

Bits Security of the Elliptic Curve Diffie–Hellman Secret Keys

Dimitar Jetchev¹ and Ramarathnam Venkatesan^{2,3}

¹ Dept. of Mathematics, University of California at Berkeley, Berkeley, CA 94720
jetchev@math.berkeley.edu

² Microsoft Research, One Microsoft Way, Redmond WA 98052

³ Microsoft Research India Private Limited, "Scientia", No:196/36,
2nd Main Road, Sadashivnagar, Bangalore – 560080, India
venkie@microsoft.com

Abstract. We show that the least significant bits (LSB) of the elliptic curve Diffie–Hellman secret keys are hardcore. More precisely, we prove that if one can efficiently predict the LSB with non-negligible advantage on a polynomial fraction of all the curves defined over a given finite field \mathbb{F}_p , then with polynomial factor overhead, one can compute the entire Diffie–Hellman secret on a polynomial fraction of all the curves over the same finite field. Our approach is based on random self-reducibility (assuming GRH) of the Diffie–Hellman problem among elliptic curves of the same order. As a part of the argument, we prove a refinement of H. W. Lenstra’s lower bounds on the sizes of the isogeny classes of elliptic curves, which may be of independent interest.

Key words: elliptic curves, Diffie–Hellman function, hardcore bits, expander graphs, isogeny graphs, random self-reducibility

1 Introduction

The Diffie–Hellman protocol for key exchange [16] is based on the hardness of computing the function $\text{DH}_g(g^u, g^v) = g^{uv}$, where g is a fixed generator of the multiplicative group of a finite field \mathbb{F}_p , and $1 \leq u, v \leq p-1$ are integers. A natural question is whether one can compute some of the bits of g^{uv} given g, g^u, g^v . It is unknown if predicting partial information with significant advantage over a random guess will lead to a compromise of the Diffie–Hellman function. Boneh and Venkatesan [2], [25] have shown that if one is able to compute (in time polynomial in $\log p$) the $5\sqrt{\log p}$ most significant bits of g^{uv} for every input (g^u, g^v) then one can compute (in polynomial time) the entire shared secret key g^{uv} . For motivation, note that g^{uv} may be 1024 bits long, but one may want to use the least significant 128 bits of g^{uv} as a block cipher key. Thus, it is important to know that partial information is not computable or predictable with any significant advantage over a random guess. Another motivation stems from the fact that the methods used in [2] suggest attacks on cryptographic systems that reveal some information about g^{uv} to the attacker [8], [10], [18], [19], [20], [24], [26].

The analogous problem for elliptic curves studies the bit security of the following function:

Diffie–Hellman function: Let E be an elliptic curve over \mathbb{F}_p and let $P \in E$ be a point of prime order q . We define the Diffie–Hellman function as

$$\text{DH}_{E,P}(uP, vP) = uvP,$$

where $1 \leq u, v \leq q$ are integers. Moreover, we refer to the triple (P, uP, vP) as a *Diffie–Hellman triple* for E .

For simplicity, we restrict ourselves to *short Weierstrass equations (models)* of E , i.e., models of the form $y^2 = x^3 + ax + b$ with $a, b \in \mathbb{F}_p$ and $4a^3 + 27b^2 \neq 0$. By abuse of common terminology, an *elliptic curve* for us will be an \mathbb{F}_p -isomorphism class of short Weierstrass equations. It is not hard to see that the complexity of the Diffie–Hellman function is independent of the choice of the short Weierstrass equation for the elliptic curve E over \mathbb{F}_p . Indeed, given two different models W and W' for E over \mathbb{F}_p and an explicit isomorphism $\varphi : W \rightarrow W'$ and its inverse $\varphi^{-1} : W' \rightarrow W$, a Diffie–Hellman triple (P, uP, vP) on W is mapped to a Diffie–Hellman triple $(\varphi(P), u\varphi(P), v\varphi(P))$ on W' and therefore, if one can compute $uv\varphi(P)$, one would know uvP . Yet, if one wants to formalize the notion of security of single bits of the Diffie–Hellman function, one needs to choose a short Weierstrass model (it is not necessarily true any more that if one knows one bit of the Diffie–Hellman secret $uv\varphi(P)$ on W' then one can compute the corresponding bit of uvP on W).

Boneh and Shparlinski [1] have reduced (in time polynomial in $\log p$) the Diffie–Hellman problem on an elliptic curve E to the problem of predicting the LSB of the secret key uvP with non-negligible advantage over a random guess on a polynomial fraction of all short Weierstrass models for E . Alternatively, if one looks for a polynomial time reduction of the Diffie–Hellman problem to the problem of predicting partial information on the *same* short Weierstrass model W , some results have been established using Gröbner bases [12].

A more general and natural situation would be to consider an oracle \mathcal{A} that predicts the least significant bit of the Diffie–Hellman secret key for short Weierstrass models W chosen from a non-negligible subset G (i.e., from a $(\log p)^{O(1)}$ -fraction) of all the short Weierstrass equations over \mathbb{F}_p and arbitrary Diffie–Hellman triples on these models. Here, one encounters extra challenges. First, the set G may be distributed arbitrarily over all (exponentially many in $\log p$) isogeny classes of short Weierstrass models, where each isogeny class contains exponentially many isomorphism classes of short Weierstrass models, with each isomorphism class containing exponentially many short Weierstrass models. Second, relating the difficulty of computing the Diffie–Hellman function within each isogeny class is itself a nontrivial task: having an explicit (computable in time polynomial in $\log p$) isogeny from an elliptic curve E to another curve E' in the same class would achieve this task. By Tate’s isogeny theorem [28], such a map exists if and only if E and E' have the same number of points (E and E' are said to be isogenous). Yet, such an isogeny can have large degree and it can take

superpolynomial number of steps to compute it. Typically, isogeny computations are used in attacks such as the Weil descent attack [3], [7].

We show that such an oracle \mathcal{A} is unlikely to exist by proving that its existence would imply the existence of a set S of polynomial (in $\log p$) fraction of all elliptic curves over \mathbb{F}_p so that one can solve the Diffie–Hellman problem for every $E \in S$ and every Diffie–Hellman triple (P, uP, vP) for E . This is based on random self-reducibility among elliptic curves, which was first studied in [13]; by Tate’s theorem achieving this via algebraic maps (isogenies) is possible only among those curves that have the same order (or trace). Thus our focus here is to identify the values of the trace for which the self-reducibility is applicable. This allows us to use Boneh-Shparlinski hard core bit result on isomorphism classes and enlarge the set of curves where it is applicable. For example, if on a specific isomorphism class their oracle algorithm does not apply, our random walk can (with a good probability) link it to another class where it applies. To show the hard core bit theorem for all the curves, one may consider the analysis based only on isomorphism classes, but the associated hardness assumption is clear and natural when restricted isogeny classes (in view of Tate’s theorem). It will be interesting to see if one can develop new attacks, similar to the ones mentioned earlier for the finite field case. We remark that hard core bit theorems for finite field Diffie–Hellman function remain open and the best in this case is computing one bit (without error) is hard, if the generator is small [2].

