

Proof of Corollary 1.3. The proof follows from Lemma 2.1 once we have verified that $\log \frac{k}{c}$ is bounded below by a constant (depending on B and n) times $\log \log q$ once q is sufficiently large. Indeed, in our setting the degree is $k = \lambda_{\text{triv}}$, and c may be taken to be the bound in (1.7). For q sufficiently large, $\log \frac{k}{c}$ is indeed bounded below by a constant times $\log \lambda_{\text{triv}} \gg B \log \log q$. \square

Remark 2.2. The main point of the Corollary is to give examples of rapid mixing over large graphs. However, for a finite number of cases when q is small, the graph Γ_x may actually be disconnected. In addition, the equidistribution is not as interesting in situations when the graph Γ_x has relatively few vertices, i.e. when the narrow ray class number of \mathfrak{m} is small. This can be computed explicitly as

$$|G| = |\Gamma_x| = 2^{r_1} \frac{h(K) |(\mathcal{O}_K/\mathfrak{m})^*|}{[U(K) : U_{\mathfrak{m}}(K)]}, \quad (2.7)$$

where r_1 is the number of real embeddings of K , $h(K)$ its class number, \mathcal{O}_K its ring of integers, $U(K)$ its unit group, and $U_{\mathfrak{m}}(K) \subset U(K)$ its subgroup of totally positive units which are congruent to 1 (*mod* \mathfrak{m}) [8, Prop. 3.2.4]. For a fixed degree n , the class number $h(K)$ is $O_{\varepsilon}(|D|^{1/2+\varepsilon})$ for any $\varepsilon > 0$, and so $|G|$ above is bounded by $O(q)$.

3 Elliptic curves

In this section we explain the connection between the GRH graphs and elliptic curves, and prove Theorem 1.5. For ease of presentation, we begin first with the case of elliptic curves defined over complex numbers, and then later explain how our results over complex numbers imply the corresponding results over finite fields.

Let \mathcal{O}_D be an imaginary quadratic order of discriminant $D < 0$. Denote by $\text{Ell}(\mathcal{O}_D)$ the set of all isomorphism classes of elliptic curves E over \mathbb{C} having \mathcal{O}_D as their full ring of complex multiplication (i.e. having $\text{End}(E) \cong \mathcal{O}_D$). It is well known that isomorphism classes of elliptic curves over \mathbb{C} correspond bijectively with homothety classes of complex lattices [46, I.1]; accordingly, we will write E_{Λ} throughout for the elliptic curve corresponding to a complex lattice $\Lambda \subset \mathbb{C}$. Moreover, fixing an embedding $\mathcal{O}_D \subset \mathbb{C}$, one can show that ideal classes $\mathfrak{a} \subset \mathcal{O}_D$ give rise to precisely those lattices representing elliptic curves in $\text{Ell}(\mathcal{O}_D)$ [9, 10.20], and that the map $\mathfrak{a} \mapsto E_{\mathfrak{a}}$ induces a bijection between the ideal class group $\text{Cl}(\mathcal{O}_D)$ of \mathcal{O}_D and $\text{Ell}(\mathcal{O}_D)$.

The above paragraph thus explains the correspondence between ideal class groups and elliptic curves over \mathbb{C} . The following proposition describes how this correspondence behaves with respect to isogenies:

Proposition 3.1.

1. *There is a well defined simply transitive action of $\text{Cl}(\mathcal{O}_D)$ on $\text{Ell}(\mathcal{O}_D)$, given by the formula*

$$\mathfrak{a} * E_\Lambda := E_{\mathfrak{a}^{-1}\Lambda},$$

valid for any nonzero fractional ideal $\mathfrak{a} \subset \mathcal{O}_D$.

2. *If \mathfrak{a} is an invertible ideal of \mathcal{O}_D , one has $\Lambda \subset \mathfrak{a}^{-1}\Lambda$, and this inclusion induces an isogeny $E_\Lambda \rightarrow \mathfrak{a} * E_\Lambda$ of degree equal to the norm $N(\mathfrak{a})$ of the ideal \mathfrak{a} .*
3. *Up to isomorphism, every isogeny between two elliptic curves $E_1, E_2 \in \text{Ell}(\mathcal{O}_D)$ arises in the above manner.*

Proof. Items 1 and 2 are proved in [46, II.1] (for the case of \mathcal{O}_D maximal) and [31] (for the general case).

