Using the Satoh-FGH algorithm in cryptography

November 6, 2000

Abstract

The use of elliptic curves in cryptography relies on the ability to count the number of points on a given curve. Before 1999, the SEA algorithm was the only efficient method known for random curves. Then T. Satoh proposed a new algorithm based on the canonical p-adic lift of the curve. Fouquet-Gaudry-Harley extended Satoh's method to the case of characteristic two. This paper presents an implementation of Satoh-FGH and its application to the problem of finding curves suitable for cryptography. By combining Satoh-FGH and an early-abort strategy based on SEA, we are able to find secure random curves in much less time than previously reported. In particular we can generate curves widely considered to be as secure as RSA-1024 in less than one minute each on a fast workstation.

Keywords: Elliptic curve cryptography, Satoh-FGH algorithm, early abort, point counting.

1 Introduction

Since elliptic curve cryptosystems were first proposed in the mid-eighties by Koblitz [Kob87] and Miller [Mil87], their efficiency and security have been the focus of intense study. In recent years, they have become widely accepted as an alternative to cryptosystems based on factorisation or discrete logarithms in finite fields, especially for constrained environments.

One of the initial steps in protocols based on elliptic curve cryptography is to generate a suitable curve defined over a finite field. To ensure that the system is secure, the curve must be chosen to have a number of points which is divisible by a large prime so that computing discrete logarithms on the curve is intractable using known attacks. Hence it is necessary to know the cardinality of the curve.

Among the elliptic curves defined over a given finite field, there are some classes of curves with particular properties that are useful for counting points or for accelerating arithmetic operations occurring in the protocols. However choosing such curves can be a risky proposition.

Perhaps the most striking example is anomalous curves. The number of points over $\mathbb{F}_q$ is simply q. However Smart [Sma99], Satoh-Araki [SA98] and Semaev [Sem98] independently discovered a polynomial-time attack.

Another attack due to Menezes-Okamoto-Vanstone [MOV91], and generalised by Frey-Rück [FR94], reduces discrete logs on supersingular and trace 2 curves to discrete logs in a small-degree extension of $\mathbb{F}_q$. This yields an algorithm that runs in sub-exponential time.

A minor weakness is known for curves with many automorphisms [vOW99], [GLV98], [DGM99] including curves defined over a small subfield, proposed by Koblitz, and some complex-multiplication curves. Attacks on these curves take less time than for generic curves, but remain in exponential time.

It has recently been shown by Gaudry-Hess-Smart [GHS00] that curves defined over composite extension fields are also weak in certain cases, using a reduction via hyperelliptic curves.

These results suggest that for maximum security one should avoid curves with special properties and instead choose a random curve whose number of points is divisible by a large prime, over a prime field or an extension of prime degree. This ideal procedure was made possible in practice by the SEA algorithm due to Schoof [Sch85], [Sch95], Elkies [Elk98], Atkin [Atk92] and others [Cou94, Cou96], [Mor95], [Ler97a], [Mül95], [Dew98], etc. With this method, counting points on one given curve is reasonably fast.

However finding a cryptographically suitable curve requires testing many curves and this takes much more time. For instance, Johnson and Menezes [JM99] recently described this process as a "complicated and cumbersome task" requiring "a few hours on a workstation" for 200 bits.

Recently, a new algorithm for counting points on curves in small characteristic $p \geq 5$ was designed by Satoh [Sat] and extended to characteristics two and three by Fouquet-Gaudry-Harley [FGH00]. An independent extension to characteristic two is described by Skjernaa [Skj].

Satoh's algorithm is asymptotically superior to SEA for fixed p, requiring $O(\log^{3+\varepsilon} q)$ deterministic time, instead of $O(\log^{4+\varepsilon} q)$ under reasonable hypotheses. As demonstrated in [FGH00], the Satoh-FGH algorithm is much faster in practice, at least in characteristic two. Indeed the authors were able to count points over much larger fields (up to 8009 bits) than had previously been possible, and could match the largest size reached with SEA (i.e. 1999 bits) in just three hours.

In the following we will describe a method for generating cryptographically suitable curves, over fields of 113 to 571 bits, using an implementation of the Satoh-FGH algorithm combined with an efficient early-abort strategy based on ideas from SEA. In this manner we reduce substantially the time required for curve-generation, finding suitable 200-bit curves in minutes rather than hours on a workstation, for instance.

In section 2, we recall some basic facts about elliptic curves defined over finite fields of characteristic two. Next we review some algorithms that can be used to compute the cardinality of a curve, and in particular we give a description of the Satoh-FGH algorithm. Section 4 gives the conditions that a curve must satisfy in order to be suitable for cryptographic applications. It also describes the early-abort strategy first used by Lercier in [Ler97a] for selecting good curves. Last but not least we describe our implementation and the results we obtained by combining a more aggressive early-abort strategy and the Satoh-FGH algorithm.

2 Elliptic curves over finite fields of characteristic two

In this section, we recall some basic facts about elliptic curves defined over $\mathbb{F}_q$ where $q = 2^d$. We will only be concerned with characteristic two. The reader can find details for the general case in [Men93, Sil86, BSS99].