2 Notation and preliminaries

Throughout, $p \geq 5$ will be a prime number and $\tilde{\varepsilon} > 0$ will be a fixed real number. We will be considering the Diffie–Hellman problem for elliptic curves E over \mathbb{F}_p and triples (P, uP, vP) , where P is a point of prime order $q > (\log p)^{2+\tilde{\varepsilon}}$ and $1 \leq u, v \leq q$ are integers. We make this assumption because an isogeny $\phi : E \rightarrow E'$ of prime degree $\ell \leq (\log p)^{2+\tilde{\varepsilon}}$ will preserve the order of P and this assumption will be necessary for what follows.

We say that an oracle \mathcal{B} computes the Diffie–Hellman function for E if for any point P of prime order $q > (\log p)^{2+\tilde{\varepsilon}}$,

$$\mathcal{B}(P, uP, vP) = uvP$$

holds with probability at least $1 - 1/p$ (here, the probability is taken over all possible choices of u and v).

Moreover, if z is a non-negative integer then $\text{LSB}(z)$ will denote the least significant bit of z . To define the least significant bit of an element $x \in \mathbb{F}_p$, we first look at the identity map $\iota : \mathbb{F}_p \rightarrow \mathbb{Z}/p\mathbb{Z}$. If $0 \leq z \leq p - 1$ is the unique integer whose image is $\iota(x)$, we define $\text{LSB}(x) = \text{LSB}(z)$. Also, if $Q \in E(\mathbb{F}_p)$ then $x(Q)$ and $y(Q)$ denote the x - and y -coordinates of Q , respectively.

Finally, let $H = \{t \in \mathbb{Z} : |t| \leq 2\sqrt{p}\}$ be the Hasse interval. For $t \in H$ one can write $t^2 - 4p$ uniquely as $d_t c_t^2$, where $d_t < 0$ is square-free and $c_t > 0$. We call c_t the conductor of t .

Advantage: Let \mathcal{A} be an algorithm that, given a short Weierstrass equation W over \mathbb{F}_p , a point $P \in W(\mathbb{F}_p)$ of prime order $q > (\log p)^{2+\varepsilon}$ and two multiples uP and vP with $1 \leq u, v \leq q-1$, outputs a single bit. We define the advantage $\text{Adv}_{W,P}(\mathcal{A})$ of \mathcal{A} as

$$\text{Adv}_{W,P}(\mathcal{A}) := \left| \Pr_{u,v} [\mathcal{A}(P, uP, vP) = \text{LSB}(x(uvP))] - \frac{1}{2} \right|.$$

We say that \mathcal{A} has an advantage ε on W if $\text{Adv}_{W,P}(\mathcal{A}) > \varepsilon$ holds for any point $P \in W(\mathbb{F}_p)$ of prime order $q > (\log p)^{2+\varepsilon}$.

3 The main result

For each prime p , let

$$\Gamma_p = \{W_{a,b} : (a, b) \in \mathbb{F}_p \times \mathbb{F}_p, 4a^3 + 27b^2 \neq 0\}$$

be the set of all short Weierstrass equations and let Ω_p be the set of all elliptic curves over \mathbb{F}_p (i.e., $\Omega_p = \Gamma_p / \cong_{\mathbb{F}_p}$). Let $\Omega_p^{(t)}$ and $\Gamma_p^{(t)}$ denote the restriction to those curves with trace t .

Theorem 3.1. *Assume the Generalized Riemann Hypothesis (GRH) and let $c > 0$ be a fixed real. (a) For almost every t in the Hasse interval, the Diffie–Hellman problem is random self reducible among the set of elliptic curves with trace t . (b) Given a subset $G \subset \Gamma_p$, such that $|G| = \delta |\Gamma_p|$ for some $0 < \delta \leq 1$ with $1/\delta = O((\log p)^c)$, assume that there exists $\varepsilon > 0$ and an algorithm \mathcal{A} running in time t that takes as input a short Weierstrass model W and a Diffie–Hellman triple (P, uP, vP) and outputs a single bit. Assume that \mathcal{A} satisfies the following property: for any $W \in G$ and any point P of prime order $q > (\log p)^{2+\varepsilon}$ on W , $\text{Adv}_{W,P}(\mathcal{A}) > \varepsilon$. Then there exists a subset $S \subseteq \Omega_p$ satisfying*

$$\frac{|\Omega_p|}{|S|} = O_c \left((\log p)^{\frac{3(c+1)}{2}} (\log \log p)^4 \right),$$

and an algorithm \mathcal{B} running in time $(\varepsilon^{-1} \log p)^{O(1)}$, such that \mathcal{B} computes the entire Diffie–Hellman secret $\text{DH}_{E,P}(uP, vP)$ for any $E \in S$ and any Diffie–Hellman triple (P, uP, vP) for E (Note that in the above displayed formula, the implied constant depends only on c). Moreover, these statements hold true with Ω_p and Γ_p replaced by $\Omega_p^{(t)}$ and $\Gamma_p^{(t)}$ for almost every value of the trace t .

Intuitively, (a) implies that an efficient algorithm for computing the Diffie–Hellman function in the average case would imply an efficient algorithm for the same function in the worst case (see Section 6 for the precise technical definition).

4 Counting elliptic curves

Let $p \geq 5$ be a prime and let $\Gamma_p = \{W_{a,b} : (a,b) \in \mathbb{F}_p \times \mathbb{F}_p, 4a^3 + 27b^2 \neq 0\}$ be the set of all short Weierstrass equations over \mathbb{F}_p . Explicitly, $W_{a,b}$ is the short Weierstrass equation $y^2 = x^3 + ax + b$. Then $|\Gamma_p| = p(p-1)$ since the number of all pairs (a,b) , such that $4a^3 + 27b^2 = 0$ is equal to p . Indeed, any such pair is parameterized by $a = -3c^2$ and $b = 2c^3$ for some $c \in \mathbb{F}_p$ and each such c is uniquely determined from (a,b) .

4.1 Isomorphism classes

Two short Weierstrass equations $W_{a,b}$ and $W_{a',b'}$ are isomorphic over \mathbb{F}_p if there exists an element $u \in \mathbb{F}_p^\times$, such that $a' = u^4a$ and $b' = u^6b$. To count the elliptic curves E over \mathbb{F}_p , we observe that the number of short Weierstrass equations $W \in \Gamma_p$ for E is exactly $\frac{p-1}{\#\text{Aut}(E)}$. In particular, this gives us the formula

$$\sum_E \frac{1}{\#\text{Aut}(E)} = p,$$

where the sum is taken over all elliptic curves E over \mathbb{F}_p .

