A Note on the Security Proof of SQIsign

Maher Mamah and David Jao

University of Waterloo, Department of Combinatorics & Optimization, Waterloo, Canada

Abstract. Aardal et al. (CRYPTO 2025) provided the first complete security proof of SQIsign; however, their reduction incurs a square-root loss in the prime characteristic due to the application of a loose bound on the min-entropy. For instance, at NIST security level I, an adversary making 2^{64} signing queries renders the security proof vacuous. In this note, we show that the min-entropy of SQIsign is optimal, namely $\mathcal{O}(1/p)$. Although this improvement does not yield full λ-bit security, we show that it preserves two-thirds of the expected bit-security. We show that this artifact comes from an information-theoretic loss in the zero-knowledge simulation of SQIsign, suggesting a new proof technique is needed to achieve full λ-bit security at the current parameters.

1 Introduction

SQIsign is a compact post-quantum signature scheme that recently advanced to round 3 in the NIST PQC call for additional digital signatures. Its security relies on the hardness of computing the endomorphism ring of supersingular elliptic curves. Aardal et al. [ABDF+25] provided the first complete EUF-CMA security proof for SQIsign using a novel Fiat-Shamir with hints framework. While their proof is an important step in establishing the foundations of SQIsign, its concrete reduction contains a quantitatively loose term arising from the analysis of the commitment distribution. This looseness stems from two factors. First, they show that the min-entropy of the commitment distribution in the identification protocol Σ_{SQI} is bounded by $\mathsf{MinEnt}(\Sigma_{\mathsf{SQI}}) \leq p^{-1/2}$ [ABDF+25, Lemma 4.1], where p is the underlying prime characteristic. Second, the bound on the weak Honest Verifier Zero Knowledge (wHVZK) distinguisher that realizes the statistical distance is of the same order [ABDF+25, Lemma 4.6]. These bounds are subsequently invoked in reducing the EUF-CMA security of SQIsign to the One Endomorphism problem OneEnd with hints [ABDF+25, Theorems 4, 6]. This indicates that the dominant term in their security bound is $\mathcal{O}(s(q+s)/\sqrt{p})$, where s and q are the number of signing and random oracle queries, $p \approx 2^{2\lambda}$ and λ is the security parameter. Furthermore, it suggests that SQIsign achieves only half the desired bit-security[1] due to the quadratic and multiplicative factors present in this bound. In particular, as noted in [ABDF+25], an adversary making $q = 2^{64}$ random oracle queries[2] renders the security proof vacuous.

In SQISign, the commitment distribution is obtained by taking the codomain of a random cyclic isogeny of prime degree N_{mix}. In the NIST submission [AAA+25], the prime degree is $N_{\text{mix}} \approx 2^{4\lambda}$, where λ is the security parameter. Thus the degree of random walk isogeny is roughly p^2 which suggests that the min-entropy of the commitment (the codomain of the walk) should be quite close to the stationary distribution rather than having a square-root loss in the aforementioned bound.

E-mail: mmamah@uwaterloo.ca (Maher Mamah), djao@uwaterloo.ca (David Jao)

[1] A cryptographic scheme has λ *bits of security* if for any running time parameter $t \leq 2^{\lambda}$, any attacker running in time t succeeds with probability $\mathcal{O}(t/2^{\lambda})$.

[2] NIST [Nat16] explicitly states that for the purpose of estimating security strengths, it may be assumed that the attacker can make 2^{64} random oracle queries

This work is licensed under a "CC BY 4.0" license.

Date of this document: 2026-09-23.

The situation in the quantum random oracle model (QROM) is even more delicate, as SQIsign achieves no provable security against adversaries making 2^{40} signing queries due to an additional square-root loss on top of the min-entropy arising from collisions in the QROM, see [ABR26, Theorem 3].
As noted in [ABDF+25], this security loss can be mitigated by increasing the commitment walk length to $N_{\text{mix}} \approx 2^{8\lambda}$. However, this roughly doubles the commitment-generation time and therefore impacts SQIsign signing performance. We note that Wesolowski [Wes26] very recently proposed a new algorithm for the isogeny path problem with time and space complexity $p^{1/3+o(1)}$. While this improves on the previous $\tilde{O}(p^{1/2})$ bound, its practical impact on concrete security remains unclear and requires further assessment by the cryptography community.

Since the distribution of sufficiently large-degree random walks in the ℓ-isogeny graph plays a fundamental role in analyzing the underlying cryptographic security of many protocols, it has become important to give a thorough analysis of this distribution.

