INVALID-CURVE ATTACKS ON (HYPER)ELLIPTIC CURVE CRYPTOSYSTEMS

Koray Karabina

Department of Combinatorics and Optimization
Faculty of Mathematics
University of Waterloo
200 University Avenue West
Waterloo, ON, Canada, N2L 3G1

Berkant Ustaoglu

NTT Information Sharing Platform Laboratories
3-9-11, Midori-cho Musashino-shi
Tokyo 180-8585, Japan

(Communicated by Renate Scheidler)

Abstract. We extend the notion of an invalid-curve attack from elliptic curves to genus 2 hyperelliptic curves. We also show that invalid singular (hyper)elliptic curves can be used in mounting invalid-curve attacks on (hyper)elliptic curve cryptosystems, and make quantitative estimates of the practicality of these attacks. We thereby show that proper key validation is necessary even in cryptosystems based on hyperelliptic curves. As a byproduct, we enumerate the isomorphism classes of genus g hyperelliptic curves over a finite field by a new counting argument that is simpler than the previous methods.

1. Introduction

The purpose of public-key validation is to verify that a public key possesses certain arithmetic properties. Public-key validation is especially important in discrete logarithm protocols where a party $\hat{B}$ combines his private key with a public key received from a second party $\hat{A}$ to form a group element σ. A dishonest party $\hat{A}$ might select an invalid public key in such a way that the subsequent use of σ in the protocol leaks information about $\hat{B}$'s private key. Lim and Lee [18] demonstrated the importance of public-key validation by presenting *small-subgroup attacks* on some discrete logarithm key agreement protocols that are effective if the receiver of a group element does not verify that the element belongs to the desired group of high order (e.g., a prime-order DSA-type subgroup of $\mathbb{F}_p^*$). In [3, 1], *invalid-curve attacks* were designed that are effective on elliptic curve protocols if the receiver of a point does not verify that the point indeed lies on the chosen elliptic curve; see also [20, 22, 21]. Chen, Cheng and Smart [6] illustrated the importance of public-key validation in identity-based key agreement protocols that use bilinear pairings.

The performance of low-genus hyperelliptic curves has been shown to be competitive with that of elliptic curves*; see [2] for a summary of recent work. We demonstrate that invalid-curve attacks can be successfully mounted on protocols

2000 *Mathematics Subject Classification:* 94A60.
Key words and phrases: Invalid-curve attacks, hyperelliptic curves.
*Elliptic curves are the genus 1 hyperelliptic curves.

based on genus 2 hyperelliptic curves if the appropriate public-key validation is not performed. We also show that *singular* curves can be used in mounting invalid-curve attacks against (hyper)elliptic curve protocols. We illustrate our attacks on two recently-proposed discrete logarithm protocols — the Twin Diffie-Hellman key agreement scheme [5] and the XCR signature scheme [14].

In order to analyze invalid-curve attacks on hyperelliptic curve cryptosystems, it is useful to know the number $N_g(q)$ of isomorphism classes of genus g hyperelliptic curves over a finite field $\mathbb{F}_q$. The isomorphism classes of genus g hyperelliptic curves over an algebraically closed field K are in 1-1 correspondence with the elements of a $(2g-1)$-dimensional irreducible subvariety H_g of the moduli space M_g over K (see [11, p.347]), suggesting that number is of the order of q^{2g-1}. This was confirmed by Nart [23], who gave a closed formula for $N_g(q)$. We give an elementary counting argument that $N_g(q) = 2q^{2g-1} + \mathcal{O}(gq^{2g-2})$.

The remainder of the paper is organized as follows. After laying some of the mathematical groundwork in §2, we present in §3 our derivation of the number of genus g hyperelliptic curves over a finite field. In §4, we extend the notion of an invalid curve from elliptic curves to genus 2 curves. We also present the notion of an invalid singular curve, and enumerate the invalid elliptic and genus 2 singular curves. Our invalid-curve attacks are demonstrated and analyzed in §5; we conclude in §6.

2. Mathematical preliminaries

Notation. The operator $[x^i]$ denotes the coefficient extraction operator when x is an indeterminate. For indeterminate x and polynomial $f(x)$ we adopt the convention $[x^i]f(x) = f_i$. The set of monic polynomials of degree d over a finite field $\mathbb{F}_q$ is denoted by $\mathcal{P}^d$; the subset of polynomials with at least one repeated root will be denoted by $\tilde{\mathcal{P}}^d$. Let $\tilde{f} = \tilde{f}_{d-1}, \tilde{f}_{d-2}, \ldots, \tilde{f}_{d-i}$ be an ordered sequence where each $\tilde{f}_j \in \mathbb{F}_q$. Then

$$\mathcal{P}^d_{\tilde{f}} := \{x^d + \tilde{f}_{d-1}x^{d-1} + \cdots + \tilde{f}_{d-i}x^{d-i} + f_{d-i-1}x^{d-i-1} + \cdots + f_0 \mid f_{d-i-1}, \ldots, f_0 \in \mathbb{F}_q\}.$$

For example, $\mathcal{P}^2_{-3}$ denotes the set of polynomials of the form $x^2 - 3x + f_0$ where $f_0 \in \mathbb{F}_q$.

A *hyperelliptic curve* $\mathcal{H}$ of *genus* g over a finite field $\mathbb{F}_q$ is defined by a non-singular *Weierstrass equation*

$$\mathcal{H}: \; y^2 + H(x)y = F(x),$$

where $F, H \in \mathbb{F}_q[x]$, F is monic, $\deg(F) = 2g+1$, and $\deg(H) \leq g$. The *Jacobian* $J_{\mathcal{H}}(\mathbb{F}_q)$ of $\mathcal{H}$ over $\mathbb{F}_q$ is the quotient group of the *degree zero divisors* defined over $\mathbb{F}_q$ by the group of *principal divisors* defined over $\mathbb{F}_q$. The divisor classes $\bar{D} \in J_{\mathcal{H}}(\mathbb{F}_q)$ are in one-to-one correspondence with the pairs of polynomials (u, v) with $u, v \in \mathbb{F}_q[x]$, $\deg(v) < \deg(u) \leq g$, u monic, and $u \mid (v^2 + Hv - F)$; we write $\bar{D} = [u, v]$. $J_{\mathcal{H}}(\mathbb{F}_q)$ is a finite abelian group with $|J_{\mathcal{H}}(\mathbb{F}_q)| \in [(\sqrt{q}-1)^{2g}, (\sqrt{q}+1)^{2g}]$ [25]. Given two divisor classes $\bar{D}_1 = [u_1, v_1]$ and $\bar{D}_2 = [u_2, v_2] \in J_{\mathcal{H}}(\mathbb{F}_q)$, Cantor's algorithm [4] can be used to find the unique divisor $\bar{D} = [u, v]$ such that $\bar{D} = \bar{D}_1 + \bar{D}_2$.

If $\mathrm{char}(\mathbb{F}_q) \notin \{2, 2g+1\}$ then the same curve $\mathcal{H}$ (up to isomorphism) can be given by the equation

$$\mathcal{H}: y^2 = f(x) = x^{2g+1} + \sum_{i=0}^{2g-1} f_i x^i. \tag{1}$$

The non-singularity requirement on the equation of $\mathcal{H}$ means that f has no repeated roots, in which case we call $\mathcal{H}$ a *non-singular curve*. If f has x_0 as a repeated root, we call $\mathcal{H}$ a *singular curve*, and (x_0, y_0), where $y_0^2 = f(x_0)$, a singular point on $\mathcal{H}$.

The remainder of this work assumes that a hyperelliptic curve $\mathcal{H}$ of genus g over a finite field $\mathbb{F}_q$ is given via (1). The set of all (non-singular) genus g hyperelliptic curves over $\mathbb{F}_q$ will be denoted by $\mathcal{H}^*$. When a hyperelliptic curve $\mathcal{H}$ is defined over $\mathbb{F}_q$ we will abbreviate $J_{\mathcal{H}}(\mathbb{F}_q)$ to $J_{\mathcal{H}}$.