For our purposes, we can choose the equation of an elliptic curve E (with non-zero j-invariant) to be:

$$E: \; y^2 + xy = x^3 + a_6 \qquad \text{where } a_6 \in \mathbb{F}_q^*.$$

Its *twist* curve is:

$$E^*: \; y^2 + xy = x^3 + a_2 x^2 + a_6$$

where a_2 is some fixed element of trace 1.

One main invariant of the curve is its j-invariant which is: $j(E) = 1/a_6$. (In the following we assume $j(E) \notin \mathbb{F}_4$.)

The set of points $E(\mathbb{F}_q)$ of the curve is:

$$E(\mathbb{F}_q) = \{(x, y) \in \mathbb{F}_q^2 \mid (x, y) \text{ satisfies the equation of } E\} \cup \{\mathcal{O}_E\},$$

where $\mathcal{O}_E$ is the *point at infinity*.

The Frobenius automorphism F is the map $x \mapsto x^q$. It can be extended to an endomorphism of E:

$$\begin{array}{rccc} F: & E & \to & E \\ & (x, y) & \mapsto & (x^q, y^q) \end{array}$$

Its characteristic equation is

$$F^2 - cF + q = 0.$$

One can deduce from this that the number of points on E is

$$N = q + 1 - c,$$

where c is the trace of Frobenius on E. Note that $4 \mid N$ since the point $(\sqrt[4]{a_6}, \sqrt{a_6})$ on E has order four. The number of points on E^* is $N^* = q + 1 + c$ and one has $2 \parallel N^*$. The following bound is due to Hasse [Has33]:

$$|c| \leq 2\sqrt{q}.$$

The little Frobenius automorphism σ is the map $x \mapsto x^2$. It can be extended to an isogeny from E to the conjugate curve $E^\sigma : y^2 + xy = x^3 + a_6^\sigma$ as follows:

$$\begin{array}{rccc} \sigma: & E & \to & E^\sigma \\ & (x, y) & \mapsto & (x^2, y^2). \end{array}$$

3 Counting the number of points

3.1 The Schoof-Elkies-Atkin algorithm

The first polynomial-time algorithm for counting points on elliptic curves over finite fields was described by Schoof in [Sch85]. The basic idea is to find the trace of the curve modulo small primes ℓ by studying the action of F on the ℓ–torsion part of E. Restricting the characteristic equation of F to the ℓ–torsion results in

$$(X^{q^2}, Y^{q^2}) - [q](X, Y) = [c_\ell](X^q, Y^q)$$

for each point (X, Y). This equality can be tested, for each candidate $c_\ell \in [0..\ell - 1]$, by doing polynomial arithmetic modulo the ℓ–division polynomial. Now, it suffices to compute c_ℓ for many small primes ℓ and then to recover the exact result using the Chinese Remainder Theorem. The time required for point-counting over $\mathbb{F}_q$ with this algorithm is $O(\log^{5+\varepsilon} q)$ using asymptotically fast methods for arithmetic (or $O(\log^8 q)$ using naïve arithmetic). The degree of the ℓ–division polynomial is $O(\ell^2)$, which grows quickly and causes this algorithm to be slow in practice.

In large characteristic, Elkies [Elk98] and Atkin [Atk92] improved Schoof's method yielding the so-called SEA algorithm (see [Sch95]) with run-time reduced to $O(\log^{4+\varepsilon} q)$ (or $O(\log^6 q)$) under reasonable hypotheses. Their idea is to construct a factor of degree $O(\ell)$ of the division polynomial and work with it instead. Such a factor can be found by factoring the modular polynomial to find eigenspaces of the Frobenius F restricted to $E[\ell]$.

Further work by Morain [Mor95] and others led to practical implementations of SEA for prime fields. Couveignes extended SEA to work in small characteristic using the formal group [Cou94] or the p-torsion [Cou96] and Lercier found an efficient method for characteristic two [Ler97a].

3.2 The Satoh-FGH algorithm

Here we present Fouquet-Gaudry-Harley's adaptation of Satoh's algorithm to the case of characteristic two. The reader can find more details, including for odd characteristic, in [Sat] and [FGH00].

The principal idea of this new algorithm is to lift E to a curve $\mathcal{E}$ over a 2–adic ring $\mathbb{Z}_q$ and to compute the trace of the Frobenius on $\mathcal{E}$.

3.2.1 Canonical lift of the curve

Just as $\mathbb{F}_q$ is obtained from $\mathbb{F}_2$ by taking an algebraic extension modulo an irreducible polynomial $f(x)$, one can obtain $\mathbb{Z}_q$ from the 2–adic integers $\mathbb{Z}_2$ by taking an extension modulo a polynomial $g(x)$ which reduces modulo 2 to $f(x)$. Thus we have $\mathbb{Z}_q = \mathbb{Z}_2[x]/(g(x))$. We represent this situation with the following figure.