Contribution. In this note, we show that the min-entropy of the SQIsign commitment, $\mathsf{MinEnt}(\Sigma_{\mathsf{SQI}})$, is in fact $\mathcal{O}(1/p)$ for the NIST parameter regime where $N_{\mathsf{mix}} \approx p^2$. This takes care of the quadratic factor s^2 associated with the number of signing queries in the previously mentioned security bound.
However, we are only able to reduce the multiplicative factor qs to $q\sqrt{s}$, which we heuristically argue is optimal (see Remark 3.2). This limitation stems from a wHVZK distinguisher that achieves a statistical distance bounded by $\mathcal{O}(p^{-1/2})$. While this demonstrates that, under the current proof, SQIsign cannot achieve full λ-bit security without increasing parameter sizes, it does not imply a fundamental security flaw in SQIsign itself. Rather, it suggests that the generic Fiat-Shamir with hints proof technique introduces an inherent artifact, highlighting the need for more tailored proof strategies. Finally, we discuss the implications of our improved bound for parameter choices. In particular, our proof shows that increasing the commitment degree to $N_{\text{mix}} \approx 2^{5\lambda}$ is sufficient to restore the expected λ-bit security.

2 Min-Entropy of SQIsign Commitment

We start by defining the min-entropy parameter of the commitment distribution. In what follows, we denote by Σ_{SQI} the identification protocol underlying SQIsign via the Fiat–Shamir transform, and by $\mathcal{S}_p$ the set of isomorphism classes of supersingular elliptic curves over $\mathbb{F}_{p^2}$. The commitment distribution of Σ_{SQI} is a curve obtained from a large enough cyclic walk which is independent of the statement and the witness, thus the definition of min-entropy in [ABDF+25, Definition 2.4] simplifies as follows.

Definition 1. Let D_{com} denote the commitment distribution of Σ_{SQI}. The min-entropy of D_{com}

$$\mathsf{MinEnt}(\Sigma_{\mathsf{SQI}}) = \max_{E \in \mathcal{S}_p} \Pr_{\mathsf{com} \leftarrow D_{\mathsf{com}}}[\mathsf{com} = E].$$

Let E_0 denote the fixed auxiliary supersingular elliptic curve used in the commitment generation algorithm. A commitment curve is generated by choosing a uniformly random cyclic subgroup

$$C \subset E_0[\ell^k], \qquad C \simeq \mathbb{Z}/\ell^k\mathbb{Z},$$

and outputting the codomain $E_C = E_0/C$. In the NIST submission, $k = 1$ and $\ell = N_{\text{mix}}$ is a large prime, however for the sake of generality, we will assume we are working with ℓ^k degree walk and then specialize our result to SQIsign's case. Thus, the commitment min-entropy reduces to

$$\mathsf{MinEnt}(\Sigma_{\mathsf{SQI}}) = \max_{F \in \mathcal{S}_p} \Pr_{C}[E_0/C \simeq F].$$

In contrast, the stationary distribution on supersingular j-invariants is

$$\pi(F) = \frac{24}{(p-1)|\operatorname{Aut}(F)|},$$

where $\operatorname{Aut}(F)$ denotes the automorphism group of F satisfying $|\operatorname{Aut}(F)| \in \{2, 4, 6\}$. Therefore, we have

$$\|\pi\|_\infty \leq \frac{12}{p-1}.$$

The tightness loss in [ABDF+25] stems from using an L_1 bound for the error distribution rather than a maximum probability mass L_∞-bound for the commitment. More precisely, Aardel et al. [ABDF+25, Lemma 4.1] bounded the min-entropy using:

$$\max_F \Pr_C[E_0/C \simeq F] \leq ||\pi||_\infty + \Delta(D_{\mathsf{com}}, \pi),$$

where

$$\Delta(D_{\mathsf{com}}, \pi) = \frac{1}{2}\sum_E |\Pr[E \leftarrow D_{\mathsf{com}}] - \Pr[E \leftarrow \pi]| = \mathcal{O}(1/\sqrt{p})$$

using [BCC+23, Theorem 11]. Although mathematically valid, this bound is overly pessimistic and results in the aforementioned loose bound, as the single maximum probability is effectively scaled by an L_1 total mass factor of $\sqrt{p}$. Our motivation, therefore, is to derive a bound for the min-entropy using the L_∞-norm rather than the L_1-norm.

