

Lossy CSI-FiSh: Efficient Signature Scheme with Tight Reduction to Decisional CSIDH-512

Ali El Kaafarani^{1,2}, Shuichi Katsumata³, Federico Pintore¹

¹Mathematical Institute, University of Oxford, UK

federico.pintore@maths.ox.ac.uk

²PQShield, UK

elkaafarani@pqshield.com

³National Institute of Advanced Industrial Science and Technology (AIST), JP

shuichi.katsumata@aist.go.jp

March 10, 2020

Abstract

Recently, Beullens, Kleinjung, and Vercauteren (Asiacrypt’19) provided the first practical isogeny-based digital signature, obtained from the Fiat-Shamir (FS) paradigm. They worked with the CSIDH-512 parameters and passed through a new record class group computation. However, as with all standard FS signatures, the security proof is highly non-tight and the concrete parameters are set under the heuristic that the only way to attack the scheme is by finding collisions for a hash function.

In this paper, we propose an FS-style signature scheme, called Lossy CSI-FiSh, constructed using the CSIDH-512 parameters and with a security proof based on the “Lossy Keys” technique introduced by Kiltz, Lyubashevsky and Schaffner (Eurocrypt’18). Lossy CSI-FiSh is *provably secure* under the same assumption which underlies the security of the key exchange protocol CSIDH (Castruck et al. (Asiacrypt’18)) and is *almost as efficient* as CSI-FiSh. For instance, aiming for small signature size, our scheme is expected to take around ≈ 800 ms to sign/verify while producing signatures of size ≈ 280 bytes. This is only twice slower than CSI-FiSh while having similar signature size for the same parameter set. As an additional benefit, our scheme is by construction secure *both* in the classical and quantum random oracle model.

1 Introduction

1.1 Background

Isogeny-based cryptography is one of the promising candidates for post-quantum cryptography. While isogeny problems offer simple and efficient solutions to encryption schemes (or equivalently, key-exchange protocols) [JD11, CLM⁺18], they turned out to be rather elusive to use for constructing signature schemes.

At the highest level, all isogeny-based signatures we know thus far are based on the Fiat-Shamir paradigm [FS87, AABN02]: prepare a hard relation \mathcal{R} based on an isogeny problem, construct an identification protocol (or sigma protocol) for \mathcal{R} , and use a cryptographic hash function to compile the identification protocol into a signature scheme in the random oracle model (ROM). Both the two central isogeny problems — the computational supersingular isogeny (CSSI) problem [DFJP14] and the group action inverse problem (GAIP) [CLM⁺18] — have been the basis for constructing signatures. Those based on CSSI, proposed in [YAJ⁺17, GPS17], produce signatures of size at least 12KB even in the most optimized variant [GPS17]. On the other hand, relying on GAIP and employing the Fiat-Shamir with aborts strategy [Lyu09], De Feo

and Galbraith introduced a compact isogeny-based signature named SeaSign [DG19]. Despite the inefficiency in the signature generation and verification, SeaSign provides signatures of a remarkably small size (less than 1 kilobyte at the 128-bit security level).

Very recently, a new record class group computation has allowed Beullens, Kleinjung and Vercauteren [BKV19] to improve SeaSign and obtain the first practical isogeny-based signature scheme, named CSI-FiSh. Their computation has shed light on the structure of the ideal class group determined by a specific set of CSIDH parameters, named CSIDH-512 [CLM⁺18]. This granted a proper uniform sampling from the ideal class group, and canonical representation of its elements, which enabled to overcome the costly remedy made by SeaSign. That is, the adoption of a redundant representation of class group elements and performing rejection sampling. The result is practical efficiency in both signature generation and verification while maintaining the short signature size offered by SeaSign. However, one important remark is that, since CSI-FiSh is specific to the special set of parameters CSIDH-512, it can offer *at most* the same security level provided by a hard problem defined over the CSIDH-512 parameters. Specifically, CSI-FiSh relies on the GAIP problem, which is believed to have 128-bits of classical and (at most) 64-bits of quantum security over the CSIDH-512 parameters [CLM⁺18, Pei19].

Tight Security. Fiat-Shamir (FS) signatures [FS87, AABN02] admit an intuitive and simple construction in the ROM, however, they are notorious for having a very loose reduction. Since a loose reduction forces for a stronger hardness assumption, and consequently a less efficient scheme, it has been the focus of several works to tighten the reduction loss, e.g., [BR96, MR02, KW03, PV05, GBL08, Seu12, FJS14].

To give a more precise perception of the security loss, assume we had a FS signature that is secure based on the hardness of a particular hard problem Π . Then, the security proof of FS signatures in the classical ROM dictates that the reduction algorithm can break the underlying problem Π with advantage $Q^{-1} \cdot \epsilon^2$, where Q is the number of hash evaluations an adversary can perform and ϵ is the advantage of an adversary breaking the security of the FS signature. Therefore, if we want to instantiate the FS signature with *provably secure* parameters, we must assume the hardness of the problem Π for a security level that is much higher than expected. For instance, if we aim for 128-bits of security for the FS signature (i.e., $\epsilon = 2^{-128}$), then assuming a modest $Q \approx 2^{40}$, we require at least 296-bits of security for the hard problem Π . Since a hard problem with a higher level of security must necessitate larger parameters, this leads to inefficient schemes.

This undesirable loss in security and efficiency is common to all standard FS signatures and CSI-FiSh is no exception. However, one large difference between CSI-FiSh and other FS signatures is that CSI-FiSh relies on a hard problem defined for a *specific* security level — the GAIP problem over the CSIDH-512 parameters. For the time being, no other parameter sets are known to provide the nice algebraic structure required for CSI-FiSh. This is in sharp contrast with FS signatures based on other hardness assumptions since most hardness assumptions can “absorb” the reduction loss by setting the parameters larger. Since GAIP over the CSIDH-512 parameters only offers 128-bits of classical security, we cannot argue any notion of *provable* security for CSI-FiSh if we aim for 128-bits of security. Concretely, if we plug in $Q \approx 2^{40}$ as above, we can only provably argue 44-bits of security for CSI-FiSh. Moreover, if we aim for quantum security, the situation is worse since the reduction algorithm can break the underlying problem Π with only advantage $Q^{-6} \cdot \epsilon^3$ [DFMS19, LZ19]. We note that the currently available resources would probably allow other record computations for bigger parameters for which GAIP is believed to have a much higher security level; however, the benefit of having a higher security level would likely be beaten by the significant slow-down in efficiency.

In practice, this inconvenient reduction loss in FS signatures is usually overlooked or simply ignored, and the parameters are set assuming that the best attack against the FS signature is (roughly) finding a collision in the hash function. In [BKV19], the parameters for CSI-FiSh are set under this simplified assumption as well. Considering this undesirable gap between practice and theory, a natural question which arises is:

Can we design an isogeny-based signature scheme as efficient as CSI-FiSh with provable secure parameters?

1.2 Our Contribution

In this work, we provide a partial answer to the above problem and propose a new signature scheme, Lossy CSI-FiSh, with the following features:

- It is *tightly secure* under a natural hardness assumption over the CSIDH-512 parameters, that is, the *decisional CSIDH* (D-CSIDH) assumption.

We note D-CSIDH is not a new assumption introduced in this paper, as it was originally defined by Stolbunov in his PhD thesis [Sto12, Problem 2.2] and implicitly underlies the security of the key exchange protocol CSIDH [CLM⁺18].¹

- It is *almost as efficient* as CSI-FiSh. Compared to CSI-FiSh, the signature size is the same, the public key is only twice as large, and the runtime of the signature generation and verification is estimated to be (at most) twice as slow. For instance, aiming for small signature size, our scheme is expected to take around ≈ 800 ms to sign/verify while producing signatures of size ≈ 280 bytes. This is still 150 times faster and around 3 times smaller than an optimized version of SeaSign for the same parameter set.
- It is secure *both* in the classical and quantum ROM (QROM). In particular, we do not require a separate construction using the Unruh transform [Unr15] to achieve security in the QROM.

We obtain our results by following the line of work that constructs *lossy* identification protocols to obtain tightly secure FS signatures [KW03, AFLT12, Unr17, KLS18]. A lossy identification protocol comes with an additional *lossy statement* generator that produces lossy statements which are computationally indistinguishable from honestly generated statements for the hard relation \mathcal{R} induced by some hardness assumption. Moreover, relative to the lossy statements, the protocol admits statistical soundness. That is, not even a computationally unbounded adversary can successfully impersonate a prover. Using the result of Kiltz, Lyubashevsky, and Schaffner [KLS18] (see Theorem 2.5), a lossy identification protocol directly provides us an FS signature with a *tight reduction* in the *classical and quantum ROM*.

The idea to use a lossy identification protocol to achieve tight security for isogeny-based FS signatures was also considered by De Feo and Galbraith for SeaSign [DG19, Section 8]. In particular, they proposed to take a very large ideal class group (determined by a big prime p) and then only a small subset as the space of possible private keys (that results in valid public keys being chosen from a set of roughly the same cardinality). The signature generation and verification processes are not altered from the *standard* SeaSign scheme. The result is that the lossy variant inherits the same inefficiency of the main scheme, with the increment of the prime p further aggravating the issue. It is evident that the above approach does not extend to the current version of CSI-FiSh, which requires the specific CSIDH-512 parameter set.

The lossy identification protocol proposed in this work — which arises from the observation that the D-CSIDH relation over the CSIDH-512 parameters naturally admits a lossy mode — appears to be much simpler and it smoothly leads to a practical signature scheme. Our identification protocol enjoys the same optimizations used in [DG19] and [BKV19]. Using D-CSIDH instead of GAIP as the underlying assumption, we encounter an obstacle that stems from the fact that D-CSIDH does not provide natural random self-reducibility properties. However, we discuss that this issue does not have much of a big impact on the concrete choice of parameters.

Related Works. There are only a handful of efficient signature schemes that are tightly and provably secure in the (Q)ROM that we are aware of. The lattice-based Gentry-Peikert-Vaikuntanathan (GPV) signature [GPV08] or its much-optimized successor FALCON [FHK⁺18] have tight security in the (Q)ROM. One notable feature is that the construction natively supports tight security in both classical and quantum ROM *without* incurring any overhead. Dilithium [DKL⁺18], which is a lattice-based FS-type signature, also has tight security in the (Q)ROM [KLS18]. To achieve tight security, they must modify the public key of

¹Roughly, this is parallel to the relation between the Diffie-Hellman (DH) protocol and the decisional DH assumption [DH76]. For a more formal discussion, we refer to Section 3.1.

their non-tightly secure scheme to obtain a lossy mode. Unfortunately, when using a lattice-based hard problem (that is, the learning with errors problem), this comes at the cost of making the public key size at least 5 times larger and the signature size at least 2 times larger, e.g., public key and signature size grows from (1472, 2701) bytes up to (7712, 5690) bytes. As we mentioned above, SeaSign [DG19] goes through the lossy argument as well. They require to use of a non-standard variant of the GAIP problem and makes it difficult to assess the increase in signature and public key sizes. We like to highlight that although we go through the same paradigm of *lossy* arguments, Lossy CSI-FiSh is based on a standard assumption and does not incur a large blow up in size; the public key is only 2 times larger and the signature size remains the same compared to the non-tight variant CSI-FiSh. Finally, the hash-based signature SPHINCS+ [BHK+19] also enjoys tight security in the (Q)ROM under several heuristic assumptions on the underlying cryptographic hash function.

Finally, recently in a different context of distributed signature, Cozzo and Smart [CS20] independently proposed a CSIDH-based identification protocol similar to that we introduce in Section 3.

Roadmap. The rest of the paper is organized as follows. In Section 2 we give a brief preliminary on identification protocols and class group actions. In Section 3 and 4 we introduce the new lossy identification protocol and we adapt it using the optimizations proposed in [DG19, BKV19] to enlarge the challenge space. In Section 5 we describe the signature scheme obtained through the Fiat-Shamir transform, and we compare it to CSI-FiSh in terms of bandwidth and computational complexity. In Section 6 we report concluding remarks.

2 Preliminaries

2.1 Identification Protocols

Given two sets X and Y , a subset $\mathcal{R} \subset X \times Y$ is a polynomially computable binary relation on $X \times Y$ if, given $(X, W) \in X \times Y$, we can check $(X, W) \in \mathcal{R}$ in time $\text{poly}(|X|)$. The language $\mathcal{L}_{\mathcal{R}}$ corresponding to \mathcal{R} is the set $\{X \in X \mid \exists W \in Y : (X, W) \in \mathcal{R}\}$, where we call W a witness for the statement $X \in \mathcal{L}_{\mathcal{R}}$.

An identification protocol ID for a relation \mathcal{R} is a three-move interactive protocol between a prover and a verifier. Informally, a prover holding a statement-witness pair $(X, W) \in \mathcal{R}$ can prove to the verifier that they indeed possess a valid witness W without revealing any more than the mere fact that they know W .