To prove item 3, let $\phi: E_1 \rightarrow E_2$ be an isogeny and choose fractional ideals $\mathfrak{a} \subset \mathfrak{b}$ of \mathcal{O}_D such that $E_{\mathfrak{a}} \cong E_1$ and $E_{\mathfrak{b}} \cong E_1/\ker(\phi) \cong E_2$. Since $\mathfrak{a} \subset \mathfrak{b}$, there exists an integral ideal $\mathfrak{c} \subset \mathcal{O}_D$ such that $\mathfrak{b}\mathfrak{c} = \mathfrak{a}$, whereupon the morphism $\psi: E_{\mathfrak{a}} \rightarrow \mathfrak{c} * E_{\mathfrak{a}}$ yields an isogeny which has the same kernel as ϕ , and hence must be isomorphic to ϕ . \square

We now state and prove an analogue of Theorem 1.5 over the complex numbers.

Theorem 3.2. *Let Γ be the graph whose vertices are elements of $\text{Ell}(\mathcal{O}_D)$ and whose edges are isogenies of prime degree less than some fixed bound $M \geq (\log |D|)^B$, for some absolute constant $B > 2$. Then, assuming GRH, the graph Γ is an expander graph satisfying the bound (1.7).*

Proof. We have already seen that the elements of $\text{Ell}(\mathcal{O}_D)$ are in bijection with the elements of the group $\text{Cl}(\mathcal{O}_D)$ [9, 10.20], and that the action of $\text{Cl}(\mathcal{O}_D)$ on $\text{Ell}(\mathcal{O}_D)$ defined in Proposition 3.1 coincides exactly with the translation action of $\text{Cl}(\mathcal{O}_D)$ on itself under this bijection. Moreover, isogenies of prime degree less than M correspond to integral ideals of prime norm less than M , and the inverses (i.e. complex conjugates) of these ideals

have the same prime norm and thus also yield such isogenies. Hence, the graph Γ is isomorphic to the Cayley graph of $\text{Cl}(\mathcal{O}_D)$ under the generating set consisting of ideals of prime norm less than $M \geq (\log |D|)^B$.

Next we relate this graph to one covered by Theorem 1.1. Let $K = \mathbb{Q}(\sqrt{D})$ and \mathfrak{m} the principal ideal generated by the conductor c of the discriminant D (i.e. the largest integer whose square divides D). Then the class group $\text{Cl}(\mathcal{O}_D)$ is a quotient of the narrow ray class group of K relative to \mathfrak{m} [9, Prop. 7.22], and Theorem 1.1 applies directly to Γ and equation (1.7) with $x = M$ gives the desired bound. \square

In order to prove Theorem 1.5 from Theorem 3.2, we require the following classical result, known as Deuring's lifting theorem [10]:

Theorem 3.3.

1. *Let E be an elliptic curve defined over \mathbb{F}_q , and let ϕ be a nontrivial endomorphism of E . There exists an elliptic curve \tilde{E} defined over a number field L , a prime ideal \mathfrak{p} of L , and an endomorphism $\tilde{\phi}$ of \tilde{E} such that \tilde{E} and $\tilde{\phi}$ reduce to E and ϕ modulo \mathfrak{p} .*
2. *When E is ordinary, the mod \mathfrak{p} reduction map induces an isomorphism $\text{End}(\tilde{E}) \cong \text{End}(E)$.*

Proof of Theorem 1.5: Since the curves in Theorem 1.5 are ordinary, there exists an imaginary quadratic order \mathcal{O}_D such that $\text{End}(E) = \mathcal{O}_D$. Observe that $(\log 4q)^B \geq (\log |D|)^B$, since $D = t^2 - 4q$ where the trace t satisfies the Hasse bound $|t| < 2\sqrt{q}$. Hence $(\log 4q)^B$ satisfies the condition for M in Theorem 3.2.

We will now show that the graph Γ in Theorem 3.2 is isomorphic to the graph defined in Theorem 1.5. The elliptic curves in $\text{Ell}(\mathcal{O}_D)$ are all defined over the ring class field H of \mathcal{O}_D . Identification of the vertices is accomplished by choosing a prime $\mathfrak{p} \subset H$ lying over the characteristic p of \mathbb{F}_q , and reducing curves in $\text{Ell}(\mathcal{O}_D)$ to obtain curves in $S_{N,q}$. Theorem 3.3 shows that this identification is surjective. To show that it is injective, consider two non-isomorphic curves $E_{\mathfrak{a}}$ and $E_{\mathfrak{b}}$ in $\text{Ell}(\mathcal{O}_D)$, meaning that \mathfrak{a} and \mathfrak{b} lie in different ideal classes in $\text{Cl}(\mathcal{O}_D)$. By the Chebotarev density theorem, there exists an unramified prime ideal \mathfrak{c} belonging to the same ideal class as $\mathfrak{a}\mathfrak{b}^{-1}$; note in particular that \mathfrak{c} is not principal. By Proposition 3.1, the ideal \mathfrak{c} induces an isogeny ϕ between $E_{\mathfrak{a}}$ and $E_{\mathfrak{b}}$ having degree equal to $N(\mathfrak{c})$. If the reductions $\bar{E}_{\mathfrak{a}}$ and $\bar{E}_{\mathfrak{b}}$ of $E_{\mathfrak{a}}$ and $E_{\mathfrak{b}}$ modulo \mathfrak{p} were to be somehow isomorphic,