4.2 Isogeny classes

Tate’s isogeny theorem states that two elliptic curves E_1, E_2 over \mathbb{F}_p are isogenous if and only if $\#E_1(\mathbb{F}_p) = \#E_2(\mathbb{F}_p)$. For any elliptic curve E/\mathbb{F}_p we have the Hasse bound $|p+1 - \#E(\mathbb{F}_p)| \leq 2\sqrt{p}$. For an integer $t \in H$ consider the isogeny class of short Weierstrass equations

$$C_t = \{W_{a,b} \in \Gamma : \#W_{a,b}(\mathbb{F}_p) = p+1-t\}.$$

Our goal is to provide upper and lower bounds on the size of C_t for any $t \in H$. We show how to do this in the next two sections. A useful definition for what follows is the *weighted cardinality*:

Definition 4.1 (Weighted cardinality). *Let U be any set of elliptic curves over \mathbb{F}_p . We define the weighted cardinality to be the sum*

$$\#^!U = \sum_{E \in U} \frac{1}{\#\text{Aut}(E)}.$$

4.3 Lenstra’s upper bound

Lemma 4.1. *Let Σ be a set of integers t satisfying $|t| \leq 2\sqrt{p}$. There exists an effectively computable constant c_u (independent of p), such that*

$$\sum_{t \in \Sigma} |C_t| \leq c_u |\Sigma| p^{3/2} (\log p) (\log \log p)^2.$$

Proof. By [15, Prop.1.9(a)], there exists an effective constant c , such that

$$\#\{W \in \Gamma_p : 1 + p - \#W(\mathbb{F}_p) \in \Sigma\} /_{\cong_{\mathbb{F}_p}} \leq c|\Sigma|p^{1/2}(\log p)(\log \log p)^2.$$

Now, the lemma is a consequence of the fact that the weight of an elliptic curve E is $(\#\text{Aut}(E))^{-1}$ (which is either $1/2$, $1/3$ or $1/6$) and that the isomorphism class corresponding to E contains $\frac{p-1}{\#\text{Aut}(E)}$ short Weierstrass equations.

4.4 Refining Lenstra's lower bound

We need a simple refinement of the lower bound established by Lenstra in [15, Prop.1.9(b)] on the size of a collection of isogeny classes.

If $|t| \leq 2\sqrt{p}$, the weighted number of elliptic curves over \mathbb{F}_p whose trace of Frobenius is t is equal to the *Kronecker class number* $H(t^2 - 4p)$ (see [4], [15, pp.654-655]). For a fixed integer $\Delta < 0$, $\Delta \equiv 0, 1 \pmod{4}$, the Kronecker class number $H(\Delta)$ is the weighted number of equivalence classes of binary quadratic forms of discriminant Δ (the weight of a quadratic form is defined to be inverse of the number of automorphisms of the form). Let Δ_0 be the fundamental discriminant associated with Δ and let χ_0 be the quadratic character associated to Δ_0 . Using an analytic class number formula, one expresses $H(\Delta)$ in terms of the special value $L(1, \chi_0)$ of the L -function of the character χ_0 and the discriminant Δ . Thus, a lower bound for $H(\Delta)$ would follow from a lower bound on the special value of the above L -function. The following result is proved in [15, Prop.1.8]:

Lemma 4.2. (i) *There exists an effectively computable positive constant c_0 , such that for each $z \in \mathbb{Z}_{>0}$, there exists $\Delta^* = \Delta^*(z)$, such that*

$$H(\Delta) \geq c_0 \frac{|\Delta|^{1/2}}{\log z},$$

for each Δ which satisfies $|\Delta| \leq z$, $\Delta < 0$, $\Delta \equiv 0, 1 \pmod{4}$ and $\Delta_0 \neq \Delta^*$.

(ii) *Assume the Generalized Riemann Hypothesis. There exists an effectively computable constant $c'_0 > 0$, such that for each $z \in \mathbb{Z}_{>0}$*

$$H(\Delta) \geq c'_0 \frac{|\Delta|^{1/2}}{\log \log z},$$

for each Δ which satisfies $|\Delta| \leq z$, $\Delta < 0$ and $\Delta \equiv 0, 1 \pmod{4}$.

The following refinement of Lenstra's Proposition 1.9(b) is necessary for our argument:

Proposition 4.1. *Let $0 < \mu < 1$ and let Σ be a set of integers t satisfying $|t| \leq 2\sqrt{p}(1 - \mu)$. Let*

$$w_\Sigma = \#\{E : E \text{ elliptic curve over } \mathbb{F}_p, 1 + p - \#E(\mathbb{F}_p) \in \Sigma\} /_{\cong_{\mathbb{F}_p}},$$

be the weighted cardinality of the short Weierstrass equations whose traces of Frobenius are in Σ .

(i) There exists an effectively computable constant $c_1 > 0$, such that

$$w_\Sigma \geq c_1(|\Sigma| - 2) \frac{\mu^{1/2} p^{1/2}}{\log p}.$$

(ii) Assume the Generalized Riemann Hypothesis. Then there exists an effectively computable constant $c'_1 > 0$, such that

$$w_\Sigma \geq c'_1 |\Sigma| \frac{\mu^{1/2} p^{1/2}}{\log \log p}.$$

Proof. One can express

$$w_\Sigma = \sum_{t \in \Sigma} H(t^2 - 4p).$$

i) We apply Lemma 4.2 with $z = 4p$ to get that there exists a constant $c_0 > 0$, such that $H(\Delta) \geq c_0 \frac{|\Delta|^{1/2}}{\log p}$ unless $\Delta_0 = \Delta^*$. As in the proof of Lenstra's Proposition 1.9(b), there are at most two values of t for which the fundamental discriminant of $t^2 - 4p$ is equal to Δ^* . Hence, it remains to estimate $|t^2 - 4p|$ from below to obtain a lower estimate on w_Σ . But if $t \in \Sigma$ then

$$|t^2 - 4p| \geq 4p - 4p(1 - \mu)^2 = 8\mu p - 4\mu^2 p > 8\mu p - 4\mu p = 4\mu p.$$

Thus, $|t^2 - 4p|^{1/2} \geq 2\mu^{1/2} p^{1/2}$. Hence, if $c_1 = c_0$ then

$$w_\Sigma \geq c_1(|\Sigma| - 2) \frac{\mu^{1/2} p^{1/2}}{\log p}.$$

ii) The second part follows similarly except that we use the lower bound under the Generalized Riemann Hypothesis for the Kronecker class number $H(\Delta)$ from Lemma 4.2(ii).

5 Isogeny Graphs

We recall a construction for isogeny graphs for ordinary elliptic curves [13]. For an integer $t \in H$ consider the isogeny class $C_t \subset \Gamma_p$ of short Weierstrass equations over \mathbb{F}_p . Let $S_t = C_t / \sim$ be the corresponding isogeny class of elliptic curves (i.e., we identify two short Weierstrass equations $W_{a,b}, W_{a',b'} \in C_t$ if they are isomorphic over \mathbb{F}_p).

Throughout the whole paper, an *isogeny* between two elliptic curves will always mean an isogeny defined over \mathbb{F}_p .