Definition 2.1 (Identification Protocol). *An identification protocol ID for a relation \mathcal{R} consists of four PPT algorithms $(\text{IGen}, \text{P} = (\text{P}_1, \text{P}_2), \text{V})$, where V is deterministic and we assume P_1 and P_2 share states. Let ComSet , ChSet , and ResSet be the commitment space, challenge space, and response space, respectively. Then, an identification protocol is defined in the following way.*

- The key generation algorithm IGen takes the security parameter 1^λ as input, and outputs a statement-witness pair $(X, W) \in \mathcal{R}$.
- The prover, on input (X, W) , first executes $\text{com} \leftarrow \text{P}_1(X, W)$, and then sends the commitment com to the verifier.
- The verifier chooses a random challenge $\text{ch} \leftarrow \text{ChSet}$ and sends ch to the prover.
- The prover, given ch , runs $\text{resp} \leftarrow \text{P}_2(X, W, \text{com}, \text{ch})$ and returns a response resp to the verifier. Finally, the verifier runs $\text{V}(X, \text{com}, \text{ch}, \text{resp})$ and outputs 1 if they accept, 0 otherwise.

The protocol transcript $(\text{com}, \text{ch}, \text{resp}) \in \text{ComSet} \times \text{ChSet} \times \text{ResSet}$ is said to be valid in case $\text{V}(X, \text{com}, \text{ch}, \text{resp})$ outputs 1.

We require the following properties from an identification protocol ID. Some of them may seem non-standard, however, they are all necessary to argue security of the Fiat-Shamir transform in the (quantum) random

oracle model. We note that some of the properties are simplified and stronger than those in [KLS18], e.g. we ignore negligible correctness errors. This is done without loss of generality, since our proposed identification protocol satisfies all the stronger properties.

Correctness. The following holds for all $(X, W) \in \mathcal{R}$:

$$\Pr \left[V(X, \text{com}, \text{ch}, \text{resp}) = 1 \mid \begin{array}{l} \text{com} \leftarrow P_1(X, W), \\ \text{ch} \leftarrow \text{ChSet}, \\ \text{resp} \leftarrow P_2(X, W, \text{com}, \text{ch}) \end{array} \right] = 1.$$

(Perfect) Honest-Verifier Zero-Knowledge (HVZK). There exists a PPT simulator algorithm Sim that takes as inputs a statement $X \in \mathcal{L}_{\mathcal{R}}$ and a challenge $\text{ch} \in \text{ChSet}$, and outputs a commitment com and a response resp such that $(\text{com}, \text{ch}, \text{resp})$ is a valid transcript for X . Moreover, the output distribution of Sim on input (X, ch) is equal to the distribution of those outputs generated via an honest execution conditioned on the verifier using ch as the challenge. We note we can consider relaxed variants of HVZK where the distributions are only required to be computationally indistinguishable.

Min-Entropy. The identification protocol ID has α bits of min-entropy if

$$\Pr_{(X, W) \leftarrow \text{IGen}(1^\lambda)} \left[\text{min-entropy}(\text{com} \mid \text{com} \leftarrow P_1(X, W)) \geq \alpha \right] \geq 1 - 2^{-\alpha}.$$

(Optional) Perfect Unique Response. With overwhelming probability over the random choice of $(X, W) \leftarrow \text{IGen}(1^\lambda)$, for any $\text{com} \in \text{ComSet}$ and $\text{ch} \in \text{ChSet}$, there exists a unique response $\text{resp} \in \text{ResSet}$ that leads to a valid transcript $(\text{com}, \text{ch}, \text{resp})$. This property is required when aiming for *strong* unforgeability (i.e., *su-cma*) of the FS signature scheme. As we will see, our identification protocol supports this property by default.

(Optional) Commitment Revocability. With overwhelming probability over the random choice of $(X, W) \leftarrow \text{IGen}(1^\lambda)$, for any $\text{ch} \in \text{ChSet}$ and $\text{resp} \in \text{ResSet}$, there exists a unique commitment $\text{com} \in \text{ComSet}$ that makes $(\text{com}, \text{ch}, \text{resp})$ a valid transcript. Such a commitment can be publicly computed by means of an algorithm taking $(X, \text{ch}, \text{resp})$ as input. This property is unnecessary from a security stand point and only allows for shorter signatures. Again, our identification protocol supports this property by default.

To achieve a *tight* security proof for Fiat-Shamir signatures (formally defined later), we further require the identification protocol to satisfy some notion of *lossiness* defined below.

Definition 2.2 (Lossy Identification Protocol). *An identification protocol ID is called lossy - and denoted by ID_{ls} - if it admits an extra PPT algorithm LossyGen , named lossy key generation algorithm, that on input 1^λ outputs $X_{\text{ls}} \in X \setminus \mathcal{L}_{\mathcal{R}}$.*

We require a lossy identification protocol ID_{ls} to satisfy the following two properties.

Indistinguishability of Lossy Statements. We ask that a statement generated with the lossy key generation algorithm is indistinguishable from a statement generated by the real key generation algorithm. Let us define the following advantage for an adversary \mathcal{A} :

$$\text{Adv}_{\mathcal{A}}^{\text{lossy}}(\lambda) := \left| \Pr[\mathcal{A}(X_{\text{ls}}) = 1 \mid X_{\text{ls}} \leftarrow \text{LossyGen}(1^\lambda)] - \Pr[\mathcal{A}(X) = 1 \mid (X, W) \leftarrow \text{IGen}(1^\lambda)] \right|$$

We say the lossy identification protocol satisfies indistinguishability of lossy statements if for any PPT (or quantum PT) adversary we have $\text{Adv}_{\mathcal{A}}^{\text{lossy}}(\lambda) = \text{negl}(\lambda)$.

Statistical Lossy Soundness. The definition of statistical lossy soundness relies on the following game, named *lossy impersonation game*, played by an adversary \mathcal{A} and a challenger.

Setup: The challenger runs $X_{\text{ls}} \leftarrow \text{LossyGen}(1^\lambda)$ and provides the adversary \mathcal{A} the lossy statement X_{ls} .

Commitment and challenge selection: On input X_{IS} the adversary \mathcal{A} selects a commitment $\text{com} \in \text{ComSet}$ and sends it to the challenger. The challenger responds by returning a random challenge $\text{ch} \in \text{ChSet}$.

Output: \mathcal{A} outputs a response $\text{resp} \in \text{ResSet}$. The adversary \mathcal{A} wins the game if $(\text{com}, \text{ch}, \text{resp})$ is a valid transcript for X_{IS} .

We say ID_{IS} is ϵ_{IS} -lossy sound if for any unbounded (possibly quantum) adversary \mathcal{A} the winning probability in the above game is less than ϵ_{IS} .

2.2 Digital Signature Schemes

Here we introduce the definition of standard signature schemes.

Definition 2.3. A signature scheme $\Pi_{\mathcal{S}}$ consists of three PPT algorithms $(\text{S.KeyGen}, \text{S.Sign}, \text{S.Vrfy})$ such that:

- $\text{S.KeyGen}(1^\lambda) \rightarrow (\text{vk}, \text{sk})$: On input a security parameter 1^λ , the key generation algorithm outputs a pair of verification and signing keys (vk, sk) ;
- $\text{S.Sign}(\text{sk}, \text{M}) \rightarrow \sigma$: On input a signing key sk and a message M , the signing algorithm outputs a signature σ ;
- $\text{S.Vrfy}(\text{vk}, \text{M}, \sigma) \rightarrow 1/0$: On input a verification key vk , a message M and a signature σ , the verification key outputs 1 (accept) or 0 (reject).

We require a signature scheme $\Pi_{\mathcal{S}}$ to satisfy the following two properties.

Correctness. For every security parameter 1^λ , with $\lambda \in \mathbb{N}$, and every message M the following holds:

$$\Pr \left[\text{S.Vrfy}(\text{vk}, \text{M}, \sigma) = 1 \mid \begin{array}{l} (\text{vk}, \text{sk}) \leftarrow \text{S.KeyGen}(1^\lambda), \\ \sigma \leftarrow \text{S.Sign}(\text{sk}, \text{M}) \end{array} \right] = 1.$$

Unforgeability. We define the *strong unforgeability under chosen message attack* su-cma by the following game played by an adversary \mathcal{A} and a challenger.

Setup: The challenger runs $(\text{vk}, \text{sk}) \leftarrow \text{S.KeyGen}(1^\lambda)$ and provides the adversary \mathcal{A} the verification key vk . It also prepares an empty set $\mathcal{S} = \emptyset$.

Signing Queries: The adversary \mathcal{A} may adaptively submit messages M to the challenger. The challenger responds by returning $\sigma \leftarrow \text{S.Sign}(\text{sk}, \text{M})$ to \mathcal{A} . It then updates the set $\mathcal{S} \leftarrow \mathcal{S} \cup \{(\text{M}, \sigma)\}$.

Output: Finally, \mathcal{A} outputs a forgery (M^*, σ^*) . We say the adversary \mathcal{A} wins if $(\text{M}^*, \sigma^*) \notin \mathcal{S}$ and $\text{S.Vrfy}(\text{vk}, \text{M}^*, \sigma^*) = 1$.

We define the advantage of \mathcal{A} as the probability it wins the above game, that is, $\text{Adv}_{\mathcal{A}}^{\text{su-cma}}(1^\lambda) := \Pr[\mathcal{A} \text{ wins}]$.

Definition 2.4 (Su-cma Security). We say a signature scheme $\Pi_{\mathcal{S}}$ is *su-cma secure* if for all PPT adversaries \mathcal{A} , we have $\text{Adv}_{\mathcal{A}}^{\text{su-cma}}(\lambda) = \text{negl}(\lambda)$.

2.3 Pseudorandom Functions

Consider a mapping $\text{PRF} : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$, where \mathcal{K} is a key space. We say PRF is a *pseudorandom function* if for all PPT (or quantum) adversaries, their advantage defined below is negligible:

$$\text{Adv}_{\mathcal{A}}^{\text{PRF}}(\lambda) := \left| \Pr[\mathcal{A}^{\text{PRF}(K, \cdot)}(1^\lambda) = 1 \mid K \leftarrow \mathcal{K}] - \Pr[\mathcal{A}^{\text{RF}(\cdot)}(1^\lambda) = 1] \right|,$$

where $\text{RF} : \mathcal{X} \rightarrow \mathcal{Y}$ is a perfect random function. In practice, any standard hash function (e.g., SHA-3) is believed to be a (quantumly) secure PRF.

2.4 Fiat-Shamir Transformation

The original Fiat-Shamir transformation [FS87, AABN02] turns a (not necessarily lossy) identification protocol ID into a digital signature scheme by means of a cryptographic hash function $H : \{0, 1\}^* \rightarrow \text{ChSet}$ modeled as a *classical* random oracle (RO). For each parallel execution of ID, the challenge is obtained as $H(\text{com}, M)$, where M is the message to sign. Then the resulting digital signature σ is a t -tuple composed by t commitments and the corresponding responses, where t is set in such a way that $|\text{ChSet}|^t$ is exponentially large. Recently, the Fiat-Shamir transformation has been extended to the *quantum* random oracle model (QROM) as well [KLS18, DFMS19, LZ19].

In this work, we will be interested in Fiat-Shamir transformations for a specific type of identification protocol (namely, *lossy* identification protocol) which admits *tight* security proofs. For a general identification protocol, it is well-known that the Fiat-Shamir signature incurs a prohibitively large reduction loss: the advantage of breaking the underlying hard problem degrades as $O(Q^{-1} \cdot \epsilon^2)$ in the classical ROM and as $O(Q^{-6} \cdot \epsilon^3)$ in the quantum ROM, where Q is the number of random oracle queries made by the adversary and ϵ is the advantage against the Fiat-Shamir signature scheme.

The following result is taken from the recent work of Kiltz, Lyubashevsky, and Schaffner [KLS18].

Theorem 2.5. *Assume the identification protocol ID is lossy, perfect HVZK, has α bits of min-entropy, has perfect unique response, and is ϵ_{ls} -lossy sound. The Fiat-Shamir transformation provides a signature scheme such that, for any quantum adversary \mathcal{A} against su-cma security that issues at most Q_H queries to the quantum random oracle, there exists quantum adversaries \mathcal{B} and \mathcal{D} such that*

$$\text{Adv}_{\mathcal{A}}^{\text{su-cma}}(\lambda) \leq \text{Adv}_{\mathcal{B}}^{\text{lossy}}(\lambda) + 8(Q_H + 1)^2 \cdot \epsilon_{\text{ls}} + 2^{-\alpha+1} + \text{Adv}_{\mathcal{D}}^{\text{PRF}}(\lambda),$$

and $\text{Time}(\mathcal{B}) = \text{Time}(\mathcal{D}) = \text{Time}(\mathcal{A}) + Q_H \approx \text{Time}(\mathcal{A})$.

In the classical setting, the only difference is that the bound depends linearly on Q_H instead of quadratically.

The above theorem is obtained by derandomizing the Fiat-Shamir signature by a pseudorandom function PRF and plugging it in Theorem 3.1 of [KLS18]. We note that some simplification to Theorem 3.1 of [KLS18] is made since our proposed lossy identification protocol achieves *perfect* HVZK and *perfect* unique response.

2.5 Class Group Actions and Hardness Assumption

The action of ideal class groups on elliptic curves was firstly proposed for cryptographic purposes by Couveignes [Cou06], and Rostovtsev and Stolbunov [RS06, Sto10]. Their approach was then revised by De Feo, Kieffer and Smith [DKS18], who were unable to turn it into practicality despite the introduction of remarkable mathematically-driven speed-ups. The efficiency issues were overcome by Castryck *et al.* [CLM⁺18], that introduced the CSIDH key-exchange protocol restricting to supersingular elliptic curves. In the following, we will give a brief background on ideal class groups and their action on supersingular curves. For a more detailed overview we suggest the consultation of [CLM⁺18] and Cox's book [Cox13].