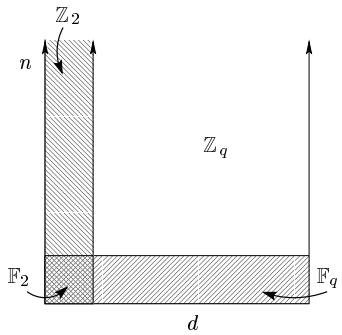

A Frobenius morphism $\mathcal{F}$ can also be defined on $\mathbb{Z}_q$. In this case it is not a simple q-th powering operation but something much more complicated. We do not define it explicitly since we will never have to compute it. Similarly, there exists a little Frobenius morphism Σ. For further details on $\mathbb{Z}_q$ and its Frobenius maps, see [Ser68].

A theorem of Lubin, Serre and Tate [LST64] guarantees the existence and uniqueness of a canonical lifted curve $\mathcal{E}$ over $\mathbb{Z}_q$ such that $\mathrm{End}(\mathcal{E}) = \mathrm{End}(E)$, via a canonical lift of the j-invariant. Indeed $J = j(\mathcal{E})$ is characterised by $J \equiv j(E) \bmod 2$ and $\Phi_2(J, \Sigma(J)) = 0$, where Φ_2 is the 2–modular polynomial.

A crucial part of Satoh's contribution is an efficient algorithm for lifting j-invariants. Instead of lifting $j(E)$ in isolation, he suggests lifting the whole cycle of conjugate j's simultaneously. He also proposes replacing the little Frobenius isogeny, Σ, with its dual, $\hat{\Sigma}$. Indeed the dual is separable and hence is determined by its kernel. As a result, for matters of computation we lift the cycle of curves related by $\hat{\Sigma}$. We get the following diagram:

$$\begin{array}{ccccccccc}
\mathcal{E}_0 & \xrightarrow{\hat{\Sigma}_0} & \mathcal{E}_1 & \xrightarrow{\hat{\Sigma}_1} & \cdots & \xrightarrow{\hat{\Sigma}_{d-2}} & \mathcal{E}_{d-1} & \xrightarrow{\hat{\Sigma}_{d-1}} & \mathcal{E}_0 \\
\downarrow{\scriptstyle \pi} & & \downarrow{\scriptstyle \pi} & & & & \downarrow{\scriptstyle \pi} & & \\
E_0 & \xrightarrow{\hat{\sigma}_0} & E_1 & \xrightarrow{\hat{\sigma}_1} & \cdots & \xrightarrow{\hat{\sigma}_{d-2}} & E_{d-1} & \xrightarrow{\hat{\sigma}_{d-1}} & E_0
\end{array}$$

where π is the reduction modulo 2 from $\mathbb{Z}_q$ to $\mathbb{F}_q$.

After having lifted the j-invariants using Satoh's method, we lift the coefficients of the curves using the algorithm of FGH. We also use FGH's deterministic method for computing the kernels of $\hat{\Sigma}$ by lifting a 2–torsion point on each conjugate curve.

3.2.2 Computing the trace in $\mathbb{Z}_q$

Since the trace of an isogeny is preserved by taking the dual and by lifting, we have the equation:

$$\mathrm{Tr}(F) = \mathrm{Tr}(\hat{F}) = \mathrm{Tr}(\hat{\mathcal{F}}).$$

Moreover $\hat{\mathcal{F}}$ can be written as the composition

$$\hat{\mathcal{F}} = \hat{\Sigma}_0 \circ \hat{\Sigma}_1 \circ \ldots \circ \hat{\Sigma}_{d-1}.$$

To find its trace we go to the formal groups of the curves. In formal groups, isogenies are represented by power series and composing isogenies is done by composing the power series. The first coefficient c_1 of the power series of $\hat{\mathcal{F}}$ is related to its trace as follows:

$$\mathrm{Tr}\,\hat{\mathcal{F}} = c_1 + \frac{q}{c_1}.$$

Therefore, computing the trace can be done by computing c_1, and the latter can be computed by composing all the power series of the $\hat{\Sigma}_i$. Only the first coefficient is needed at the end, so only the first coefficients g_i of the $\hat{\Sigma}_i$ have to be determined, and this can be done with Vélu's formulae [Vél71]. More precisely, the square of g_i is given by an explicit formula involving the lifted curves and 2–torsion. Taking one of the square roots of Πg_i^2 produces the desired trace.

3.3 Description of the algorithm

In this section, we give a synthetic description of the algorithm. For a more detailed one, we refer the reader to [FGH00]. The general procedure is:

Procedure MainAlgorithm
Input: An elliptic curve E defined over $\mathbb{F}_q$, with $j(E) \notin \mathbb{F}_4$.
Output: The trace of the curve.

1. Compute the cycle of d curves E_i and their j-invariants j_i.
2. Lift all the j_i's simultaneously, yielding J_i.
3. Lift each curve by lifting its a_6 coefficient.
4. Lift each curve's 2–torsion subgroup.
5. Compute the trace from the lifted data.

In this procedure, points 2, 3 and 4 concern the lifting of the cycle of curves and of the 2–torsion. We will detail these first. An essential ingredient is *Newton's iteration* for improving the (2–adic) precision of a root of a function.

Procedure LiftCurvesAnd2Torsion
Input: A cycle of d conjugate curves, and their j-invariants.
Output: The canonical lift of this cycle over $\mathbb{Z}_q$.