5.1 Ordinary isogeny classes and isogeny volcanoes

1. *Ordinary isogeny classes.* Suppose that S_t is an isogeny class of ordinary elliptic curves. To understand the structure of S_t one looks at the endomorphism rings of the elliptic curves inside S_t . For any curve $E \in S_t$, the endomorphism ring $\text{End}(E)$ is an order in a quadratic imaginary field [27, §III.9]. Let $\pi : E \rightarrow E$ be the Frobenius endomorphism. The characteristic polynomial of π is $X^2 - tX + p = 0$, so we can regard π as an algebraic integer. It only depends on the class S_t . The following theorem is proved in [14, §4.2] (see also [13, Thm.2.1])

Theorem 5.1 (Kohel). *Let E and E' be two elliptic curves over \mathbb{F}_p that are isogenous over \mathbb{F}_p , let K be the quadratic imaginary field $\text{End}(E) \otimes \mathbb{Q}$ and \mathcal{O}_K be the maximal order of K .*

1. *We have $\mathbb{Z}[\pi] \subseteq \text{End}(E) \subseteq \mathcal{O}_K$ and $\mathbb{Z}[\pi] \subseteq \text{End}(E') \subseteq \mathcal{O}_K$.*
2. *The following are equivalent:*
 - (a) $\text{End}(E) = \text{End}(E')$
 - (b) *There exist isogenies $\phi : E \rightarrow E'$ and $\psi : E' \rightarrow E$ of relatively prime degree.*
 - (c) $(\mathcal{O}_K : \text{End}(E)) = (\mathcal{O}_K : \text{End}(E'))$.
3. *Let $\phi : E \rightarrow E'$ be an isogeny from E to E' of prime degree ℓ defined over \mathbb{F}_p . Then one of the three cases occurs: i) $\text{End}(E)$ contains $\text{End}(E')$ with index ℓ ; ii) $\text{End}(E')$ contains $\text{End}(E)$ with index ℓ ; iii) $\text{End}(E') = \text{End}(E)$.*
4. *Let ℓ be a prime that divides exactly one of $(\mathcal{O}_K : \text{End}(E))$ and $(\mathcal{O}_K : \text{End}(E'))$. Then every isogeny $\phi : E \rightarrow E'$ has degree a multiple of ℓ .*

2. *Isogeny volcanoes.* A convenient visualization of the elliptic curves in an isogeny class in the ordinary case together with the isogenies between them is given by *isogeny volcanoes* [5], [14]. The curves are represented in levels according to their endomorphism rings. Two curves E_1 and E_2 are in the same level if and only if $\text{End}(E_1) \cong \text{End}(E_2)$. Thus, every level corresponds to an order \mathcal{O} in a fixed quadratic imaginary field K . The level corresponding to an order \mathcal{O} is above the level corresponding to an order \mathcal{O}' if $\mathcal{O} \supsetneq \mathcal{O}'$.

Following [5], [6] and [14], we distinguish among three types of isogenies $\phi : E \rightarrow E'$ of prime degree ℓ over \mathbb{F}_p :

1. ϕ is *horizontal* if $\text{End}(E) = \text{End}(E')$;
2. ϕ is *up* if $(\text{End}(E') : \text{End}(E)) = \ell$;
3. ϕ is *down* if $(\text{End}(E) : \text{End}(E')) = \ell$.

One can compute the number of horizontal, up and down isogenies of a given prime degree coming out of a particular ordinary elliptic curve E in terms of the degree and the Legendre symbol. The result (see [5, §2.1], [6, Thm.4] and [14, Ch.4, Prop.23]) is summarized in the following

Proposition 5.1. *Let E be an ordinary elliptic curve over \mathbb{F}_p , with endomorphism ring $\text{End}(E)$ contained in the quadratic imaginary field K with fundamental discriminant $-D < 0$. Let ℓ be a prime different from p and let $c_\pi = (\mathcal{O}_K : \mathbb{Z}[\pi])$ and $c_E = (\mathcal{O}_K : \text{End}(E))$. Then*

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

Finally, we say that two isomorphism classes of elliptic curves E_1 and E_2 in the same isogeny class belong to the same level of the isogeny volcano if and only if $c_{E_1} = c_{E_2}$.

5.2 Expander graphs and a rapid mixing lemma

Let k be a positive integer and let \mathcal{I} be an infinite set of positive integers. Consider a sequence of graphs $\{G_h\}_{h \in \mathcal{I}}$, each of which is k -regular and connected, such that G_h has h vertices. Let A_h be the adjacency matrix of G_h . Since G_h is k -regular, the vector v_h consisting of 1's in each coordinate is an eigenvector for A_h with eigenvalue $\lambda_{\text{triv}} = k$ and any other eigenvalue λ of A_h satisfies $|\lambda| \leq k$. We refer to the eigenvalue λ_{triv} as the trivial eigenvalue. Furthermore, since G_h is connected, the eigenvalue k has multiplicity one.

Definition 5.1. *The sequence $\{G_h\}_{h \in \mathcal{I}}$ is called a sequence of expander graphs if there exists a constant $0 < \nu < 1$, such that for any h and any eigenvalue $\lambda \neq \lambda_{\text{triv}}$ of A_h , $|\lambda| \leq \nu \lambda_{\text{triv}}$.*

The main application of expander graphs is to prove the rapid mixing of random walks provided we have a good upper bound on the spectral gap ν . The property is summarized in the following proposition which will be used in our particular application (see [13, Prop.3.1] for a proof):

Proposition 5.2. *Let G be a k -regular graph with h vertices. Assume that every eigenvalue $\lambda \neq \lambda_{\text{triv}}$ of G satisfies the bound $|\lambda| \leq \nu \lambda_{\text{triv}}$ for some $0 < \nu < 1$. Let S be a set of vertices of G and let x be any vertex of G . Then a random walk of length at least $\frac{\log\left(\frac{2h}{|S|^{1/2}}\right)}{\log(\nu^{-1})}$ starting at x will land in S with probability at least $\frac{|S|}{2h}$.*

5.3 Isogeny graphs in the ordinary case

Fix an isogeny class C_t of short Weierstrass equations for some $t \in H$ and the corresponding set S_t of elliptic curves. Following [13, §2.1] we define an isogeny graph to be a graph \mathcal{G} whose vertices are all the elements of S_t that belong to a fixed level of the isogeny volcano for S_t .

Let $E_1, E_2 \in S_t$. Two isogenies $\phi : E_1 \rightarrow E_2$ and $\phi' : E_1 \rightarrow E_2$ are said to be equivalent if there exists an automorphism $\alpha \in \text{Aut}(E_2)$, such that $\phi' = \alpha\phi$ (see also [9, Prop.2.3]). The edges of the graph are equivalence classes of *horizontal* isogenies that have prime degrees at most $(\log p)^{2+\varepsilon}$. The degree bound is chosen in such a way that it is small enough to allow the isogenies to be computed and large enough to allow the graph to be connected and to have rapid mixing properties.