3. The number of hyperelliptic curves

In this section we estimate the number of non-isomorphic genus g hyperelliptic curves given by (1). First we need the following formula for the number of monic polynomials with no repeated roots and fixed second leading coefficient.

Theorem 1. *Let $g \geq 1$ be an integer, $\mathbb{F}_q$ be a finite field, and $f_{2g} \in \mathbb{F}_q$. Then*

$$|\mathcal{P}^{2g+1}_{f_{2g}} \setminus \tilde{\mathcal{P}}^{2g+1}_{f_{2g}}| = q^{2g} - q^{2g-1}.$$

Proof. Let $\bar{\mathbb{F}}_q$ denote the closure of $\mathbb{F}_q$. The argument proceeds by induction on g.

For $g = 1$, consider $f(x) = x^3 + f_2x^2 + f_1x + f_0 \in \mathcal{P}^3_{f_2}$. Since f has degree three, it can have at most one repeated root, say α. If $\alpha \in \bar{\mathbb{F}}_q \setminus \mathbb{F}_q$ then the conjugates of α are also repeated roots of f contradicting the fact that f has at most one repeated root. Therefore $\alpha \in \mathbb{F}_q$. Factoring f gives $f(x) = (x-\alpha)^2(x-\beta)$, where $\beta = -f_2 - 2\alpha$. Hence, the only degree of freedom is α and so $|\tilde{\mathcal{P}}^3_{f_2}| = q$. And, since $|\mathcal{P}^3_{f_2}| = q^2$, we have $|\mathcal{P}^3_{f_2} \setminus \tilde{\mathcal{P}}^3_{f_2}| = q^2 - q$.

Assume the result holds for all integers $1, 2, \ldots, (g-1)$, where $g \geq 2$. We will show the result holds for g. Let $f(x) = x^{2g+1} + \sum_{j=0}^{2g} f_j x^j \in \mathcal{P}^{2g+1}_{f_{2g}}$. To each repeated root α of f with multiplicity $k \geq 2$, we associate $\lfloor k/2 \rfloor$ pairs (α, α); we call each such pair a *paired repeated root* of f corresponding to α. Note that f can have at most g paired repeated roots. Since $f \in \mathbb{F}_q[x]$, if f has exactly i paired repeated roots then it can be written as $f = a^2b$, where $a, b \in \mathbb{F}_q[x]$,

$$\begin{aligned} a(x) &= x^i + a_{i-1}x^{i-1} + \cdots + a_0, \\ b(x) &= x^{2(g-i)+1} + \beta x^{2(g-i)} + b_{2(g-i)-1}x^{2(g-i)-1} + \cdots + b_0, \end{aligned}$$

and b has no repeated roots in $\bar{\mathbb{F}}_q$. Since the coefficient β in $b(x)$ satisfies $\beta + [x^{2i-1}]a(x)^2 = f_{2g}$, β is determined by a and f_{2g}. For a fixed $i \in [1, g-1]$ the number of polynomials a is q^i. By induction, the number of polynomials b of degree $2(g-i)+1$ that have no repeated roots and have fixed β is $q^{2(g-i)} - q^{2(g-i)-1}$. Therefore, the number of polynomials f with exactly i paired repeated roots is $q^i \cdot (q^{2(g-i)} - q^{2(g-i)-1}) = q^{2g-i} - q^{2g-i-1}$. For $i = g$ the polynomial f factors as $a^2(x-\beta)$ and as before β is determined by a and f_{2g}. Then the number of polynomials f with exactly g paired repeated roots equals the number of choices for a, which is q^g.

Hence, the number of polynomials f with at least one paired repeated root is

$$\begin{aligned} |\tilde{\mathcal{P}}^{2g+1}_{f_{2g}}| &= q^g + \sum_{i=1}^{g-1}(q^{2g-i} - q^{2g-i-1}) \\ &= q^g + q^{2g-1} + \sum_{i=2}^{g-1} q^{2g-i} - \sum_{i=1}^{g-1} q^{2g-i-1} \;=\; q^{2g-1}. \end{aligned}$$

Finally, since $|\mathcal{P}_{f_{2g}}^{2g+1}| = q^{2g}$, we have

$$|\mathcal{P}_{f_{2g}}^{2g+1} \setminus \tilde{\mathcal{P}}_{f_{2g}}^{2g+1}| = |\mathcal{P}_{f_{2g}}^{2g+1}| - |\tilde{\mathcal{P}}_{f_{2g}}^{2g+1}| = q^{2g} - q^{2g-1},$$

which completes the argument. □

Setting $f_{2g} = 0$ in Theorem 1 determines the number of polynomials that define a genus g hyperelliptic curve over $\mathbb{F}_q$, where $\mathrm{char}(\mathbb{F}_q) \notin \{2, 2g+1\}$. However, such curves can have more than one representation, and we do not wish to distinguish between isomorphic curves. The following result, due to Lockhart [19], gives a one-to-one correspondence between isomorphism classes of curves and equivalence classes of Weierstrass equations.

Theorem 2. [19, Proposition 1.2] *If $\mathcal{H}_1$ and $\mathcal{H}_2$ are two genus g hyperelliptic curves defined over $\mathbb{F}_q$, then $\mathcal{H}_1$ and $\mathcal{H}_2$ are isomorphic over $\mathbb{F}_q$ if and only if there exists $\alpha \in \mathbb{F}_q^*$, $\beta \in \mathbb{F}_q$, and $t \in \mathbb{F}_q[x]$ with $deg(t) \leq g$, such that the change of coordinates $(x, y) \to (\alpha^2 x + \beta, \alpha^{2g+1} y + t(x))$, transforms the equation of $\mathcal{H}_1$ to the equation of $\mathcal{H}_2$.*

Invalid curve attacks are based on the curve representation (and the explicit group law). Hence we are interested in isomorphisms that preserve the representation of curves. Lockhart's theorem can be specialized to suit our needs as follows.

Suppose that $\mathrm{char}(\mathbb{F}_q)$ is odd and $\mathrm{char}(\mathbb{F}_q) \nmid (2g+1)$. Let $\mathcal{H}$ be a genus g hyperelliptic curve over $\mathbb{F}_q$, and let $\tau : (x, y) \to (\alpha^2 x + \beta, \alpha^{2g+1} y + t(x))$ be an isomorphism to another genus g hyperelliptic curve over $\mathbb{F}_q$. We claim that if τ preserves the form of equation (1), it must be the case that $\beta = 0$ and $t(x) = 0$. Indeed, if $t(x) \neq 0$ then applying τ to $\mathcal{H}$ will result in a linear term in y, which is not present in (1). Now, applying the transformation $(x, y) \to (\alpha^2 x + \beta, \alpha^{2g+1} y)$ to (1), the coefficient of x^{2g} is $(2g+1)\beta\alpha^{4g}$ which has to be zero as in (1). Since $\alpha \neq 0$ and $\mathrm{char}(\mathbb{F}_q) \nmid (2g+1)$ it must be the case that $\beta = 0$. Therefore, the class of transformations that correspond to isomorphisms of a hyperelliptic curve that preserve (1) are of the form $(x, y) \to (\alpha^2 x, \alpha^{2g+1} y)$, where $\alpha \in \mathbb{F}_q^*$. From now on, when we talk about isomorphisms we will identify them with the element α.