1. Lift the j-invariants simultanously using an adaptation of the Newton iteration to the multivariate case. The function that we consider acts on a vector of size d : $\Theta(x_0, \ldots, x_{d-1}) = (\Phi_2(x_0, x_1), \Phi_2(x_1, x_2), \cdots, \Phi_2(x_{d-1}, x_0))$, where the initial approximation of the root is the vector $(j_0, j_1, \ldots, j_d)$.

2. Lift each curve E_i by lifting its a_6 coefficient, yielding A_i, using a Newton iteration with the function $f(x) = 1 + J(x + 432x^2)$ and the initial approximation $-1/J_i$ modulo 16.

3. Lift the 2–torsion point of each curve yielding a 2–torsion point X_i of $\mathcal{E}_i$ with a Newton iteration based on the function $f(x) = 8x^3 + x^2 + A_i$ and with initial approximation $1/J_{i+1}$ modulo 4.

With these algorithms, one can perform the lifting efficiently. Once this is done, it remains to compute the trace of $\hat{\mathcal{F}}$. The equations in the following algorithm are derived from Vélu's formulae.

Procedure ComputeTrace
Input: A cycle of d curves, given by A_i, and 2–torsion abscissae X_i.
Output: The trace of $\hat{\mathcal{F}}$.

1. Compute the square of the first coefficient of the expansion of each $\hat{\Sigma}_i$ in the formal group of $\mathcal{E}_i$ using Vélu's formulae. The result is:

$$g_i^2 = \frac{1 - 252X_i + 19008A_i}{\left(1 + 120(X_i + 6X_i^2)\right)\left(1 + 864A_{i+1}\right)}.$$

2. Compute $c^2 = \prod g_i^2$.

3. Compute c by computing a square root of c^2 and by determining the sign using $c \equiv 1 \mod 4$.

4 Good elliptic curves in cryptography

The security of elliptic curve cryptosystems depends on the difficulty of solving the elliptic curve discrete logarithm (ECDL) problem. As mentioned in the introduction, there are several attacks against curves with special properties such as the anomalous attack, MOV reduction for supersingular curves, etc.

For random curves, the chance that one of these methods can apply is vanishingly small. However there are other attacks that work for generic abelian finite groups.

The first is Pohlig-Hellman reduction [PH78]. When the group order N has all its prime factors small, discrete logs can be computed quickly by working in small subgroups. Thus for good security it is essential to pick a group whose order is divisible by a large prime.

The other attacks are algorithms that run in time $O(\sqrt{N})$. They include Shanks' baby-step giant-step algorithm (see [Coh96]) and Pollard's ρ method [Pol78]. In practice, the most difficult ECDL that has been computed [Har] is on a Koblitz curve over $\mathbb{F}_{2^{109}}$ using a distributed version of Pollard–ρ.

By extrapolating the work required to larger sizes and allowing safety margins for future increases in computing power, it is generally believed (see [FIPS186, LV00, P1363, Sil00]) that a random curve whose order is divisible by a prime of at least 160 bits will offer reasonable security, comparable to 80-bit symmetric systems or 1024-bit RSA. For applications with the highest security requirements, one may take larger safety margins.

To find a secure curve, Lercier [Ler97a] proposed an early-abort strategy to use when computing the cardinality of the curve using SEA. The idea is to test on the fly if $q + 1 - c \bmod \ell = 0$. If the test is true, then we throw away the curve and try again with another one. Since SEA computes $c \bmod \ell$, this test is easy to implement and costs no extra run-time.

A difficulty that arises when designing an early-abort strategy to use with the Satoh-FGH algorithm is that $c \bmod \ell$ is not readily available (except for ℓ a power of two). Our solution is to

Field size	163 bits	193 bits	239 bits	409 bits	571 bits
Cost of a multiplication in $\mathbb{F}_q$	0.488 μs	0.639 μs	0.917 μs	2.632 μs	4.685 μs

Table 1: Cost of a multiplication in $\mathbb{F}_q$ on a 750 MHz Alpha.

implement a simplified version of SEA using just the first few primes ℓ, as a preliminary step before launching Satoh-FGH. There is a trade-off to be made between the extra cost of these calculations and the benefit to be gained by avoiding an entire cardinality computation. In practice we found this strategy to be very worthwhile and obtained run-times lower than those previously reported in the literature.

5 Implementation and results

5.1 Implementation details

We wrote optimised implementations of the early-abort strategy and the Satoh-FGH algorithm, in C. For multiplication in $\mathbb{F}_q$ we used Karatsuba's algorithm; in $\mathbb{Z}_q$ we used Toom's algorithm. To ensure that modular reduction took very little time, we chose the irreducible polynomial to be a trinomial or pentanomial. For division we used the binary Euclidean algorithm in $\mathbb{F}_q$, and inversion by Newton iterations in $\mathbb{Z}_q$.

Our code was written for machines with either 32-bit or 64-bit words. Most of our timing tests were run on a 750 MHz EV6 Alpha. In order to compare results with [Ler97a], we also ran some tests on a 266 MHz EV4 Alpha identical to the one Lercier used. Finally we timed curve generation for one small field on a 275 MHz StrongARM chip.

