

Do All Elliptic Curves of the Same Order Have the Same Difficulty of Discrete Log?

David Jao¹, Stephen Miller^{*,2,3}, and Ramarathnam Venkatesan¹

¹ Microsoft Research, 1 Microsoft Way, Redmond WA 98052
{davidjao,venkie}@microsoft.com

² Department of Mathematics, Rutgers University
110 Frelinghuysen Rd, Piscataway, NJ 08854-8019

³ Einstein Institute of Mathematics, Edmond J. Safra Campus, Givat Ram
The Hebrew University of Jerusalem, Jerusalem 91904 Israel
miller@math.huji.ac.il

Abstract. The aim of this paper is to justify the common cryptographic practice of selecting elliptic curves using their order as the primary criterion. We can formalize this issue by asking whether the discrete log problem (DLOG) has the same difficulty for all curves over a given finite field with the same order. We prove that this is essentially true by showing polynomial time random reducibility of DLOG among such curves, assuming the Generalized Riemann Hypothesis (GRH). We do so by constructing certain expander graphs (of Ramanujan type), with elliptic curves as nodes and low degree isogenies as edges, and utilizing the rapid mixing of random walks on this graph. Our proof works only for curves with (nearly) the same endomorphism rings. Without this technical restriction such DLOG equivalence might be false; however, in practice the restriction may be moot, because all known polynomial time techniques for constructing equal order curves produce only curves with nearly equal endomorphism rings.

Keywords: random reducibility, discrete log, elliptic curves, isogenies, modular forms, L -functions, generalized Riemann hypothesis, Ramanujan graphs, expanders, rapid mixing.

1 Introduction

Public key cryptosystems based on the elliptic curve discrete logarithm (DLOG) problem [22, 34] have received considerable attention because they are currently the most widely used systems whose underlying mathematical problem has yet to admit subexponential attacks (see [3, 31, 46]). Hence it is important to formally understand how the choice of elliptic curve affects the difficulty of the resulting DLOG problem. This turns out to be more intricate than the corresponding problem of DLOG over finite fields and their selection.

* Partially supported by NSF grant DMS-0301172 and an Alfred P. Sloan Foundation Fellowship.

To motivate the questions in this paper, we begin with two observations. First, we note that one typically picks an elliptic curve at random, and examines its group order (e.g. to check if it is smooth) to decide whether to keep it, or discard it and pick another one. It is therefore a natural question whether or not DLOG is of the same difficulty on curves over the same field with the same number of points. Indeed, it is a theorem of Tate that curves E_1 and E_2 defined over the same finite field \mathbb{F}_q have the same number of points if and only if they are *isogenous*, i.e., there exists a nontrivial algebraic group homomorphism $\phi: E_1 \rightarrow E_2$ between them. If this ϕ is efficiently computable and has a small kernel over \mathbb{F}_q , we can solve DLOG on E_1 , given a DLOG oracle for E_2 .

Secondly, we recall the observation that DLOG on $(\mathbb{Z}/p\mathbb{Z})^*$ has *random self-reducibility*: given any efficient algorithm $A(g^x) = x$ that solves DLOG on a polynomial fraction of inputs, one can solve *any* instance $y = g^x$ by an expected polynomial number of calls to A with *random* inputs of the form $A(g^r y)$. Thus, if DLOG on $(\mathbb{Z}/p\mathbb{Z})^*$ is hard in a sense suitable for cryptography at all (e.g., has no polynomial on average attack), then all but a negligible fraction of instances of DLOG on $(\mathbb{Z}/p\mathbb{Z})^*$ must necessarily be hard. This result is comforting since for cryptographic use we need the DLOG problem to be hard with overwhelming probability when we pick inputs at random. The same random self-reduction statement also holds true for DLOG on any abelian group, and in particular for DLOG on a *fixed* elliptic curve. We consider instead the following question: given a polynomial time algorithm to solve DLOG on some positive (or non-negligible) fraction of isogenous elliptic curves over \mathbb{F}_q , can we solve DLOG for *all* curves in the same isogeny class in polynomial time? In this paper we show that the answer to this question is essentially yes, by proving (assuming GRH) the mixing properties of random walks of isogenies on elliptic curves. It follows that if DLOG is hard at all in an isogeny class, then DLOG is hard for all but a negligible fraction of elliptic curves in that isogeny class. This result therefore justifies, in an average case sense, the cryptographic practice of selecting curves at random within an isogeny class.

1.1 Summary of our results

The conventional wisdom is that if two elliptic curves over the same finite field have the same order, then their discrete logarithm problems are equally hard. Indeed, this philosophy is embodied in the way one picks curves in practice. However, such a widely relied upon assertion merits formal justification. Our work shows that this simplified belief is essentially true for all elliptic curves which are constructible using present techniques, but with an important qualification which we shall now describe.

Specifically, let $S_{N,q}$ denote the set of elliptic curves defined over a given finite field \mathbb{F}_q , up to \mathbb{F}_q -isomorphism, that have the same order N over \mathbb{F}_q . We split $S_{N,q}$ into *levels* (as in Kohel [23]), where each level represents all elliptic curves having a particular endomorphism ring over \mathbb{F}_q . The curves in each level form the vertices of an *isogeny graph* [10, 11, 33], whose edges represent prime degree isogenies between curves of degree less than some specified bound m .

Theorem 1.1. *(Assuming GRH) There exists a polynomial $p(x)$, independent of N and q , such that for $m = p(\log q)$ the isogeny graph \mathcal{G} on each level is an expander graph, in the sense that any random walk on \mathcal{G} will reach a subset of size h with probability at least $\frac{h}{2|\mathcal{G}|}$ after $\text{polylog}(q)$ steps (where the implicit polynomial is again independent of N and q).*

Corollary 1.2 *(Assuming GRH) The DLOG problem on elliptic curves is random reducible in the following sense: given any algorithm A that solves DLOG on some fixed positive proportion of curves in fixed level, one can probabilistically solve DLOG on any given curve in that same level with $\text{polylog}(q)$ expected queries to A with random inputs.*

The proofs are given at the end of Section 4. These results constitute the first formulation of a polynomial time random reducibility result for the elliptic curve DLOG problem which is general enough to apply to typical curves that one ordinarily encounters in practice. An essential tool in our proof is the *nearly Ramanujan* property of Section 3, which we use to prove the expansion properties of our isogeny graphs. The expansion property in turn allows us to prove the rapid mixing of random walks given by compositions of small degree isogenies within a fixed level. Our method uses GRH to prove eigenvalue separation for these graphs, and provides a new technique for constructing expander graphs.

The results stated above concern a fixed level. One might therefore object that our work does not adequately address the issue of DLOG reduction in the case where two isogenous elliptic curves belong to different levels. If an attack is *balanced*, i.e., successful on each level on a polynomial fraction of curves, then our results apply. However, if only unbalanced attacks exist, then a more general equivalence may be false for more fundamental reasons. Nevertheless, at present this omission is not of much practical importance. First of all, most random curves over \mathbb{F}_q belong to sets $S_{N,q}$ consisting of only one level (see Section 6); for example, in Figure 1, we find that 10 out of the 11 randomly generated curves appearing in international standards documents have only one level. Second, if the endomorphism rings corresponding to two levels have conductors whose prime factorizations differ by quantities which are polynomially smooth, then one can use the algorithms of [11, 23] to navigate to a common level in polynomial time, and then apply Corollary 1.2 within that level to conclude that DLOG is polynomial-time random reducible between the two levels. This situation always arises in practice, because no polynomial time algorithm is known which even produces a pair of curves lying on levels whose conductor difference is not polynomially smooth. It is an open problem if such an algorithm exists.

Since our results link the DLOG complexity of all curves with the same conductor and most curves have small (or smooth) c_π , one may interpret those curves with small values of c_π as being generic. We emphasize that such genericity does not indicate whether DLOG on generic curves is easier or harder than on non-generic curves in general.