The graph \mathcal{G} is known to be isomorphic to a graph \mathcal{H} whose vertices are elliptic curves \mathbb{C}/\mathfrak{a} with complex multiplication by the order \mathcal{O} corresponding to the level for the graph \mathcal{G} in the isogeny volcano (here, $\mathfrak{a} \subset \mathcal{O}$ is an ideal) and whose edges are isogenies of the form $\mathbb{C}/\mathfrak{a} \rightarrow \mathbb{C}/\mathfrak{a}\mathfrak{l}^{-1}$, where $\mathfrak{l} \subset \mathcal{O}$ is an invertible prime ideal satisfying $N(\mathfrak{l}) \leq (\log p)^{2+\varepsilon}$ [6, §3], [7], [13, §2.1]. Equivalently, \mathcal{H} is the Cayley graph of the Picard group $\text{Pic}(\mathcal{O})$ of the order \mathcal{O} with respect to the generators $[\mathfrak{l}] \in \text{Pic}(\mathcal{O})$, where \mathfrak{l} ranges over the invertible prime ideals of \mathcal{O} whose norm is at most $(\log p)^{2+\varepsilon}$.

5.4 The spectral gap of an isogeny graph

For a particular isogeny graph \mathcal{G} of ordinary elliptic curves, one can bound the nontrivial eigenvalues via character sum estimates under the Generalized Riemann Hypothesis. This is done via spectral analysis of the corresponding Cayley graph \mathcal{H} . For what follows, it will be convenient to view the eigenvectors of the adjacency matrix of \mathcal{H} as functions on the corresponding ideal classes of the Picard group. The following proposition is proven in [13, §4]:

Proposition 5.3. *Let $m = (\log p)^{2+\varepsilon}$ and let $e = \#\mathcal{O}^\times$.*

(i) *The graph \mathcal{H} has eigenfunctions equal to the characters χ of $\text{Pic}(\mathcal{O})$ with corresponding eigenvalues the character sums*

$$\lambda_\chi = \sum_{p \leq m} \sum_{\substack{\mathfrak{a} \subset \mathcal{O}, \\ N\mathfrak{a} = p}} \chi(\mathfrak{a}).$$

(ii) *Let $D < 0$ and let \mathcal{O} be an order of discriminant D . The trivial eigenvalue λ_{triv} is equal to the number of ideal classes of the form $[\mathfrak{l}]$ where \mathfrak{l} invertible prime ideal of \mathcal{O} of norm at most m (note that λ_{triv} is asymptotically equal to $\frac{m}{e \log m}$ where $e = \#\mathcal{O}^\times$). If χ is a nontrivial character of the Picard group $\text{Pic}(\mathcal{O})$, then under the Generalized Riemann Hypothesis,*

$$\lambda_\chi = O(m^{1/2} \log |mD|).$$

Remark 5.1. Propositions 5.2 and 5.3 show the following: suppose that S is a set of elliptic curves belonging to the same level of the isogeny volcano, such that $|\mathcal{G}|/|S| = (\log p)^{O(1)}$ and such that one can efficiently compute $\text{DH}_{E,P}(uP, vP)$ for every $E \in S$ and every Diffie–Hellman triple (P, uP, vP) for E . Then there is a random polynomial time reduction of the computation of the Diffie–Hellman function on an arbitrary curve $E \in V(\mathcal{G})$ to the Diffie–Hellman function on a curve in S . Hence, one can compute the Diffie–Hellman secret on any curve E in $V(\mathcal{G})$ with high probability in time polynomial in $\log p$.

Indeed, a random walk of length polynomial in $\log p$ will connect E to a curve in S with high probability (high probability means $1 - O(p^{-r})$ for some $r > 0$). Since any step in this random walk is an isogeny that is computable in time polynomial in $\log p$, the resulting composition of isogenies and their duals are computable in time polynomial in $\log p$ (even if the degree of the composition is large). Finally, if (P, uP, vP) is a Diffie–Hellman triple for E and $\phi : E \rightarrow E'$ is an isogeny to an elliptic curve $E' \in S$, one can consider the Diffie–Hellman triple $(\phi(P), u\phi(P), v\phi(P))$ on E' and compute the Diffie–Hellman function for that triple to obtain $uv\phi(P)$. After applying the dual isogeny, we obtain the point $duvP$, where d is the degree of the composition (note that the degree is polynomial in $\log p$). Finally, since we are in a prime-order subgroup, we compute e , such that de is congruent to 1 modulo the group order. The point $ed(uvP) = uvP$ is then the desired point.

Remark 5.2. There exist isogeny graphs for supersingular elliptic curves as well. These supersingular graphs were first considered in [11] and [17]. Their expansion properties were shown much later by Pizer [21], [22]. Given a prime p , the supersingular elliptic curves are always defined over \mathbb{F}_{p^2} . According to [17], all isomorphism classes of supersingular elliptic curves belong to the same isogeny class. In practice, we ignore supersingular curves in our argument for the main theorem. Yet, the corresponding isogeny graph is still an expander graphs.

6 Random self-reducibility

We define random self-reducibility. Intuitively, we would like to prove that an efficient algorithm for the Diffie–Hellman function in the average case would imply an efficient algorithm in the worst case.

6.1 Smooth isogeny classes and random self-reducibility

Let R be a fixed polynomial. Consider the following properties of a set S of elliptic curves over \mathbb{F}_p :

1. There exists a subset $S' \subseteq S$ with $|S'|/|S| \geq R(\log p)^{-1}$.
2. There exists an algorithm \mathcal{A} , such that: i) \mathcal{A} computes the Diffie–Hellman function on any elliptic curve $E \in S'$; ii) \mathcal{A} produces random output whenever one feeds in a Diffie–Hellman triple for an elliptic curve $E \notin S'$.

Definition 6.1. *Let S be a set of elliptic curves that satisfies conditions 1. and 2. We call S random self-reducible with respect to R if given an elliptic curve $E \in S$, one can compute the Diffie–Hellman function for any triple (Q, uQ, vQ) on E with expected $(\log p)^{O(1)}$ queries to \mathcal{A} on elliptic curves $E' \in S$ that are randomly distributed among all classes in S .*

6.2 Random self-reducibility for single levels in the isogeny volcanoes

We first show that horizontal levels in the isogeny volcanoes with sufficiently many curves on which the Diffie–Hellman problem is solvable are random self-reducible:

Lemma 6.1. *Let \mathcal{G} be the graph corresponding to a particular level of the isogeny volcano for some isogeny class of elliptic curves. Assume that the set of vertices $V(\mathcal{G})$ of \mathcal{G} satisfies 1. and 2. for some polynomial R . Then $V(\mathcal{G})$ is random self-reducible with respect to R .*

Proof. Let E be any elliptic curve in $V(\mathcal{G})$ and (P, uP, vP) be any Diffie–Hellman triple for E . We will show how to connect this input to the Diffie–Hellman function to an input on a *random* elliptic curve E' from $V(\mathcal{G})$ via a sequence of isogenies that are computable in polynomial time. Let $S' \subset V(\mathcal{G})$ be the distinguished set from item 1 above. and let $\mu = |S'|/|V(\mathcal{G})|$. Let $E_0 = E$. We

will use the fact that \mathcal{G} is an expander graph. Let $\tau = \left\lceil \frac{\log \left(\frac{2|V(\mathcal{G})|}{|S'|^{1/2}} \right)}{\log(\nu^{-1})} \right\rceil + 1$, where ν is the spectral gap for \mathcal{G} . Using the upper bound for ν from Proposition 5.2, we obtain that τ is polynomial in $\log p$, i.e., $\tau = (\log p)^{O(1)}$.