Let B denote the adjacency matrix of the associated supersingular ℓ-isogeny graph, and let B_{ℓ^k} denote the Brandt matrix whose (E, F)-entry counts the non-backtracking, cyclic, degree-ℓ^k isogenies from E to F. Then,

$$\Pr[E_0/C \simeq F] = \frac{B_{\ell^k}(E_0, F)}{\psi(\ell^k)},$$

where $\psi(\ell^k) = \ell^k + \ell^{k-1}$ is the number of cyclic subgroups of order ℓ^k. Consequently, the min-entropy is given by:

$$\mathsf{MinEnt}(\Sigma_{\mathsf{SQI}}) = \max_{F \in \mathrm{SS}p} \frac{B_{\ell^k}(E_0, F)}{\psi(\ell^k)}.$$

We are now ready to state our main theorem about long random walks in ℓ-isogeny graph.

Theorem 1. *For all supersingular curves $E, F \in \mathcal{S}_p$,*

$$\left|\frac{B_{\ell^k}(E, F)}{\psi(\ell^k)} - \pi(F)\right| \leq \frac{\sqrt{3}(k+1)}{\ell^{k/2}}.$$

Proof. The non-backtracking matrices B_{ℓ^k} satisfy the Hecke recurrence [Mur03, section 6]

$$B_{\ell^{k+1}} = B_\ell B_{\ell^k} - \ell B_{\ell^{k-1}}.$$

Hence every eigenvalue α_k of B_{ℓ^k} satisfies

$$\alpha_{k+1} = \alpha \cdot \alpha_k - \ell\alpha_{k-1},$$

where α is the corresponding eigenvalue of B_ℓ.

The supersingular ℓ-isogeny graph is Ramanujan, and Deligne's bound [Del80] yields

$$|\alpha| \leq 2\sqrt{\ell},$$

for every eigenvalue α different from the trivial eigenvalue $\ell + 1$. Writing

$$\alpha = 2\sqrt{\ell}\cos\theta,$$

and defining $\beta_k = \ell^{-k/2}\alpha_k$, the recurrence becomes

$$\beta_{k+1} = 2\cos\theta \cdot \beta_k - \beta_{k-1},$$

with $\beta_0 = 1$. Thus

$$\beta_k = U_k(\cos\theta),$$

where U_k denotes the Chebyshev polynomial of the second kind. Since

$$|U_k(x)| \leq k + 1 \qquad \text{for } x \in [-1, 1],$$

we obtain

$$|\alpha_k| \leq (k+1)\ell^{k/2}.$$

Note that the trivial eigenvalue of B_{ℓ^k} equals $\psi(\ell^k) = \ell^k + \ell^{k-1}$. Define the normalized operator

$$P_k = \frac{1}{\psi(\ell^k)} B_{\ell^k},$$

which has same eigenvectors as B_ℓ and a trivial eigenvalue of 1. This P_k can be thought of as a Markov operator for the non-backtracking walks. Therefore every nontrivial eigenvalue of P_k is bounded by

$$|\lambda| \leq \frac{(k+1)\ell^{k/2}}{\ell^k + \ell^{k-1}} \leq \frac{k+1}{\ell^{k/2}}.$$

We now convert the spectral bound $\rho_k = \dfrac{k+1}{\ell^{k/2}}$ into an entrywise estimate for the distributions. By the dual-isogeny symmetry of Brandt matrices, the operator P_k satisfies the reversibility relation

$$\pi(E)P_k(E, F) = \pi(F)P_k(F, E).$$

Thus, P_k is self-adjoint with respect to the weighted inner product

$$\langle f, g\rangle_\pi = \sum_{X \in \mathcal{S}_p} \pi(X) f(X) g(X).$$

Consequently, on the space orthogonal to the trivial eigenvector $\mathbf{1} = (1, 1, \ldots, 1)$, namely the mean-zero subspace $L_0^2(\pi) := \{f : \sum_X \pi(X)f(X) = 0\}$, the spectral bound implies:

$$\|P_k f\|_{2,\pi} \leq \rho_k \|f\|_{2,\pi}. \tag{1}$$

Fix $F \in \mathcal{S}_p$ and define the test function $f_F(X) = \frac{\mathbf{1}_{X=F}}{\pi(F)}$. Since $\sum_{X \in \mathcal{S}_p} \pi(X) f_F(X) = 1$, the shifted function $f_F - 1$ lies in $L_0^2(\pi)$. Applying P_k to this function yields:

$$P_k(f_F - 1)(E) = \frac{P_k(E, F)}{\pi(F)} - 1.$$

For any h on $\mathcal{S}_p$, the standard point-evaluation bound gives $|h(E)| \leq \|h\|_{2,\pi}/\sqrt{\pi(E)}$. Applying this to $h = P_k(f_F - 1)$ and utilizing the spectral bound, we obtain:

$$\left| \frac{P_k(E, F)}{\pi(F)} - 1 \right| \leq \frac{\|P_k(f_F - 1)\|_{2,\pi}}{\sqrt{\pi(E)}} \leq \frac{\rho_k \|f_F - 1\|_{2,\pi}}{\sqrt{\pi(E)}}.$$