Curve	c_π (maximal conductor gap in isogeny class)	$P(c_\pi)$ = largest prime factor of c_π
NIST P-192	1	1
NIST P-256	3	3
NIST P-384	1	1
NIST P-521	1	1
NIST K-163	$45641 \cdot 82153 \cdot 56498081 \cdot P(c_\pi)$	86110311
NIST K-233	$5610641 \cdot 85310626991 \cdot P(c_\pi)$	150532234816721999
NIST K-283	$1697 \cdot 162254089 \cdot P(c_\pi)$	1779143207551652584836995286271
NIST K-409	$21262439877311 \cdot 22431439539154506863 \cdot P(c_\pi)$	57030553306655053533734286593
NIST K-571	$3952463 \cdot P(c_\pi)$	9021184135396238924389891(<i>contd</i>) 9451926768145189936450898(<i>contd</i>) 07769277009849103733654828039
NIST B-163	1	1
NIST B-233	1	1
NIST B-283	1	1
NIST B-409	1	1
NIST B-571	1	1
IPSec $3^{rd} OG, F_{2^{155}}$	1	1
IPSec $4^{th} OG, F_{2^{185}}$	1	1

Fig. 1. A table of curves recommended as international standards [16, 36]. Note that the value of c_π for each of the standards curves is small (at most 3), except for the curves in the NIST K (Koblitz curve) family. These phenomena are to be expected and are explained in Section 6. Any curve with $c_\pi = 1$ has the property that its isogeny class consists of only one level. It follows from the results of Section 1.1 that randomly generated elliptic curves with $c_\pi = 1$ (or, more generally, with smooth c_π) will have discrete logarithm problems of typical difficulty amongst all elliptic curves in their isogeny class.

Our use of random walks to reach large subsets of the isogeny graph is crucial, since constructing an isogeny between two specific curves⁴ is believed to be inherently hard, whereas constructing an isogeny from a fixed curve to a subset constituting a positive (or polynomial) fraction of the isogeny graph is proved in this paper to be easy. Kohel [23] and Galbraith [11] present exponential time algorithms (and thus exponential time reductions) for navigating between two nodes in the isogeny graph, some of which are based on random walk *heuristics* which we prove here rigorously. Subsequent papers on Weil descent attacks [12, 32] and elliptic curve trapdoor systems [45] also use isogeny random walks in order to extend the GHS Weil descent attack [13] to elliptic curves which are not themselves directly vulnerable to the GHS attack. Our work does not imply any changes to the deductions of these papers, since they also rely on the above heuristic assumptions involving exponentially long random walks. In our case, we achieve polynomial time instead of exponential time reductions; this is possible since we keep one curve fixed, and the random reducibility requires only that the other curve be randomly distributed.

⁴ If one uses polynomial size circuits (i.e., polynomial time algorithms with exponential time pre-processing) for reductions, then one can relate DLOG on two given curves. This claim follows using the smallness of diameter of our graphs and the smoothness of the degrees of isogenies involved. We omit the details.

2 Preliminaries

Let E_1 and E_2 be elliptic curves defined over a finite field \mathbb{F}_q of characteristic p . An isogeny $\phi: E_1 \rightarrow E_2$ defined over \mathbb{F}_q is a non-constant rational map defined over \mathbb{F}_q which is also a group homomorphism from $E_1(\mathbb{F}_q)$ to $E_2(\mathbb{F}_q)$ [42, §III.4]. The degree of an isogeny is its degree as a rational map. For any elliptic curve $E: y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6$ defined over \mathbb{F}_q , the Frobenius endomorphism is the isogeny $\pi: E \rightarrow E$ of degree q given by the equation $\pi(x, y) = (x^q, y^q)$. It satisfies the equation

$$\pi^2 - \text{Trace}(E)\pi + q = 0,$$

where $\text{Trace}(E) = q + 1 - \#E(\mathbb{F}_q)$ is the trace of the Frobenius endomorphism of E over \mathbb{F}_q . The polynomial $p(X) := X^2 - \text{Trace}(E)X + q$ is called the characteristic polynomial of E .

An endomorphism of E is an isogeny $E \rightarrow E$ defined over the algebraic closure $\overline{\mathbb{F}_q}$ of \mathbb{F}_q . The set of endomorphisms of E together with the zero map forms a ring under the operations of pointwise addition and composition; this ring is called the endomorphism ring of E and denoted $\text{End}(E)$. The ring $\text{End}(E)$ is isomorphic either to an order in a quaternion algebra or to an order in an imaginary quadratic field [42, V.3.1]; in the first case we say E is supersingular and in the second case we say E is ordinary. In the latter situation, the Frobenius endomorphism π can be regarded as an algebraic integer which is a root of the characteristic polynomial.

Two elliptic curves E_1 and E_2 defined over \mathbb{F}_q are said to be isogenous over \mathbb{F}_q if there exists an isogeny $\phi: E_1 \rightarrow E_2$ defined over \mathbb{F}_q . A theorem of Tate states that two curves E_1 and E_2 are isogenous over \mathbb{F}_q if and only if $\#E_1(\mathbb{F}_q) = \#E_2(\mathbb{F}_q)$ [43, §3]. Since every isogeny has a dual isogeny [42, III.6.1], the property of being isogenous over \mathbb{F}_q is an equivalence relation on the finite set of $\overline{\mathbb{F}_q}$ -isomorphism classes of elliptic curves defined over \mathbb{F}_q . We define an isogeny class to be an equivalence class of elliptic curves, up to $\overline{\mathbb{F}_q}$ -isomorphism, under this equivalence relation; the set $S_{N,q}$ of Section 1.1 is thus equal to the isogeny class of elliptic curves over \mathbb{F}_q having cardinality N .

Curves in the same isogeny class are either all supersingular or all ordinary. We assume for the remainder of this paper that we are in the **ordinary case**, which is the more interesting case from the point of view of cryptography in light of the MOV attack [30]. Theorem 1.1 in the supersingular case was known earlier by results of Pizer [37, 38], and a proof has been included for completeness in Appendix A.

The following theorem describes the structure of elliptic curves within an isogeny class from the point of view of their endomorphism rings.

Theorem 2.1. *Let E and E' be ordinary elliptic curves defined over \mathbb{F}_q which are isogenous over \mathbb{F}_q . Let K denote the imaginary quadratic field containing $\text{End}(E)$, and write \mathcal{O}_K for the maximal order (i.e., ring of integers) of K .*

1. *The order $\text{End}(E)$ satisfies the property $\mathbb{Z}[\pi] \subseteq \text{End}(E) \subseteq \mathcal{O}_K$.*

2. The order $\text{End}(E')$ also satisfies $\text{End}(E') \subset K$ and $\mathbb{Z}[\pi] \subseteq \text{End}(E') \subseteq \mathcal{O}_K$.
3. The following are equivalent:
 - (a) $\text{End}(E) = \text{End}(E')$.
 - (b) There exist two isogenies $\phi: E \rightarrow E'$ and $\psi: E \rightarrow E'$ of relatively prime degree, both defined over \mathbb{F}_q .
 - (c) $[\mathcal{O}_K : \text{End}(E)] = [\mathcal{O}_K : \text{End}(E')]$.
 - (d) $[\text{End}(E) : \mathbb{Z}[\pi]] = [\text{End}(E') : \mathbb{Z}[\pi]]$.
4. Let $\phi: E \rightarrow E'$ be an isogeny from E to E' of prime degree ℓ , defined over \mathbb{F}_q . Then either $\text{End}(E)$ contains $\text{End}(E')$ or $\text{End}(E')$ contains $\text{End}(E)$, and the index of the smaller in the larger divides ℓ .
5. Suppose ℓ is a prime that divides one of $[\mathcal{O}_K : \text{End}(E)]$ and $[\mathcal{O}_K : \text{End}(E')]$, but not the other. Then every isogeny $\phi: E \rightarrow E'$ defined over \mathbb{F}_q has degree equal to a multiple of ℓ .

Proof. [23, §4.2].

For any order $\mathcal{O} \subseteq \mathcal{O}_K$, the conductor of \mathcal{O} is defined to be the integer $[\mathcal{O}_K : \mathcal{O}]$. The field K is called the CM field of E . We write c_E for the conductor of $\text{End}(E)$ and c_π for the conductor of $\mathbb{Z}[\pi]$. Note that this is not the same thing as the arithmetic conductor of an elliptic curve [42, §C.16], nor is it related to the conductance of an expander graph [21]. It follows from [4, (7.2) and (7.3)] that $\text{End}(E) = \mathbb{Z} + c_E \mathcal{O}_K$ and $D = c_E^2 d_K$, where D (respectively, d_K) is the discriminant of the order $\text{End}(E)$ (respectively, \mathcal{O}_K). Furthermore, the characteristic polynomial $p(X)$ has discriminant $d_\pi = \text{disc}(p(X)) = \text{Trace}(E)^2 - 4q = \text{disc}(\mathbb{Z}[\pi]) = c_\pi^2 d_K$, with $c_\pi = c_E \cdot [\text{End}(E) : \mathbb{Z}[\pi]]$.