We repeat the following procedure $m \geq \frac{2}{\mu} \log p$ times:

1. Consider a random walk E_0, E_1, \dots, E_τ on \mathcal{G} of length τ . Let ϕ be the composition of the isogenies along the walk, $\widehat{\phi}$ be the dual isogeny of ϕ and d be their degree. Compute $e = d^{-1}$ modulo q (recall that q is the prime order of the original point P).
2. If $E' = E_\tau$, query the oracle on the elliptic curve E' and the Diffie–Hellman triple $(\phi(P), u\phi(P), v\phi(P))$ under ϕ .
3. If the oracle returns the point Q on E' , compute and return $e\widehat{\phi}(Q) \in E(\mathbb{F}_p)$.

Since the computation of a single isogeny of degree ℓ takes $O(\ell^4)$ time (see [14]), each of the above steps runs in time $O((\log p)^{8+4\epsilon}\tau)$ which is polynomial in $\log p$ (as do all other steps below).

By Proposition 5.2, the probability that $E_\tau \notin S'$ is at most $1 - \frac{\mu}{2}$. Thus, if we repeat the above steps m times, the probability that none of the end points of the random walk is in S' is at most

$$\left(1 - \frac{\mu}{2}\right)^m \leq e^{-\frac{\mu m}{2}} \leq e^{-\frac{\mu \cdot 2/\mu \log p}{2}} = O(p^{-1}).$$

Therefore, the above procedure will produce a list $A = L(P, uP, vP)$ of points that contains the desired point uvP with high probability. To obtain the desired solution, we compute the list $B = L(P, (u+r)P, vP)$ for a random $r \in [1, q-1]$ as in the method of Shoup [23]. We check if A and $-rvP + B$ have a unique common element, and if so, we output it. Otherwise, we report a failure. The analysis of this last step is the same as in [23].

6.3 Random self-reducibility for multiple levels in the isogeny volcanoes

Owing to space limitations we will only outline how one can apply the methods of the single level case to solve the case of multiple levels in the isogeny volcano. Outside of this section, we restrict our discussion to the case of a single level.

Definition 6.2. *Let B be a positive real number. An isogeny class S_t of elliptic curves is called B -smooth if its conductor c_t is B -smooth, i.e., if any prime divisor of c_t is at most B .*

The next lemma proves reducibility of the Diffie–Hellman problem for a whole isogeny class (not just a single level).

Lemma 6.2. *Let $r > 0$ be any real constant and assume that S_t satisfies i) and ii) for some polynomial R , and that S_t is $(\log p)^r$ -smooth. Assuming the Generalized Riemann Hypothesis, any instance of the Diffie–Hellman problem on any elliptic curve $E \in S_t$ can be computed in time polynomial in $\log p$.*

The next lemma guarantees that the conductor c_t will have $O(\log \log p)$ distinct prime factors for almost all traces t in the Hasse interval. Let m be a positive integer such that $\log \log m > 1$ and let N_m be the number of traces $t \in H$, such that c_t has less than m distinct prime factors.

Lemma 6.3. *There exists a constant C (independent of m and p) such that*

$$N_m \geq (1 - e^{-Cm \log m})|H|.$$

Proof omitted.

Remark 6.1. Suppose that $c > 0$ is fixed. By choosing k large enough (independent of p) and applying the above lemma for $m = k \log \log p$, we can guarantee that $N_m = (1 - O((\log p)^{-c}))|H|$. This means that for most of the traces $t \in H$, c_t will have $O(\log \log p)$ distinct prime divisors.

For the classes S_t for which the volcano has multiple levels, we may not be able to exploit random self-reducibility in some of them. We can bound c_t to be small enough and having $O(\log \log p)$ prime divisors, so that starting from an arbitrary elliptic curve, we can reach the appropriate random self-reducible level in time polynomial in $\log p$ by searching through the levels via vertical isogenies and testing each level for random self-reducibility.

7 Proof of Theorem 3.1

7.1 Notation

Let \mathcal{A} be the oracle from Theorem 3.1 and ε be the corresponding advantage. A short Weierstrass equation W is called *LSB-predictable*, if for any point $P \in W(\mathbb{F}_p)$ of prime order $q > (\log p)^{2+\varepsilon}$, $\text{Adv}_{W,P}(uP, vP) > \varepsilon$ (in other words, \mathcal{A} predicts the least significant bit of the Diffie–Hellman function for W and the generator P with advantage ε).

More generally, if T is any set of short Weierstrass equations over \mathbb{F}_p and $0 < \delta' < 1$ is a real number, we refer to T as δ' -*predictable* if at least $\delta'|T|$ elliptic curves in T are LSB-predictable.

7.2 Most of the isogeny classes are smooth

Let B be an arbitrary integer. The following lemma shows that almost all of the isogeny classes S_t of elliptic curves over \mathbb{F}_p are B -smooth. The latter will be useful in applying the tunneling argument and Lemma 6.2.

Lemma 7.1. *The number of traces $t \in H$, such that the isogeny class S_t corresponding to t is B -smooth is at least $\left(1 - \frac{2}{B}\right)|H|$.*

Proof. Fix a prime ℓ , such that $B < \ell < \sqrt{p}$ and consider the solutions of the congruence

$$t^2 \equiv 4p \pmod{\ell^2}$$

for $t \in H$. First, the congruence $t^2 \equiv 4p \pmod{\ell}$ has exactly $1 + \left(\frac{4p}{\ell}\right)$ solutions.

Each such solution t lifts uniquely to a solutions \tilde{t} modulo ℓ^2 by Hensel's lemma since the derivative of $f(x) = x^2 - 4p$ does not vanish modulo $\ell > 2$ at any such t . Thus,

$$\begin{aligned} \Pr_{t \in H} [c_t \text{ is not } B\text{-smooth}] &\leq \sum_{B < \ell < \sqrt{p}} \frac{1}{\ell^2} \left[1 + \left(\frac{4p}{\ell}\right) \right] < \\ &< \sum_{B < \ell < \sqrt{p}} \frac{2}{\ell^2} < \int_B^\infty \frac{2}{u^2} du = \frac{2}{B}. \end{aligned}$$

7.3 Lower bound on smooth, predictable isogeny classes

Here, we show that there is a polynomial fraction of traces $t \in H$ such that S_t is smooth and C_t contains sufficiently many LSB-predictable short Weierstrass equations.