In Theorem 3 we estimate the number of isomorphism classes of genus g hyperelliptic curves over finite fields $\mathbb{F}_q$ of odd characteristic not dividing $2g+1$. Note that, by Theorem 1, the number of polynomials $f \in \mathbb{F}_q[x]$ that give rise to a genus g hyperelliptic curve is $q^{2g} - q^{2g-1}$, and since the number of isomorphisms (nonzero field elements) is $q - 1$ we expect that the number of non-isomorphic curves to be about q^{2g-1}. A slightly stronger result is proved in [23, Theorem 3.3] that appeared after our paper was submitted. The proof given here is simpler and is presented for the reader's convenience.

Theorem 3. *Let g be fixed. Let $\mathbb{F}_q$ be a finite field of odd characteristic p, and suppose that $p \nmid (2g+1)$ and $q > 4g+2$. Then the number of non-isomorphic genus g hyperelliptic curves over $\mathbb{F}_q$ is*

$$N_g(q) = 2q^{2g-1} + \mathcal{O}(gq^{2g-2}).$$

Proof. Let $f \in \mathcal{P}_0^{2g+1} \setminus \tilde{\mathcal{P}}_0^{2g+1}$. Define $z_i(f) = 0$ if $f_i = 0$, and $z_i(f) = 1$ otherwise. We will abbreviate $z_i(f)$ to z_i in case f is clear from context. We call the sequence $z(f) = (z_0, z_1, \ldots, z_{2g-1})$ the *characteristic* sequence of f. We will use $ab\tilde{z}$ to denote the sequence z where $z_0 = a$, $z_1 = b$, and the remaining entries are given by $\tilde{z}$. Let $\mathcal{H}_{ab}^z$ be the set of polynomials f with characteristic sequence z such that

$(z_0, z_1) = (a, b)$. Let $|z|$ denote the number of nonzero entries in a sequence z. Then $|\mathcal{H}^z_{10}| \le (q-1)^{|z|}$ and $|\mathcal{H}^z_{01}| \le (q-1)^{|z|}$.

An isomorphism α acting on the curve $y^2 = f(x)$ preserves the characteristic sequence z. Indeed, applying an isomorphism $\alpha : (x, y) \to (\alpha^2 x, \alpha^{2g+1} y)$ to (1) results in

$$\alpha^{4g+2} y^2 = \alpha^{4g+2} x^{2g+1} + \alpha^{4g-2} f_{2g-1} x^{2g-1} + \cdots + \alpha^2 f_1 x + f_0,$$

which can be rewritten as

$$y^2 = x^{2g+1} + \alpha^{-4} f_{2g-1} x^{2g-1} + \cdots + \alpha^{-4g} f_1 x + \alpha^{-4g-2} f_0.$$

Thus the curves $y^2 = f(x)$ for $f \in \mathcal{H}^z_{ab}$ have the same automorphism group; we denote this group by $\mathtt{Aut}\mathcal{H}^z_{ab}$. Observe also that if the mapping α is an automorphism then the order of α in $\mathbb{F}_q$ must divide $4g+2$. Since $q > 4g+2$, $|\mathtt{Aut}\mathcal{H}^z_{ab}|$ is no larger than $4g+2$ which yields

$$\begin{aligned}
\sum_z |\mathcal{H}^z_{01}|(|\mathtt{Aut}\mathcal{H}^z_{01}| - 2) &\le 4g \sum_z |\mathcal{H}^z_{01}| \le 4g \sum_z (q-1)^{|z|} \\
&= 4g \sum_{|z|=1}^{2g-1} \binom{2g-2}{|z|-1} (q-1)^{|z|} \\
&= 4g(q-1) \sum_{i=1}^{2g-1} \binom{2g-2}{i-1} (q-1)^{i-1} \\
&= 4g(q-1) \sum_{i=0}^{2g-2} \binom{2g-2}{i} (q-1)^{i} \\
&= 4g(q-1) q^{2g-2} \\
&= \mathcal{O}(g q^{2g-1}).
\end{aligned}$$

Note that in the above equations, the sequence z has its first two coordinates fixed, and therefore the remaining $|z|-1$ nonzero entries are chosen from $2g-2$ possible indices. Similarly, we have

$$\sum_z |\mathcal{H}^z_{10}|(|\mathtt{Aut}\mathcal{H}^z_{10}| - 2) = \mathcal{O}(g q^{2g-1}).$$

If both z_0 and z_1 are equal to zero, then 0 is a repeated root of f. So if f has no repeated roots then at least one of z_0 or z_1 is nonzero. By Theorem 1 we have

$$\sum_z (|\mathcal{H}^z_{11}| + |\mathcal{H}^z_{01}| + |\mathcal{H}^z_{10}|) = q^{2g} - q^{2g-1}.$$

An automorphism α that fixes f when f_1 and f_0 are simultaneously nonzero must satisfy $\alpha^2 = 1$. There are two such automorphisms, so $|\mathtt{Aut}\mathcal{H}^z_{11}| = 2$. Therefore,

$$\begin{aligned}
N_g(q) &= \sum_z \frac{|\mathcal{H}^z_{ab}|}{[\mathbb{F}_q^* : \mathtt{Aut}\mathcal{H}^z_{ab}]} = \sum_z \frac{|\mathcal{H}^z_{ab}||\mathtt{Aut}\mathcal{H}^z_{ab}|}{|\mathbb{F}_q^*|} = \sum_z \frac{|\mathcal{H}^z_{ab}||\mathtt{Aut}\mathcal{H}^z_{ab}|}{q-1} \\
&= \sum_{01\tilde{z}} \frac{|\mathcal{H}^z_{01}||\mathtt{Aut}\mathcal{H}^z_{01}|}{q-1} + \sum_{10\tilde{z}} \frac{|\mathcal{H}^z_{10}||\mathtt{Aut}\mathcal{H}^z_{10}|}{q-1} + \sum_{11\tilde{z}} \frac{|\mathcal{H}^z_{11}||\mathtt{Aut}\mathcal{H}^z_{11}|}{q-1}
\end{aligned}$$

$$
\begin{aligned}
&= \frac{1}{q-1}\left(\sum_{01\tilde{z}}|\mathcal{H}^z_{01}||\mathtt{Aut}\mathcal{H}^z_{01}| + \sum_{10\tilde{z}}|\mathcal{H}^z_{10}||\mathtt{Aut}\mathcal{H}^z_{10}| + \sum_{11\tilde{z}}2|\mathcal{H}^z_{11}|\right)\\
&= \frac{1}{q-1}\left(\sum_{01\tilde{z}}|\mathcal{H}^z_{01}||\mathtt{Aut}\mathcal{H}^z_{01}| + \sum_{10\tilde{z}}|\mathcal{H}^z_{10}||\mathtt{Aut}\mathcal{H}^z_{10}|\right.\\
&\qquad\left.+2\left(q^{2g}-q^{2g-1}-\sum_{01\tilde{z}}|\mathcal{H}^z_{01}|-\sum_{10\tilde{z}}|\mathcal{H}^z_{10}|\right)\right)\\
&= \frac{1}{q-1}\left(\sum_{01\tilde{z}}|\mathcal{H}^z_{01}|\left(|\mathtt{Aut}\mathcal{H}^z_{01}|-2\right) + \sum_{10\tilde{z}}|\mathcal{H}^z_{10}|\left(|\mathtt{Aut}\mathcal{H}^z_{10}|-2\right)\right.\\
&\qquad\left.+2\left(q^{2g}-q^{2g-1}\right)\right)\\
&= \frac{1}{q-1}\left(\mathcal{O}(gq^{2g-1}) + \mathcal{O}(gq^{2g-1}) + 2q^{2g} - 2q^{2g-1}\right)\\
&= \frac{1}{q-1}\left(2q^{2g} + \mathcal{O}(gq^{2g-1})\right)\\
&= 2q^{2g-1} + \mathcal{O}(gq^{2g-2}),
\end{aligned}
$$

as required. □

4. Invalid and singular curves

We now extend the notion of invalid elliptic curves, proposed by Antipa et al. [1], to genus 2 curves. We emphasize that invalid curves are defined with respect to a specific curve representation and explicit formulae for the group law. That is, given a curve representation and formulae for the group operations that do not make use of a specific coefficient in the selected curve representation, one can define an invalid curve with respect to that representation-formulae pair. If the explicit formulae utilize all the coefficients in the curve representation then invalid curves in this context do not exist. However, for curves of genus 1 and genus 2, which are widely considered for cryptographic applications, the notion of invalid curves is indeed relevant and important.