A direct computation gives $\|f_F - 1\|_{2,\pi}^2 = \frac{1}{\pi(F)} - 1 < \frac{1}{\pi(F)}$. Substituting this bound into the inequality above and multiplying through by $\pi(F)$, we arrive at the final entrywise estimate:

$$|P_k(E,F) - \pi(F)| \leq \rho_k \sqrt{\frac{\pi(F)}{\pi(E)}}.$$

For every supersingular curve E, we have $|\operatorname{Aut}(E)| \in \{2,4,6\}$, and therefore

$$\sqrt{\frac{\pi(F)}{\pi(E)}} \leq \sqrt{3}.$$

This completes the proof upon recalling that $P_k(E,F) = \dfrac{B_{\ell^k}(E,F)}{\psi(\ell^k)}$. □

Corollary 1. *The commitment min-entropy of* Σ_{SQI} *satisfies*

$$\mathsf{MinEnt}(\Sigma_{\mathsf{SQI}}) \leq \frac{16}{p}, \qquad p > 500$$

where the degree of the commitment isogeny N_{mix} *is the smallest prime greater than* p^2*.*

Proof. Applying Theorem 1 with $E = E_0$ gives

$$\Pr[E_0/C \simeq F] \leq \pi(F) + \sqrt{3}\frac{k+1}{\ell^{k/2}}.$$

Taking the maximum over all $F \in \mathcal{S}_p$ yields

$$\mathsf{MinEnt}(\Sigma_{\mathsf{SQI}}) \leq \|\pi\|_\infty + \sqrt{3}\frac{k+1}{\ell^{k/2}}.$$

Using $\|\pi\|_\infty \leq \dfrac{12}{p-1}$ and the fact that $N_{\mathrm{mix}} = \ell > p^2$ is a prime, i.e. $k = 1$, we obtain the desired result by standard algebraic manipulation for $p > 500$. □

3 Refining the Security Proof of SQIsign

The preceding section improves the commitment min-entropy parameter by replacing the L^1-based estimate of the commitment distribution with an L_∞-bound. This removes the square-root loss from the min-entropy term. However, this is not the only place where the commitment distribution enters the security proof of SQIsign. From this point on, we assume some familiarity with the structure of the security proof in [ABDF+25], particularly sections 4 and 5.

3.1 A tighter bound for the distribution marginals

A separate statistical loss appears in the weak honest-verifier zero-knowledge (wHVZK) simulation [ABDF+25, Lemma 4.6]. In the real protocol, the commitment curve is distributed according to D_{com}, whereas in the simulator used in the proof, it is sampled from the stationary distribution π. For q simulated signing transcripts, their proof requires a tight bound for

$$\Delta(D_{\mathsf{com}}^{\otimes q}, \pi^{\otimes q}) = \frac{1}{2}\sum_{\mathbf{j}\in\mathcal{S}_p^q} \left|\Pr[\mathbf{j} \leftarrow D_{\mathsf{com}}^{\otimes q}] - \Pr[\mathbf{j} \leftarrow \pi^{\otimes q}]\right|.$$

The original proof bounds this mismatch using a linear hybrid argument, yielding a term of the form

$$2q\Delta(D_{\mathsf{com}}, S_j).$$

This bound is valid but loses the product structure of the commitment distribution. In this section, we isolate this purely statistical step and replace the linear hybrid estimate with a product χ^2-bound. We note that this gain is modest, as it replaces the q-loss with a $\sqrt{q}$-loss. We start by defining the χ^2-divergence and follow with a general lemma for the multiplicativity of the χ^2 under products.