In the early-abort part, as explained below, the most time consuming parts are lazy factorizations of small-degree polynomials over $\mathbb{F}_q$. The most frequent operation is multiplication in $\mathbb{F}_q$. We give relevant timings obtained on the 750 MHz Alpha in Table 1.

The most frequent operation in the point-counting part is multiplication in $\mathbb{Z}_q$. In Table 2, we give the time for one such operation at the highest 2–adic precision required i.e., $\lceil d/2 \rceil + 1$ bits, for various field sizes d. These measurements were also done on the 750 MHz Alpha.

5.2 Counting the number of points on one curve

When computing the cardinality of a curve, one has to decide whether to use SEA or Satoh-FGH. Two cases have to be dealt with differently: the case of large characteristic and the case of small characteristic and in particular, characteristic two.

The complexity of Satoh's algorithm has a bad dependency in the characteristic p of the base field and when p is large, it is not efficient at all. This is due to the use of the modular equation Φ_p for the lifting of the curves. This equation has $O(p^2)$ coefficients that have to be known at least

Base field size	163 bits	193 bits	239 bits	409 bits	571 bits
Maximal precision	83	98	121	206	287
Cost of a multiplication in $\mathbb{Z}_q$	0.19 ms	0.24 ms	0.36 ms	4.8 ms	9.3 ms

Table 2: Cost of a multiplication in $\mathbb{Z}_q$ on a 750 MHz Alpha.

	SEA (timings from [Ler97b])			Satoh-FGH
Field size	Min	Max	Avg	
155 bits	58.8 s	132 s	86.5 s	36.5 s
196 bits	212 s	1029 s	308 s	68.8 s
300 bits	1519 s	3686 s	2434 s	443.9 s

Table 3: Times for point-counting on a 266 MHz Alpha

Field size	Satoh-FGH	Field size	Satoh-FGH	Field size	Satoh-FGH
157 bits	2.37 s	197 bits	4.52 s	283 bits	29.5 s
163 bits	2.74 s	233 bits	6.61 s	409 bits	82.5 s
193 bits	4.12 s	239 bits	6.93 s	571 bits	295 s

Table 4: Times for point-counting on a 750 MHz Alpha

modulo $p^{(d/2)+O(1)}$. Hence a complexity which is exponential in p appears to be unavoidable. On the other hand, the SEA algorithm is polynomial-time independently of p. For instance, Morain succeeded in counting the number of points of a curve over a field of size $10^{499} + 153$ [Mor95].

In characteristic two, we no longer encounter this problem. To illustrate the difference in speed between the two algorithms, we compare Lercier's results [Ler97b] with the statistics we get over the same fields, using an identical 266 MHz Alpha. The results are given in Table 3. We do not give minimal or maximal times in the case of Satoh-FGH since this algorithm is deterministic and thus the runtime is almost constant when treating different curves over the same field. These results show that Satoh-FGH is more efficient than SEA and that the bigger the field, the greater the advantage.

We give timings for point-counting on the 750 MHz Alpha in Table 4. Most of the field sizes that we chose are recommended in cryptographic standards (ANSI X9.63, IEEE P1363, IPSec, NIST, WAP).

Remark: In some cases, the SEA and Satoh-FGH algorithms can be combined to speed-up point-counting. This works particularly well when the field size is such that the maximum precision required in Satoh-FGH is a little more than a multiple of the machine word-size. A good example is $q = 2^{257}$: the maximum precision in the lifting calculations is $\lceil \frac{257}{2} \rceil + 1 = 130$ bits. In this case, computing the trace modulo 5 with the SEA algorithm takes very little time, but allows the precision to be reduced to 128 bits which fits perfectly in a whole number of words.

5.3 Finding a good curve

The naïve strategy to find a curve suitable for cryptographic use is to count the number of points for many curves, until one with almost prime order is found. As mentioned before, if the SEA algorithm is used then many bad curves can be detected early; this nice property does not hold for the Satoh-FGH algorithm.

Hence, for small to medium sizes, the naïve strategy using Satoh-FGH does not give compelling results. For instance over $\mathbb{F}_{2^{155}}$, Lercier [Ler97b] was able to select the good curves among a set of 1000 random ones in 14112 seconds. On the same computer, counting points with Satoh-FGH takes 36.5 seconds per curve so that selecting the good curves among these 1000 and their 1000 twists would take 36500 seconds.

	$q = 2^{163}$		$q = 2^{239}$	
ℓ	Root finding of $\Phi_\ell^c(X, j)$	Average total time	Root finding of $\Phi_\ell^c(X, j)$	Average total time
3	0.17 ms	0.17 ms	0.28 ms	0.28 ms
5	0.34 ms	0.38 ms	0.61 ms	0.68 ms
7	0.34 ms	1.18 ms	0.56 ms	2.18 ms
11	4.42 ms	6.93 ms	9.14 ms	14.1 ms
13	1.07 ms	4.19 ms	1.94 ms	8.36 ms
17	3.71 ms	8.63 ms	7.34 ms	17.9 ms
19	4.97 ms	11.6 ms	10.1 ms	23.7 ms

Table 5: Average runtime for checking if E is ℓ-good (750 MHz)

To improve matters, we take advantage of both methods: we first eliminate many candidates by an early-abort strategy based on SEA's techniques, and then run Satoh-FGH on the remaining curves.