Following [10] and [11], we say that an isogeny $\phi: E \rightarrow E'$ of prime degree ℓ defined over \mathbb{F}_q is “down” if $[\text{End}(E) : \text{End}(E')] = \ell$, “up” if $[\text{End}(E') : \text{End}(E)] = \ell$, and “horizontal” if $\text{End}(E) = \text{End}(E')$. The following theorem classifies the number of degree ℓ isogenies of each type in terms of the Legendre symbol $\left(\frac{D}{\ell}\right)$.

Theorem 2.2. *Let E be an ordinary elliptic curve over \mathbb{F}_q , with endomorphism ring $\text{End}(E)$ of discriminant D . Let ℓ be a prime different from the characteristic of \mathbb{F}_q .*

- Assume $\ell \nmid c_E$. Then there are exactly $1 + \left(\frac{D}{\ell}\right)$ horizontal isogenies $\phi: E \rightarrow E'$ of degree ℓ .
 - If $\ell \nmid c_\pi$, there are no other isogenies $E \rightarrow E'$ of degree ℓ over \mathbb{F}_q .
 - If $\ell \mid c_\pi$, there are $\ell - \left(\frac{D}{\ell}\right)$ down isogenies of degree ℓ .
- Assume $\ell \mid c_E$. Then there is one up isogeny $E \rightarrow E'$ of degree ℓ .
 - If $\ell \nmid \frac{c_\pi}{c_E}$, there are no other isogenies $E \rightarrow E'$ of degree ℓ over \mathbb{F}_q .
 - If $\ell \mid \frac{c_\pi}{c_E}$, there are ℓ down isogenies of degree ℓ .

Proof. [10, §2.1] or [11, §11.5].

It follows that the maximal conductor difference between levels in an isogeny class is achieved between a curve at the top level (with $\text{End}(E) = \mathcal{O}_K$) and a curve at the bottom level (with $\text{End}(E) = \mathbb{Z}[\pi]$).

2.1 Isogeny Graphs

We define two curves E_1 and E_2 in an isogeny class $S_{N,q}$ to have the same level if $\text{End}(E_1) = \text{End}(E_2)$. An *isogeny graph* is a graph whose nodes consist of all elements in $S_{N,q}$ belonging to a fixed level. Note that a horizontal isogeny always goes between two curves of the same level; likewise, an up isogeny enlarges the size of the endomorphism ring and a down isogeny reduces the size. Since there are fewer elliptic curves at higher levels than at lower levels, the collection of isogeny graphs under the level interpretation visually resembles a “pyramid” or a “volcano” [10], with up isogenies ascending the structure and down isogenies descending.

As in [15, Prop. 2.3], we define two isogenies $\phi: E_1 \rightarrow E_2$ and $\phi': E_1 \rightarrow E_2$ to be equivalent if there exists an automorphism $\alpha \in \text{Aut}(E_2)$ (i.e., an invertible endomorphism) such that $\phi' = \alpha\phi$. The edges of the graph consist of equivalence classes of isogenies over \mathbb{F}_q between elliptic curve representatives of nodes in the graph, which have prime degree less than the bound $(\log q)^{2+\delta}$ for some fixed constant $\delta > 0$. The degree bound must be small enough to permit the isogenies to be computed, but large enough to allow the graph to be connected and to have the rapid mixing properties that we want. We will show in Section 4 that there exists a constant $\delta > 0$ for which a bound of $(\log q)^{2+\delta}$ satisfies all the requirements, provided that we restrict the isogenies to a single level.

Accordingly, fix a level of the isogeny class, and let $\text{End}(E) = \mathcal{O}$ be the common endomorphism ring of all of the elliptic curves in this level. Denote by \mathcal{G} the regular graph whose vertices are elements of $S_{N,q}$ with endomorphism ring \mathcal{O} , and whose edges are equivalence classes of horizontal isogenies defined over \mathbb{F}_q of prime degree $\leq (\log q)^{2+\delta}$. By standard facts from the theory of complex multiplication [4, §10], each invertible ideal $\mathfrak{a} \subset \mathcal{O}$ produces an elliptic curve \mathbb{C}/\mathfrak{a} defined over some number field $L \subset \mathbb{C}$ (called the ring class field of \mathcal{O}) [4, §11]. The curve \mathbb{C}/\mathfrak{a} has complex multiplication by \mathcal{O} , and two different ideals yield isomorphic curves if and only if they belong to the same ideal class. Likewise, each invertible ideal $\mathfrak{b} \subset \mathcal{O}$ defines an isogeny $\mathbb{C}/\mathfrak{a} \rightarrow \mathbb{C}/\mathfrak{ab}^{-1}$, and the degree of this isogeny is the norm $N(\mathfrak{b})$ of the ideal \mathfrak{b} . Moreover, for any prime ideal \mathfrak{P} in L lying over p , the reductions mod \mathfrak{P} of the above elliptic curves and isogenies are defined over \mathbb{F}_q , and every elliptic curve and every horizontal isogeny in \mathcal{G} arises in this way (see [11, §3] for the $p > 3$ case, and [12] for the small characteristic case). Therefore, the isogeny graph \mathcal{G} is isomorphic to the corresponding graph \mathcal{H} whose nodes are elliptic curves \mathbb{C}/\mathfrak{a} with complex multiplication by \mathcal{O} , and whose edges are complex analytic isogenies represented by ideals $\mathfrak{b} \subset \mathcal{O}$ and subject to the same degree bound as before. This isomorphism preserves the degrees of isogenies, in the sense that the degree of any isogeny in \mathcal{G} is equal to the norm of its corresponding ideal \mathfrak{b} in \mathcal{H} .

The graph \mathcal{H} has an alternate description as a Cayley graph on the ideal class group $\text{Cl}(\mathcal{O})$ of \mathcal{O} . Indeed, each node of \mathcal{H} is an ideal class of \mathcal{O} , and two ideal classes $[\mathfrak{a}_1]$ and $[\mathfrak{a}_2]$ are connected by an edge if and only if there exists a prime ideal \mathfrak{b} of norm $\leq (\log q)^{2+\delta}$ such that $[\mathfrak{a}_1\mathfrak{b}] = [\mathfrak{a}_2]$. Therefore, the graph \mathcal{H} (and hence the graph \mathcal{G}) is isomorphic to the Cayley graph of the group $\text{Cl}(\mathcal{O})$

with respect to the generators $[\mathfrak{b}] \in \text{Cl}(\mathcal{O})$, as \mathfrak{b} ranges over all prime ideals of \mathcal{O} of norm $\leq (\log q)^{2+\delta}$.

Remark 2.1. The isogeny graph \mathcal{G} consists of objects defined over the finite field \mathbb{F}_q , whereas the objects in the graph \mathcal{H} are defined over the number field L . One passes from \mathcal{H} to \mathcal{G} by taking reductions mod \mathfrak{P} , and from \mathcal{G} to \mathcal{H} by using Deuring's Lifting Theorem [8, 11, 24]. There is no known polynomial time or even subexponential time algorithm for computing the isomorphism between \mathcal{G} and \mathcal{H} [11, §3]. For our purposes, such an explicit algorithm is not necessary, since we only use the complex analytic theory to prove abstract graph-theoretic properties of \mathcal{G} .

Remark 2.2. The isogeny graph \mathcal{G} is typically a symmetric graph, since each isogeny ϕ has a unique dual isogeny $\hat{\phi}: E_2 \rightarrow E_1$ of the same degree as ϕ in the opposite direction [42, §III.6]. (From the viewpoint of \mathcal{H} , an isogeny represented by an ideal $\mathfrak{b} \subset \mathcal{O}$ has its dual isogeny represented simply by the complex conjugate $\bar{\mathfrak{b}}$.) However, the definition of equivalence of isogenies from [15] given in 2.1 contains a subtle asymmetry which can sometimes render the graph \mathcal{G} asymmetric in the supersingular case (Appendix A). Namely, if $\text{Aut}(E_1)$ is not equal to $\text{Aut}(E_2)$, then two isogenies $E_1 \rightarrow E_2$ can sometimes be equivalent even when their dual isogenies are not. For ordinary elliptic curves within a common level, the equation $\text{End}(E_1) = \text{End}(E_2)$ automatically implies $\text{Aut}(E_1) = \text{Aut}(E_2)$, so the graph \mathcal{G} is always symmetric in this case. Hence, we may regard \mathcal{G} as undirected and apply known results about undirected expander graphs (as in the following section) to \mathcal{G} .