Let \mathbb{F}_p denote a prime field, with p being an odd prime. Given two elliptic curves E, E' defined over \mathbb{F}_p , an isogeny $\varphi : E \rightarrow E'$ is a non-constant morphism mapping 0_E to $0_{E'}$. Hence each coordinate of $\varphi(x, y)$ can be expressed as a fraction of two polynomials belonging to $\overline{\mathbb{F}}_p[x, y]$. If their coefficients are contained in \mathbb{F}_p , then we say that φ is defined over \mathbb{F}_p . A separable isogeny (it induces a separable extension of function fields) having $\{0_E\}$ as kernel is an isomorphism; an isogeny having the same domain and range is an endomorphism.

The set of all endomorphisms of an elliptic curve E , together with the zero map, form a ring under pointwise addition and composition. Such a ring is called the *endomorphism ring of E* and it is denoted by $\text{End}(E)$. If $\text{End}(E)$ is abelian, the curve is said to be ordinary, otherwise it is said to be supersingular. The restriction $\text{End}_p(E)$ to the endomorphisms defined over \mathbb{F}_p constitutes a subring, which is isomorphic to an order in the quadratic field $\mathbb{K} = \mathbb{Q}(\sqrt{-p})$. An order is a subring of $\mathbb{Q}(\sqrt{-p})$ which is also a finitely-generated \mathbb{Z} -module

containing a basis of \mathbb{K} as a \mathbb{Q} -vector space. The set $\mathbb{Z}[\sqrt{-p}] = \{m + n\sqrt{-p} \mid m, n \in \mathbb{Z}\}$ satisfies the above three conditions and we will denote it by \mathcal{O} . We then consider the set $\mathcal{E}ll_p(\mathcal{O}, \pi)$ containing all supersingular curves E defined over \mathbb{F}_p - modulo isomorphisms defined over \mathbb{F}_p - such that there exists an isomorphism between \mathcal{O} and $\text{End}_p(E)$ mapping $\sqrt{-p} \in \mathcal{O}$ into the Frobenius endomorphism $(x, y) \mapsto (x^p, y^p)$. As shown in [CLM⁺18], each isomorphism class in $\mathcal{E}ll_p(\mathcal{O}, \pi)$ can be uniquely represented by a single element of \mathbb{F}_p if $p \geq 5$ is a prime such that $p \equiv 3 \pmod{8}$.

A fractional ideal \mathfrak{a} of \mathcal{O} is a finitely generated \mathcal{O} -submodule of \mathbb{K} . When \mathfrak{a} is contained in \mathcal{O} , it is said to be integral; when $\mathfrak{a} = \alpha\mathcal{O}$ for some $\alpha \in \mathbb{K}$, it is said to be principal; when there exists another fractional ideal \mathfrak{b} such that $\mathfrak{a}\mathfrak{b} = \mathcal{O}$, it is called invertible. The invertible fractional ideals of \mathcal{O} form an abelian group. Its quotient by the subgroup composed by principal fractional ideals is a finite group called *ideal class group* of \mathcal{O} , usually denoted by $\mathcal{C}l(\mathcal{O})$. Its cardinality is the class number of \mathcal{O} .

The ideal class group $\mathcal{C}l(\mathcal{O})$ acts freely and transitively on the set $\mathcal{E}ll_p(\mathcal{O}, \pi)$ via the group action \star :

$$\begin{aligned} \star : \mathcal{C}l(\mathcal{O}) \times \mathcal{E}ll_p(\mathcal{O}, \pi) &\rightarrow \mathcal{E}ll_p(\mathcal{O}, \pi) \\ (\mathfrak{a}, E) &\mapsto \mathfrak{a} \star E. \end{aligned}$$

For simplicity, we will use representatives instead of equivalence classes to denote elements of $\mathcal{C}l(\mathcal{O})$ and $\mathcal{E}ll_p(\mathcal{O}, \pi)$. When p is of the form $4\ell_1\ell_2 \cdots \ell_s - 1$, with ℓ_1, \dots, ℓ_s small odd primes, a special integral ideal $\mathcal{I}_{\ell_i} \subset \mathcal{O}$ corresponds to each prime ℓ_i . These ideals allow an easy computation of the group action. In particular, the action of \mathcal{I}_{ℓ_i} on a curve $E \in \mathcal{E}ll_p(\mathcal{O}, \pi)$ is determined by an isogeny having as kernel the unique rational ℓ_i -torsion subgroup of E .

The general variant of the CSIDH key-exchange scheme relies on the heuristic that the equivalence classes of the ideals $\mathcal{I}_{\ell_1}, \dots, \mathcal{I}_{\ell_s}$, together with their inverses, generate the entire ideal class group $\mathcal{C}l(\mathcal{O})$. In [CLM⁺18], Castryck *et al.* propose different sets of parameters for CSIDH, each of them supposedly achieving a specific quantum security level. For the *smallest*² set of parameters, named CSIDH-512 since $p \simeq 2^{512}$, the class group structure of $\mathcal{C}l(\mathcal{O})$ has been recently computed by Beullens *et al.* [BKV19]. They showed that $\mathcal{C}l(\mathcal{O})$ is a cyclic group of odd order N , where $N \simeq 2^{257.1}$ and $\mathcal{C}l(\mathcal{O}) = \langle \mathcal{I}_3 \rangle$. As a consequence, this group admits a canonical representation (as \mathbb{Z}_N) and an efficient uniform sampling of its elements. For simplicity, in the following we will denote by \mathfrak{g} the generator \mathcal{I}_3 .

Hardness Assumption. The group action inverse problem (GAIP) is the hardness assumption originally introduced by [CLM⁺18], which underlies the security of both SeaSign [DG19] and CSI-FiSh [BKV19]. Although we will not directly use GAIP in our construction, we provide it as a base point to compare the assumption we introduce.

Definition 2.6 (Group Action Inverse Problem (GAIP)). *Given two supersingular elliptic curves, $E, E_1 \in \mathcal{E}ll_p(\mathcal{O}, \pi)$, find an element $\mathfrak{a} \in \mathcal{C}l(\mathcal{O})$ such that $\mathfrak{a} \star E = E_1$.*

3 Base Lossy Identification Protocol from CSIDH-512

The CSI-FiSh signature is obtained by applying the Fiat-Shamir transformation to an identification protocol originally sketched by Couveignes [Cou06] and Stolbunov [Sto12]. In this section, we introduce our base *lossy* identification protocol for any set of CSIDH parameters for which the ideal class group $\mathcal{C}l(\mathcal{O})$ is cyclic, with a known order N and generator \mathfrak{g} . We further discuss the corresponding hardness assumption on which its security relies. Such a scheme considers an exponent $a \in \mathbb{Z}_N$, the private key, and two pairs of curves, where the second pair, the public key, is determined by the action of \mathfrak{g}^a on the first pair. For the concrete instantiation in Section 5, we use the CSIDH-512 parameters.

²The parameter set having the smallest value for the prime p .

3.1 Hardness Assumption: Decisional CSIDH

We construct a lossy identification protocol based on the *decisional* CSIDH (D-CSIDH) problem, originally defined by Stolbunov in his PhD thesis [Sto12, Problem 2.2].

Definition 3.1 (Decisional CSIDH Problem). *Given the set $\mathcal{E}\ell_p(\mathcal{O}, \pi)$ and the ideal class group $\mathcal{C}\ell(\mathcal{O})$, the decisional CSIDH (D-CSIDH) problem asks to distinguish between the following two distributions:*

- $(E, H, \mathfrak{a} \star E, \mathfrak{a} \star H)$, where the supersingular elliptic curves E and H are sampled uniformly from $\mathcal{E}\ell_p(\mathcal{O}, \pi)$, while \mathfrak{a} is sampled uniformly from $\mathcal{C}\ell(\mathcal{O})$;
- (E, H, E', H') where E, H, E', H' are supersingular elliptic curves sampled uniformly from $\mathcal{E}\ell_p(\mathcal{O}, \pi)$.

We denote by $\text{Adv}_{\mathcal{A}}^{\text{D-CSIDH}}(\lambda)$ the advantage of an adversary \mathcal{A} distinguishing the two distributions. We say that the D-CSIDH assumption holds if for every PPT (or possibly quantum) adversary \mathcal{A} , $\text{Adv}_{\mathcal{A}}^{\text{D-CSIDH}}(\lambda)$ is negligible.

The D-CSIDH assumption forms the foundation of the security of the key exchange protocol proposed by [CLM⁺18], called CSIDH. However, to be completely accurate, the security of CSIDH *not* always is equivalent to the D-CSIDH problem we defined above. The reason for this is that when the structure of the ideal class group is not known, we cannot properly sample a uniform ideal from $\mathcal{C}\ell(\mathcal{O})$ (and hence a uniform elliptic curve from the set $\mathcal{E}\ell_p(\mathcal{O}, \pi)$). Namely, in that case, a party will sample an ideal that is *heuristically* shown to be close to uniformly random over $\mathcal{C}\ell(\mathcal{O})$. Then, to show security of CSIDH, we must assume the hardness of D-CSIDH for that particular heuristically uniform distribution. Notably, we do not get a reduction from the above D-CSIDH assumption defined for truly uniform samples over $\mathcal{C}\ell(\mathcal{O})$. Hence, for the D-CSIDH assumption to be useful both in a theoretical *and* practical sense, it is desirable to have an *efficient* uniform sampler from the ideal class group $\mathcal{C}\ell(\mathcal{O})$. In this case, the security of CSIDH will indeed be equivalent to the D-CSIDH assumption.

As for the definition of D-CSIDH, we would like to simply keep it agnostic to the existence of an efficient sampler from the ideal class group $\mathcal{C}\ell(\mathcal{O})$. However, throughout the paper, we will always consider a cyclic class group $\mathcal{C}\ell(\mathcal{O})$ with known order and generator (i.e., the one derived from the CSIDH-512 parameters) so as to be able to efficiently sample uniformly over $\mathcal{C}\ell(\mathcal{O})$.

3.2 Construction of Base Lossy Identification Protocol

The base lossy identification protocol we are going to describe requires $\mathcal{C}\ell(\mathcal{O})$ to be efficiently sampleable. As anticipated, we will restrict to the case where $\mathcal{C}\ell(\mathcal{O})$ is cyclic, with a known order N and generator \mathfrak{g} . This reduces sampling from $\mathcal{C}\ell(\mathcal{O})$ to uniformly sampling from \mathbb{Z}_N , and considering the corresponding power of \mathfrak{g} .

Let the set X be composed by pairs $((E_1^{(0)}, E_2^{(0)}), (E_1^{(1)}, E_2^{(1)}))$, where $E_1^{(0)}, E_2^{(0)}, E_1^{(1)}, E_2^{(1)}$ belong to $\mathcal{E}\ell_p(\mathcal{O}, \pi)$. By Y we denote the set of witnesses $\{a \in \mathbb{Z}_N\}$, with N being the cardinality of $\mathcal{C}\ell(\mathcal{O})$. We consider the following binary relation \mathcal{R} on $X \times Y$:


$$\mathcal{R} = \{(((E_1^{(0)}, E_2^{(0)}), (E_1^{(1)}, E_2^{(1)})), a) \mid E_1^{(1)} = \mathfrak{g}^a \star E_1^{(0)}, E_2^{(1)} = \mathfrak{g}^a \star E_2^{(0)}\} \quad (1)$$

We note that the language $\mathcal{L}_{\mathcal{R}}$ is strictly contained in X , i.e. X contains *lossy* statements. On the other hand, each statement in X is a valid instance of the D-CSIDH problem.

The lossy identification protocol $\text{ID}_{\mathfrak{g}}^{\text{base}}$ deduced from relation \mathcal{R} consists of a challenge set $\text{ChSet} = \{0, 1\}$ and five algorithms $(\text{IGen}, \text{LossyGen}, \text{P}_1, \text{P}_2, \text{V})$, detailed in the following. We note that $E_0 \in \mathcal{E}\ell_p(\mathcal{O}, \pi)$ is the base curve, specified by the system parameters, and defined by the equation $y^2 = x^3 + x$ over \mathbb{F}_p .

- Algorithm IGen uniformly samples $a, b, c \in \mathbb{Z}_N$ and outputs a statement-witness pair $(X, W) \in \mathcal{R}$, where $X = ((E_1^{(0)} = \mathfrak{g}^b \star E_0, E_2^{(0)} = \mathfrak{g}^c \star E_0), (E_1^{(1)} = \mathfrak{g}^a \star E_1^{(0)}, E_2^{(1)} = \mathfrak{g}^a \star E_2^{(0)}))$, and $W = a$.

- Algorithm `LossyGen` uniformly samples $a, a', b, c \in \mathbb{Z}_N$ and outputs a lossy statement $\mathbf{X}_{\text{ls}} = ((E_1^{(0)} = \mathbf{g}^b \star E_0, E_2^{(0)} = \mathbf{g}^c \star E_0), (E_1^{(1)} = \mathbf{g}^a \star E_1^{(0)}, E_2^{(1)} = \mathbf{g}^{a'} \star E_2^{(0)}))$.
- On input (X, W) , P_1 generates a random integer $r \in \mathbb{Z}_N$ and returns the commitment $\text{com} = (F_1 = \mathbf{g}^r \star E_1^{(0)}, F_2 = \mathbf{g}^r \star E_2^{(0)})$.