In genus 1 setting, we use the affine formulae for the group law as described in [2, §13.2.1], and refer to these formulae as $\mathcal{F}_{1a}$ throughout the paper. The explicit computations in $\mathcal{F}_{1a}$ require only the coefficient f_1. In our definition of invalid elliptic curves, we include singular elliptic curves as well.

Definition 1. Let $\mathcal{E}$ be an elliptic curve defined over $\mathbb{F}_q$ with equation

$$\mathcal{E} : y^2 = x^3 + f_1 x + f_0.$$

An *invalid curve* relative to $\mathcal{E}$ and $\mathcal{F}_{1a}$ is an elliptic curve over $\mathbb{F}_q$ with equation

$$\mathcal{IE} : y^2 = x^3 + f_1 x + \tilde{f}_0,$$

where $\tilde{f}_0 \neq f_0$ and $\mathcal{IE}$ is not isomorphic to $\mathcal{E}$. In addition, if the polynomial $\tilde{f}(x) = x^3 + f_1 x + \tilde{f}_0$ has a repeated root then $\mathcal{IE}$ is called an *invalid singular curve* relative to $\mathcal{E}$ and $\mathcal{F}_{1a}$.

In genus 2 setting, we will use the affine formulae for the group law as described in [2, §14.3.2], and refer to these formulae as $\mathcal{F}_{2a}$ throughout the paper. The

formulae $\mathcal{F}_{2a}$ depends only on f_2 and f_3. In our definition of invalid hyperelliptic curves, we include singular hyperelliptic curves as well.

Definition 2. Let $\mathcal{H}$ be a genus 2 hyperelliptic curve defined over $\mathbb{F}_q$ with equation

$$\mathcal{H} : y^2 = x^5 + f_3x^3 + f_2x^2 + f_1x + f_0.$$

An *invalid curve* relative to $\mathcal{H}$ and $\mathcal{F}_{2a}$ is a hyperelliptic curve over $\mathbb{F}_q$ with equation

$$\mathcal{IH} : y^2 = x^5 + f_3x^3 + f_2x^2 + \tilde{f}_1x + \tilde{f}_0,$$

where $(\tilde{f}_1, \tilde{f}_0) \neq (f_1, f_0)$ and $\mathcal{IH}$ is not isomorphic to $\mathcal{H}$. In addition, if the polynomial $\tilde{f}(x) = x^5 + f_3x^3 + f_2x^2 + \tilde{f}_1x + \tilde{f}_0$ has a repeated root then $\mathcal{IH}$ is called an *invalid singular curve* relative to $\mathcal{H}$ and $\mathcal{F}_{2a}$.

Invalid singular curves are very interesting in the genus 1 case. If $\mathcal{SE}$ is an invalid singular curve over $\mathbb{F}_q$ relative to the elliptic curve $\mathcal{E}$ over $\mathbb{F}_q$ and $\mathcal{F}_{1a}$, then $\mathcal{SE}$ has exactly one singular point $P = (x_0, y_0)$. Applying the isomorphism $(x, y) \to (x + x_0, y + y_0)$ to $\mathcal{SE}$, we can assume that $\mathcal{SE}$ is given by the equation

$$\mathcal{SE} : y^2 = x^3 + a_2x^2,\ a_2 \in \mathbb{F}_q,$$

and $P = (0, 0)$ is the singular point of $\mathcal{SE}$. Now, let $y^2 - a_2x^2 = (y - \alpha x)(y - \beta x)$ where $\alpha, \beta \in \bar{\mathbb{F}}_q$. If $a_2 = 0$ then $\alpha = \beta = 0$, and P is called a *cusp* singularity of $\mathcal{SE}$. If $a_2 \neq 0$ then $\alpha = -\beta$, and $\alpha^2 = a_2$; P is called a *node* singularity of $\mathcal{SE}$. In this case, $\alpha, \beta \in \mathbb{F}_q$ if a_2 is a quadratic residue in $\mathbb{F}_q$; and $\alpha, \beta \in \mathbb{F}_{q^2} \setminus \mathbb{F}_q$, otherwise. It is well known that the set $\mathcal{SE}_{ns}(\mathbb{F}_q)$ of non-singular $\mathbb{F}_q$-points on $\mathcal{SE}$ together with the point at infinity forms a group and in fact the group law $\mathcal{F}_{1a}$ for $\mathcal{E}$ is also the group law for $\mathcal{SE}$. Moreover, if P is a cusp singularity of $\mathcal{SE}$ then $\mathcal{SE}_{ns}(\mathbb{F}_q)$ is isomorphic to the additive group of $\mathbb{F}_q$. If P is a node singularity of $\mathcal{SE}$ and $\alpha \in \mathbb{F}_q$ then $\mathcal{SE}_{ns}(\mathbb{F}_q)$ is isomorphic to $\mathbb{F}_q^*$; and if P is a node singularity of $\mathcal{SE}$ and $\alpha \in \mathbb{F}_{q^2} \setminus \mathbb{F}_q$ then $\mathcal{SE}_{ns}(\mathbb{F}_q)$ is isomorphic to the order-$(q+1)$ multiplicative subgroup of $\mathbb{F}_{q^2}^*$. In all cases, the isomorphisms are efficiently computable (see [12, §7.2] for more details). The key point in applying singular invalid-curve attacks is that the group law for non-singular elliptic curves can readily be used for the non-singular part of singular elliptic curves.

The next theorem assumes the setting in Definition 1 and establishes the existence and the number of invalid singular elliptic curves.

Theorem 4. *Over $\mathbb{F}_q$, where $gcd(q, 6) = 1$, the number of invalid singular curves relative to $y^2 = x^3 + f_1x + f_0$ and $\mathcal{F}_{1a}$ is*

(i) 1, *if $f_1 = 0$;*
(ii) 2, *if $-\frac{f_1}{3}$ is a quadratic residue in $\mathbb{F}_q$;*
(iii) 0, *if $-\frac{f_1}{3}$ is a quadratic non-residue in $\mathbb{F}_q$.*

Proof. Let $f_1 \in \mathbb{F}_q$ be fixed and consider the set of polynomials

$$\mathcal{P}^3_{0,f_1} = \{P_{f_0'}(x) = x^3 + f_1x + f_0' : f_0' \in \mathbb{F}_q\}.$$

If $P_{f_0'}(x)$ has a repeated root in $\bar{\mathbb{F}}_q$ then we must have

$$P_{f_0'}(x) = x^3 + f_1x + f_0' = (x + a)^2(x + k_0),\ a, k_0 \in \mathbb{F}_q,$$

or equivalently,

$$k_0 = -2a \tag{2}$$
$$f_1 = -3a^2 \tag{3}$$
$$f_0' = -2a^3. \tag{4}$$

For a fixed f_1, we consider the solutions (a, k_0, f_0') to (2)–(4).
Case (i). If $f_1 = 0$ then $(0,0,0)$ is the only solution to (2)–(4) and so $\mathcal{P}^3_{0,f_1}$ has exactly one polynomial that has a repeated root, namely $P(x) = x^3$. In this case, the curve defined by $y^2 = P(x)$ has a cusp singularity at $S = (0,0)$.
Case (ii). If $-f_1/3 = a_1^2$ for some $a_1 \in \mathbb{F}_q^*$ then $(a_1, -2a_1, -2a_1^3)$ and $(-a_1, 2a_1, 2a_1^3)$ are the only two solutions to (2)–(4). Hence, $\mathcal{P}^3_{0,f_1}$ has exactly two polynomials with repeated roots: $P(x) = x^3 + f_1 x - 2a_1^3$ in which case the curve defined by $y^2 = P(x)$ has a node singularity at $S = (-a_1, 0)$; and $P(x) = x^3 + f_1 x + 2a_1^3$ in which case the curve defined by $y^2 = P(x)$ has a node singularity at $S = (a_1, 0)$.
Case (iii). If $-f_1/3$ is a quadratic non-residue in $\mathbb{F}_q$ then the system defined by (2)–(4) has no solutions and so $\mathcal{P}^3_{0,f_1}$ does not contain any polynomial with repeated roots. □

In the attacks described in §5, the adversary will need curves with small-order subgroups. In the genus 1 case the adversary has the ability to choose $\tilde{f}_0$, and in [1, §4.4] it was already argued that an adversary can efficiently find suitable invalid curves, essentially by picking curves at random. To extend that argument for genus 2 we need the following result.