3 Expander Graphs

Let $G = (\mathcal{V}, E)$ be a finite graph on h vertices \mathcal{V} with undirected edges E . Suppose G is a regular graph of degree k , i.e., exactly k edges meet at each vertex. Given a labeling of the vertices $\mathcal{V} = \{v_1, \dots, v_h\}$, the adjacency matrix of G is the symmetric $h \times h$ matrix A whose ij -th entry $A_{ij} = 1$ if an edge exists between v_i and v_j , and 0 otherwise.

It is convenient to identify functions on \mathcal{V} with vectors in \mathbb{R}^h via this labeling, and therefore also think of A as a self-adjoint operator on $L^2(\mathcal{V})$. All of the eigenvalues of A satisfy the bound $|\lambda| \leq k$. Constant vectors are eigenfunctions of A with eigenvalue k , which for obvious reasons is called the trivial eigenvalue λ_{triv} . A family of such graphs G with $h \rightarrow \infty$ is said to be a sequence of *expander graphs* if all other eigenvalues of their adjacency matrices are bounded away from $\lambda_{\text{triv}} = k$ by a fixed amount.⁵ In particular, no other eigenvalue is equal to k ; this implies the graph is connected. A *Ramanujan graph* [29] is a special type of expander which has $|\lambda| \leq 2\sqrt{k-1}$ for any nontrivial eigenvalue

⁵ Expansion is usually phrased in terms of the number of neighbors of subsets of G , but the spectral condition here is equivalent for k -regular graphs and also more useful for our purposes.

which is not equal to $-k$ (this last possibility happens if and only if the graph is bipartite). The supersingular isogeny graphs in Appendix A are sometimes Ramanujan, while the ordinary isogeny graphs in Section 2.1 do not qualify, partly because their degree is not bounded. Nevertheless, they still share the most important properties of expanders as far as our applications are concerned. In particular their degree k grows slowly (as a polynomial in $\log |\mathcal{V}|$), and they share a qualitatively similar eigenvalue separation: instead the nontrivial eigenvalues λ can be arranged to be $O(k^{1/2+\varepsilon})$ for any desired value of $\varepsilon > 0$. Since our goal is to establish a polynomial time reduction, this enlarged degree bound is natural, and in fact necessary for obtaining expanders from *abelian* Cayley graphs [1]. Obtaining *any* nontrivial *exponent* $\beta < 1$ satisfying $\lambda = O(k^\beta)$ is a key challenge for many applications, and accordingly we shall focus on a type of graphs we call “nearly Ramanujan” graphs: families of graphs whose nontrivial eigenvalues λ satisfy that bound.

A fundamental use of expanders is to prove the rapid mixing of the random walk on \mathcal{V} along the edges \mathcal{E} . The following rapid mixing result is standard but we present it below for convenience. For more information, see [5, 28, 40].

Proposition 3.1 *Let G be a regular graph of degree k on h vertices. Suppose that the eigenvalue λ of any nonconstant eigenvector satisfies the bound $|\lambda| \leq c$ for some $c < k$. Let S be any subset of the vertices of G , and x be any vertex in G . Then a random walk of any length at least $\frac{\log 2h/|S|^{1/2}}{\log k/c}$ starting from x will land in S with probability at least $\frac{|S|}{2h} = \frac{|S|}{2|G|}$.*

Proof. There are k^r random walks of length r starting from x . One would expect in a truly random situation that roughly $\frac{|S|}{h}k^r$ of these land in S . The lemma asserts that for $r \geq \frac{\log 2h/|S|^{1/2}}{\log k/c}$ at least half that number of walks in fact do. Denoting the characteristic functions of S and $\{x\}$ as χ_S and $\chi_{\{x\}}$, respectively, we count that

$$\#\{\text{walks of length } r \text{ starting at } x \text{ and landing in } S\} = \langle \chi_S, A^r \chi_{\{x\}} \rangle, \quad (3.1)$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product of functions in $L^2(\mathcal{V})$. We estimate this as follows. Write the orthogonal decompositions of χ_S and $\chi_{\{x\}}$ as

$$\chi_S = \frac{|S|}{h} \mathbf{1} + u \quad \text{and} \quad \chi_{\{x\}} = \frac{1}{h} \mathbf{1} + w, \quad (3.2)$$

where $\mathbf{1}$ is the constant vector and $\langle u, \mathbf{1} \rangle = \langle w, \mathbf{1} \rangle = 0$. Then (3.1) equals the expected value of $\frac{|S|}{h}k^r$, plus the additional term $\langle u, A^r w \rangle$, which is bounded by $\|u\| \|A^r w\|$. Because $w \perp \mathbf{1}$ and the symmetric matrix A^r has spectrum bounded by c^r on the span of such vectors,

$$\|u\| \|A^r w\| \leq c^r \|u\| \|w\| \leq c^r \|\chi_S\| \|\chi_{\{x\}}\| = c^r |S|^{1/2}. \quad (3.3)$$

For our values of r this is at most half of $\frac{|S|}{h}k^r$, so indeed at least $\frac{1}{2} \frac{|S|}{h}k^r$ of the paths terminate in S as was required.

In our application the quantities k , $\frac{k}{k-c}$, and $\frac{h}{|S|}$ will all be bounded by polynomials in $\log(h)$. Under these hypotheses, the probability is at least $1/2$ that some $\text{polylog}(h)$ trials of random walks of $\text{polylog}(h)$ length starting from x will reach S at least once. This mixing estimate is the source of our polynomial time random reducibility (Corollary 1.2).

4 Spectral Properties of the Isogeny Graph

4.1 Navigating the Isogeny Graph

Let \mathcal{G} be as in Section 2.1. The isogeny graph \mathcal{G} has exponentially many nodes and thus is too large to be stored. However, given a curve E and a prime ℓ , it is possible to efficiently compute the curves which are connected to E by an isogeny of degree ℓ . These curves E' have j -invariants which can be found by solving the modular polynomial relation $\Phi_\ell(j(E), j(E')) = 0$; the cost of this step is $O(\ell^3)$ field operations [11, 11.6]. Given the j -invariants, the isogenies themselves can then be obtained using the algorithms of [10] (or [26, 27] when the characteristic of the field is small). In this way, it is possible to navigate the isogeny graph locally without computing the entire graph. We shall see that it suffices to have the degree of the isogenies in the graph be bounded by $(\log q)^{2+\delta}$ to assure the Ramanujan properties required for \mathcal{G} to be an expander.

4.2 θ -Functions and Graph Eigenvalues

The graph \mathcal{H} (and therefore also the isomorphic graph \mathcal{G}) has one node for each ideal class of \mathcal{O} . Therefore, the total number of nodes in the graph \mathcal{G} is the ideal class number of the order \mathcal{O} , and the vertices \mathcal{V} can be identified with ideal class representatives $\{\alpha_1, \dots, \alpha_h\}$. Using the isomorphism between \mathcal{G} and \mathcal{H} , we see that the generating function $\sum M_{\alpha_i, \alpha_j}(n)q^n$ for degree n isogenies between the vertices α_i and α_j of \mathcal{G} is given by

$$\sum_{n=1}^{\infty} M_{\alpha_i, \alpha_j}(n)q^n := \frac{1}{e} \sum_{z \in \alpha_i^{-1}\alpha_j} q^{N(z)/N(\alpha_i^{-1}\alpha_j)}, \quad (4.1)$$

where e is the number of units in \mathcal{O} (which always equals 2 for $\text{disc}(\mathcal{O}) > 4$). The sum on the righthand side depends only on the ideal class of the fractional ideal $\alpha_i^{-1}\alpha_j$; by viewing the latter as a lattice in \mathbb{C} , we see that $N(z)/N(\alpha_i^{-1}\alpha_j)$ is a quadratic form of discriminant D where $D := \text{disc}(\mathcal{O})$ [4, p. 142]. That means this sum is a θ -series, accordingly denoted as $\theta_{\alpha_i^{-1}\alpha_j}(q)$. It is a holomorphic modular form of weight 1 for the congruence subgroup $\Gamma_0(|D|)$ of $SL(2, \mathbb{Z})$, transforming according to the character $\left(\frac{D}{\cdot}\right)$ (see [19, Theorem 10.9]).