then ϕ would represent an endomorphism of \bar{E}_a , of degree $N(\mathfrak{c})$. However, we know the endomorphism ring of \bar{E}_a is equal to \mathcal{O}_D , and no element of \mathcal{O}_D has norm equal to $N(\mathfrak{c})$ (this is because $\mathbb{Q}(\sqrt{D})$ is an imaginary quadratic number field). Thus the endomorphism ring \mathcal{O}_D cannot contain any endomorphism of degree equal to $N(\mathfrak{c})$.

Likewise, for each prime $\ell < (\log 4q)^B$, the reduction map modulo \mathfrak{p} sends every isogeny of degree ℓ in characteristic 0 to an isogeny of degree ℓ in characteristic p . All isogenies in characteristic p are obtained in this way, since isogenies of degree ℓ are given by the roots of the modular polynomial $\Phi_\ell(x, y)$, and this polynomial does not have more roots over the algebraic closure in characteristic p than in characteristic 0. \square

4 Relationship with discrete logarithms

Given a generator g of a cyclic group G of order n , the discrete logarithm of an element h of G is defined to be the residue class x of integers mod n such that $g^x = h$. The elliptic curve discrete logarithm problem is the problem of computing discrete logarithms when G is the group of points on an elliptic curve defined over a finite field \mathbb{F}_q . Determining the difficulty of this problem is important because much of elliptic curve cryptography is based, at least conjecturally, on the infeasibility of computing discrete logarithms on elliptic curves defined over a finite field.

Galbraith [15] has observed that given an efficiently computable isogeny $\phi: E \rightarrow E'$, one can compute discrete logarithms on E by computing discrete logarithms on E' . The procedure is as follows: given $P, Q \in E$, compute $\phi(P)$ and $\phi(Q)$, and determine the discrete logarithm x of $\phi(Q)$ on E' with respect to the generator $\phi(P)$. The equation $x \cdot \phi(P) = \phi(Q)$ determines the solution for x modulo the kernel of ϕ . When ϕ is furthermore a low-degree isogeny, it is both efficiently computable and has small kernel (which itself can be efficiently enumerated). Such an isogeny provides a reduction between the discrete logarithm problems on E and E' , in time polynomial in $\log q$ and the degree. Moreover, a theorem of Tate [49] states that two elliptic curves E and E' defined over \mathbb{F}_q have the same number of points if and only if they are isogenous. Tate's theorem guarantees the existence of an isogeny defined over \mathbb{F}_q between curves in their equivalence classes, which computationally amounts to one between the curves themselves (see footnote 2). However, this isogeny usually is difficult to compute and has enormous degree.

We now use the above observation to give a proof of Theorem 1.6. Our proof consists of showing that, for curves of the same level, a composition of low-degree isogenies between them exists. Indeed, though the degree of such a composition may be very large, it can be computed efficiently; furthermore, it gives efficient reductions between all curves it connects.

Proof of Theorem 1.6: Returning to the isogeny graph of Theorem 1.5, let S denote the subset consisting of the μ -fraction of elliptic curves to which the algorithm \mathcal{A} applies. Let E be any curve of the same level as the curves in S . Because of the effective upper bounds on class numbers, one has that $\log |S_{N,q}| \leq c' \log q$, for some $c' > 0$. Construct a random walk of length $Cc'(\log q)/\log \log q$ starting at E , where C is the constant in Corollary 1.3. Let ϕ denote the isogeny equal to the composition of the isogenies represented by the edges comprising the random walk. Then ϕ can be evaluated in polynomial time, and hence the discrete logarithm problem on E can be solved efficiently by querying \mathcal{A} , as long as the random walk above lands in S . By Corollary 1.3, the probability that the random walk lands in S is at least $\frac{\mu}{2}$, so by repeating this process until the walk lands in S , we can solve discrete logarithms on E in probabilistic polynomial time using an expected number of queries to \mathcal{A} bounded by $\frac{2}{\mu}$. \square

5 Reductions between different levels

It is natural to ask whether the equivalence of discrete logarithms holds for elliptic curves in different levels. We begin by observing that the CM field $\text{End}(E) \otimes \mathbb{Q}$ is the same for all curves $E \in S_{N,q}$ regardless of level. Moreover, two curves E, E' have the same level if and only if the conductors of their endomorphism rings in $\text{End}(E) \otimes \mathbb{Q}$ are equal. It is thus natural to define the *conductor gap* to be the value of the largest prime factor at which the prime factorizations of the conductors of $\text{End}(E)$ and $\text{End}(E')$ differ; in addition, for a single curve E we define the conductor gap of E to be the maximal possible conductor gap over all possible pairs of isogenous E, E' . The conductor gap provides a rough measurement of how much the levels of E and E' differ.