Let E be a curve over $\mathbb{F}_q$. For a small prime ℓ, E is called ℓ-good if its order is coprime to ℓ, and ℓ-bad otherwise. Early-abort works as follows for each ℓ:

1. Compute the number of roots of $\Phi_\ell(X, j(E))$. It can be 0, 1, 2 or $\ell + 1$. (The cases 1 or $\ell + 1$ cannot occur unless q is a square modulo ℓ.)

2. If there are no roots, E is ℓ-good.

3. Otherwise, for each root of Φ_ℓ, build the corresponding factor of the ℓ-division polynomial and search for a root x of the factor. If there is one such x in $\mathbb{F}_q$ and if the corresponding y's are also in $\mathbb{F}_q$, then E is ℓ-bad.

4. Otherwise E is ℓ-good.

The major cost in step 1 is that of computing X^q modulo $\Phi_\ell(X, j(E))$, which has degree $\ell + 1$. To accelerate the calculation, we replace Φ_ℓ by the *canonical* modular polynomial Φ_ℓ^c, which has the same degree but is sparser and involves lower powers of j. We refer to [Mor95] for the construction and the properties of these equations.

Heuristically, in half of the cases there will be no root (in such a case ℓ is called an Atkin prime) and we are done. Otherwise, we have to continue to step 3. The factor of the division polynomial corresponding to a root of the modular polynomial is calculated using a system of formulae due to Lercier [Ler97a]. For small ℓ the solution to this system can be written explicitly, and the factor is obtained at almost no cost. (For larger ℓ the system could be solved efficiently by an algorithm also due to Lercier.) The cost of searching for a root is dominated by the computation of X^q modulo the factor, which has degree $(\ell - 1)/2$.

In Table 5 we give the run-time for this procedure, measured on the 750 MHz Alpha.

It is necessary to bound the maximum size of ℓ in order to balance the cost of early-abort against the gain obtained by avoiding point-counting. In theory, it would be beneficial to increase ℓ until the above early-abort procedure took approximately one ℓ-th of the time required for point-counting. However, almost all of the advantage to be gained comes from using the first few primes, and in practice we found $\ell \leq 19$ to be a good trade-off.

Field size	SEA (from [Ler97b])	Satoh-FGH + early-abort
155 bits	14112 s	4490 s
196 bits	30254 s	7850 s

Table 6: Time to select good curves among 1000 (266 MHz)

Field size (in bits)	Time for e.-a. on 1000 curves	Number of remaining curves	Time to count remaining curves	Number of good curves	Average time to find a good curve
157	1.9 s	43	102 s	5	20.8 s
163	1.9 s	37	101 s	6	17.2 s
193	2.9 s	57	235 s	4	59.4 s
197	2.9 s	45	204 s	2	103 s
233	4.2 s	50	331 s	1	335 s
239	4.2 s	36	249 s	3	84.6 s
283	11.7 s	45	22 min	1	22 min
409	26.3 s	52	72 min	1	72 min
571	46.6 s	35	172 min	1	172 min

Table 7: Average time to find a good curve (750 MHz)

For comparison with Lercier's results reported in [Ler97b], we ran some further tests on the 266 MHz Alpha. We initially chose a similar early-abort strategy, searching for a good curve with order $4p$ without considering the twist curves at all (but see below). As a first step in the early-abort, we determine whether the order is divisible by 8. This can be decided very quickly by computing $\operatorname{Tr} a_6$. The results can be found in Table 6. Note that we measured our timings for 157 and 197 bits instead of 155 and 196 because composite extension fields are no longer considered secure, as mentioned in the introduction.

Next, in order to maximise the performance of curve generation we decided to search simultaneously for twist curves with order $2p$ and this allowed us roughly to double the speed. As is clear from section 2, the cardinality of the twist can be found immediately from that of the curve itself. Furthermore, the early-abort strategy can easily be adapted to take the twist into account since it has the same j-invariant and the same division polynomials since the curve and its twist are isomorphic over an algebraic closure and the isomorphism preserves the abscissae.

One possibility would be to reject a pair consisting of a curve and its twist only when the early-abort strategy determines that both curves are cryptographically unsuitable. Alternatively one may pursue a more aggressive strategy by rejecting them both as soon as either one is found to be unsuitable, and immediately moving on to a new pair. Using the latter method for 1000 random curve pairs on the 750 MHz Alpha, we measured the timing results shown in Table 7. Note that the values given in the last column (average time to find a good curve) are not precise. To obtain accurate timings, it would be necessary to test many more curves[1].

Although the $O(d^3)$ space complexity of Satoh's algorithm grows quickly, the tricks described by FGH keep the constant factor small. With these tricks, the largest key size we dealt with (571 bits) requires under 10 megabytes. For moderate key sizes, the memory usage was only a few hundred kilobytes so we chose a different trade-off, using more memory in exchange for slightly higher speed.