Before discussing exactly which degrees of isogenies to admit into our isogeny graph \mathcal{G} , let us first make some remarks about the simpler graph on $\mathcal{V} = \{\alpha_1, \dots, \alpha_h\}$ whose edges represent isogenies of degree exactly equal to n . Its adjacency matrix is of course the $h \times h$ matrix $M(n) = [M_{\alpha_i, \alpha_j}(n)]_{\{1 \leq i, j \leq h\}}$

defined by series coefficients in (4.1). It can be naturally viewed as an operator which acts on functions on $\mathcal{V} = \{\alpha_1, \dots, \alpha_h\}$, by identifying them with h -vectors according to this labeling. We will now simultaneously diagonalize all $M(n)$, or what amounts to the same, diagonalize the matrix $A_q = \sum_{n \geq 1} M(n)q^n$ for any value of $q < 1$ (where the sum converges absolutely). The primary reason this is possible is that for each fixed n this graph is an abelian Cayley graph on the ideal class group $\text{Cl}(\mathcal{O})$, with generating set equal to those classes α_i which represent an n -isogeny. The eigenfunctions of the adjacency matrix of an abelian Cayley graph are always given by characters of the group (viewed as functions on the graph), and their respective eigenvalues are sums of these characters over the generating set. This can be seen directly in our circumstance as follows. The ij -th entry of A_q is $\frac{1}{e}\theta_{\alpha_i^{-1}\alpha_j}(q)$, which we recall depends only on the ideal class of the fractional ideal $\alpha_i^{-1}\alpha_j$. If χ is any character of $\text{Cl}(\mathcal{O})$, viewed as the h -vector whose i -th entry is $\chi(\alpha_i)$, then the i -th entry of the vector $A_q\chi$ may be evaluated through matrix multiplication as

$$(A_q\chi)(\alpha_i) = \frac{1}{e} \sum_{\alpha_j \in \text{Cl}(\mathcal{O})} \theta_{\alpha_i^{-1}\alpha_j}(q) \chi(\alpha_j) = \frac{1}{e} \left(\sum_{\alpha_j \in \text{Cl}(\mathcal{O})} \chi(\alpha_j) \theta_{\alpha_j}(q) \right) \chi(\alpha_i), \quad (4.2)$$

where in the last equality we have reindexed $\alpha_j \mapsto \alpha_i \alpha_j$ using the group structure of $\text{Cl}(\mathcal{O})$. Therefore χ is in fact an eigenvector of the matrix eA_q , with eigenvalue equal to the sum of θ -functions enclosed in parentheses, known as a *Hecke θ -function* (see [19, §12]). These, which we shall denote $\theta_\chi(q)$, form a more natural basis of modular forms than the ideal class θ -functions θ_{α_j} because they are in fact Hecke eigenforms. Using (4.1), the L -functions of these Hecke characters can be written as

$$\begin{aligned} L(s, \chi) &= L(s, \theta_\chi) = \sum_{\substack{\text{integral ideals } \mathfrak{a} \subset K}} \chi(\mathfrak{a}) (N\mathfrak{a})^{-s} = \sum_{n=1}^{\infty} a_n(\chi) n^{-s}, \\ \text{where } a_n(\chi) &= \sum_{\substack{\text{integral ideals } \mathfrak{a} \subset K \\ N\mathfrak{a}=n}} \chi(\mathfrak{a}) \end{aligned} \quad (4.3)$$

is in fact simply the eigenvalue of $eM(n)$ for the eigenvector formed from the character χ as above, which can be seen by isolating the coefficient of q^n in the sum on the righthand side of (4.2).

4.3 Eigenvalue Separation under the Generalized Riemann Hypothesis

Our isogeny graph is a superposition of the previous graphs $M(n)$, where n is a prime bounded by a parameter m (which we recall is $(\log q)^{2+\delta}$ for some fixed $\delta > 0$). This corresponds to a graph on the elliptic curves represented by ideal classes in an order \mathcal{O} of $K = \mathbb{Q}(\sqrt{d})$, whose edges represent isogenies of prime degree $\leq m$. The graphs with adjacency matrices $\{M(p) \mid p \leq m\}$ above share

common eigenfunctions (the characters χ of $\text{Cl}(\mathcal{O})$), and so their eigenvalues are

$$\lambda_\chi = \frac{1}{e} \sum_{p \leq m} a_p(\chi) = \frac{1}{e} \sum_{p \leq m} \sum_{\substack{\text{integral ideals } \mathfrak{a} \subset K \\ N\mathfrak{a} = p}} \chi(\mathfrak{a}). \quad (4.4)$$

When χ is the trivial character, λ_{triv} equals the degree of the regular graph \mathcal{G} . Since roughly half of rational primes p split in K , and those which do split into two ideals of norm p , λ_{triv} is roughly $\frac{\pi(m)}{e} \sim \frac{m}{e \log m}$ by the prime number theorem. This eigenvalue is always the largest in absolute value, as can be deduced from (4.4), because $|\chi(\mathfrak{a})|$ always equals 1 when χ is the trivial character. For the polynomial mixing of the random walk in Theorem 1.1 we will require a separation between the trivial and nontrivial eigenvalues of size $1/\text{polylog}(q)$. This would be the case, for example, if for each nontrivial character χ there merely exists one ideal \mathfrak{a} of prime norm $\leq m$ with $\text{Re } \chi(\mathfrak{a}) \leq 1 - \frac{1}{\text{polylog}(q)}$. This is analogous to the problem of finding a small prime nonresidue modulo, say, a large prime Q , where one merely needs to find any cancellation at all in the character sum $\sum_{p \leq m} \left(\frac{p}{Q}\right)$. However, the latter requires a strong assumption from analytic number theory, such as the Generalized Riemann Hypothesis (GRH). In the next section we will accordingly derive such bounds for λ_χ , under the assumption of GRH. As a consequence of the more general Lemma 5.3 we will show the following.

Lemma 4.1 *Let $D < 0$ and let \mathcal{O} be the quadratic order of discriminant D . If χ is a nontrivial ideal class character of \mathcal{O} , then the Generalized Riemann Hypothesis for $L(s, \chi)$ implies that the sum (4.4) is bounded by $O(m^{1/2} \log |mD|)$ with an absolute implied constant.*

Proof (of Theorem 1.1). There are only finitely many levels for q less than any given bound, so it suffices to prove the theorem for q large and $p(x) = x^{2+\delta}$, where $\delta > 0$ is fixed. The eigenvalues of the adjacency matrix for a given level are given by (4.4). Recall that $|D| \leq 4q$ and $\lambda_{\text{triv}} \sim \frac{m}{e \log m}$. With our choice of $m = p(\log q)$, the bound for the nontrivial eigenvalues in Lemma 4.1 is $\lambda_\chi = O(\lambda_{\text{triv}}^\beta)$ for any $\beta > \frac{1}{2} + \frac{1}{\delta+2}$. That means indeed our isogeny graphs are expanders for q large; the random walk assertion follows from this bound and Proposition 3.1.

Proof (of Corollary 1.2). The Theorem shows that a random walk from any fixed curve E probabilistically reaches the proportion where the algorithm A succeeds, in at most $\text{polylog}(q)$ steps. Since each step is a low degree isogeny, their composition can be computed in $\text{polylog}(q)$ steps. Even though the degree of this isogeny might be large, the degrees of each step are small. This provides the random polynomial time reduction of DLOG along successive curves in the random walk, and hence from E to a curve for which the algorithm A succeeds.

5 The Prime Number Theorem for Modular Form L -functions

In this section we prove Lemma 4.1, assuming the Generalized Riemann Hypothesis (GRH) for the L -functions (4.3). Our argument is more general, and in fact gives estimates for sums of the form $\sum_{p \leq m} a_p$, where a_p are the prime coefficients of any L -function. This can be thought of as an analog of the Prime Number Theorem because for the simplest L -function, $\zeta(s)$, $a_p = 1$ and this sum is in fact exactly $\pi(m)$. As a compromise between readability and generality, we will restrict the presentation here to the case of modular form L -functions (including (4.3)). Background references for this section include [19, 20, 35]; for information about more general L -functions see also [14, 39].