- On input $(X, W, \text{com}, \text{ch})$, where $\text{ch} \in \text{ChSet}$, P_2 outputs the response resp which is r if $\text{ch} = 0$, $r - a$ if $\text{ch} = 1$.
- On input $(X, \text{com}, \text{ch}, \text{resp})$, the verification algorithm V checks that

$$\begin{cases} (\mathbf{g}^{\text{resp}} \star E_1^{(0)} = F_1, \mathbf{g}^{\text{resp}} \star E_2^{(0)} = F_2) & \text{if } \text{ch} = 0 \\ (\mathbf{g}^{\text{resp}} \star E_1^{(1)} = F_1, \mathbf{g}^{\text{resp}} \star E_2^{(1)} = F_2) & \text{if } \text{ch} = 1 \end{cases} \quad (2)$$

The interaction between a prover and a verifier within the identification protocol is summarised in Figure 1.




Figure 1: The base lossy identification protocol and its transcript $(\text{com}, \text{ch}, \text{resp})$.

3.3 Security of Base Lossy Identification Protocol $\text{ID}_{\text{ls}}^{\text{base}}$

We show that the proposed lossy identification protocol $\text{ID}_{\text{ls}}^{\text{base}}$ satisfies all the desired properties presented in Section 2.1. Properties for standard identification protocols - namely, correctness, perfect unique response, and commitment revocability - are straightforward to prove, with the last two verified by noticing that the group action \star is transitive and free. Moreover, for the Honest-Verifier Zero-Knowledge property, consider a simulator Sim defined as follows:

$\text{Sim}(X, \text{ch})$: on input a statement $\mathbf{X} = ((E_1^{(0)}, E_2^{(0)}), (E_1^{(1)}, E_2^{(1)})) \in \mathcal{L}_{\mathcal{R}}$ and a challenge bit $\text{ch} \in \{0, 1\}$, the simulator samples a random $u \in \mathbb{Z}_N$ and outputs either of the following tuples, depending on whether $\text{ch} = 0$ or $\text{ch} = 1$:

$$((\mathbf{g}^u \star E_1^{(0)}, \mathbf{g}^u \star E_2^{(0)}), \text{ch} = 0, u), \quad ((\mathbf{g}^u \star E_1^{(1)}, \mathbf{g}^u \star E_2^{(1)}), \text{ch} = 1, u).$$

It can be checked that the transcripts output by the simulator Sim are indistinguishable from honest transcripts, since both have uniformly random distributed values as responses. Finally, by construction, we have $\log N$ bits of min-entropy.

The remaining issue is showing that $\text{ID}_{\text{ls}}^{\text{base}}$ satisfies the lossy properties (see Definition 2.2). Specifically, it has indistinguishability of lossy statements and statistical lossy soundness.

Lemma 3.2. *Our lossy identification protocol $\text{ID}_{\text{Is}}^{\text{base}}$ satisfies indistinguishability of lossy statements assuming the hardness of the D-CSIDH problem. Specifically, an adversary \mathcal{A} with advantage $\text{Adv}_{\mathcal{A}}^{\text{lossy}}(\lambda)$ can be turned into an adversary \mathcal{B} against the D-CSIDH problem with advantage $\text{Adv}_{\mathcal{B}}^{\text{D-CSIDH}}(\lambda) = \text{Adv}_{\mathcal{A}}^{\text{lossy}}(\lambda)$ and the same running time.*

Proof. The statement is an immediate consequence of the D-CSIDH problem. In particular, the distribution induced by LGen corresponds to valid D-CSIDH instances and that of LossyLGen corresponds to random D-CSIDH instances. \square

Lemma 3.3. *Our lossy identification protocol $\text{ID}_{\text{Is}}^{\text{base}}$ satisfies statistical ϵ_{Is} -lossy soundness for $\epsilon_{\text{Is}} = 1/2 + 1/2N$, where $N = |\mathcal{Cl}(\mathcal{C})|$.*

Proof. First of all, a simple calculation shows that the set of valid statements $\mathcal{L}_{\mathcal{R}}$ has size N^3 . Therefore, since LossyLGen outputs a uniformly random image in the set X , which has size N^4 , we have $\Pr[\mathbf{X}_{\text{Is}} \leftarrow \text{LossyLGen}(1^\lambda) : \mathbf{X}_{\text{Is}} \in \mathcal{L}_{\mathcal{R}}] = 1/N$. Furthermore, for an adversary \mathcal{A} against the lossy impersonation game, the following holds:

$$\begin{aligned} \Pr[\mathcal{A} \text{ wins}] &= \Pr[\mathcal{A} \text{ wins} \mid \mathbf{X}_{\text{Is}} \notin \mathcal{L}_{\mathcal{R}}] \Pr[\mathbf{X}_{\text{Is}} \notin \mathcal{L}_{\mathcal{R}}] + \\ &\quad \Pr[\mathcal{A} \text{ wins} \mid \mathbf{X}_{\text{Is}} \in \mathcal{L}_{\mathcal{R}}] \Pr[\mathbf{X}_{\text{Is}} \in \mathcal{L}_{\mathcal{R}}] \\ &\leq \Pr[\mathcal{A} \text{ wins} \mid \mathbf{X}_{\text{Is}} \notin \mathcal{L}_{\mathcal{R}}] \cdot \left(1 - \frac{1}{N}\right) + \frac{1}{N}. \end{aligned}$$

We show that for any statement $\mathbf{X}_{\text{Is}} \notin \mathcal{L}_{\mathcal{R}}$ and commitment $\text{com} \in \text{ComSet}$, there exists at most one challenge $\text{ch} \in \text{ChSet}$ that admits a *valid* response $\text{resp} \in \text{ResSet}$. Since this implies $\Pr[\mathcal{A} \text{ wins} \mid \mathbf{X}_{\text{Is}} \notin \mathcal{L}_{\mathcal{R}}] \leq 1/|\text{ChSet}| = 1/2$, we obtain $(1/2 + 1/2N)$ -lossy soundness as desired.

Given a statement $\mathbf{X}_{\text{Is}} = ((E_1^{(0)}, E_2^{(0)}), (E_1^{(1)}, E_2^{(1)})) \notin \mathcal{L}_{\mathcal{R}}$, let us assume there exist two valid transcripts for \mathbf{X}_{Is} . Namely, consider $(\text{com}, \text{ch}, \text{resp})$ and $(\text{com}, \text{ch}', \text{resp}')$, with $\text{ch} \neq \text{ch}'$ and $\text{com} = (F_1, F_2)$. Then, it is possible to extract a witness W such that $(\mathbf{X}_{\text{Is}}, W) \in \mathcal{L}_{\mathcal{R}}$. Indeed, assuming $\text{ch} = 0$, the responses $\text{resp}, \text{resp}'$ must satisfy

$$\begin{cases} \mathbf{g}^{\text{resp}} \star E_1^{(0)} = F_1, & \mathbf{g}^{\text{resp}} \star E_2^{(0)} = F_2, \\ \mathbf{g}^{\text{resp}'} \star E_1^{(1)} = F_1, & \mathbf{g}^{\text{resp}'} \star E_2^{(1)} = F_2. \end{cases} \quad (3)$$

Therefore, $\text{resp} - \text{resp}'$ is the desired witness, that is, $E_1^{(1)} = g^{\text{resp} - \text{resp}'} \star E_1^{(0)}$ and $E_2^{(1)} = g^{\text{resp} - \text{resp}'} \star E_2^{(0)}$. However, this is a contradiction to $\mathbf{X}_{\text{Is}} \notin \mathcal{L}_{\mathcal{R}}$. Therefore, there can exist at most one challenge that possesses a valid response. This concludes the proof. \square

3.4 Lossy Soundness Amplification of $\text{ID}_{\text{Is}}^{\text{base}}$

As typically done, we use standard parallel repetition of the base lossy identification protocol $\text{ID}_{\text{Is}}^{\text{base}}$ to make the lossy soundness ϵ_{Is} negligibly small, as required when setting the concrete parameters for the relative FS signature according to Theorem 2.5. Specifically, on input (\mathbf{X}, W) , the prover runs parallel execution of the protocol with the verifier, where the verifier uses independent challenges in each execution.

We make this standard procedure explicit since, unlike sigma-protocols with 2-special soundness, lossy soundness is not closed under parallel repetition. That is, even if we run t parallel instances of our base protocol $\text{ID}_{\text{Is}}^{\text{base}}$, this will not result in a protocol with $(\epsilon_{\text{Is}})^t$ -lossy soundness. Namely, we have the following result.

Lemma 3.4. *Consider running t parallel rounds of the base lossy identification protocol $\text{ID}_{\text{Is}}^{\text{base}}$ (with the same statement-witness pair). Then it satisfies statistical ϵ_{Is} -lossy soundness for $\epsilon_{\text{Is}} = 1/2^t \cdot (1 - 1/N) + 1/N$, where $N = |\mathcal{Cl}(\mathcal{C})|$. In particular, we have $\epsilon_{\text{Is}} \leq 1/2^t + 1/N$.*

Proof. The proof is straightforward. In case $\mathbf{X}_{\text{Is}} \notin \mathcal{L}_{\mathcal{R}}$, we can argue that the adversary has at most $1/2^t$ probability in winning the lossy impersonation game. Recalling that $\mathbf{X}_{\text{Is}} \in \mathcal{L}_{\mathcal{R}}$ happens with probability $1/N$ over the random choice of LossyLGen , we can upper bound the advantage of \mathcal{A} by $\epsilon_{\text{Is}} = 1/2^t(1 - 1/N) + 1/N$. This concludes the proof. \square

All other properties are closed under parallel repetition and inherited directly from $\text{ID}_{\text{IS}}^{\text{base}}$.

4 Optimized Lossy Identification Protocol from CSIDH-512

We show several methods to optimize our base lossy identification protocol, following closely the work of [DG19, BKV19]. We first prepare a slight variant of the D-CSIDH assumption, which will form the basis of our optimized schemes.

4.1 Hardness Assumption: Fixed-Curve Multi-Decisional CSIDH

We consider a slight variant of D-CSIDH, where we are given many D-CSIDH tuples, with the first two elliptic curves of each tuple being fixed. Formally, we consider the following problem, which is equivalent to D-CSIDH when $S = 1$.

Definition 4.1 (Fixed-Curve Multi-Decisional CSIDH Problem). *Let S be a positive integer. Given the ideal class group $\mathcal{Cl}(\mathcal{O})$ and the set $\mathcal{Ell}_p(\mathcal{O}, \pi)$, the fixed-curve multi-decisional CSIDH (FCMD-CSIDH) problem with parameter S asks to distinguish between the following two distributions³:*

- $(E, H, (\mathbf{a}_i \star E, \mathbf{a}_i \star H)_{i \in [S]})$, where the supersingular elliptic curves E and H are sampled uniformly from $\mathcal{Ell}_p(\mathcal{O}, \pi)$, and \mathbf{a}_i for $i \in [S]$ are sampled uniformly from $\mathcal{Cl}(\mathcal{O})$;
- $(E, H, (E'_i, H'_i)_{i \in [S]})$ where E, H, E'_i, H'_i for $i \in [S]$ are supersingular elliptic curves sampled uniformly from $\mathcal{Ell}_p(\mathcal{O}, \pi)$.

We denote by $\text{Adv}_{\mathcal{A}, S}^{\text{FCMD-CSIDH}}(\lambda)$ the advantage of an adversary \mathcal{A} distinguishing the two distributions. We say that the FCMD-CSIDH assumption with parameter S holds if for any PPT (or possibly quantum) adversary \mathcal{A} , $\text{Adv}_{\mathcal{A}, S}^{\text{FCMD-CSIDH}}(\lambda)$ is negligible.

A tight reduction from the (one-instance) decisional CSIDH problem to the fixed-curve multi-decisional CSIDH problem with parameter S would have been desirable, however, this seems to be highly challenging (as long as we view the group action \star as a black box). This is in sharp contrast with the classical decisional DH problem, which admits a nice random self-reducibility property. The main reason why D-CSIDH does not possess this property seems to stem from the fact that the group action only allows to add a known constant to the exponent of \mathfrak{g} when considering a curve $\mathfrak{g}^a \star E$. In other words, we do not have an analogous of the mapping $g^a \mapsto (g^a)^r$ exploited in the classical DH setting.

Therefore, we only have a trivial non-tight reduction from the D-CSIDH problem to the FCMD-CSIDH problem with parameter S . This is formally stated in the following lemma.

Lemma 4.2 (D-CSIDH to FCMD-CSIDH). *Let S be a positive integer. Let $\mathcal{Cl}(\mathcal{O})$ be the ideal class group of an order \mathcal{O} in $\mathbb{Q}(\sqrt{-p})$, with p a prime, and $\mathcal{Ell}_p(\mathcal{O}, \pi)$ be the corresponding set of supersingular elliptic curves. Then, for any adversary \mathcal{A} for the FCMD-CSIDH problem with parameter S , there exists an adversary \mathcal{B} for the D-CSIDH problem such that*

$$\text{Adv}_{\mathcal{A}, S}^{\text{FCMD-CSIDH}} \leq S \cdot \text{Adv}_{\mathcal{B}}^{\text{D-CSIDH}},$$

and $\text{Time}(\mathcal{B}) \approx \text{Time}(\mathcal{A})$.