Given any curve E whose endomorphism ring has conductor c , it is possible to compute a curve E' with conductor $c\ell$ together with an isogeny $E \rightarrow E'$ of degree ℓ in time $O(\ell^3)$; the reverse, starting from E' of conductor $c\ell$ and ending up with E of conductor c , is also possible in the same amount of time

([11, 15, 30]). Consider a union of any number of levels which collectively have conductor gap bounded polynomially in $\log q$. Though the individual sizes of each level may be difficult to compute, formula (2.7) or [9, Cor. 7.28] allows one to compute their relative sizes efficiently. By weighing these sizes it is possible to select a level at random with probability proportional to its total size amongst this union. This level can be reached by appropriate low degree isogenies. Thus it is possible to reach a random curve through walks of low degree isogenies, and it follows that Theorem 1.6 holds for the union of any number of levels which collectively have conductor gap bounded polynomially in $\log q$.

Large conductor gaps do pose an obstacle in the statement of Theorem 1.6, but they rarely arise in practice. Indeed, every curve $E \in S_{N,q}$ has at least the endomorphisms $\mathbb{Z} \subset \text{End}(E)$ and $\pi_q \in \text{End}(E)$, with π_q denoting the Frobenius endomorphism. The discriminant of the quadratic order $\mathbb{Z}[\pi_q]$ is equal to $t^2 - 4q$ where $t = q + 1 - N$, and the conductor of any curve in $S_{N,q}$ must be an integer c satisfying $c^2 \mid (t^2 - 4q)$. Thus, if $t^2 - 4q$ is square free, then all curves in $S_{N,q}$ are of the same level, and in this case the level restriction in Theorem 1.6 is vacuous. More generally, as long as $t^2 - 4q$ has no large repeated prime factors, the statement of Theorem 1.6 holds for all of $S_{N,q}$, by the previous paragraph.

We can analyze the expected frequency of large conductor gaps as follows. The Hasse-Weil bound on t implies $-4q \leq t^2 - 4q \leq 0$. A random integer within this interval has probability $1 - \prod_{p > \beta}^{2\sqrt{q}} (1 - p^{-2})$ of admitting a repeated prime factor $p > \beta$. Since this probability is bounded above by $O(1/\beta)$, we expect as a heuristic that, for any positive $\beta < p$, random choices of (N, q) will admit repeated prime factors exceeding β with probability $1/\beta$. In fact, [37, Theorem 1] rigorously proves the probability estimate $\frac{(\log p)^2}{\beta}$ for $\beta \ll p^{1/6}$, where $q = p$ is odd. Therefore, in most cases, conductor gaps between elliptic curves are quite small and we can ignore the effects of differing endomorphism rings in our discrete logarithm comparisons. For example, an investigation of nine randomly generated curves listed in international standards documents reveals that all of them satisfy $c_{N,q} \leq 3$ (cf. Section 6). In fact, a somewhat surprising observation is that there is currently no efficient algorithm to construct pairs of elliptic curves with conductor gaps that are *not* small, even though such pairs are known to exist in abundance (cf. Section 7).

6 Government standards for curves

In the previous section we showed that all curves in an isogeny class have identical security on average whenever the conductor gap is small. However, determining the conductor gap of a curve requires factoring a large integer and hence is a nontrivial computation. In this section we provide the computation of the conductor gap for a family of randomly selected curves which appear as part of a US government standard.

In 2000 the National Institute of Standards and Technology (NIST), a branch of the United States Department of Commerce, introduced a family of elliptic curves as standards for cryptographic applications [43]. The selection of these curves was the outcome of several years of testing. The NIST curves are generated by the values of secure hash functions applied to publicly-revealed seeds, making it plausible that they were not excessively manipulated before their public release. However, the user cannot be totally confident that there is not a backdoor or weakness in the published curve.