We shall now consider a classical holomorphic modular form f , with Fourier expansion $f(z) = \sum_{n=0}^{\infty} c_n e^{2\pi i n z}$. We will assume that f is a Hecke eigenform, since this condition is met in the situation of Lemma 4.1 (see the comments between (4.2) and (4.3)). It is natural to study the renormalized coefficients $a_n = n^{-(k-1)/2} c_n$, where $k \geq 1$ is the weight of f (in Section 4.2 $k = 1$, so $a_n = c_n$). The L -function of such a modular form can be written as the Dirichlet series $L(s, f) = \sum_{n=1}^{\infty} a_n n^{-s} = \prod_p (1 - \alpha_p p^{-s})^{-1} (1 - \beta_p p^{-s})^{-1}$, the last equality using the fact that f is a Hecke eigenform. The L -function $L(s, f)$ is entire when f is a cusp form (e.g. $a_0 = 0$). The Ramanujan conjecture (in this case a theorem of [6] and [7]) asserts that $|\alpha_p|, |\beta_p| \leq 1$.

Lemma 4.1 is concerned with estimates for the sums

$$S(m, f) := \sum_{p \leq m} a_p. \quad (5.1)$$

As with the prime number theorem, it is more convenient to instead analyze the weighted sum

$$\psi(m, f) := \sum_{p^k} b_{p^k} \log p \quad (5.2)$$

over prime powers, where the coefficients b_n are those appearing in the Dirichlet series for $-\frac{L'}{L}(s)$:

$$-\frac{L'}{L}(s) = \sum_{n=1}^{\infty} b_n \Lambda(n) n^{-s} = \sum_{p, k} b_{p^k} \log(p) p^{-k s},$$

i.e., $b_{p^k} = \alpha_p^k + \beta_p^k$.

Lemma 5.1 *For a holomorphic modular form f one has*

$$\psi(m, f) = \sum_{p \leq m} a_p \log p + O(m^{1/2}).$$

Proof. The error term represents the contribution of proper prime powers. Since $|b_{p^k}| \leq 2$, it is bounded by twice

$$\sum_{\substack{p^k \leq m \\ k \geq 2}} \log p = \sum_{\substack{p \leq m^{1/2} \\ 2 \leq k \leq \frac{\log m}{\log p}}} \log p \leq \sum_{p \leq m^{1/2}} \log p \frac{\log m}{\log p} \leq \pi(m^{1/2}) \log m, \quad (5.3)$$

which is $O(m^{1/2})$ by the Prime Number Theorem.

Lemma 5.2 (*Iwaniec* [20, p. 114]) *Assume that f is a holomorphic modular cusp form of level⁶ N and that $L(s, f)$ satisfies GRH. Then $\psi(m, f) = O(m^{1/2} \log(m) \log(mN))$.*

We deduce that $S'(m, f) := \sum_{p \leq m} a_p \log p = O(m^{1/2} \log(m) \log(mN))$. Finally we shall estimate the sums $S(m, f)$ from (5.1) by removing the $\log(m)$ using a standard partial summation argument.

Lemma 5.3 *Suppose that f is a holomorphic modular cusp form of level N and $L(s, f)$ satisfies GRH. Then $S(m, f) = O(m^{1/2} \log(mN))$.*

Proof. First define \tilde{a}_p to be a_p , if p is prime, and 0 otherwise. Then

$$\sum_{p \leq m} a_p = \sum_{p \leq m} [\tilde{a}_p \log p] \frac{1}{\log p} = \sum_{n \leq m} [\tilde{a}_n \log n] \frac{1}{\log n}.$$

By partial summation over $2 \leq n \leq m$, we then find

$$\begin{aligned} \sum_{p \leq m} a_p &= \sum_{n < m} S'(n, f) \left(\frac{1}{\log(n)} - \frac{1}{\log(n+1)} \right) + \frac{S'(m, f)}{\log m} \\ &\ll \sum_{n < m} \left(n^{1/2} \log(n) \log(nN) \right) \left| \frac{d}{dn} ((\log n)^{-1}) \right| + m^{1/2} \log(mN) \\ &\ll \sum_{n < m} n^{1/2} \log(n) \log(nN) \frac{1}{n(\log n)^2} + m^{1/2} \log(mN), \end{aligned}$$

so in fact $S(m, f) = \sum_{p \leq m} a_p = O(m^{1/2} \log(mN))$.

All the implied constants in these 3 lemmas are absolute. Some useful estimates for them may be found in [2].

5.1 Subexponential Reductions via Lindelöf Hypothesis

In the previous lemma we have assumed GRH. It seems very difficult to get a corresponding unconditional bound for $S(m, f)$. However, a slightly weaker

⁶ Actually in [20] N equals the conductor of the L -function, which in general may be smaller than the level. The lemma is of course nevertheless valid.

statement can be proven by assuming only the Lindelöf hypothesis (which is a consequence of GRH). Namely, one has that $\sum_{n \leq m} a_n = O_\varepsilon(m^{1/2+\varepsilon}N^\varepsilon)$, for any $\varepsilon > 0$ ([19, (5.61)]). The fact that this last sum is over all $n \leq m$, not just primes, is not of crucial importance for our application. However, the significant difference here is that the dependence on N is not polynomial in $\log N$, but merely subexponential. This observation can be used to weaken the hypothesis in Theorem 1.1 and Corollary 1.2 from GRH to the Lindelöf hypothesis, at the expense of replacing “polynomial” by “subexponential.”

6 Distribution of c_π

Theorem 1.1 and Corollary 1.2 are statements about individual levels. As we mentioned in Section 1.1, our random reducibility result extends between two levels as long as the levels satisfy the requirement that their conductors differ by polynomially smooth amounts. In this section we explore this extension in more detail, and explain why the above requirement is typically satisfied.

It was mentioned after Theorem 2.2 that the largest possible conductor difference is c_π , which is the largest square factor of $d_\pi = \text{Trace}(E)^2 - 4q$. In principle this factor could be as large as $2\sqrt{q}$, though statistically speaking most integers (a proportion of $\frac{6}{\pi^2} \approx .61$) are square-free, explaining why c_π is very often 1 or at least fairly small [44]. This means, for example, that most randomly selected elliptic curves have an isogeny class consisting of only one level.

When an isogeny class consists of multiple levels, we need to be able to construct vertical isogenies between levels in order to conclude that DLOG instances between the levels are randomly reducible to each other. The fastest known algorithm for constructing vertical isogenies between two levels, due to Kohel [23], has runtime $O(\ell^4)$, where ℓ is the largest prime dividing the conductor of one of the levels, but not the other. Any two levels which can be efficiently bridged via Kohel’s algorithm can be considered as one unit for the purposes of random reducibility. Accordingly, polynomial time random reducibility holds within an isogeny class if c_π for that isogeny class is polynomially smooth.

With this in mind, we will now determine a heuristic estimate for the expected size of the largest prime factor $P(c_\pi)$ of c_π , i.e., the largest prime which divides d_π to order at least 2. The trace $t = \text{Trace}(E)$, when sampled over random elliptic curves, is thought to have a fairly uniform distribution over most of the Hasse interval. This serves to predict the useful heuristic that $-d_\pi = 4q - t^2$ is typically of size q (see for example [25, 41]). Assuming that, the probability that $P(c_\pi)$ exceeds β can be loosely estimated as $O(1/\beta)$. This is because roughly a fraction of $\rho = \prod_{p > \beta}^{\sqrt{q}} (1 - p^{-2})$ integers of size q have no repeated prime factor $p > \beta$. It is easy to see that $\log(\rho) = O(\sum_{n > \beta} n^{-2}) = O(1/\beta)$, so that $1 - \rho = O(1/\beta)$ as suggested.

It follows that a randomly selected elliptic curve is extremely likely to have a small enough value of $P(c_\pi)$ to allow for random reducibility throughout its entire isogeny class. This explains why in Figure 1 all of the randomly generated curves have $P(c_\pi) = 1$, except for one curve which has $P(c_\pi) = 3$.