Proof. The proof is elementary. We consider $S + 1$ hybrid games where, in the j -th game⁴, an adversary is given $(E, H, (E'_i, H'_i)_{i \in [S]})$, where $(E'_i, H'_i)_{i \in [j]}$ is random over $\mathcal{Ell}_p(\mathcal{O}, \pi)^2$ and $(E'_i, H'_i)_{i \in [S] \setminus [j]}$ is of the form $(\mathbf{a}_i \star E, \mathbf{a}_i \star H)$ for a random $\mathbf{a}_i \in \mathcal{Cl}(\mathcal{O})$. We then simply show that each game is indistinguishable using the D-CSIDH problem to conclude the proof. However, one thing we remark is that in order for the D-CSIDH adversary \mathcal{B} to simulate the view to the FCMD-CSIDH adversary \mathcal{A} , \mathcal{B} must be able to sample uniformly from $\mathcal{Cl}(\mathcal{O})$. This justifies once more our restriction to cyclic ideal class groups $\mathcal{Cl}(\mathcal{O})$ having known order and generator. \square

³With $[S]$ we denote the set $\{1, \dots, S\}$.

⁴ j varies from 0 to S , and with $[0]$ we denote the set $\{0\}$.

We leave it as an interesting open problem to achieve a tight reduction. We believe a technique which allows such a reduction will most likely have applications elsewhere.

Impact on Signature Scheme (and Identification Protocol). Although this loose reduction is not desirable, fortunately, the integer S will not have a tremendous impact on the concrete choice of parameters for our signature scheme (and identification protocol). This is because S is only a parameter chosen at the setup of the scheme, which is in particular *independent* of the adversary. This should be compared to standard non-tight Fiat-Shamir signatures which incurs a reduction loss of $Q^{-1} \cdot \epsilon^2$ in the classical ROM and $Q^{-6} \cdot \epsilon^3$ in the quantum ROM, where Q is an *adversarially dependent* parameter denoting the number of RO queries. In particular, in the original paper of CSI-FiSh [BKV19], S is a constant set between 1 to $2^{18} - 1$. Depending on the value of S , we have a tradeoff between the runtimes of several algorithms and size of public keys and signatures. We refer to Section 5 for more details.

4.2 Enlarging Challenge Space of Base Lossy Identification Protocol

We show a variant of our base lossy identification protocol which is obtained adapting the idea from [DG19, BKV19] to enlarge the challenge space. In particular, we will use the FCMD-CSIDH problem with parameter S instead of the D-CSIDH problem to define the language used in the identification protocol. Formally, the set of (possibly non-valid) statements is:

$$X = \{((E_1^{(0)}, E_2^{(0)}), (E_1^{(1)}, E_2^{(1)}), \dots, (E_1^{(S)}, E_2^{(S)})) \mid E_1^{(i)}, E_2^{(i)} \in \mathcal{Ell}_p(\mathcal{O})\},$$

while the set of witnesses is $Y = \{(a_1, \dots, a_S) \mid a_1, \dots, a_S \in \mathbb{Z}_N\}$. We then consider the following binary relation on $X \times Y$:

$$\mathcal{R} = \{(((E_1^{(0)}, E_2^{(0)}), (E_1^{(1)}, E_2^{(1)}), \dots, (E_1^{(S)}, E_2^{(S)})), (a_1, \dots, a_S)) \in X \times Y \mid \\ \mathfrak{g}^{a_i} \star E_1^{(0)} = E_1^{(i)}, \mathfrak{g}^{a_i} \star E_2^{(0)} = E_2^{(i)} \text{ for } i \in [S]\}.$$

The lossy identification protocol with enlarged challenge space $\text{ID}_{\text{ls}}^{\text{enCh}}$ deduced from the above relation \mathcal{R} is a simple adaptation of the base scheme $\text{ID}_{\text{ls}}^{\text{base}}$. We provide the details below for completeness, where the challenge space is enlarged to $\text{ChSet} = \{0, 1, \dots, S\}$. Note that S is a parameter chosen by the scheme. Our base scheme is obtained by setting $S = 1$.

- Algorithm IGen uniformly samples $(a_i)_{i \in [S]}, b, c \in \mathbb{Z}_N$ and outputs a statement-witness pair $(X, W) \in \mathcal{R}$, where

$$X = \left((E_1^{(0)} = \mathfrak{g}^b \star E_0, E_2^{(0)} = \mathfrak{g}^c \star E_0), (E_1^{(i)} = \mathfrak{g}^{a_i} \star E_1^{(0)}, E_2^{(i)} = \mathfrak{g}^{a_i} \star E_2^{(0)})_{i \in [S]} \right),$$

and $W = (a_i)_{i \in [N]}$.

- Algorithm LossyIGen uniformly samples $(a_i, a'_i)_{i \in [S]}, b, c \in \mathbb{Z}_N$ and outputs a lossy statement

$$X = \left((E_1^{(0)} = \mathfrak{g}^b \star E_0, E_2^{(0)} = \mathfrak{g}^c \star E_0), (E_1^{(i)} = \mathfrak{g}^{a_i} \star E_1^{(0)}, E_2^{(i)} = \mathfrak{g}^{a'_i} \star E_2^{(0)})_{i \in [S]} \right),$$

- On input (X, W) , P_1 generates a random integer $r \in \mathbb{Z}_N$ and returns the commitment $\text{com} = (F_1 = \mathfrak{g}^r \star E_1^{(0)}, F_2 = \mathfrak{g}^r \star E_2^{(0)})$.
- On input $(X, W, \text{com}, \text{ch})$, where $\text{ch} \in \text{ChSet}$, P_2 outputs the response resp which is r if $\text{ch} = 0$, $r - a_{\text{ch}}$ if $\text{ch} > 0$.
- On input $(X, \text{com}, \text{ch}, \text{resp})$, the verification algorithm V checks that

$$\mathfrak{g}^{\text{resp}} \star E_1^{(\text{ch})} = F_1, \quad \mathfrak{g}^{\text{resp}} \star E_2^{(\text{ch})} = F_2$$

Security of Lossy Identification Protocol $ID_{\mathcal{I}_S}^{\text{enCh}}$. The proposed lossy identification protocol $ID_{\mathcal{I}_S}^{\text{enCh}}$ inherits most of the desired standard properties presented in Section 2.1 from the base lossy identification protocol $ID_{\mathcal{I}_S}^{\text{base}}$. Namely, correctness, min-entropy, perfect unique response, and commitment revocability trivially follow from those of $ID_{\mathcal{I}_S}^{\text{base}}$. Moreover, the Honest-Verifier Zero-Knowledge property holds similarly as well. Simply consider a simulator Sim which, on input $X \in \mathcal{L}_{\mathcal{R}}$ and $\text{ch} \in \{0, 1, \dots, S\}$, outputs $((g^u \star E_1^{(\text{ch})}, g^u \star E_2^{(\text{ch})}), \text{ch}, u)$, where u is randomly sampled from \mathbb{Z}_N .

We next show that $ID_{\mathcal{I}_S}^{\text{enCh}}$ satisfies the lossy properties (see Definition 2.2). Specifically, it has indistinguishability of lossy statements and statistical lossy soundness.

Lemma 4.3. *Our lossy identification protocol $ID_{\mathcal{I}_S}^{\text{enCh}}$ satisfies indistinguishability of lossy statements assuming the hardness of the FCMD-CSIDH problem with parameter S . Specifically, an adversary \mathcal{A} with advantage $\text{Adv}_{\mathcal{A}}^{\text{lossy}}(\lambda)$ can be turned into an adversary \mathcal{B} against the FCMD-CSIDH problem with advantage $\text{Adv}_{\mathcal{B}, S}^{\text{FCMD-CSIDH}}(\lambda) = \text{Adv}_{\mathcal{A}}^{\text{lossy}}(\lambda)$ and same running time.*

Proof. The proof is analogous to that of Lemma 3.2. \square

Lemma 4.4. *The lossy identification protocol $ID_{\mathcal{I}_S}^{\text{enCh}}$ satisfies statistical $\epsilon_{\mathcal{I}_S}$ -lossy soundness for $\epsilon_{\mathcal{I}_S} = (1/(S+1)) \prod_{i=1}^S ((N-i)/N) + (1 - \prod_{i=1}^S ((N-i)/N))$, where $N = |\mathcal{C}\ell(\mathcal{O})|$.*

Proof. The general strategy is similar to that used for proving Lemma 3.4. We separate the set X in such a way that in one of the subsets the adversary \mathcal{A} has exactly $1/(S+1)$ probability in winning the lossy impersonation game. We then argue that LossyGen outputs a statement belonging to this subset with overwhelming probability. However, unlike the proof in Lemma 3.4, we will not be able to simply use $X \setminus \mathcal{L}_{\mathcal{R}}$ as such a subset. This is because a computationally unbounded adversary may be able, for some of the instances in $X \setminus \mathcal{L}_{\mathcal{R}}$, to forge a response for any $\text{ch} \in \text{ChSet}$.

Recall the set X we consider is of the following form:

$$\left((E_1^{(0)}, E_2^{(0)}), (E_1^{(i)} = \mathfrak{g}^{a_i} \star E_1^{(0)}, E_2^{(i)} = \mathfrak{g}^{a'_i} \star E_2^{(0)})_{i \in [S]} \right),$$

where $(E_1^{(0)}, E_2^{(0)})$ are arbitrary elements in $\mathcal{E}\ell_p(\mathcal{O}, \pi)$, and a_i, a'_i are arbitrary elements in \mathbb{Z}_N . We define the set X_{BAD} as the subset of X which satisfies the following conditions for all distinct $i, j \in [S]$:

$$\begin{cases} a_i \neq a'_i, \\ a_j - a_i \neq a'_j - a'_i. \end{cases} \quad (4)$$

Below, we first compute $|X_{\text{BAD}}|$ and then show that $\Pr[\mathcal{A} \text{ wins} \mid X_{\mathcal{I}_S} \in X_{\text{BAD}}]$ is at most $1/(S+1)$.

First, fix arbitrary $(E_1^{(0)}, E_2^{(0)})$. Then, let us consider fixing arbitrary $(a_1, a'_1) \in (\mathbb{Z}_N)^2$, conditioned on conditions (4). Then, there exist at most $N(N-1)$ choices of such pairs. Let us further consider fixing arbitrary $(a_2, a'_2) \in (\mathbb{Z}_N)^2$, conditioned on conditions (4). Then, since we have to also satisfy $a_2 - a_1 \neq a'_2 - a'_1$, there exist at most $N(N-2)$ choices of such pairs. Continuing this procedure, each pair $(a_i, a'_i) \in (\mathbb{Z}_N)^2$, with $i \in [S]$, has exactly $N(N-i)$ freedom. Therefore, we have $|X_{\text{BAD}}| = N^{2+S}(N-1) \cdots (N-S)$ and $\Pr[X_{\mathcal{I}_S} \leftarrow \text{LossyGen} : X_{\mathcal{I}_S} \in X_{\text{BAD}}]$ equal to $(N-1) \cdots (N-S)/N^S$.

Let us now compute $\Pr[\mathcal{A} \text{ wins} \mid X_{\mathcal{I}_S} \in X_{\text{BAD}}]$. Assume there exist two valid transcripts for $X_{\mathcal{I}_S}$. Namely, consider $(\text{com}, \text{ch}, \text{resp})$ and $(\text{com}, \text{ch}', \text{resp}')$, with $\text{ch} \neq \text{ch}'$ and $\text{com} = (F_1, F_2)$. Then, we have

$$\begin{cases} \mathfrak{g}^{\text{resp}} \star E_1^{(\text{ch})} = F_1, & \mathfrak{g}^{\text{resp}} \star E_2^{(\text{ch})} = F_2, \\ \mathfrak{g}^{\text{resp}'} \star E_1^{(\text{ch}')} = F_1, & \mathfrak{g}^{\text{resp}'} \star E_2^{(\text{ch}')} = F_2. \end{cases}$$

Therefore, we can deduce

$$\mathfrak{g}^{\text{resp} - \text{resp}'} \star E_1^{(\text{ch})} = E_1^{(\text{ch}')} \quad \text{and} \quad \mathfrak{g}^{\text{resp} - \text{resp}'} \star E_2^{(\text{ch})} = E_2^{(\text{ch}')}.$$

However, this clearly contradicts conditions (4). Therefore, there can exist at most one challenge that admits a valid response in case $X_{\text{Is}} \in X_{\text{BAD}}$. In particular, this proves $\Pr[\mathcal{A} \text{ wins} \mid X_{\text{Is}} \in X_{\text{BAD}}] \leq 1/(S+1)$. Combining everything together, we conclude.

$$\begin{aligned} & \Pr[\mathcal{A} \text{ wins}] \\ &= \Pr[\mathcal{A} \text{ wins} \mid X_{\text{Is}} \in X_{\text{BAD}}] \Pr[X_{\text{Is}} \in X_{\text{BAD}}] + \Pr[\mathcal{A} \text{ wins} \mid X_{\text{Is}} \notin X_{\text{BAD}}] \Pr[X_{\text{Is}} \notin X_{\text{BAD}}] \\ &\leq \frac{1}{S+1} \cdot \frac{(N-1) \cdots (N-S)}{N^S} + \left(1 - \frac{(N-1) \cdots (N-S)}{N^S}\right). \end{aligned}$$

□

4.3 (Almost) Doubling Challenge Space of Lossy Identification Scheme $\text{ID}_{\text{Is}}^{\text{enCh}}$

Following the work of [BKV19] and their exploitation of quadratic twists, we show a simple method to almost double the challenge space of the previous scheme $\text{ID}_{\text{Is}}^{\text{enCh}}$. The new scheme $\text{ID}_{\text{Is}}^{\text{denCh}}$ (with a doubly-enlarged challenge set) has statement-witness pairs almost identical to those of $\text{ID}_{\text{Is}}^{\text{enCh}}$. The statement remains the same, while the witness contains two extra-coordinates, namely $b, c \in \mathbb{Z}_N$ such that $\mathbf{g}^b \star E_0 = E_1^{(0)}$, $\mathbf{g}^c \star E_0 = E_2^{(0)}$. The algorithm IGen is adjusted according to this modification, while the lossy key generation algorithm LossyIGen and prover's first move P_1 are defined exactly the same.