Though it is hard to imagine arguing directly that discrete logarithms on a *specific* elliptic curve do not have good attacks, our results can be used to give some assurance that the NIST curves are not weaker than comparable elliptic curves. Namely, Theorem 1.6 and the comment immediately following it show that the discrete logarithm problem has roughly equivalent difficulty as one ranges over curves defined over the same field, and whose endomorphism rings have small conductor gap.

Some of the NIST curves are *Koblitz* curves [29], which are not expected to have small conductor gaps. However, for the remaining NIST curves, some lengthy computations showed that the conductor gap is very small: all but one curve had a conductor gap of 1, and the only exception had a conductor gap of 3. That means that in the former cases, the isogeny class consists of only one level, and Theorem 1.6 provides a full equivalence of discrete logarithms. Only in the exceptional case with conductor gap 3 must one navigate between levels (the topic of Section 5); this can easily be done by constructing a degree 3 isogeny between them. Therefore we may conclude that these curves have typical difficulty among all elliptic curves defined over the same field and having the same number of points.

As an example, consider the NIST curve B-571, which is given by the Weierstrass equation $y^2 + xy = x^3 + x^2 + b$ over the field $\mathbb{F}_{2^{571}}$. Here b is an element of $\mathbb{F}_{2^{571}}$ which is cumbersome to describe but can be found on p. 47

of [43]. It has discriminant

$$d = -210092063841005638410400838462812964562253124135523060955333 \backslash \\ 767330638498791801056156659734237518468659692798673383993911 \backslash \\ 78057790576859207002963481895511008772786625592941143 \quad (6.1)$$

and prime factorization

$$= -137 * 1502689 * 5608493523058319 * 3563521804312876303 \\ * 46393104672338327566438581332776443577 \\ * 1100628851017477373738489717699925956411395060089467152067605 \backslash \\ 28637300688225399301632484625559$$

(we have written out the decimal expansion of d over several lines owing to its length). One can determine the conductor gap knowing this factorization: it is the largest square factor, which in this example is 1.

We wish to mention that finding the above factorization was far from trivial, taking about 5 days on a dedicated cluster in the Netherlands which utilized specialized factoring software. Although determining the conductor gap is useful in assuring that a given elliptic curve is not cryptographically weak, clearly this is not a test which the average user can perform. It may be good practice for standards bodies to publish the factorization of the discriminants along with their recommended curves so that users have this information.

7 Open problems

In this section we address two shortcomings of Theorem 1.6. The first is that the Theorem, as stated, applies only to individual levels of curves. As noted just after its statement and further in Sections 5 and 6, curves whose levels differ by a ratio composed of small primes can be bridged by random isogenies; the issue is when the conductor gap has a large prime factor. The second is the strong analytic assumption of the Generalized Riemann Hypothesis. We conclude by discussing some related cryptographic problems.

7.1 Large conductor gaps

The equivalence result of Theorem 1.6 is incomplete in the sense that it does not apply to curves having a large conductor gap. Pairs of such curves certainly exist, but no efficient method is known for finding them, and indeed no explicit example is known at the present time. A curve chosen at random will have conductor greater than ℓ with probability heuristically equal to $1/\ell$ (see Section 5). As we mentioned in Section 5, it is possible to produce an explicit isogeny between two curves with conductor gap ℓ in time $O(\ell^3)$, which for large ℓ is far slower than solving discrete logarithms themselves. Additionally, it was recently shown in [12] how to create special pairs of curves with conductor gap ℓ in time $O(\ell^2)$, without finding an explicit isogeny between them. All of these methods are too slow for large values of ℓ , but leave an intermediate range of conductor gaps which presently cannot be bridged by computable isogenies.

The conductor gap question is especially pertinent for certain special classes of curves in cryptography such as pairing friendly curves (see [14]). All constructible examples of such curves are presently restricted to small discriminants, with the exception of certain families of curves having conductor gaps which fall within the abovementioned intermediate range [6]; note, however, that these conductor gaps are still small enough that improvements such as Moore's law affect the boundaries of this range. There is some concern (although no proof) that discrete logarithms on such curves are weaker than on general pairing friendly curves. Achieving large conductor gaps for pairing friendly curves would help alleviate this concern, since our work then implies that pairing friendly curves with large discriminant are provably as secure as random pairing friendly curves.

7.2 The assumption of GRH

The theorems in this paper all assume the Generalized Riemann Hypothesis, which is used to obtain the error estimate in (2.6). Lighter analytic assumptions still imply nontrivial error estimates; for example the Generalized Lindelöf Hypothesis instead implies a bound of $O_{\varepsilon,K}(x^{1/2+\varepsilon}Q^\varepsilon)$ for any $\varepsilon > 0$ [24]. This corresponds to a subexponential time algorithm in Theorem 1.6, as opposed to a polynomial time one.