Finally, let us consider the situation where a *non-random* curve is deliberately selected so as to have a large value of c_π . Currently the only known methods for constructing such curves is to use complex multiplication methods [3, Ch. VIII] to construct curves with a predetermined number of points chosen to ensure that c_π is almost as large as $\sqrt{d_\pi}$. Some convenient examples of such curves are the Koblitz curves listed in the NIST FIPS 186-2 document [36], which we have also tabulated in Figure 1. Since these curves all have complex multiplication by the field $K = \mathbb{Q}(\sqrt{-7})$, the discriminants of these curves are of the form $d_\pi = -7c_\pi^2$. If we assume that c_π behaves as a random integer of size $\sqrt{d_\pi}$, which is roughly \sqrt{q} , then the distribution of $P(c_\pi)$ is governed by the usual smoothness bounds for large integers [44], and hence is typically too large to permit efficient application of Kohel’s algorithm for navigating between levels. Thus we cannot prove random reducibility from a theoretical standpoint for all of the elliptic curves within the isogeny class $S_{N,q}$ of such a specially constructed curve. However, in practice only a small subset of the elliptic curves in $S_{N,q}$ are efficiently constructible using the complex multiplication method (or any other presently known method), and this subset coincides exactly with the subcollection of levels in $S_{N,q}$ which are accessible from the top level (where $\text{End}(E) = \mathcal{O}_K$) using Kohel’s algorithm. Pending future developments, it therefore remains true that all of the special curves that we can construct within an isogeny class have equivalent DLOG problems in the random reducible sense.

Acknowledgments: It is a pleasure to thank William Aiello, Michael Ben-Or, Dan Boneh, Brian Conrad, Henryk Iwaniec, Dimitar Jetchev, Neal Koblitz, Alexander Lubotzky, Peter Sarnak, Adi Shamir, and Yacov Yacobi for their discussions and helpful comments. We are also indebted to Peter Montgomery for his factoring assistance in producing Figure 1.

References

1. Noga Alon and Yuval Roichman, *Random Cayley graphs and expanders*, Random Structures Algorithms **5** (1994), no. 2, 271–284.
2. Eric Bach and Jonathan Sorenson, *Explicit bounds for primes in residue classes*, Math. Comp. **65** (1996), no. 216, 1717–1735.
3. 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.
4. David A. Cox, *Primes of the form $x^2 + ny^2$* , A Wiley-Interscience Publication, John Wiley & Sons Inc., New York, 1989.
5. Giuliana Davidoff, Peter Sarnak, and Alain Valette, *Elementary number theory, group theory, and Ramanujan graphs*, London Mathematical Society Student Texts, vol. 55, Cambridge University Press, Cambridge, 2003.
6. Pierre Deligne, *La conjecture de Weil. I*, Inst. Hautes Études Sci. Publ. Math. (1974), no. 43, 273–307 (French).
7. Pierre Deligne and Jean-Pierre Serre, *Formes modulaires de poids 1*, Ann. Sci. École Norm. Sup. (4) **7** (1974), 507–530 (1975) (French).
8. Max Deuring, *Die Typen der Multiplikatorenringe elliptischer Funktionenkörper*, Abh. Math. Sem. Hansischen Univ. **14** (1941), 197–272 (German).

9. Martin Eichler, *Quaternäre quadratische Formen und die Riemannsche Vermutung für die Kongruenzzetafunktion*, Arch. Math. **5** (1954), 355–366 (German).
10. Mireille Fouquet and François Morain, *Isogeny volcanoes and the SEA algorithm*, Algorithmic number theory (Sydney, 2002), 2002, pp. 276–291.
11. Steven D. Galbraith, *Constructing isogenies between elliptic curves over finite fields*, LMS J. Comput. Math. **2** (1999), 118–138 (electronic).
12. 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.
13. P. Gaudry, F. Hess, and N. P. Smart, *Constructive and destructive facets of Weil descent on elliptic curves*, J. Cryptology **15** (2002), no. 1, 19–46.
14. Stephen S. Gelbart and Stephen D. Miller, *Riemann’s zeta function and beyond*, Bull. Amer. Math. Soc. (N.S.) **41** (2004), no. 1, 59–112 (electronic).
15. Benedict H. Gross, *Heights and the special values of L -series*, Number theory (Montreal, Que., 1985), 1987, pp. 115–187.
16. D. Harkins and D. Carrel, *The Internet key exchange (IKE)*, Technical Report IETF RFC 2409, November 1998. <http://www.ietf.org/rfc/rfc2409.txt>.
17. Jun-ichi Igusa, *Fibre systems of Jacobian varieties. III. Fibre systems of elliptic curves*, Amer. J. Math. **81** (1959), 453–476.
18. Yasutaka Ihara, *Discrete subgroups of $PL(2, k_\varphi)$* , Algebraic Groups and Discontinuous Subgroups (Proc. Sympos. Pure Math., Boulder, Colo., 1965), 1966, pp. 272–278.
19. Henryk Iwaniec, *Topics in classical automorphic forms*, Graduate Studies in Mathematics, vol. 17, American Mathematical Society, Providence, RI, 1997.
20. Henryk Iwaniec and Emmanuel Kowalski, *Analytic number theory*, American Mathematical Society Colloquium Publications, vol. 53, American Mathematical Society, Providence, RI, 2004.
21. M. Jerrum and A. Sinclair, *Conductance and the rapid mixing property for Markov chains: the approximation of permanent resolved*, ACM Symposium on Theory of Computing (May 1988), 235–243.
22. Neal Koblitz, *Elliptic curve cryptosystems*, Math. Comp. **48** (1987), no. 177, 203–209.
23. David Kohel, *Endomorphism rings of elliptic curves over finite fields*, University of California, Berkeley, 1996, Ph.D thesis.
24. Serge Lang, *Elliptic functions*, 2nd ed., Graduate Texts in Mathematics, vol. 112, Springer-Verlag, New York, 1987. With an appendix by J. Tate.
25. H. W. Lenstra Jr., *Factoring integers with elliptic curves*, Ann. of Math. (2) **126** (1987), no. 3, 649–673.
26. Reynald Lercier, *Computing isogenies in \mathbf{F}_{2^n}* , Algorithmic number theory (Talence, 1996), 1996, pp. 197–212.
27. R. Lercier and F. Morain, *Algorithms for computing isogenies between elliptic curves*, Computational perspectives on number theory (Chicago, IL, 1995), 1998, pp. 77–96.
28. Alexander Lubotzky, *Discrete groups, expanding graphs and invariant measures*, Progress in Mathematics, vol. 125, Birkhäuser Verlag, Basel, 1994.
29. A. Lubotzky, R. Phillips, and P. Sarnak, *Ramanujan graphs*, Combinatorica **8** (1988), no. 3, 261–277.
30. 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.

31. Alfred J. Menezes, Paul C. van Oorschot, and Scott A. Vanstone, *Handbook of applied cryptography*, CRC Press Series on Discrete Mathematics and its Applications, CRC Press, Boca Raton, FL, 1997. With a foreword by Ronald L. Rivest.
32. Alfred Menezes, Edlyn Teske, and Annegret Weng, *Weak fields for ECC*, Topics in cryptography—CT-RSA 2004, 2004, pp. 366–386.
33. J.-F. Mestre, *La méthode des graphes. Exemples et applications*, Proceedings of the international conference on class numbers and fundamental units of algebraic number fields (Katata, 1986), 1986, pp. 217–242 (French).
34. Victor S. Miller, *Use of elliptic curves in cryptography*, Advances in cryptography—CRYPTO '85 (Santa Barbara, Calif., 1985), 1986, pp. 417–426.
35. M. Ram Murty, *Problems in analytic number theory*, Graduate Texts in Mathematics, vol. 206, Springer-Verlag, New York, 2001. Readings in Mathematics.
36. National Institute of Standards and Technology, *Digital Signature Standard (DSS)*, Technical Report FIPS PUB 186–2, January 2000.
<http://csrc.nist.gov/publications/fips/>.
37. Arnold K. Pizer, *Ramanujan graphs and Hecke operators*, Bull. Amer. Math. Soc. (N.S.) **23** (1990), no. 1, 127–137.
38. ———, *Ramanujan graphs*, Computational perspectives on number theory (Chicago, IL, 1995), 1998, pp. 159–178.
39. Zeév Rudnick and Peter Sarnak, *Zeros of principal L -functions and random matrix theory*, Duke Math. J. **81** (1996), no. 2, 269–322. A celebration of John F. Nash, Jr.
40. Peter Sarnak, *Some applications of modular forms*, Cambridge Tracts in Mathematics, vol. 99, Cambridge University Press, Cambridge, 1990.
41. Jean-Pierre Serre, *Abelian l -adic representations and elliptic curves*, McGill University lecture notes written with the collaboration of Willem Kuyk and John Labute, W. A. Benjamin, Inc., New York-Amsterdam, 1968.
42. Joseph H. Silverman, *The arithmetic of elliptic curves*, Graduate Texts in Mathematics, vol. 106, Springer-Verlag, New York, 1994.
43. John Tate, *Endomorphisms of abelian varieties over finite fields*, Invent. Math. **2** (1966), 134–144.
44. Gérard Tenenbaum, *Introduction to analytic and probabilistic number theory*, Cambridge Studies in Advanced Mathematics, vol. 46, Cambridge University Press, Cambridge, 1995. Translated from the second French edition (1995) by C. B. Thomas.
45. Edlyn Teske, *An elliptic curve trapdoor system (extended abstract)*, High primes and misdemeanours: lectures in honour of the 60th birthday of Hugh Cowie Williams, 2004, pp. 341–352.
46. Lawrence C. Washington, *Elliptic curves*, Discrete Mathematics and its Applications (Boca Raton), Chapman & Hall/CRC, Boca Raton, FL, 2003.