Theorem 5. *Suppose that* $gcd(q, 30) = 1$. *The number of invalid singular curves relative to* $y^2 = x^5 + f_3x^3 + f_2x^2 + f_1x + f_0$ *and* $\mathcal{F}_{2a}$ *over* $\mathbb{F}_q$ *is between* $q - 3$ *and* $q + 3$.

Proof. Let $f_2, f_3 \in \mathbb{F}_q$ be fixed, and consider the set of polynomials

$$\mathcal{P}^5_{0,f_3,f_2} = \{P_{f_1',f_0'}(x) = x^5 + f_3x^3 + f_2x^2 + f_1'x + f_0' : \ f_1', f_0' \in \mathbb{F}_q\}.$$

If $P_{f_1',f_0'}(x)$ has repeated roots in $\bar{\mathbb{F}}_q$, then there are two (not mutually exclusive) possibilities:
Case 1. There exist $a, b, k_0 \in \mathbb{F}_q$ such that

$$P_{f_1',f_0'}(x) = x^5 + f_3x^3 + f_2x^2 + f_1'x + f_0' = (x^2 + ax + b)^2(x + k_0). \tag{5}$$

Comparing the coefficients of the same degree terms, we have

$$k_0 = -2a \tag{6}$$
$$f_3 = 2b - 3a^2 \tag{7}$$
$$f_2 = -(2a^3 + 2ab) \tag{8}$$
$$f_1' = b^2 - 4a^2b \tag{9}$$
$$f_0' = -2b^2a. \tag{10}$$

We obtain from (7) and (8) that a must satisfy $5a^3 + f_3a + f_2 = 0$. Moreover, since f_2 and f_3 are already fixed, any choice of a fixes k_0 and b by (6) and (7). Therefore, the number of polynomials $P_{f_1',f_0'}(x)$ of the form (5) is at most three.

Case 2. There exist $a, k_0, k_1, k_2 \in \mathbb{F}_q$ such that

$$P_{f_1',f_0'}(x) = x^5 + f_3x^3 + f_2x^2 + f_1'x + f_0' = (x+a)^2(x^3 + k_2x^2 + k_1x + k_0).$$

Comparing the coefficients of the same degree terms, we have

$$k_2 = -2a \tag{11}$$
$$k_1 = f_3 + 3a^2 \tag{12}$$
$$k_0 = f_2 - 2a(f_3 + 2a^2) \tag{13}$$
$$f_1' = 2af_2 - 3a^2f_3 - 5a^4 \tag{14}$$
$$f_0' = a^2f_2 - 2a^3f_3 - 4a^5. \tag{15}$$

For a fixed pair $(f_2, f_3) \in \mathbb{F}_q \times \mathbb{F}_q$, we first count the number of solutions $(a, k_0, k_1, k_2, f_0', f_1')$ to (11)–(15). Note that if $a = 0$ then we must have $k_0 = f_2$, $k_1 = f_3$, $k_2 = f_0' = f_1' = 0$. On the other hand, if $a \neq 0$, we have $k_1 \neq f_3$ and $(k_1 - f_3)/3$ is a quadratic residue in $\mathbb{F}_q$ for exactly $(q-1)/2$ elements $k_1 \in \mathbb{F}_q$. For each such k_1, we obtain two solutions determined by setting a equal to the square roots of $(k_1 - f_3)/3$. That is, there are $1 + 2((q-1)/2) = q$ solutions in total and each solution $(a, k_0, k_1, k_2, f_0', f_1')$ leads to a polynomial $P_{f_1',f_0'}(x)$ which has either one or two paired repeated roots in $\mathbb{F}_q$. We note that different solutions may lead to the same polynomial $P_{f_1',f_0'}(x)$. It is easy to see that this occurs only if $P_{f_1',f_0'}(x)$ has exactly two paired repeated roots in $\mathbb{F}_q$ and in this case we show that, for a fixed $P_{f_1',f_0'}(x)$, there could exist at most two different solutions that yield this polynomial. The proof is as follows. Suppose there are three different solutions $S := (a, k_0, k_1, k_2, f_0', f_1')$, $\tilde{S} := (\tilde{a}, \tilde{k}_0, \tilde{k}_1, \tilde{k}_2, f_0', f_1')$ and $\hat{S} := (\hat{a}, \hat{k}_0, \hat{k}_1, \hat{k}_2, f_0', f_1')$ that yield the polynomial $P_{f_1',f_0'}(x)$. If $a = \tilde{a}$ then $S = \tilde{S}$ by (11), (12) and (13). Therefore, we will assume that $a, \tilde{a}$ and $\hat{a}$ are pairwise different; and in this case one can see that $(x+\tilde{a})^2(x+\hat{a})^2 \mid (x^3 + k_2x^2 + k_1x + k_0)$, which is impossible.

We are now ready to prove the theorem. Suppose that the number of polynomials $P_{f_1',f_0'}(x)$ that have exactly two paired repeated roots both of which are in $\mathbb{F}_q$ is β, and the number of polynomials $P_{f_1',f_0'}(x)$ that have exactly two paired repeated roots both of which are in $\mathbb{F}_{q^2} \setminus \mathbb{F}_q$ is γ. Then, by our above argument, the number of solutions $(a, k_0, k_1, k_2, f_0', f_1')$ to (11)–(15), such that each solution is leading to a polynomial $P_{f_1',f_0'}(x)$ that have exactly two paired repeated roots both of which are in $\mathbb{F}_q$, is at most 2β. Hence, there are at least $q - 2\beta$ polynomials $P_{f_1',f_0'}(x)$ having exactly one paired repeated root in $\mathbb{F}_q$, and there are at least $(q - 2\beta) + \beta + \gamma$ polynomials $P_{f_1',f_0'}(x)$ with at least one repeated root in $\bar{\mathbb{F}}_q$. By Case 1, $\beta + \gamma \leq 3$, and we can see that $q - \beta + \gamma \geq q - 3$.