An unconditional proof of expansion seems out of reach at present. In the introduction it was explained why expansion bounds for λ_χ imply bounds

on the least quadratic nonresidue, and thus at present require an analytic assumption. The recent preprint [33] considers cancellation in the sums λ_χ defined in (1.4) for other characters.

Intriguingly, it has been widely speculated that the GRH implication of $B > 2$ in Theorem 1.1 is not sharp, and that $B > 1$ is in fact expected. This feature dates back to the suggestion of Littlewood that the Euler product for $L(1, \chi)$ could be approximated by the partial Euler product over primes smaller than $(\log Q)^B$, for any $B > 1$. This approximation is consistent with the best known constructions of lower bounds for the error terms in the sums (2.6), and for related problems such as the least nonresidue problem [19, 20, 34]. Recent work of [21, 42, 50] supports the validity of the wider range $B > 1$. This bound is also sharp from the point of view of the Alon-Roichman Theorem [2], which asserts that expanders must have at least logarithmic degree in the size of the graph.

Finally, the constants in (2.6) are effective and numerical values for them have been obtained in [3, 4].

7.3 Generalizations to other cryptographic problems

The elliptic curve discrete logarithm problem can be generalized to Jacobians of hyperelliptic curves or other curves of higher genus, and recently there has been some progress in obtaining efficiently computable isogenies between such abelian varieties [48]. At present, not enough such isogenies are known to enable any statement about reducibility of discrete logarithms between such Jacobians, but further developments could likely yield new results in this area.

In a different vein, one can consider alternative cryptographic problems such as the Diffie-Hellman problem instead of the discrete logarithm problem. For example, the recent paper [26] shows that, for curves over a prime field, computing the least significant bit of a Diffie-Hellman secret with greater than 50% probability over a non-negligible fraction of curves is almost always equivalent to solving the full Diffie-Hellman problem itself (assuming GRH). The proof relies heavily on the rapid mixing properties of isogeny graphs for ordinary elliptic curves.

Acknowledgements: It is a pleasure to thank Noga Alon, R. Balasubramanian, Dan Boneh, Ehud de Shalit, Noam Elkies, Andrew Granville,

Dimitar Jetchev, Nati Linial, Alexander Lubotzky, Peter Montgomery, Kumar Murty, Ravi Ramakrishna, Ze’ev Rudnick, Peter Sarnak, Adi Shamir, Igor Shparlinksi, and Martin Weissman for their helpful comments and suggestions. In particular, we wish to thank Peter Montgomery for his assistance and advice in factoring the discriminants of NIST curves in Section 6, and for making computational resources available to us.

References

- [1] Noga Alon, *Tough Ramsey graphs without short cycles*, J. Algebraic Combin. **4** (1995), no. 3, 189–195.
- [2] Noga Alon and Yuval Roichman, *Random Cayley graphs and expanders*, Random Structures Algorithms **5** (1994), no. 2, 271–284.
www.math.tau.ac.il/~nogaa/PDFS/exp1.pdf.
- [3] Eric Bach, *Explicit bounds for primality testing and related problems*, Math. Comp. **55** (1990), no. 191, 355–380.
- [4] Eric Bach and Jonathan Sorenson, *Explicit bounds for primes in residue classes*, Math. Comp. **65** (1996), no. 216, 1717–1735.
- [5] N. L. Biggs and A. G. Boshier, *Note on the girth of Ramanujan graphs*, J. Combin. Theory Ser. B **49** (1990), no. 2, 190–194.
- [6] Gaetan Bisson and Takakazu Satoh, *More Discriminants with the Brezing-Weng Method*, Progress in Cryptology - INDOCRYPT 2008, 9th International Conference on Cryptology in India, Kharagpur, India, December 14-17, 2008. Proceedings, 2008, pp. 389–399.
- [7] I. F. Blake, G. Seroussi, and N. P. Smart, *Elliptic curves in cryptography*, London Mathematical Society Lecture Note Series, vol. 265, Cambridge University Press, Cambridge, 2000.
- [8] Henri Cohen, *Advanced topics in computational number theory*, Graduate Texts in Mathematics, vol. 193, Springer-Verlag, New York, 2000.
- [9] David A. Cox, *Primes of the form $x^2 + ny^2$* , A Wiley-Interscience Publication, John Wiley & Sons Inc., New York, 1989.
- [10] Max Deuring, *Die Typen der Multiplikatorenringe elliptischer Funktionenkörper*, Abh. Math. Sem. Hansischen Univ. **14** (1941), 197–272 (German).
- [11] Andreas Enge, *Computing modular polynomials in quasi-linear time*. To appear in Mathematics of Computation.
<http://arxiv.org/abs/0704.3177>.