A Supersingular Case

In this appendix we discuss the isogeny graphs for supersingular elliptic curves and prove Theorem 1.1 in this setting. The isogeny graphs were first considered by Mestre [33], and were shown by Pizer [37, 38] to have the Ramanujan property. Curiously, the actual graphs were first described by Ihara [18] in 1965, but not noticed to be examples of expander graphs until much later. We have decided

to give an account here for completeness, mainly following Pizer's arguments. The isogeny graphs we will present here differ from those in the ordinary case in that they are *directed*. This will cause no serious practical consequences, because one can arrange that only a bounded number of edges in these graphs will be unaccompanied by a reverse edge. Also, the implication about rapid mixing used for Theorem 1.1 carries over as well in the directed setting with almost no modification. It is instructive to compare the proofs for the ordinary and supersingular cases, in order to see how GRH plays a role analogous to the Ramanujan conjectures.

Every $\overline{\mathbb{F}}_q$ -isomorphism class of supersingular elliptic curves in characteristic p is defined over either \mathbb{F}_p or \mathbb{F}_{p^2} [42], so it suffices to fix $\mathbb{F}_q = \mathbb{F}_{p^2}$ as the field of definition for this discussion. Thus, in contrast to ordinary curves, there is a finite bound g on the number of isomorphism classes that can belong to any given isogeny class (this bound is in fact the genus of the modular curve $X_0(p)$, which is roughly $\frac{p+1}{12}$). It turns out that all isomorphism classes of supersingular curves defined over \mathbb{F}_{p^2} belong to the same isogeny class [33]. Because the number of supersingular curves up to isomorphism is so much smaller than the number of ordinary curves up to isomorphism, correspondingly fewer of the edges need to be included in order to form a Ramanujan graph. For a fixed prime value of $\ell \neq p$, we define the vertices of the supersingular isogeny graph \mathcal{G} to consist of these g isomorphism classes, with directed edges indexed by equivalence classes of degree- ℓ isogenies as defined below. In fact, we will prove that \mathcal{G} is a directed $k = \ell + 1$ -regular graph satisfying the Ramanujan bound of $|\lambda| \leq 2\sqrt{\ell} = 2\sqrt{k-1}$ for the nontrivial eigenvalues of its adjacency matrix. The degree ℓ in particular may be taken to be as small as 2 or 3.

For the definition of the equivalence classes of isogenies — as well as later for the proofs — we now need to recall the structure of the endomorphism rings of supersingular elliptic curves. In contrast to the ordinary setting (Section 2), the endomorphism ring $\text{End}(E)$ is a maximal order in the quaternion algebra $R = \mathbb{Q}_{p,\infty}$ ramified at p and ∞ . Moreover, isomorphism classes of supersingular curves E_i isogenous to E are in 1-1 correspondence with the left ideal classes $I_i := \text{Hom}(E_i, E)$ of R . As in Section 2.1, call two isogenies $\phi_1, \phi_2: E_i \rightarrow E_j$ equivalent if there exists an automorphism α of E_j such that $\phi_2 = \alpha\phi_1$. Under this relation, the set of equivalence classes of isogenies from E_i to E_j is equal to $I_j^{-1}I_i$ modulo the units of I_j . This correspondence is degree preserving, in the sense that the degree of an isogeny equals the reduced norm of the corresponding element in $I_j^{-1}I_i$, normalized by the norm of $I_j^{-1}I_i$ itself. This is the notion of equivalence class of isogenies referred to in the definition of \mathcal{G} in the previous paragraph. Thus, for any integer n , the generating function for the number $M_{ij}(n)$ of equivalence classes of degree n isogenies from E_i to E_j (i.e., the number of edges between vertices representing elliptic curves E_i and E_j) is given by

$$\sum_{n=0}^{\infty} M_{ij}(n) q^n := \frac{1}{e_j} \sum_{\alpha \in I_j^{-1}I_i} q^{N(\alpha)/N(I_j^{-1}I_i)}, \quad (\text{A.1})$$

where e_j is the number of units in I_j (equivalently, the number of automorphisms of E_j). One knows that $e_j \leq 6$, and in fact $e_j = 2$ except for at most two values of j – see the further remarks at the end of this appendix. Proofs for the statements in this paragraph can be found in [15, 38].

The θ -series on the righthand side of (A.1) is a weight 2 modular form for the congruence subgroup $\Gamma_0(p)$, and the matrices

$$B(n) := \begin{pmatrix} M_{11}(n) & \cdots & M_{1g}(n) \\ \vdots & \ddots & \vdots \\ M_{g1}(n) & \cdots & M_{gg}(n) \end{pmatrix}$$

(called Brandt matrices) are simultaneously both the n -th Fourier coefficients of various modular forms, as well the adjacency matrices for the graph \mathcal{G} . A fundamental property of the Brandt matrices $B(n)$ is that they represent the action of the n^{th} Hecke operator $T(n)$ on a certain basis of modular forms of weight 2 for $\Gamma_0(p)$ (see [37]). Thus the eigenvalues of $B(n)$ are given by the n^{th} coefficients of the weight-2 Hecke eigenforms for $\Gamma_0(p)$. These eigenforms include a single Eisenstein series, with the rest being cusp forms. Now we suppose that $n = \ell$ is prime (mainly in order to simplify the following statements). The n^{th} Hecke eigenvalue of the Eisenstein series is $n+1$, while those of the cusp forms are bounded in absolute value by $2\sqrt{n}$ according to the Ramanujan conjectures (in this case a theorem of Eichler [9] and Igusa [17]). Thus the adjacency matrix of \mathcal{G} has trivial eigenvalue equal to $\ell+1$ (the degree k), and its nontrivial eigenvalues indeed satisfy the Ramanujan bound $|\lambda| \leq 2\sqrt{k} - 1$.

Finally, we conclude with some comments about the potential asymmetry of the matrix $B(n)$. This is due to the asymmetry in the definition of equivalence classes of isogenies. Indeed, if $\text{Aut}(E_1)$ and $\text{Aut}(E_2)$ are different, then two isogenies $E_1 \rightarrow E_2$ can sometimes be equivalent even when their dual isogenies are not equivalent. This problem arises only if one of the curves E_i has complex multiplication by either $\sqrt{-1}$ or $e^{2\pi i/3}$, since otherwise the only possible automorphisms of E_i are the scalar multiplication maps ± 1 [42, §III.10]. In the supersingular setting, one can avoid curves with such unusually rich automorphism groups by choosing a characteristic p which splits in both $\mathbb{Z}[\sqrt{-1}]$ and $\mathbb{Z}[e^{2\pi i/3}]$, i.e., $p \equiv 1 \pmod{12}$ (see [37, Prop. 4.6]). In the case of ordinary curves, however, the quadratic orders $\mathbb{Z}[\sqrt{-1}]$ and $\mathbb{Z}[e^{2\pi i/3}]$ both have class number 1, which then renders the issue moot because the isogeny graphs corresponding to these levels each have only one node.