[1] This will be done for the final version

Field size	Frequency	Time for one point-counting	Average time to find a good curve	RAM used
113 bits	275 MHz	5.9 s	49.5 s	380 KB

Table 8: Time to find a good WAP curve on an ARM chip

To investigate the possibility of generating curves in constrained environments, we ran some tests at 113 bits on an ARM chip. This small key size is recommended in the Wireless Application Forum's WTLS standard (WAP) and can be used for short-term security at a level comparable to DES. The results can be seen in table 8.

6 Conclusion

The Satoh-FGH algorithm has proven to be the method of choice whenever one wants to compute the cardinality of a random elliptic curve defined over a finite field of characteristic two. But in spite of Satoh-FGH's excellent performance (see Table 4), the SEA algorithm should not be abandoned too quickly. In the case of large characteristic it is the only practical method available. Moreover the early-abort strategy which is closely related to it is valuable when looking for a curve for cryptographic use, even in small characteristic. By combining this technique and the Satoh-FGH algorithm, we obtain an efficient way of computing secure curves (see Table 7). We conclude that it is no longer necessary, nor even advisable, to use precomputed curves in cryptography as one can easily compute new curves as desired. Finding a curve with a security level comparable with RSA-1024 takes a matter of minutes. Curve generation for short-term security, with a level equivalent to DES, is doable on a low-power chip. Finally, very high security levels similar to the highest AES level are now possible albeit in several hours.

References

[FIPS186] FIPS 186-2. Digital Signature Standard. Federal Information Processing Standards publication, January 2000. U.S. Departement of Commerce/National Institute of Standards and Technology. Available at `http://csrc.nist.gov/cryptval/dss.htm`.

[P1363] IEEE P1363. Standard specifications for public key cryptography. Available at `http://www.manta.ieee.org/groups/1363/`.

[Atk92] A. O. L. Atkin. The number of points on an elliptic curve modulo a prime. Series of e-mails to the `NMBRTHRY` mailing list, 1992.

[BSS99] I. Blake, G. Seroussi, and N. Smart. *Elliptic curves in cryptography*, volume 265 of *London Math. Soc. Lecture Note Ser.* Cambridge University Press, 1999.

[Coh96] H. Cohen. *A course in algorithmic algebraic number theory*, volume 138 of *Graduate Texts in Mathematics*. Springer–Verlag, 1996. Third printing.

[Cou94] J.-M. Couveignes. *Quelques calculs en théorie des nombres*. Thèse, Université de Bordeaux I, July 1994.

[Cou96] J.-M. Couveignes. Computing ℓ-isogenies using the p-torsion. In H. Cohen, editor, *Algorithmic Number Theory*, volume 1122 of *Lecture Notes in Comput. Sci.*, pages 59–65. Springer Verlag, 1996. Second International Symposium, ANTS-II, Talence, France, May 1996, Proceedings.

[Dew98] L. Dewaghe. Remarks on the Schoof-Elkies-Atkin algorithm. *Math. Comp.*, 67(223):1247–1252, July 1998.

[DGM99] I. Duursma, P. Gaudry, and F. Morain. Speeding up the discrete log computation on curves with automorphisms. In Kwok Yan Lam, Eiji Okamoto, and Chaoping Xing, editors, *Advances in Cryptology – ASIACRYPT '99*, volume 1716 of *Lecture Notes in Comput. Sci.*, pages 103–121. Springer-Verlag, 1999. International Conference on the Theory and Applications of Cryptology and Information Security, Singapore, November 1999, Proceedings.

[Elk98] N. Elkies. Elliptic and modular curves over finite fields and related computational issues. In D.A. Buell and eds. J.T. Teitelbaum, editors, *Computational Perspectives on Number Theory*, pages 21–76. AMS/International Press, 1998. Proceedings of a Conference in Honor of A.O.L. Atkin.

[FGH00] M. Fouquet, P. Gaudry, and R. Harley. On Satoh's algorithm and its implementation. Technical report, LIX, École polytechnique, June 2000. Available at `http://ultralix.polytechnique.fr/Labo/Mireille.Fouquet/elliptic.html`.

[FR94] G. Frey and H.-G. Rück. A remark concerning m-divisibility and the discrete logarithm in the divisor class group of curves. *Math. Comp.*, 62(206):865–874, April 1994.

[GHS00] P. Gaudry, F. Hess, and N. Smart. Constructive and destructive facets of Weil descent on elliptic curves, 2000. Submitted to J. Crypt. and available at `http://www.cs.bris.ac.uk/~nigel/weil_descent.html`.

[GLV98] R. Gallant, R. Lambert, and S. Vanstone. Improving the parallelized Pollard lambda search on binary anomalous curves, 1998. Available at `http://www.certicom.com/chal/download/paper.ps`.

[Har] R. Harley. `http://cristal.inria.fr/~harley/ecdl/`.