The challenge set ChSet now admits also negative values, in particular it is the set $\{0, \pm 1, \dots, \pm S\}$. The third move P_2 and the Verification algorithm V are hence converted to deal with these new challenge values:

- On input $(X, W, \text{com}, \text{ch})$, where $\text{ch} \in \text{ChSet}$, P_2 outputs the response resp which is r if $\text{ch} = 0$, $r - a_{\text{ch}}$ if $\text{ch} > 0$ and $r + b + c + a_{|\text{ch}|}$ if $\text{ch} < 0$.
- On input $(X, \text{com}, \text{ch}, \text{resp})$, the verification algorithm V checks that $\mathbf{g}^{\text{resp}} \star E_1^{(\text{ch})} = F_1$, $\mathbf{g}^{\text{resp}} \star E_2^{(\text{ch})} = F_2$ if $\text{ch} \geq 0$, and

$$\mathbf{g}^{\text{resp}} \star E_1^{(\text{ch}), \text{tw}} = F_2, \quad \mathbf{g}^{\text{resp}} \star E_2^{(\text{ch}), \text{tw}} = F_1$$

if $\text{ch} < 0$.

We note that the symbols $E_1^{(\text{ch}), \text{tw}}$, $E_2^{(\text{ch}), \text{tw}}$ denote the quadratic twists of the curve $E_1^{(\text{ch})}$ and $E_2^{(\text{ch})}$, respectively. In particular $E_1^{(\text{ch}), \text{tw}} = \mathbf{g}^{-a_{|\text{ch}|} - b} \star E_0$, and $E_2^{(\text{ch}), \text{tw}} = \mathbf{g}^{-a_{|\text{ch}|} - c} \star E_0$.

Remark 4.5. We exploit the quadratic twist in a slightly different way compared to [BKV19]. This has the effect of allowing us to base security on the FCMD-CSIDH assumption rather than the more restricted FCMD-CSIDH assumption where $E_1^{(0)}$ is fixed to be the special elliptic curve E_0 . The variant proposed in [BKV19, Section 2.5] in order to extend the challenge set to negative values relies on the fact that the public key and the commitment are computed starting from the specific elliptic curve E_0 . Consequently, the security of their derived sigma protocol requires the GAIP problem to be hard for this specific E_0 as the base point. This is in contrast to all other schemes provided in [BKV19] which only need the standard GAIP problem.

Security of Lossy Identification Scheme $\text{ID}_{\text{Is}}^{\text{denCh}}$. The proposed lossy identification protocol $\text{ID}_{\text{Is}}^{\text{denCh}}$ inherits all the standard properties of a lossy identification protocol (see Definition 2.1) from the previous scheme $\text{ID}_{\text{Is}}^{\text{enCh}}$. Moreover, since the statement output by IGen and LossyIGen is identical to $\text{ID}_{\text{Is}}^{\text{enCh}}$, the protocol $\text{ID}_{\text{Is}}^{\text{denCh}}$ satisfies indistinguishability of lossy statements assuming the hardness of the FCMD-CSIDH problem.

Finally, the statistical lossy soundness is addressed in the following lemma. As it can be seen, the shape of ϵ_{Is} remains unchanged with respect to Lemma 4.4.

Lemma 4.6. *Our lossy identification protocol $\text{ID}_{\text{Is}}^{\text{denCh}}$ satisfies statistical ϵ_{Is} -lossy soundness for $\epsilon_{\text{Is}} = (1/(2S+1)) \cdot \prod_{i=1}^S ((N-i)/N) + (1 - \prod_{i=1}^S ((N-i)/N))$, where $N = |\mathcal{C}(\mathcal{O})|$.*

Proof. The proof is almost identical to that of Lemma 4.4. We consider exactly the same partition X_{BAD} , $X \setminus X_{\text{BAD}}$ for the set of statements X which was introduced in Lemma 4.4. The only difference is that three extra-cases arise from the extension of the challenge space when computing $\Pr[\mathcal{A} \text{ wins} \mid X_{\text{is}} \in X_{\text{BAD}}]$. Namely, consider $(\text{com}, \text{ch}, \text{resp})$ and $(\text{com}, \text{ch}', \text{resp}')$, with $\text{ch} \neq \text{ch}'$ and $\text{com} = (F_1, F_2)$, as valid transcripts for X_{is} . If ch and ch' are both negative, we have that $\text{resp} - \text{resp}'$ satisfies

$$\begin{cases} \mathbf{g}^{\text{resp}-\text{resp}'} \star E_1^{(|\text{ch}|), \text{tw}} = E_1^{(|\text{ch}'|), \text{tw}} \\ \mathbf{g}^{\text{resp}-\text{resp}'} \star E_2^{(|\text{ch}|), \text{tw}} = E_2^{(|\text{ch}'|), \text{tw}} \end{cases}$$

i.e. $a_{|\text{ch}|} - a_{|\text{ch}'|} = a'_{|\text{ch}|} - a'_{|\text{ch}'|}$. When $\text{ch} > 0$ and $\text{ch}' < 0$, for the value $\text{resp} - \text{resp}'$ it holds

$$\begin{cases} \mathbf{g}^{\text{resp}-\text{resp}'} \star E_1^{(\text{ch})} = E_2^{(|\text{ch}'|), \text{tw}} \\ \mathbf{g}^{\text{resp}-\text{resp}'} \star E_2^{(\text{ch})} = E_1^{(|\text{ch}'|), \text{tw}} \end{cases}$$

which implies the analogous relation $a_{\text{ch}} - a_{|\text{ch}'|} = a'_{\text{ch}} - a'_{|\text{ch}'|}$. The last case to be taken into account has $\text{ch} = 0$ and $\text{ch}' < 0$, for which we deduce

$$\begin{cases} \mathbf{g}^{\text{resp}-\text{resp}'} \star E_1^{(0)} = E_2^{(|\text{ch}'|), \text{tw}} \\ \mathbf{g}^{\text{resp}-\text{resp}'} \star E_2^{(0)} = E_1^{(|\text{ch}'|), \text{tw}} \end{cases}$$

and then the relation $a_{|\text{ch}'|} = a'_{|\text{ch}'|}$.

Therefore, combining this with conditions (4) in Lemma 4.4, we conclude that in case $X_{\text{is}} \in X_{\text{BAD}}$, there can exist at most one $\text{ch} \in \{0, \pm 1, \dots, \pm S\}$ which leads to a valid response resp . This concludes the proof. \square

4.4 Lossy Soundness Amplification of $\text{ID}_{\text{is}}^{\text{denCh}}$

For completeness, we provide the following lemma.

Lemma 4.7. *Consider running t parallel rounds of the lossy identification protocol $\text{ID}_{\text{is}}^{\text{denCh}}$ (with the same statement-witness pair). Then it satisfies statistical ϵ_{is} -lossy soundness for $\epsilon_{\text{is}} = (1/(2S+1)^t) \cdot \prod_{i=1}^S ((N-i)/N) + (1 - \prod_{i=1}^S ((N-i)/N))$, where $N = |\mathcal{C}(\mathcal{O})|$.*

Proof. The proof is analogous to Lemma 3.4. \square

5 Lossy CSI-FiSh: Tightly Secure Signature from CSIDH-512

5.1 Construction of Lossy CSI-FiSh

We depict our Lossy CSI-FiSh signature scheme, whose security is based on the FCMD-CSIDH assumption with parameter S , in Algorithms 1 to 3. It is obtained by applying the Fiat-Shamir transformation on the (soundness-amplified) lossy identification protocol $\text{ID}_{\text{is}}^{\text{denCh}}$ introduced in Section 4.3. We note that we use a (quantumly secure) PRF to derandomize the signature generation, to comply with the hypothesis of Theorem 2.5. In practice, one can simply use any standard hash function (e.g., SHA-3).⁵ Moreover, we use the extra property of commitment revocability (see Definition 2.1) of our lossy identification protocol $\text{ID}_{\text{is}}^{\text{denCh}}$ and let the verifier recover com from resp and ch . This allows us to send t -hash values rather than $2t$ -elliptic curves over $\mathcal{E}\ell_p(\mathcal{O}, \pi)$, and greatly reduces the signature size.

The values S and t are parameters of the signature scheme and can be chosen by the user allowing for different tradeoffs between security, efficiency and signature size. Roughly, the only condition which S and

⁵We note that assuming that a standard cryptographic hash function acts as a PRF does not add to our set of assumptions, since we are already working in the ROM.

t must satisfy is $t \cdot \log_2 S \approx \lambda$ in the classical setting, where λ is the desired security level. In the quantum setting, we will require $t \cdot \log_2 S \approx \lambda + \log_2 Q_H$, where Q_H is the number of hash evaluations an adversary can make. For fixed S and t , the resulting signature size is $t \cdot (\lceil \log_2 N \rceil + \lceil \log_2 S \rceil)$. A selection of candidate parameters is provided in Section 5.2.

The following asserts the tight security of Lossy CSI-FiSh based on the FCMD-CSIDH assumption. Observe that the computational advantages appear with a constant factor (one). Moreover, viewing S as a constant parameter, Lossy CSI-FiSh admits tight security based on the D-CSIDH assumption as well.

Theorem 5.1. *Let Lossy CSI-FiSh be the signature scheme depicted in Algorithms 1, 2, and 3. Then, for any quantum adversary \mathcal{A} against su-cma security of Lossy CSI-FiSh that issues at most Q_H queries to the quantum random oracle, there exists a quantum adversary \mathcal{B} against the FCMD-CSIDH problem with parameter S and an quantum adversary \mathcal{D} against the PRF such that*

$$\begin{aligned} \text{Adv}_{\mathcal{A}}^{\text{su-cma}}(\lambda) &\leq \text{Adv}_{\mathcal{B},S}^{\text{FCMD-CSIDH}}(\lambda) + \text{Adv}_{\mathcal{D}}^{\text{PRF}}(\lambda) + \frac{2}{N} + \\ &\quad + 8(Q_H + 1)^2 \cdot \left(\frac{1}{(2S + 1)^t} \cdot \prod_{i \in [S]} \frac{N - i}{N} + \left(1 - \prod_{i \in [S]} \frac{N - i}{N} \right) \right) \end{aligned}$$

and $\text{Time}(\mathcal{B}) = \text{Time}(\mathcal{D}) = \text{Time}(\mathcal{A}) + Q_H \approx \text{Time}(\mathcal{A})$. Moreover, we can replace \mathcal{B} by a quantum adversary \mathcal{B}' against the D-CSIDH problem such that

$$\text{Adv}_{\mathcal{B},S}^{\text{FCMD-CSIDH}}(\lambda) \leq S \cdot \text{Adv}_{\mathcal{B}'}^{\text{D-CSIDH}}(\lambda)$$

and $\text{Time}(\mathcal{B}) \approx \text{Time}(\mathcal{B}')$.

In the classical setting, the only difference is that the above bound depends linearly on Q_H instead of quadratically. That is, we can replace $8(Q_H + 1)^2$ with $Q_H + 1$.⁶

Proof. The theorem is a consequence of Theorem 2.5, Lemma 4.2, and Lemma 4.7, along with the additional security claims made in Section 4. Note that the lossy identification protocol $\text{ID}_{\text{is}}^{\text{denCh}}$ has N bits of min entropy, where N is the cardinality of $\mathcal{C}(\mathcal{O})$. \square

Remark 5.2 (Shorter Secret Key). Since the secret key sk is composed of random values, we can use standard tricks to derive them from the PRF key. In particular, we only require one PRF key, e.g., a 16-byte seed for SHA-3, as the secret key. This modification has (almost) no effect on the overall concrete security. In order to simplify the readability, in Algorithm 1 we do not make the use of the PRF explicit while uniformly sampling in \mathbb{Z}_N .