From Case 1 and Case 2 there are at most $q + 3$ polynomials $P_{f_1',f_0'}(x)$ with at least one repeated root in $\bar{\mathbb{F}}_q$. □

Remark 1. Let $\mathcal{H} : y^2 = x^5 + f_3x^3 + f_2x^2 + f_1x + f_0$ be a genus 2 hyperelliptic curve defined over $\mathbb{F}_q$. According to Theorem 5, there are at least $\chi = q^2 - q - 3$ invalid curves relative to $\mathcal{H}$ and $\mathcal{F}_{2a}$. We now argue that if $f_2 \neq 0$ or $f_3 \neq 0$ then at least $\frac{\chi}{6}$ of these curves are pairwise non-isomorphic. Let $\mathcal{H}'$ and $\mathcal{H}''$ be two genus 2 hyperelliptic curves over $\mathbb{F}_q$ such that

$$\mathcal{H}' : y^2 = x^5 + f_3x^3 + f_2x^2 + f_1'x + f_0'$$
$$\mathcal{H}'' : y^2 = x^5 + f_3x^3 + f_2x^2 + f_1''x + f_0''.$$

The curves $\mathcal{H}'$ and $\mathcal{H}''$ are isomorphic if there is an isomorphism α

$$\alpha : (x, y) \to (\alpha^2 x, \alpha^5 y)$$

that applied to $\mathcal{H}'$ gives the equation of $\mathcal{H}''$. Applying α to $\mathcal{H}'$ gives

$$\alpha\mathcal{H}' : y^2 = x^5 + \alpha^{-4} f_3 x^3 + \alpha^{-6} f_2 x^2 + \alpha^{-8} f_1' x + \alpha^{-10} f_0'.$$

If $f_2 \neq 0$ or $f_3 \neq 0$ then $\mathcal{H}'$ is isomorphic to $\mathcal{H}''$ only if $\alpha^4 = 1$ or $\alpha^6 = 1$. This proves that at least $\frac{\chi}{6}$ of the curves relative to $\mathcal{H}$ and $\mathcal{F}_{2a}$ are pairwise non-isomorphic. Hence, there are at least $(\frac{\chi}{6} - 1)$ isomorphism classes of invalid curves relative to $\mathcal{H}$ and $\mathcal{F}_{2a}$, when $f_2 \neq 0$ or $f_3 \neq 0$. We are justified in omitting the case $f_2 = f_3 = 0$ in our argument since there are at most 10 isomorphism classes of such genus 2 hyperelliptic curves defined over $\mathbb{F}_q$ (see [7]).

5. Invalid-curve attacks

Suppose that a discrete logarithm cryptographic protocol requires party $\hat{B}$ to use his static (long-term) private key b by computing $\sigma = bP$ for some incoming $P \in J_{\mathcal{H}}$, where $\mathcal{H}$ is a genus 1 or genus 2 hyperelliptic curve and P has order n. If $\hat{B}$ does not verify that $P \in J_{\mathcal{H}}$ then an adversary $\mathcal{M}$ could launch an invalid-curve attack to learn b by selecting $P \in J_{\mathcal{I}}$ where $\mathcal{I}$ is an invalid curve or an invalid singular curve with respect to $\mathcal{H}$ (and either $\mathcal{F}_{1a}$ or $\mathcal{F}_{2a}$). Invalid-curve attacks come in two flavors.

In a *small subgroup attack* [18], $\mathcal{M}$ selects $\mathcal{I}$ so that $J_{\mathcal{I}}$ has an element P of small order r. The order r is small enough so that the discrete logarithm problem in $\langle P \rangle$ is feasible via exhaustive search; typically r is a small prime. Small subgroup attacks can be used in situations where $\mathcal{M}$ is able to obtain a quantity that is derived from bP, for example $k = H(bP)$ where H is a cryptographic hash function. In this case, $\mathcal{M}$ would compute $k' = H(iP)$ for all $i \in [0, r-1]$ until $k' = k$, after which $\mathcal{M}$ concludes that $b \equiv i \pmod{r}$. By repeating the procedure for different curves $\mathcal{I}$ (and different primes r) the value b can be recovered via the Chinese Remainder Theorem.

In a *large subgroup attack* [20, p.58], $\mathcal{M}$ selects an invalid curve $\mathcal{I}$ so that $J_{\mathcal{I}}$ has an element P of large order $t \approx n$ and so that the discrete logarithm problem in $\langle P \rangle$ can be efficiently solved; this is the case if t is smooth, or if there exists an efficiently computable mapping from $\langle P \rangle$ to another group where efficient DLP algorithms are known. Large subgroup attacks can be used in situations where $\mathcal{M}$ is able to obtain the group element bP itself. In this case, $\mathcal{M}$ would compute $b \bmod t$, and thereafter efficiently determine b.

There are two main reasons that invalid-curve attacks can be used by $\mathcal{M}$ in the above setting. First, the representation of elements in the valid group $J_{\mathcal{H}}$ is the same as the representation of elements in the invalid group $J_{\mathcal{I}}$. Second, the implementation of the group operation in the valid group $J_{\mathcal{H}}$ is applicable to the invalid group $J_{\mathcal{I}}$ without any modification. We illustrate our invalid-curve attacks on two recently-proposed discrete logarithm protocols — the Twin Diffie-Hellman key agreement scheme and the XCR signature scheme — that are successful if public-key validation is not performed. We emphasize that our attacks do not illustrate any weaknesses in these protocols, but rather serve to emphasize the importance of public-key validation. More precisely, the attacks we describe do not require the adversary to tamper with hardware or modify registers; the attacks

exploit only omission of public key validation. Since validation in the case of curve-based cryptography does not require full exponentiation but only a constant number of field multiplications, it introduces a negligible efficiency overhead. Therefore validation should always be performed in conjunction with curve-based protocols.

5.1. TWIN DIFFIE-HELLMAN (TDH). Cash, Kiltz and Shoup [5, Section 4] proposed and analyzed a simple Diffie-Hellman type protocol depicted in Figure 1. The security of TDH relies on the twin Diffie-Hellman assumption which is equivalent to the computational Diffie-Hellman assumption. This is in contrast with many other key agreement protocols (e.g. [15, 24]), where security has only been proven with respect to the gap Diffie-Hellman assumption – this is the assumption that the computational Diffie-Hellman problem is hard even if the solver is given an oracle for the decisional Diffie-Hellman problem.

We extend the small subgroup attacks on the static Diffie-Hellman protocol to the TDH protocol in the genus 2 setting. We show how, even in the restricted security model used in [5], an adversary can successfully break the protocol should honest parties fail to obtain assurances that the static public keys of their peers were validated. This demonstrates the importance of requiring that all elements belong to the correct group.

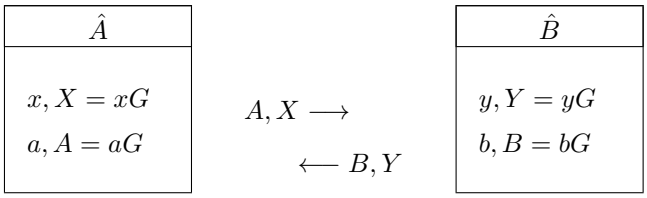

$$\kappa = H\left(\hat{A}, \hat{B}, \mathrm{CDH}(A, B), \mathrm{CDH}(A, Y), \mathrm{CDH}(X, B), \mathrm{CDH}(X, Y)\right)$$

FIGURE 1. The twin Diffie-Hellman protocol. The underlying group is $\langle G \rangle$. Party $\hat{A}$'s static key pairs are (A, a) and (X, x), while party $\hat{B}$'s static key pairs are (B, b) and (Y, y).

Informally, the security model in [5] allows the adversary $\mathcal{M}$ to observe the interaction between honest parties, obtain session keys computed by honest parties, register corrupt parties with $\mathcal{M}$'s choice of static public key, and interact with honest parties on behalf of corrupted parties. The model did not (explicitly) require any checks on the static public keys chosen by the adversary when registering corrupt parties. The implication is that the adversary can register invalid static keys.

We now describe an attack that allows $\mathcal{M}$ to recover the static private key of an honest party. Suppose that the underlying group is a prime-order subgroup of $J_{\mathcal{H}}$, where $\mathcal{H}$ is the hyperelliptic curve defined by the polynomial $y^2 = x^5 + f_3x^3 + f_2x^2 + f_1x + f_0$ over $\mathbb{F}_q$. Suppose also that honest parties use group addition formula $\mathcal{F}_{2a}$ (recall that $\mathcal{F}_{2a}$ does not explicitly use the coefficients f_1 and f_0). In this case $\mathcal{M}$ chooses invalid curves $\mathcal{IH}_1$ and $\mathcal{IH}_2$ relative to $\mathcal{H}$ and $\mathcal{F}_{2a}$ such that the invalid curves have points of small orders u and v, respectively, with $\gcd(u, v) = 1$.