- [12] ———, *The complexity of class polynomial computation via floating point approximations*. To appear in Mathematics of Computation.
<http://arxiv.org/abs/cs.CC/0601104>.
- [13] Mireille Fouquet and François Morain, *Isogeny volcanoes and the SEA algorithm*, Algorithmic number theory (Sydney, 2002), 2002, pp. 276–291.
- [14] David Freeman, Michael Scott, and Edlyn Teske, *A taxonomy of pairing-friendly elliptic curves* (2006). <http://eprint.iacr.org/2006/372>.
- [15] Steven D. Galbraith, *Constructing isogenies between elliptic curves over finite fields*, LMS J. Comput. Math. **2** (1999), 118–138 (electronic).
- [16] Steven D. Galbraith, Florian Hess, and Nigel P. Smart, *Extending the GHS Weil descent attack*, Advances in cryptology—EUROCRYPT 2002 (Amsterdam), 2002, pp. 29–44.
- [17] M. Z. Garaev, *Character sums in short intervals and the multiplication table modulo a large prime*, Monatsh. Math. **148** (2006), no. 2, 127–138. MR **2235360** (2007c:11094)
- [18] Oded Goldreich, *Candidate One-Way Functions Based on Expander Graphs*, Electronic Colloquium on Computational Complexity (ECCC) **90** (2000).
<http://citeseer.ist.psu.edu/article/goldreich00candidate.html>.
- [19] S. W. Graham and C. J. Ringrose, *Lower bounds for least quadratic nonresidues*, Analytic number theory (Allerton Park, IL, 1989), Progr. Math., vol. 85, Birkhäuser Boston, Boston, MA, 1990, pp. 269–309.
- [20] A. Granville and K. Soundararajan, *The distribution of values of $L(1, \chi_d)$* , Geom. Funct. Anal. **13** (2003), no. 5, 992–1028.
- [21] Andrew Granville and K. Soundararajan, *Extreme values of $|\zeta(1+it)|$* , The Riemann zeta function and related themes: papers in honour of Professor K. Ramachandra, Ramanujan Math. Soc. Lect. Notes Ser., vol. 2, Ramanujan Math. Soc., Mysore, 2006, pp. 65–80.
- [22] Jeremy Horwitz and Ramarathnam Venkatesan, *Random Cayley digraphs and the discrete logarithm*, Algorithmic number theory (Sydney, 2002), 2002, pp. 416–430.
- [23] Henryk Iwaniec, *Topics in classical automorphic forms*, Graduate Studies in Mathematics, vol. 17, American Mathematical Society, Providence, RI, 1997.
- [24] Henryk Iwaniec and Emmanuel Kowalski, *Analytic number theory*, American Mathematical Society Colloquium Publications, vol. 53, American Mathematical Society, Providence, RI, 2004.
- [25] David Jao, Stephen D. Miller, and Ramarathnam Venkatesan, *Do all elliptic curves of the same order have the same difficulty of discrete log?*, Advances in cryptology—ASIACRYPT 2005, Lecture Notes in Comput. Sci., vol. 3788, Springer, Berlin, 2005, pp. 21–40.