Algorithm 1 KeyGen

Input: E_0 , class number $N = |\mathcal{C}(\mathcal{O})|$

Output: (pk, sk)

- 1: $b \leftarrow \mathbb{Z}_N, c \leftarrow \mathbb{Z}_N$
 - 2: $E_1^{(0)} = \mathbf{g}^b \star E_0, E_2^{(0)} = \mathbf{g}^c \star E_0$
 - 3: **for** $i \in \{1, \dots, S\}$ **do**
 - 4: $a_i \leftarrow \mathbb{Z}_N$
 - 5: $E_1^{(i)} = \mathbf{g}^{a_i} \star E_1^{(0)}, E_2^{(i)} = \mathbf{g}^{a_i} \star E_2^{(0)}$
 - 6: $\text{pk} = [(E_1^{(j)}, E_2^{(j)}) : j \in \{0, \dots, S\}]$
 - 7: $\text{K} \leftarrow \mathcal{K}$ \triangleright Sample key for PRF.
 - 8: $\text{sk} = [b, c, a_i : i \in \{1, \dots, S\}, \text{K}]$
- return:** (pk, sk)
-

⁶We can get rid of the constant 8 in the classical setting since it is due to the reduction from the generic *quantum* search problem. See [Zha12, HRS16] for example.

Algorithm 2 Sign

Input: (pk, sk, message M)**Output:** σ

```
1: for  $k \in \{1, \dots, t\}$  do
2:    $r_k \leftarrow \mathbb{Z}_N$  ▷ Derive randomness using PRF(K, M||k).
3:    $F_1^{(k)} = \mathbf{g}^{r_k} \star E_1^{(0)}, F_2^{(k)} = \mathbf{g}^{r_k} \star E_2^{(0)}$ 
4:  $(\text{ch}_1, \dots, \text{ch}_t) = H(F_1^{(1)} \parallel F_2^{(1)} \parallel \dots \parallel F_1^{(t)} \parallel F_2^{(t)} \parallel M)$ 
5: for  $k \in \{1, \dots, t\}$  do ▷ Define  $\text{sign}(0) := 0$ .
6:    $\text{resp}_k = r_k - \text{sign}(\text{ch}_k) a_{|\text{ch}_k|} - \frac{\text{sign}(\text{ch}_k) - |\text{sign}(\text{ch}_k)|}{2} (b + c) \pmod{N}$ 
7:  $\sigma = (\text{resp}_1, \dots, \text{resp}_t, \text{ch}_1, \dots, \text{ch}_t)$ 
return:  $\sigma$ 
```

Algorithm 3 Verify

Input: (pk, signature σ , message M)**Output:** Valid / Invalid

```
1: Parse  $\sigma$  as  $(\text{resp}_1, \dots, \text{resp}_t, \text{ch}_1, \dots, \text{ch}_t)$ 
2: for  $i \in \{1, \dots, S\}$  do
3:    $E_1^{(-i)} = E_1^{(i), \text{tw}}, E_2^{(-i)} = E_2^{(i), \text{tw}}$  ▷ Compute quadratic twists.
4: for  $k \in \{1, \dots, t\}$  do
5:   if  $\text{ch}_k \geq 0$  then
6:      $F_1^{(k)} = \mathbf{g}^{\text{resp}_k} \star E_1^{(\text{ch}_k)}, F_2^{(k)} = \mathbf{g}^{\text{resp}_k} \star E_2^{(\text{ch}_k)}$ 
7:   else
8:      $F_1^{(k)} = \mathbf{g}^{\text{resp}_k} \star E_2^{(\text{ch}_k)}, F_2^{(k)} = \mathbf{g}^{\text{resp}_k} \star E_1^{(\text{ch}_k)}$ 
9:  $(\text{ch}'_1, \dots, \text{ch}'_t) = H(F_1^{(1)} \parallel F_2^{(1)} \parallel \dots \parallel F_1^{(t)} \parallel F_2^{(t)} \parallel M)$ 
10: if  $(\text{ch}_1, \dots, \text{ch}_t) == (\text{ch}'_1, \dots, \text{ch}'_t)$  then
11:   return: Valid
12: else
13:   return: Invalid
```

5.2 Instantiations and Comparison to CSI-FiSh

In this section, we specialise the Lossy CSI-FiSh to the CSIDH-512 parameters, and we consider distinct possible values for t and S both in the classical and quantum setting. For each choice of (S, t) , Theorem 5.1 dictates how many bits of classical/quantum security the scheme guarantees. Clearly, different choices for (S, t) will lead to different bandwidth and computational efficiency.

Here, the term γ -bit of security for a cryptographic scheme is defined as the non-existence of an adversary that breaks the scheme with a success ratio bigger than $2^{-\gamma}$, where the success ratio is the quotient between the adversary’s success probability and its running time [BR96]. In the light of Theorem 5.1, the number of bits of security guaranteed by the signature scheme Lossy CSI-FiSh is upper bounded by the security of the FCMD-CSIDH problem. In line with [CLM⁺18], in the following we assume that the best methodology to solve the D-CSIDH problem (and hence FCMD-CSIDH) is solving one of the corresponding GAIP instances.

Aligning with [BKV19], we consider a hash function that is a factor 2^u slower than a standard hash function (as, for example, SHA3) and vary u to obtain tradeoffs between security and efficiency. Moreover, for the sake of easy comparison, we consider the same values for S and u that are used in [BKV19]. Below, we first provide discussions on the size of the public key and signature size of Lossy CSI-FiSh, both in the classical and quantum setting. We then discuss the efficiency of our scheme with respect to the running times of signature generation and verification. The analysis on runtime will be the same for both the classical and quantum setting.

Classical Setting. The best known classical algorithm to solve the GAIP problem applies the meet-in-the-middle strategy, and hence has a time complexity $O(\sqrt{N})$, where N is the cardinality of $\mathcal{C}\ell(\mathcal{O})$. The class group computation executed in [BKV19] has shown that $N \simeq 2^{257.1}$ for CSIDH-512 parameters. This means that the D-CSIDH problem guarantees at most 128 bits of classical security and then, in turn, the FCMD-CSIDH problem guarantees at most 128-bits when $S = 1$, and at most $128/\log_2 S$ bits when $S > 1$ (see Lemma 4.2).

By Theorem 5.1, for all classical adversaries running in time at most 2^{128} and making at most 2^{128} (random) queries Q_H , it holds:

$$\begin{aligned} \frac{\text{Adv}_{\mathcal{A}}^{\text{su-cma}}(\lambda)}{\text{Time}(\mathcal{A})} &\leq S \cdot \frac{\text{Adv}_{\mathcal{B}'}^{\text{D-CSIDH}}(\lambda)}{\text{Time}(\mathcal{B}')} + \frac{\text{Adv}_{\mathcal{D}}^{\text{PRF}}(\lambda)}{\text{Time}(\mathcal{D})} + \\ &\quad + 2^{-u} \cdot \left(\frac{1}{(2S+1)^t} \cdot \prod_{i \in [S]} \frac{N-i}{N} + \left(1 - \prod_{i \in [S]} \frac{N-i}{N} \right) \right) \\ &\simeq S \cdot 2^{-128} + 2^{-256} + 2^{-u} \cdot (2S+1)^{-t}, \end{aligned}$$

where we ignore the min-entropy since it does not give any significant contribution, being smaller than 2^{256} . Furthermore, $1 - \prod_{i \in [S]} (N-i)/N$ is less than 2^{-242} even for the biggest value of S considered in the following, i.e. $2^{15} - 1$. Hence, the last term can be safely approximated as $2^{-u} \cdot (2S+1)^{-t}$. Now, since each of the values of S is of the form $2^w - 1$, we deduce that $2^{-u} \cdot (2S+1)^{-t}$ must be bounded by 2^{-129} to reach $-128 + w$ bits of security. For a fixed value of u , the smallest value of t for which the above inequality is satisfied is uniquely defined.

In the following Table 1 we report: for each choice of S and u , the minimum value of t for which we obtain the maximal security guaranteed by Lossy CSI-FiSh, the number of bits of such security level, the sizes of signatures and the sizes of public keys for Lossy CSI-FiSh and CSI-FiSh. The column “bits of security” is dismissed for CSI-FiSh as it does not provide provable concrete security. We highlight that for a fixed triple (S, t, u) , the signatures produced with our scheme Lossy CSI-FiSh have exactly the same size as those produced with CSI-FiSh. Finally, we note that the values for CSI-FiSh reported in Table 1 slightly differ from those of [BKV19, Table 3], where some approximations were made (e.g., $2S - 1$ was approximated with $2S$), while our parameters are chosen without any approximation.

Table 1: Comparison between Lossy CSI-FiSh and CSI-FiSh.

				Lossy CSI-FiSh		CSI-FiSh
S	t	u	$ \sigma $	$ \text{pk} $	Bits of security	$ \text{pk} $
1	74	16	2405B	256B	127	64B
3	43	14	1403B	512B	126	192B
7	30	16	983B	1024B	125	448B
15	25	13	822B	2048B	124	960B
$2^6 - 1$	17	16	564B	8.2KB	122	4KB
$2^8 - 1$	14	11	468B	32.8KB	120	16.3KB
$2^{10} - 1$	12	7	404B	131KB	118	65.5KB
$2^{12} - 1$	10	11	339B	524KB	116	262KB
$2^{15} - 1$	8	16	274B	4MB	113	2MB

The differences on the public key sizes between Lossy CSI-FiSh and CSI-FiSh have a double cause:

- in Lossy CSI-FiSh the *starting curves* $E_1^{(0)}, E_2^{(0)}$ are computed by each user and are part of the public key, while in CSI-FiSh the starting curve E_0 is part of the public parameters of the scheme;
- for each coordinate a_i of the private key, with $i \in [S]$, Algorithm 1 computes two curves that will become part of the public key, while in CSI-FiSh only $\mathbf{g}^{a_i} \star E_0$ is appended to the public key.

Recalling that each curve in $\mathcal{E}\ell_p(\mathcal{O}, \pi)$ can be uniquely represented by an element of \mathbb{F}_p , with $p \simeq 2^{512}$, for a given S the size of a CSI-FiSh's public key is $S \cdot 512$ while the size of a public key produced with Lossy CSI-FiSh has length equal to $(S + 2) \cdot 512$, with the increment given by the extra term more visible for small values of S .

Quantum setting. The best known quantum algorithm for the GAIP problem is Kuperberg's algorithm for the hidden shift problem [Kup05b, Kup05a], which has a subexponential complexity. The concrete security estimates, however, are still an active area of research [BLMP19, Pei19, BS18]. In the following we will consider 56 bits of quantum security as a conservative choice, and 64 bits as a more optimistic choice for the D-CSIDH problem. Consequently, we consider quantum adversaries running in time at most 2^{56} in the conservative variant, and 2^{64} in the more optimist one. Analogously, we upper bound the number of possible queries Q_H by 2^{56} in the former case, and by 2^{64} in the latter. In both cases, the upper bound on the security of Lossy CSI-FiSh depends quadratically in Q_H .

Considering the optimistic variant, the following inequality holds due to Theorem 5.1:

$$\begin{aligned} \frac{\text{Adv}_{\mathcal{A}}^{\text{su-cma}}(\lambda)}{\text{Time}(\mathcal{A})} &\leq S \cdot \frac{\text{Adv}_{\mathcal{B}'}^{\text{D-CSIDH}}(\lambda)}{\text{Time}(\mathcal{B}')} + \frac{\text{Adv}_{\mathcal{D}}^{\text{PRF}}(\lambda)}{\text{Time}(\mathcal{D})} + \\ &\quad + 8 \cdot (Q_H + 1) \cdot 2^{-u} \cdot \left(\frac{1}{(2S + 1)^t} \cdot \prod_{i \in [S]} \frac{N - i}{N} + \left(1 - \prod_{i \in [S]} \frac{N - i}{N} \right) \right) \\ &\simeq S \cdot 2^{-64} + 2^{-256} + 2^{67-u} \cdot (2S + 1)^{-t}, \end{aligned}$$

where the approximation is validated by the same argument as in the classical setting. We require $2^{67-u} \cdot (2S + 1)^{-t}$ to be bounded by 2^{-65} in order to reach $-64 + w$ bits of quantum security, with $S = 2^w - 1$. Analogously, in the conservative variant, we require $2^{59-u} \cdot (2S + 1)^{-t}$ to be bounded by 2^{-57} in order to reach $-56 + w$ bits of quantum security, with $S = 2^w - 1$.

In the following Table 2 we differentiate the Conservative and Optimistic variants, reporting the values of t for each choice of S and u , the security levels guaranteed in the two cases, and signatures and public keys sizes. We note that the size of the public key only depends on S , hence it achieves the same size as in the classical setting (see Table 1).

Table 2: Parameters and achieved quantum security level for Lossy CSI-FiSh.

S	u	$ \mathbf{pk} $	Conservative variant			Optimistic variant		
			t	$ \sigma $	Bits of security	t	$ \sigma $	Bits of security
1	16	256B	64	2080B	55	74	2405B	63
3	14	512B	37	1208B	54	43	1403B	62
7	16	1024B	26	852B	53	30	983B	61
15	13	2048B	21	691B	52	25	822B	60
$2^6 - 1$	16	8.2KB	15	497B	50	17	564B	58
$2^8 - 1$	11	32.8KB	12	401B	48	14	468B	56
$2^{10} - 1$	7	131KB	10	337B	46	12	404B	54
$2^{12} - 1$	11	524KB	9	305B	44	10	339B	52
$2^{15} - 1$	16	4MB	7	240B	41	8	274B	49

Estimated performance. The costs of key generation, signing and verifying are dominated by the class group actions to be executed in each algorithm. For fixed S and t , the number of actions for each of them is as follows:

- key generation (Algorithm 1) requires $2S + 2$ actions, while S of them are those also computed by the key generation algorithm of CSI-FiSh;
- both signing (Algorithm 2) and verifying (Algorithm 3) need $2t$ actions, exactly twice as many as required by the corresponding algorithms of CSI-FiSh.