Using the notation in Figure 1, assume that $\hat{A}$ is corrupt and $\hat{B}$ is honest. Suppose that when $\mathcal{M}$ registers $\hat{A}$, $\mathcal{M}$ picked $X \in \mathcal{IH}_1$ of order u and $A \in \mathcal{IH}_2$ of order v. Then $\mathcal{M}$ initiates a session with $\hat{B}$ and obtains the session key $\kappa =$

$H(\hat{A}, \hat{B}, bA, yA, bX, yX)$ that $\hat{B}$ computes. Now $\mathcal{M}$ computes

$$\kappa' = H(\hat{A}, \hat{B}, i_1 A, i_2 A, i_3 X, i_4 X),$$

where i_1, i_2 range over $\mathbb{Z}_v$ and i_3, i_4 range over $\mathbb{Z}_u$, until $\kappa' = \kappa$ in which case $\mathcal{M}$ learns that $b \equiv i_1 \pmod{v}$, $b \equiv i_3 \pmod{u}$, $y \equiv i_2 \pmod{v}$ and $y \equiv i_4 \pmod{u}$. For concreteness, suppose that $u, v \approx 2^{10}$. Note that $\mathcal{M}$ has to perform roughly 2^{40} steps before $\kappa' = \kappa$, which is a feasible amount of computation. After repeating the procedure for other orders (u, v), $\mathcal{M}$ can recover $\hat{B}$'s static private key (b, y).

Recall from Remark 1 that there are at least $(q^2 - q - 9)/6$ isomorphism classes of invalid curves relative to $\mathcal{H}$ and $\mathcal{F}_{2a}$, where $f_2 \neq 0$ or $f_3 \neq 0$ in the equation of $\mathcal{H}$. We can reasonably assume that the distribution of orders of these invalid curves follows the distribution of orders of all genus 2 curves over $\mathbb{F}_q$. Under this assumption, the invalid curves have group orders that are almost uniformly distributed over the Hasse interval $[(\sqrt{q} - 1)^4, (\sqrt{q} + 1)^4]$ (see [16, Proposition 1.9] and [17, Theorems 1.1 and 11.5] for results on the distribution of genus 1 and 2 curves over the Hasse interval). Thus, by randomly selecting invalid curves $\mathcal{H}'$, the adversary will quickly find one with group order divisible by a small prime. The bottleneck in the search is the time for computing the cardinality of $J_{\mathcal{H}'}(\mathbb{F}_q)$. With current technology, computing the cardinality of $J_{\mathcal{H}'}(\mathbb{F}_q)$ is feasible for 80-bit fields $\mathbb{F}_q$ [9], but not quite feasible for 128-bit fields $\mathbb{F}_q$.

5.2. SINGULAR ELLIPTIC CURVES. Krawczyk [14, §4.1 Definition 2] proposed the exponential challenge-response (XCR) signature scheme for a cryptographically strong group $\mathcal{G} = \langle G \rangle$ of prime-order n. In the XCR signature scheme the signer $\hat{B}$ with static key pair (B, b) signs a message for a verifier $\hat{A}$ who submits a challenge $X = xG$. The signature on a message m with a challenge X is (σ, Y), where $Y = yG$ is a random group element of $\hat{B}$'s choice and $\sigma = (y + H(Y, m)b)X$. $\hat{A}$ can verify the signature (σ, Y) using her knowledge of x via $\sigma = x(Y + H(Y, m)B)$. In the description of XCR in [14] there is an additional check which is not applicable for elliptic curves and hence is omitted here.

As before the goal of the adversary $\mathcal{M}$ is to learn the static private key b of the signer $\hat{B}$. The large subgroup attack assumes that the adversary can learn the ephemeral private key y that $\hat{B}$ chooses to sign a message m. The attack proceeds as follows. $\mathcal{M}$ selects an arbitrary message m and an invalid curve $\mathcal{H}'$ having an element X of large smooth order $t \approx n$. The signer $\hat{B}$ signs m and returns (σ, Y) to $\mathcal{M}$. The adversary learns the ephemeral private key y and can thereafter deduce $b \bmod t$ by using the Pohlig-Hellman algorithm to compute the logarithm of $H(Y, m)^{-1}(\sigma - yX) = bX$ to the base X; b can then be efficiently determined. Alternatively, $\mathcal{M}$ can mount a small subgroup attack. In that case, assuming that $\mathcal{M}$'s computational power is $O(2^{50})$, $\mathcal{M}$ will need to select invalid curves $\mathcal{H}'$ having elements X of 100-bit order t. To recover the secret b, $\mathcal{M}$ will need roughly $\lfloor 521/100 \rfloor = 5$ interactions with the signer.

For concreteness, suppose now that the signer uses an elliptic curve of the form $y^2 = x^3 - 3x + f_0$ over the prime field $\mathbb{F}_p$ where $p = 2^{521} - 1$. Such an elliptic curve has been specified by NIST in the FIPS 186-2 Standard [8]. Consider

$$\mathcal{E}_- : y^2 = x^3 - 3x + 2 = (x + 2)(x - 1)^2$$

and

$$\mathcal{E}_+ : y^2 = x^3 - 3x - 2 = (x - 2)(x + 1)^2$$

over $\mathbb{F}_p$, where $\gcd(p, 6) = 1$. According to Theorem 4*(ii)*, $\mathcal{E}_+$ and $\mathcal{E}_-$ are invalid singular curves relative to any curve $\mathcal{E}$ defined via $y^2 = x^3 - 3x + f_0$ over $\mathbb{F}_p$. In particular, $\mathcal{E}_+$ and $\mathcal{E}_-$ are invalid singular curves relative to all the five NIST curves defined over prime fields [8]. The sets of non-singular points on $\mathcal{E}_+$ and $\mathcal{E}_-$ over $\mathbb{F}_p$ (together with the point at infinity) form groups isomorphic to $\mathbb{F}_p^*$ that share the same group law with $\mathcal{E}$. Moreover, the isomorphisms are efficiently computable (see the discussion after Definition 2). Using $\mathcal{E}_+$, $\mathcal{M}$ can mount the following large subgroup attack. $\mathcal{M}$ picks a point X of order $p - 1$ on $\mathcal{E}_+$. As before $\mathcal{M}$ interacts with the signer and obtains a signature (σ, Y) on a message m. Suppose $\mathcal{M}$ is able to learn $\hat{B}$'s ephemeral private key y. Then $\mathcal{M}$ can map $H(Y, m)^{-1}(\sigma - yX)$ to $\mu \in \mathbb{F}_p^*$ and X to $g \in \mathbb{F}_p^*$, and then use subexponential discrete logarithm algorithms [10, 13] to compute $\log_g \mu$ thus obtaining $b \bmod (p - 1)$. Note that the discrete logarithm algorithm of [13] has a first phase that can be precomputed before the attack is launched. Subsequently, the second phase quickly computes individual logarithms and can reuse the precomputations in multiple applications of the attack.

We note that in the case of the NIST prime $p = 2^{521} - 1$, discrete logarithms in $\mathbb{F}_p^*$ can be more efficiently computed using the Pohlig-Hellman algorithm since the largest prime factor of the order of $\mathbb{F}_p^*$ is only 88 bits in length. We also note that an easier way to launch a large subgroup attack is to use the supersingular curve $E : y^2 = x^3 - 3x$ over $\mathbb{F}_p$ of order $p+1 = 2^{521}$. For this curve, discrete logarithms can be computed extremely efficiently using the Pohlig-Hellman algorithm. However, this speedup is not applicable for all primes. For example, the supersingular curve $E : \ y^2 = x^3 - 3x$ over $\mathbb{F}_p$ where $p = 2^{384} - 2^{128} - 2^{96} + 2^{32} - 1$ is the 384-bit NIST prime has a subgroup of 188-bit prime order

$$r = 213458640757090462592068633975230757544124954331352889061,$$

and hence the large subgroup attack cannot be immediately applied. However, the embedding degree of E is two, and the discrete logarithm problem can be efficiently mapped via pairings to $\mathbb{F}_{p^2}^*$.