- [26] Dimitar Jetchev and Ramarathnam Venkatesan, *Bit Secuirty of Diffie-Hellman Secret Keys*, Advances in cryptology—Crypto 2008, Lecture Notes in Comput. Sci., vol. 5157, Springer, Berlin, 2008, pp. 75–92.
- [27] Nathan Keller, Stephen D. Miller, Ilya Mironov, and Ramarathnam Venkatesan, *MV3: A New Stream Cipher Based On Random Walks and Revolving Buffers*, Topics in Cryptology — Proceedings of CT-RSA 2007, 2007, pp. 1–19.
- [28] Neal Koblitz, *Elliptic curve cryptosystems*, Math. Comp. **48** (1987), no. 177, 203–209.
- [29] ———, *CM-curves with good cryptographic properties*, Advances in cryptology—CRYPTO '91 (Santa Barbara, CA, 1991), 1992, pp. 279–287.
- [30] David Kohel, *Endomorphism rings of elliptic curves over finite fields*, University of California, Berkeley, 1996, Ph.D thesis.
- [31] Serge Lang, *Elliptic functions*, 2nd ed., Graduate Texts in Mathematics, vol. 112, Springer-Verlag, New York, 1987. With an appendix by J. Tate.
- [32] ———, *Algebraic number theory*, 2nd ed., Graduate Texts in Mathematics, vol. 110, Springer-Verlag, New York, 1994.
- [33] Y.-K. Lau and J. Wu, *On the least quadratic non-residue* (2007). preprint, <http://hkumath.hku.hk/~imr/IMRPreprintSeries/2007/IMR2007-4.pdf>.
- [34] J.E. Littlewood, *On the class number of the corpus $P(\sqrt{-k})$* , Proc. London Math. Soc. **27** (1928), 358–372.
- [35] A. Lubotzky, R. Phillips, and P. Sarnak, *Ramanujan graphs*, Combinatorica **8** (1988), no. 3, 261–277.
- [36] A. Lubotzky and B. Weiss, *Groups and expanders*, Expanding graphs (Princeton, NJ, 1992), DIMACS Ser. Discrete Math. Theoret. Comput. Sci., vol. 10, Amer. Math. Soc., Providence, RI, 1993, pp. 95–109.
- [37] Florian Luca and Igor Shparlinski, *Discriminants of Complex Multiplication Fields of Elliptic Curves over Finite Fields*. preprint.
- [38] G. A. Margulis, *Explicit constructions of expanders*, Problemy Peredači Informacii **9** (1973), no. 4, 71–80 (Russian).
- [39] ———, *Explicit group-theoretic constructions of combinatorial schemes and their applications in the construction of expanders and concentrators*, Problemy Peredachi Informatsii **24** (1988), no. 1, 51–60 (Russian).
- [40] Alfred J. Menezes, Tatsuaki Okamoto, and Scott A. Vanstone, *Reducing elliptic curve logarithms to logarithms in a finite field*, IEEE Trans. Inform. Theory **39** (1993), no. 5, 1639–1646.
- [41] Victor S. Miller, *Use of elliptic curves in cryptography*, Advances in cryptology—CRYPTO '85 (Santa Barbara, Calif., 1985), 1986, pp. 417–426.

- [42] H. L. Montgomery and R. C. Vaughan, *Extreme values of Dirichlet L -functions at 1*, Number theory in progress, Vol. 2 (Zakopane-Kościelisko, 1997), de Gruyter, Berlin, 1999, pp. 1039–1052.
- [43] National Institute of Standards and Technology, *Digital Signature Standard (DSS)*, Technical Report FIPS PUB 186-2, January 2000.
<http://www.csrc.nist.gov/publications/fips/fips186-2/fips186-2-change1.pdf>.
- [44] Stephen C. Pohlig and Martin E. Hellman, *An improved algorithm for computing logarithms over $\text{GF}(p)$ and its cryptographic significance*, IEEE Trans. Information Theory **IT-24** (1978), no. 1, 106–110.
- [45] René Schoof, *Counting points on elliptic curves over finite fields*, J. Théor. Nombres Bordeaux **7** (1995), no. 1, 219–254. Les Dix-huitièmes Journées Arithmétiques (Bordeaux, 1993).
- [46] Joseph H. Silverman, *Advanced Topics in the Arithmetic of Elliptic Curves*, Graduate Texts in Mathematics, vol. 151, Springer-Verlag, New York, 1994.
- [47] N. P. Smart, *The discrete logarithm problem on elliptic curves of trace one*, J. Cryptology **12** (1999), no. 3, 193–196.
- [48] Benjamin Smith, *Isogenies and the Discrete Logarithm Problem in Jacobians of Genus 3 Hyperelliptic Curves*, Advances in cryptology—EUROCRYPT 2008, Lecture Notes in Comput. Sci., vol. 4965, Springer, Berlin, 2008, pp. 163–180.
- [49] John Tate, *Endomorphisms of abelian varieties over finite fields*, Invent. Math. **2** (1966), 134–144.
- [50] R. C. Vaughan, *Small values of Dirichlet L -functions at 1*, Analytic number theory, Vol. 2 (Allerton Park, IL, 1995), Progr. Math., vol. 139, Birkhäuser Boston, Boston, MA, 1996, pp. 755–766.

Addresses:

David Jao
 Department of Combinatorics and Optimization
 University of Waterloo
 Waterloo, ON N2L 3G1
 Canada
djao@math.uwaterloo.ca

Stephen D. Miller
 Department of Mathematics
 110 Frelinghuysen Road
 Rutgers, The State University of New Jersey
 Piscataway, NJ 08854
miller@math.rutgers.edu

Ramarathnam Venkatesan
Microsoft Research Cryptography and Anti-Piracy Group
1 Microsoft Way
Redmond, WA 98052
and
Cryptography, Security and Applied Mathematics Research Group, Microsoft
Research India
Scientia - 196/36 2nd Main, Sadashivnagar, Bangalore 560 080, India
`venkie@microsoft.com`