Lemma 7.2. *Let δ and c be as in the statement of Theorem 3.1. There exists a constant c_1 (independent of p), such that the number of traces $t \in H$ for which S_t is $(\log p)^{c+2}$ -smooth and C_t is $\delta/2$ -predictable is at least $c_1 \frac{|H|}{(\log p)^{c+1}(\log \log p)^2}$*

Proof. Let

$$S_{\delta/2} = \{t \in H : C_t \text{ is } \delta/2\text{-predictable}\}$$

and

$$U = \{t \in H : S_t \text{ is } (\log p)^{c+2}\text{-smooth}\}.$$

By Lemma 7.1, $|U| \geq \left(1 - \frac{2}{(\log p)^{c+2}}\right) |H|$. We would like to estimate $|U \cap S_{\delta/2}|$. First, we need an estimate on $|S_{\delta/2}|$. For each $t \in S_{\delta/2}$, C_t contains at most $|C_t|$ LSB-predictable curves. For each $t \notin S_{\delta/2}$, C_t contains at most $(\delta/2)|C_t|$ LSB-predictable curves. Thus, we get the inequality

$$\sum_{t \in S_{\delta/2}} |C_t| + \sum_{t \notin S_{\delta/2}} \frac{\delta}{2} |C_t| \geq |G| = \delta |G_p|$$

We combine this with Lemma 4.1 to obtain

$$\begin{aligned} \delta |G_p| &\leq \sum_{t \in S_{\delta/2}} |C_t| + \sum_{t \notin S_{\delta/2}} \frac{\delta}{2} |C_t| = \sum_t \frac{\delta}{2} |C_t| + \sum_{t \in S_{\delta/2}} \left(1 - \frac{\delta}{2}\right) |C_t| \leq \\ &\leq \frac{\delta}{2} |G_p| + \left(1 - \frac{\delta}{2}\right) c_u |S_{\delta/2}| p^{3/2} (\log p) (\log \log p)^2. \end{aligned}$$

Thus,

$$|S_{\delta/2}| \geq \left(\frac{\delta/2}{1 - \delta/2}\right) \frac{|G_p|}{c_u p^{3/2} (\log p) (\log \log p)^2} \geq c'_1 \frac{|H|}{(\log p)^{c+1} (\log \log p)^2},$$

for some constant $c'_1 > 0$ (since $\delta = O((\log p)^e)$). Hence,

$$\begin{aligned} |U \cap S_{\delta/2}| &= |U| + |S_{\delta/2}| - |U \cup S_{\delta/2}| \geq \left(1 - \frac{2}{(\log p)^{c+2}}\right) |H| + \\ &+ c'_1 \frac{|H|}{(\log p)^{c+1} (\log \log p)^2} - |H| \geq c_1 \frac{|H|}{(\log p)^{c+1} (\log \log p)^2}, \end{aligned}$$

for some constant c_1 independent of p . This proves the lemma.

7.4 Predicting LSB within an isomorphism class

It was shown in [1] that within an isomorphism class of short Weierstrass equations, predicting the least significant bit on a non-negligible fraction of the short Weierstrass equations is at least as hard as computing the entire Diffie–Hellman secret key for the elliptic curve corresponding to this class.

For any short Weierstrass equation $W : y^2 = x^3 + ax + b$ and any $\lambda \in \mathbb{F}_p^\times$ we denote by W_λ the isomorphic curve $y^2 = x^3 + a\lambda^4 x + b\lambda^6$ and by $\phi_\lambda : W \rightarrow W_\lambda$ the isomorphism $(x, y) \mapsto (\lambda^2 x, \lambda^3 y)$. The result is summarized as follows:

Theorem 7.1 (Boneh–Shparlinski). *Let $0 < \varepsilon, \delta < 1$. Let p be a prime and W be a short Weierstrass equation over \mathbb{F}_p . Let $P \in W(\mathbb{F}_p)$ be a point of prime order. Suppose that there is a τ -time algorithm \mathcal{A} , such that $\text{Adv}_{W_\lambda, \phi_\lambda(P)}(\mathcal{A}) > \varepsilon$ for at least δ -fraction of all $\lambda \in \mathbb{F}_p^\times$. Then the Diffie–Hellman function for W with respect to the generator P can be computed in expected time $\tau \cdot (\varepsilon^{-1} \delta^{-1} \log p)^{O(1)}$.*

7.5 Predictable isomorphism classes within a predictable isogeny class

Lemma 7.3. *Let $0 < \beta < 1$, such that $1/\beta = O((\log p)^c)$, let $t \in H$ be a trace, such that C_t be a β -predictable isogeny class of short Weierstrass equations. There exists a constant $0 < c_2 < 1$, such that the number of $\beta/2$ -predictable isomorphism classes of elliptic curve inside C_t is at least $c_2 \frac{|S_t|}{(\log p)^c}$.*

Proof. Let $T_{\beta/2}$ be the set of $\beta/2$ -predictable isomorphism classes of short Weierstrass models contained C_t . Each isomorphism class $I \subset C_t$, $I \in T_{\beta/2}$ contains at most $|I|$ LSB-predictable elliptic curves and each isomorphism class $I \notin T_{\beta/2}$ contains at most $\frac{\beta}{2}|I|$ LSB-predictable elliptic curves. Thus,

$$\begin{aligned} \beta|C_t| &\leq \sum_{\substack{I \subset C_t, \\ I \in T_{\beta/2}}} |I| + \sum_{\substack{I \subset C_t, \\ I \notin T_{\beta/2}}} \frac{\beta}{2}|I| = \sum_{I \subset C_t} \frac{\beta}{2}|I| + \sum_{\substack{I \subset C_t, \\ I \in T_{\beta/2}}} \left(1 - \frac{\beta}{2}\right)|I| \leq \\ &\leq \frac{\beta}{2}|C_t| + \frac{(2-\beta)p}{4}|T_{\beta/2}|. \end{aligned}$$

Therefore,

$$|T_{\beta/2}| \geq 2 \left(\frac{\beta}{1-\beta/2} \right) \frac{|C_t|}{p} > c_2 \frac{|S|}{(\log p)^c},$$

for some constant $c_2 > 0$ independent of p (since $1/\beta = O((\log p)^c)$).