6. Concluding remarks

We have demonstrated that invalid-curve attacks can be extended to hyperelliptic curve cryptosystems that use genus 2 curves defined over prime fields and the addition formula $\mathcal{F}_{2a}$. The attacks can be extended in various ways. First, the attacks are applicable when genus 2 curves defined over binary fields are used together with the addition formulas given in [2, §14.3.2]. Furthermore, the attacks can be extended to genus 3 hyperelliptic curves using the addition formulas in [2, §14.6.1, §14.6.2]. More interestingly, it is also possible to use invalid singular hyperelliptic curves to mount attacks analogously as was done in §5.2 with singular elliptic curves. For example, suppose $\mathcal{H} : \ y^2 = x^5 + f_3x^3 + f_2x^2 + f_1x + f_0$ is a non-singular genus 2 hyperelliptic curve defined over $\mathbb{F}_q$. By the proof of Theorem 5, there exist at least $q - 6$ invalid singular curves $\mathcal{IH}$ relative to $\mathcal{H}$ and $\mathcal{F}_{2a}$ such that $\mathcal{IH} : \ y^2 = (x + a)^2(x^3 + k_2x^2 + k_1x + k_0)$ and $x^3 + k_2x^2 + k_1x + k_0$ has no repeated roots; let $\mathcal{E} : y^2 = x^3 + k_2x^2 + k_1x + k_0$ be the corresponding elliptic curve. Let $J_{ns}(\mathbb{F}_q)$ denote the group of degree zero divisor classes of $\mathcal{IH}$ over $\mathbb{F}_q$ such that the support of the divisors does not contain the singular point $(-a, 0)$. An explicit isomorphism from $J_{ns}(\mathbb{F}_q)$ to the group $\mathcal{E}(\mathbb{F}_q)$, which is induced by $\rho : \mathcal{IH} \setminus \{(-a, 0)\} \to \mathcal{E}$ such that $\rho(x, y) = (x, y/(x + a))$ and $\rho(\infty) = \infty$, can be used to map the discrete logarithm problem in $J_{ns}(\mathbb{F}_q)$ to $\mathcal{E}(\mathbb{F}_q)$. Hence, if the

validation check is omitted in a cryptographic protocol employing $\mathcal{H}$ that requires a party $\hat{B}$ to compute $b\bar{D}$ for some incoming $\bar{D} \in J_{\mathcal{H}}$, then an adversary $\mathcal{M}$ can recover the static private key b of $\hat{B}$ in time $\mathcal{O}(\sqrt{q})$ rather than $\mathcal{O}(q)$.

Altogether, these attacks emphasize the importance of validating public keys in discrete logarithm cryptosystems.

Acknowledgements

The authors thank Alfred Menezes for his careful and critical reading of the draft, and also for his suggestions and corrections. The authors also thank the two anonymous referees for their detailed and very helpful comments, and for bringing reference [23] to our attention.

References

[1] A. Antipa, D. Brown, A. Menezes, R. Struik and S. Vanstone, *Validation of elliptic curve public keys*, in "Public Key Cryptography – PKC 2003," (2003), 211–223.
[2] R. Avanzi, H. Cohen, C. Docke, G. Frey, T. Lange, K. Nguyen and F. Vercauteren, "Handbook of Elliptic and Hyperelliptic Curve Cryptography," Chapman & Hall/CRC, Boca Raton, FL, USA, (2005).
[3] I. Biehl, B. Meyer and V. Müller, *Differential fault analysis on elliptic curve cryptosystems*, in "Advances in Cryptology – CRYPTO 2000," (2000), 131–146.
[4] D. G. Cantor, *Computing the jacobian of a hyperelliptitc curve*, Math. Comput., **48** (1987), 95–101.
[5] D. Cash, E. Kiltz and V. Shoup, *The twin Diffie-Hellman problem and applications*, in "Advances in Cryptology – EUROCRYPT 2008," (2008), 127–145; available online at http://eprint.iacr.org/2008/067.
[6] L. Q. Chen, Z. H. Cheng and N. P. Smart, *Identity-based key agreement protocols from pairings*, Internat. J. Inform. Security, **6** (2007), 213–241.
[7] L. H. Encinas, A. Menezes and J. M. Masqué, *Isomorphism classes of genus 2 hyperelliptic curves over finite fields*, Appl. Algebra Engin. Commun. Comput., **13** (2002), 57–65.
[8] FIPS 186-2, *Digital signature standard (DSS)*, Federal Information Processing Standards Publication 186-2, National Institute of Standards and Technology, (2000).
[9] P. Gaudry and É. Schost, *Construction of secure random curves of genus 2 over prime fields*, in "Advances in Cryptology – EUROCRYPT 2004," (2004), 239–256.
[10] D. Gordon, *Discrete logarithms in GF(p) using the number field sieve*, SIAM J. Disc. Math., **6** (1993), 124–138.
[11] R. Hartshorne, "Algebraic Geometry," Springer-Verlag, New York, 1977.
[12] D. Husemöller, "Elliptic Curves," Springer-Verlag, New York, 1987.
[13] A. Joux and R. Lercier, *Improvements to the general number field sieve for discrete logarithms in prime fields – A comparison with the Gaussian integer method*, Math. Comput., **72** (2003), 953–967.
[14] H. Krawczyk, *HMQV: A high-performance secure Diffie-Hellman protocol*, in "Advances in Cryptology – CRYPTO 2005," (2005), 546–566; available online at http://eprint.iacr.org/2005/176.
[15] B. LaMacchia, K. Lauter and A. Mityagin, *Stronger security of authenticated key exchange*, in "Provable Security: First International Conference – ProvSec 2007," (2007), 1–16.
[16] H. Lenstra, *Factoring integers with elliptic curves*, Annals Math., **126** (1987), 649–673.
[17] H. Lenstra, J. Pila and C. Pomerance, *A hyperelliptic smoothness test, II*, Proc. London Math. Soc., **84** (2002), 105–146.
[18] C. H. Lim and P. J. Lee, *A key recovery attack on discrete log-based schemes using a prime order subgroup*, in "Advances in Cryptology – CRYPTO'97," (1997), 249–263.
[19] P. Lockhart, *On the discriminant of a hyperelliptic curve*, Trans. Amer. Math. Soc., **342** (1994), 729–752.
[20] A. Menezes, *Another look at HMQV*, J. Math. Crypt., **1** (2007), 47–64.

[21] A. Menezes and B. Ustaoglu, *On the importance of public-key validation in the MQV and HMQV key agreement protocols*, in "Progress in Cryptology – INDOCRYPT 2006," (2006), 133–147.
[22] A. Menezes and B. Ustaoglu, *On reusing ephemeral keys in Diffie-Hellman key agreement protocols*, Intern. J. Appl. Crypt., to appear.
[23] E. Nart, *Counting hyperelliptic curves*, Adv. Math., **221** (2009), 774–787.
[24] B. Ustaoglu, *Obtaining a secure and efficient key agreement protocol from (H)MQV and NAXOS*, Des. Codes Crypt., **46** (2008), 329–342.
[25] A. Weil, *L'arithmétique sur les courbes algébriques* (in French), Acta Math., **52** (1929), 281–315.

Received May 2009; revised January 2010.

E-mail address: kkarabin@uwaterloo.ca
E-mail address: bustaoglu@cryptolounge.net