As it can be seen, the key generation would be slighter slower than twice the key generation of CSI-FiSh, while the signature generation and verification would be twice that of CSI-FiSh. To provide a concrete benchmark, we estimate the running times using the two triples $(2^{15} - 1, 7, 16)$ and $(2^3 - 1, 28, 16)$ reporting the values of S , t and u for two instances from [BKV19, Table 3]. These two parameter settings are chosen in order to achieve a small signature size and a small sum of signature and public key size, respectively. For the first (resp. second) triple, CSI-FiSh takes the following: 28m (resp. 400ms) for key generation, 395ms (resp. 1.48s) for signature generation, and 393 ms (resp. 1.48s) for signature verification⁷. Therefore, we can estimate that for Lossy CSI-FiSh it will take the following for the respective tuples: ~ 56 m (resp. ~ 920 ms) for key generation, ~ 800 ms (resp. 3s) for signature generation and verification. Here for estimating the runtime of key generation, we simply scaled the runtime of CSI-FiSh by a factor $(2S + 2) \cdot S^{-1}$.

Finally, we provide one potential optimization for lowering the computation time required by the signing and verifying algorithms of Lossy CSI-FiSh. We recall that, in order to efficiently compute the action of \mathbf{g}^a on a given curve, with $a \in \mathbb{Z}_N$, it is necessary to find an equivalent representation of \mathbf{g}^a as a product of small powers of the special ideals \mathcal{J}_{ℓ_i} (see Section 2.5). In [BKV19], an algorithm solving an approximate Closest Vector Problem (CVP) has been proposed to this task. Therefore, the computation of a class group action consists of two steps: finding the equivalent representation and computing the isogenies corresponding to the ideals' powers. Here, we observe that in Lossy CSI-FiSh most of the group actions are pairwise coupled, i.e. they use the same exponent. The result is that the signing and verifying algorithms do not need to execute the finding-equivalent-representation step for each of the class actions. Therefore, this may potentially lead to more efficient algorithms depending on the exact runtime of finding the equivalent representation. We leave it as future work to implement and verify the validity of this observation.

6 Conclusions and Open Problems

In this work, we construct a new signature scheme based on the CSIDH-512 parameters, called Lossy CSI-FiSh. It is *provably secure* and *tightly reduces* to the D-CSIDH (or FCMD-CSIDH) assumption. Lossy

⁷Their benchmarking experiments were performed on a Dell OptiPlex 3050 machine with Intel Core i5-7500T CPU @ 2.70 GHz.

CSI-FiSh inherits most of the efficiency of CSI-FiSh and shows that a slight modification to CSI-FiSh allows to set the concrete parameters in a provably secure manner with minimal cost. In particular, the signature size is as small as CSI-FiSh while the signature generation and verification are around a factor of two slower. We hope that further research will allow to improve the efficiency. Optimisations may be specialized for the scheme (like, for example, halving the number of approximate CVP-problems to be solved in the key generation) or, more generally, be designed for CSI-FiSh. Indeed, the latter would likely have an impact also on our scheme.

One of the biggest open problems is to devise a (lossy or non-lossy) identification protocol that allows for the challenge set to be \mathbb{Z}_N rather than the small set $\{-S, \dots, S\}$, as also mentioned in [BKV19]. This will allow for an analogue of the highly efficient Schnorr signature [Sch90] based on the discrete logarithm problem. Another challenging yet interesting open problem is to show any type of random self-reducibility property for the D-CSIDH problem. We believe such a technique will lend hands to other tightly-secure primitives (e.g., tightly-secure key exchange protocols) and perhaps shed light to Cramer-Shoup-like techniques [CS98] in the isogeny setting.

Acknowledgement. The second author was supported by JST CREST Grant Number JPMJCR19F6.

References

- [AABN02] Michel Abdalla, Jee Hea An, Mihir Bellare, and Chanathip Namprempre. From identification to signatures via the Fiat-Shamir transform: Minimizing assumptions for security and forward-security. In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 418–433. Springer, Heidelberg, April / May 2002. 1, 2, 7
- [AFLT12] Michel Abdalla, Pierre-Alain Fouque, Vadim Lyubashevsky, and Mehdi Tibouchi. Tightly-secure signatures from lossy identification schemes. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 572–590. Springer, Heidelberg, April 2012. 3
- [BHK⁺19] Daniel J Bernstein, Andreas Hülsing, Stefan Kölbl, Ruben Niederhagen, Joost Rijneveld, and Peter Schwabe. The SPHINCS signature framework. In *ACM-CCS*, pages 17–43, 2019. Submission to the NIST PQC project. 4
- [BKV19] Ward Beullens, Thorsten Kleinjung, and Frederik Vercauteren. CSI-FiSh: Efficient isogeny based signatures through class group computations. In *ASIACRYPT 2019, Part I*, *LNCS*, pages 227–247. Springer, Heidelberg, December 2019. 2, 3, 4, 8, 12, 13, 15, 19, 21, 22
- [BLMP19] Daniel J. Bernstein, Tanja Lange, Chloe Martindale, and Lorenz Panny. Quantum circuits for the CSIDH: Optimizing quantum evaluation of isogenies. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part II*, volume 11477 of *LNCS*, pages 409–441. Springer, Heidelberg, May 2019. 20
- [BR96] Mihir Bellare and Phillip Rogaway. The exact security of digital signatures: How to sign with RSA and Rabin. In Ueli M. Maurer, editor, *EUROCRYPT’96*, volume 1070 of *LNCS*, pages 399–416. Springer, Heidelberg, May 1996. 2, 19
- [BS18] Xavier Bonnetain and André Schrottenloher. Quantum security analysis of CSIDH and ordinary isogeny-based schemes. Cryptology ePrint Archive, Report 2018/537, 2018. <https://eprint.iacr.org/2018/537>. 20
- [CLM⁺18] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. CSIDH: An efficient post-quantum commutative group action. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part III*, volume 11274 of *LNCS*, pages 395–427. Springer, Heidelberg, December 2018. 1, 2, 3, 7, 8, 9, 19

- [Cou06] Jean-Marc Couveignes. Hard homogeneous spaces. Cryptology ePrint Archive, Report 2006/291, 2006. <http://eprint.iacr.org/2006/291>. 7, 8
- [Cox13] David A. Cox. Primes of the form $x^2 + ny^2$, 2013. Wiley, 2nd edition. 7
- [CS98] Ronald Cramer and Victor Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In Hugo Krawczyk, editor, *CRYPTO'98*, volume 1462 of *LNCS*, pages 13–25. Springer, Heidelberg, August 1998. 22
- [CS20] Daniele Cozzo and Nigel P. Smart. Sashimi: Cutting up CSI-FiSh secret keys to produce an actively secure distributed signing protocol. In *PQCrypto 2020 (to appear)*. Springer, 2020. 4
- [DFJP14] Luca De Feo, David Jao, and Jerome Plüt. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. In *Journal of Mathematical Cryptology*, volume 8 (3), pages 209–247, 2014. 1
- [DFMS19] Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Security of the Fiat-Shamir transformation in the quantum random-oracle model. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 356–383. Springer, Heidelberg, August 2019. 2, 7
- [DG19] Luca De Feo and Steven D. Galbraith. SeaSign: Compact isogeny signatures from class group actions. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part III*, volume 11478 of *LNCS*, pages 759–789. Springer, Heidelberg, May 2019. 2, 3, 4, 8, 12, 13
- [DH76] Whitfield Diffie and Martin Hellman. New directions in cryptography. *IEEE transactions on Information Theory*, 22(6):644–654, 1976. 3
- [DKL⁺18] Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS-Dilithium: A lattice-based digital signature scheme. *IACR TCHES*, 2018(1):238–268, 2018. <https://tches.iacr.org/index.php/TCHES/article/view/839>. 3
- [DKS18] Luca De Feo, Jean Kieffer, and Benjamin Smith. Towards practical key exchange from ordinary isogeny graphs. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part III*, volume 11274 of *LNCS*, pages 365–394. Springer, Heidelberg, December 2018. 7
- [FHK⁺18] Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Prest, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. Falcon: Fast-Fourier lattice-based compact signatures over NTRU, 2018. Submission to the NIST PQC project. 3
- [FJS14] Nils Fleischhacker, Tibor Jager, and Dominique Schröder. On tight security proofs for Schnorr signatures. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part I*, volume 8873 of *LNCS*, pages 512–531. Springer, Heidelberg, December 2014. 2
- [FS87] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *CRYPTO'86*, volume 263 of *LNCS*, pages 186–194. Springer, Heidelberg, August 1987. 1, 2, 7
- [GBL08] Sanjam Garg, Raghav Bhaskar, and Satyanarayana V. Lokam. Improved bounds on security reductions for discrete log based signatures. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 93–107. Springer, Heidelberg, August 2008. 2
- [GPS17] Steven D. Galbraith, Christophe Petit, and Javier Silva. Identification protocols and signature schemes based on supersingular isogeny problems. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 3–33. Springer, Heidelberg, December 2017. 1

- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008. [3](#)
- [HRS16] Andreas Hülsing, Joost Rijneveld, and Fang Song. Mitigating multi-target attacks in hash-based signatures. In Chen-Mou Cheng, Kai-Min Chung, Giuseppe Persiano, and Bo-Yin Yang, editors, *PKC 2016, Part I*, volume 9614 of *LNCS*, pages 387–416. Springer, Heidelberg, March 2016. [17](#)
- [JD11] David Jao and Luca De Feo. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. In Bo-Yin Yang, editor, *Post-Quantum Cryptography - 4th International Workshop, PQCrypto 2011*, pages 19–34. Springer, Heidelberg, November / December 2011. [1](#)
- [KLS18] Eike Kiltz, Vadim Lyubashevsky, and Christian Schaffner. A concrete treatment of Fiat-Shamir signatures in the quantum random-oracle model. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 552–586. Springer, Heidelberg, April / May 2018. [3](#), [5](#), [7](#)
- [Kup05a] Greg Kuperberg. Another subexponential-time quantum algorithm for the dihedral hidden subgroup problem. *TQC*, 22:20–34, 2005. [20](#)
- [Kup05b] Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. *SIAM J. Comput.*, 35(1):170–188, 2005. [20](#)
- [KW03] Jonathan Katz and Nan Wang. Efficiency improvements for signature schemes with tight security reductions. In Sushil Jajodia, Vijayalakshmi Atluri, and Trent Jaeger, editors, *ACM CCS 2003*, pages 155–164. ACM Press, October 2003. [2](#), [3](#)
- [Lyu09] Vadim Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 598–616. Springer, Heidelberg, December 2009. [1](#)
- [LZ19] Qipeng Liu and Mark Zhandry. Revisiting post-quantum Fiat-Shamir. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 326–355. Springer, Heidelberg, August 2019. [2](#), [7](#)
- [MR02] Silvio Micali and Leonid Reyzin. Improving the exact security of digital signature schemes. *Journal of Cryptology*, 15(1):1–18, 2002. [2](#)
- [Pei19] Chris Peikert. He gives c-sieves on the CSIDH. Cryptology ePrint Archive: Report 2019/725, 2019. [2](#), [20](#)
- [PV05] Pascal Paillier and Damien Vergnaud. Discrete-log-based signatures may not be equivalent to discrete log. In Bimal K. Roy, editor, *ASIACRYPT 2005*, volume 3788 of *LNCS*, pages 1–20. Springer, Heidelberg, December 2005. [2](#)
- [RS06] Alexander Rostovtsev and Anton Stolbunov. Public-key cryptosystem based on isogenies. Cryptology ePrint Archive: Report 2006/145, 2006. [7](#)
- [Sch90] Claus-Peter Schnorr. Efficient identification and signatures for smart cards. In Gilles Brassard, editor, *CRYPTO’89*, volume 435 of *LNCS*, pages 239–252. Springer, Heidelberg, August 1990. [22](#)
- [Seu12] Yannick Seurin. On the exact security of Schnorr-type signatures in the random oracle model. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 554–571. Springer, Heidelberg, April 2012. [2](#)

- [Sto10] Anton Stolbunov. Constructing public-key cryptographic schemes based on class group action on a set of isogenous elliptic curves. *Adv. Math. Commun.*, 4(2), 2010. [7](#)
- [Sto12] Anton Stolbunov. Cryptographic schemes based on isogenies, 2012. Ph.D. thesis, Norwegian University of Science and Technology. [3](#), [8](#), [9](#)
- [Unr15] Dominique Unruh. Non-interactive zero-knowledge proofs in the quantum random oracle model. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 755–784. Springer, Heidelberg, April 2015. [3](#)
- [Unr17] Dominique Unruh. Post-quantum security of Fiat-Shamir. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 65–95. Springer, Heidelberg, December 2017. [3](#)
- [YAJ+17] Youngho Yoo, Reza Azarderakhsh, Amir Jalali, David Jao, and Vladimir Soukharev. A post-quantum digital signature scheme based on supersingular isogenies. In Aggelos Kiayias, editor, *FC 2017*, volume 10322 of *LNCS*, pages 163–181. Springer, Heidelberg, April 2017. [1](#)
- [Zha12] Mark Zhandry. How to construct quantum random functions. In *53rd FOCS*, pages 679–687. IEEE Computer Society Press, October 2012. [17](#)