7.6 Proof of Theorem 3.1

Proof (Proof of Theorem 3.1). According to Lemma 7.2, there exists a constant c_1 (independent of p), such that for at least $c_1 \frac{|H|}{(\log p)^{c+1}(\log \log p)^2}$ traces $t \in H$, S_t is $(\log p)^{c+2}$ -smooth and C_t is $\delta/2$ -predictable. Let $0 < \mu < 1$ be the real number defined by $2\sqrt{p}\mu = \frac{c_1}{4} \cdot \frac{|H|}{(\log p)^{c+1}(\log \log p)^2}$. We will apply our refinement of Lenstra's lemma with this particular μ . Indeed, let Σ be the set of all traces $t \in H$ which satisfy $|t| \leq 2\sqrt{p}(1-\mu)$ and such that S_t is $(\log p)^{c+2}$ -smooth and C_t is $\delta/2$ -predictable. Then

$$|\Sigma| \geq \left\lceil \frac{c_1}{2} \cdot \frac{|H|}{(\log p)^{c+1}(\log \log p)^2} \right\rceil.$$

Since we have assumed the Generalized Riemann Hypothesis, Proposition 4.1(ii) implies that

$$\#\{W \in C_t : t \in \Sigma\}_{/\cong_{\mathbb{F}_p}} \geq |\Sigma| \frac{\mu^{1/2} p^{1/2}}{\log \log p} \geq \tilde{c} \frac{p}{(\log p)^{\frac{3}{2}(c+1)} (\log \log p)^4},$$

for some constant \tilde{c} independent of p . Let

$$S := \{W \in C_t : t \in \Sigma\}_{/\cong_{\mathbb{F}_p}}$$

Since the weighted cardinality of each isogeny class is $p/2$, $p/4$ or $p/6$, we obtain that there exists a constant \tilde{c}' (independent of p), such that

$$|S| \geq \tilde{c}' \frac{|\Omega_p|}{(\log p)^{\frac{3}{2}(c+1)} (\log \log p)^4}.$$

We claim that S satisfies the properties of the theorem. Indeed, by Lemma 7.3 applied to $\beta = \delta/2$ we obtain that for each $t \in \Sigma$, C_t contains a polynomial fraction of $\delta/4$ -predictable isomorphism classes. The result of Boneh and Shparlinski then implies that one can compute the Diffie–Hellman function on each of these isomorphism classes in time $\tau(\log p)^{O(1)}$ (since $1/\delta$ is polynomial in $\log p$). Finally, applying Lemma 6.2 we obtain that one can solve the Diffie–Hellman problem on any $E \in S$ in time $\tau(\log p)^{O(1)}$. That completes the proof.

Acknowledgements: We thank Dan Boneh, David Jao, Steve Miller, Bjorn Poonen and Ken Ribet for discussions.

References

1. D. Boneh and I. Shparlinski, *On the unpredictability of bits of elliptic curve Diffie–Hellman scheme*, Proceedings of CRYPTO (2001), 201–212.
2. D. Boneh and R. Venkatesan, *Hardness of computing the most significant bits of secret keys in Diffie–Hellman and related schemes*, Lecture Notes in Computer Science **1109** (1996), 129–142.
3. H. Cohen, G. Frey, and Eds., *Handbook of elliptic and hyperelliptic curve cryptography*, Theory and Practice (2005).
4. M. Deuring, *Die Typen der Multiplikatorenringe elliptischer Funktionenkörpern*, Abh. Math. Sem. Hansischen Univ., vol. 14, 1941, pp. 197–272.
5. M. Fouquet and F. Morain, *Isogeny volcanoes and the SEA algorithm*, Algorithmic number theory (Sydney, 2002), Lecture Notes in Comput. Sci., vol. 2369, Springer, Berlin, 2002, pp. 276–291.
6. S. D. Galbraith, *Constructing isogenies between elliptic curves over finite fields*, LMS J. Comput. Math. **2** (1999), 118–138 (electronic).
7. S. D. Galbraith, F. Hess, and N. P. Smart, *Extending the GHS Weil descent attack*, Advances in cryptology—EUROCRYPT 2002 (Amsterdam), Lecture Notes in Comput. Sci., vol. 2332, Springer, Berlin, 2002, pp. 29–44.
8. M. I. Gonzalez Vasco and I. Shparlinski, *Security of the most significant bits of the Shamir message passing scheme.*, Math. Comput. **71** (2002), no. 237, 333–342.
9. B. H. Gross, *Heights and the special values of L-series*, Number theory (Montreal, Que., 1985), CMS Conf. Proc., vol. 7, Amer. Math. Soc., Providence, RI, 1987, pp. 115–187.
10. N. Howgrave-Graham, P. Q. Nguyen, and I. Shparlinski, *Hidden number problem with hidden multipliers, timed-release crypto, and noisy exponentiation.*, Math. Comput. **72** (2003), no. 243, 1473–1485.

11. Y. Ihara, *Discrete subgroups of $PL(2, k_\varphi)$* , Algebraic Groups and Discontinuous Subgroups (Proc. Sympos. Pure Math., Boulder, Colo., 1965), vol. IX, Amer. Math. Soc., Providence, R.I., 1966, pp. 272–278.
12. D. Jao, D. Jetchev, and R. Venkatesan, *On the security of certain partial Diffie–Hellman secrets*, INDOCRYPT (2007).
13. D. Jao, S. D. Miller, and R. Venkatesan, *Do all elliptic curves of the same order have the same difficulty of discrete log?*, ASIACRYPT, 2005, pp. 21–40.
14. D. Kohel, *Endomorphism rings of elliptic curves over finite fields.*, University of California, Berkeley, Ph.D. thesis (1996).
15. H. W. Lenstra, *Factoring integers with elliptic curves*, Ann. of Math **126** (1987), no. 2, 649–673.
16. A. J. Menezes, P. C. van Oorschot, and S. A. Vanstone, *Handbook of applied cryptography*, CRC Press, Inc., Boca Raton, FL, USA, 1996.
17. 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), 217–242.
18. P. Q. Nguyen, *The dark side of the hidden number problem: Lattice attacks on DSA*, Proc. Workshop on Cryptography and Computational Number Theory (2001), 321–330.
19. P. Q. Nguyen and I. Shparlinski, *The insecurity of the digital signature algorithm with partially known nonces.*, J. Cryptology **15** (2002), no. 3, 151–176.
20. ———, *The insecurity of the elliptic curve digital signature algorithm with partially known nonces.*, Des. Codes Cryptography **30** (2003), no. 2, 201–217.
21. A. K. Pizer, *Ramanujan graphs and Hecke operators*, Bull. Amer. Math. Soc. (N.S.) **23** (1990), no. 1, 127–137.
22. ———, *Ramanujan graphs*, Computational perspectives on number theory (Chicago, IL, 1995), AMS/IP Stud. Adv. Math., vol. 7, Amer. Math. Soc., Providence, RI, 1998, pp. 159–178.
23. V. Shoup, *Lower bounds for discrete logarithms and related problems*, Lect. Notes in Comp. Sci, vol. 1233, Springer-Verlag, Berlin, 1997, pp. 256–266.
24. I. Shparlinski, *On the generalized hidden number problem and bit security of XTR.*, AAECC, 2001, pp. 268–277.
25. ———, *Cryptographic applications of analytic number theory: Complexity lower bounds and pseudorandomness*, PCS, vol. 22, Birkhuser, 2003.
26. I. Shparlinski and A. Winterhof, *A hidden number problem in small subgroups*, Math. Comp. **74** (2005), 2073–2080.
27. J. H. Silverman, *The arithmetic of elliptic curves*, Springer-Verlag, New York, 1992, Corrected reprint of the 1986 original.
28. J. Tate, *Endomorphisms of abelian varieties over finite fields*, Invent. Math. **2** (1966), 134–144.