[Has33] H. Hasse. Beweis des Analogons der Riemannschen Vermutung für die Artinschen und F. K. Smidtschen Kongruenzzetafunktionen in gewissen elliptischen Fällen. *Ges. d. Wiss. Narichten. Math.-Phys. Klasse*, pages 253–262, 1933.

[JM99] D. Johnson and A. Menezes. The elliptic curve digital signature algorithm (ECDSA). Technical Report CORR 99-34, U. Waterloo, 1999. Available at `http://www.cacr.math.uwaterloo.ca/`.

[Kob87] N. Koblitz. Elliptic curve cryptosystems. *Math. Comp.*, 48(177):203–209, January 1987.

[Ler97a] R. Lercier. *Algorithmique des courbes elliptiques dans les corps finis.* Thèse, École polytechnique, June 1997.

[Ler97b] R. Lercier. Finding good random elliptic curves for cryptosystems defined over $\mathbb{F}_{2^n}$. In W. Fumy, editor, *Advances in Cryptology – EUROCRYPT '97*, volume 1233 of *Lecture Notes in Comput. Sci.*, pages 379–392. Springer-Verlag, 1997. International Conference on the Theory and Application of Cryptographic Techniques, Konstanz, Germany, May 1997, Proceedings.

[LST64] J. Lubin, J. P. Serre, and J. Tate. Elliptic curves and formal groups. In *Lecture notes prepared in connection with the seminars held at the Summer Institute on Algebraic Geometry, Whitney Estate, Woods Hole, Massachusetts, July 6-July 31, 1964*, 1964. Scanned copies available at `http://www.ma.utexas.edu/users/voloch/lst.html`.

[LV00] A. Lenstra and E. Verheul. Selecting cryptographic key sizes, January 2000. Presented in PKC2000.

[Men93] A. J. Menezes. *Elliptic curve public key cryptosystems.* Kluwer Academic Publishers, 1993.

[Mil87] V. Miller. Use of elliptic curves in cryptography. In A. M. Odlyzko, editor, *Advances in Cryptology – CRYPTO '86*, volume 263 of *Lecture Notes in Comput. Sci.*, pages 417–426. Springer-Verlag, 1987. Proceedings, Santa Barbara (USA), August 11–15, 1986.

[Mor95] F. Morain. Calcul du nombre de points sur une courbe elliptique dans un corps fini : aspects algorithmiques. *J. Théor. Nombres Bordeaux*, 7:255–282, 1995.

[MOV91] A. Menezes, T. Okamoto, and S. A. Vanstone. Reducing elliptic curves logarithms to logarithms in a finite field. In *Proceedings 23rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 80–89. ACM Press, 1991. May 6–8, New Orleans, Louisiana.

[Mül95] V. Müller. *Ein Algorithmus zur Bestimmung der Punktanzahl elliptischer Kurven über endlichen Körpern der Charakteristik größer drei.* PhD thesis, University of Saarland, 1995.

[PH78] S. Pohlig and M. Hellman. An improved algorithm for computing logarithms over GF(p) and its cryptographic significance. *IEEE Trans. Inform. Theory*, IT–24:106–110, 1978.

[Pol78] J. M. Pollard. Monte Carlo methods for index computation (mod p). *Math. Comp.*, 32(143):918–924, July 1978.

[SA98] T. Satoh and K. Araki. Fermat quotients and the polynomial time discrete log algorithm for anomalous elliptic curves. *Comment. Math. Univ. St. Paul.*, 47:81–92, 1998.

[Sat] T. Satoh. The canonical lift of an ordinary elliptic curve over a finite field and its point counting. Available at `http://www.rimath.saitama-u.ac.jp/lab.en/TkkzSatoh/`.

[Sch85] R. Schoof. Elliptic curves over finite fields and the computation of square roots mod p. *Math. Comp.*, 44:483–494, 1985.

[Sch95] R. Schoof. Counting points on elliptic curves over finite fields. *J. Théor. Nombres Bordeaux*, 7:219–254, 1995.

[Sem98] I. A. Semaev. Evaluation of discrete logarithms in a group of p-torsion points of an elliptic curves in characteristic p. *Math. Comp.*, 67(221):353–356, January 1998.

[Ser68] J. P. Serre. *Corps locaux*. Hermann, 1968.

[Sil86] J. H. Silverman. *The arithmetic of elliptic curves*, volume 106 of *Graduate Texts in Mathematics*. Springer–Verlag, 1986.

[Sil00] R. Silverman. A cost-based security analysis of symmetric and assymetric key lengths. Bulletin Number 13 of RSA Security, April 2000.

[Skj] B. Skjernaa. Satoh's algorithm in characteristic 2. Copies available at `http://www.imf.au.dk/~skjernaa/`.

[Sma99] N. Smart. The discrete logarithm problem on elliptic curves of trace one. *J. Cryptology*, 12:193–196, 1999.

[Vél71] J. Vélu. Isogénies entre courbes elliptiques. *C. R. Acad. Sci. Paris Sér. I Math.*, 273:238–241, July 1971. Série A.

[vOW99] P. C. van Oorschot and M. J. Wiener. Parallel collision search with cryptanalytic applications. *J. of Cryptology*, 12:1–28, 1999.