Definition 2 (χ^2-divergence)**.** Let $P \ll Q$ be probability distributions on a finite set V. The χ^2-divergence of P from Q is

$$\chi^2(P\|Q) := \sum_{x\in V} \frac{(P(x)-Q(x))^2}{Q(x)} = \mathbb{E}_{x\sim Q}\left[\left(\frac{P(x)}{Q(x)}-1\right)^2\right].$$

Lemma 1. *Let D and Π be distributions on a finite set V, with $D \ll \Pi$. Then, for every integer $q \geq 1$,*

$$\Delta(D^{\otimes q}, \Pi^{\otimes q}) \leq \frac{1}{2}\sqrt{(1+\chi^2(D\|\Pi))^q - 1}.$$

Proof. Let $P := D^{\otimes q}$ and $Q := \Pi^{\otimes q}$. By Cauchy–Schwarz,

$$\|P-Q\|_1 = \sum_{\mathbf{x}\in V^q} Q(\mathbf{x})\left|\frac{P(\mathbf{x})}{Q(\mathbf{x})}-1\right| \leq \left(\sum_{\mathbf{x}\in V^q} Q(\mathbf{x})\left(\frac{P(\mathbf{x})}{Q(\mathbf{x})}-1\right)^2\right)^{1/2}.$$

The right-hand side is

$$\sqrt{\chi^2(P\|Q)}.$$

Since P and Q are product distributions,

$$1+\chi^2(P\|Q) = \sum_{\mathbf{x}\in V^q} Q(\mathbf{x})\left(\frac{P(\mathbf{x})}{Q(\mathbf{x})}\right)^2 = \left(\sum_{x\in V}\Pi(x)\left(\frac{D(x)}{\Pi(x)}\right)^2\right)^q.$$

Thus

$$1+\chi^2(P\|Q) = \left(1+\chi^2(D\|\Pi)\right)^q,$$

and therefore

$$\|D^{\otimes q}-\Pi^{\otimes q}\|_1 \leq \sqrt{(1+\chi^2(D\|\Pi))^q - 1}.$$

The total variation statement follows from $\Delta(P,Q) = \frac{1}{2}\|P-Q\|_1$. □

It remains to estimate the quantity $\chi^2(D_{\mathsf{com}}\|S_j)$. This estimate uses our bound from Theorem 1, but now in $L^2(\pi)$ rather than in point-wise norm.

Proposition 1. *Let D_{com} be the SQIsign commitment distribution obtained from a uniform cyclic prime degree-N_{mix} isogeny from E_0, and let π be the stationary distribution on supersingular j-invariants. Then*

$$\chi^2(D_{\mathsf{com}}\|\pi) \leq \frac{2}{3p}$$

Proof. Let $P_k = \dfrac{B_{\ell^k}}{\psi(\ell^k)}$; then $D_{\mathsf{com}}(F) = P_k(E_0, F)$. Define

$$f_{E_0}(F) = \frac{\mathbf{1}_{F=E_0}}{\pi(E_0)}.$$

This is the density of the point mass at E_0 with respect to π. Following the proof of Theorem 1 and by reversibility of P_k with respect to π,

$$\frac{D_{\mathsf{com}}}{\pi} - 1 = P_k(f_{E_0} - 1).$$

Hence

$$\chi^2(D_{\mathsf{com}}\|\pi) = \left\| \frac{D_{\mathsf{com}}}{\pi} - 1 \right\|_{2,\pi}^2 = \|P_k(f_{E_0} - 1)\|_{2,\pi}^2.$$

On the mean-zero subspace, the bound in (1) yields

$$\|P_k h\|_{2,\pi} \le \rho_k \|h\|_{2,\pi}, \qquad \rho_k \le \frac{k+1}{\ell^{k/2}}.$$

Since $f_{E_0} - 1$ has mean zero,

$$\chi^2(D_{\mathsf{com}}\|\pi) \le \rho_k^2 \|f_{E_0} - 1\|_{2,\pi}^2.$$

Since SQIsign uses $j(E_0) = 1728$, we have $|\mathrm{Aut}(E_0)| = 4$. We thus obtain

$$\chi^2(D_{\mathsf{com}}\|\pi) \le \frac{(k+1)^2}{\ell^k} \cdot \frac{p}{6}$$

upon recalling that $||f_{E_0} - 1||_{2,\pi}^2 \le 1/\pi(E_0) \le p/6$ from the proof of Theorem 1. Setting $N_{\mathrm{mix}} = \ell > p^2$ concludes the proof. □

Corollary 2. *Let D_{com} be the SQIsign commitment distribution obtained from a cyclic walk of prime degree $N_{\mathrm{mix}} > p^2$ and let π be the stationary distribution. Then, for $q \le 3p/2$, we have*

$$\Delta(D_{\mathsf{com}}^{\otimes q}, \pi^{\otimes q}) \le \sqrt{\frac{q}{3p}}$$

Proof. By Lemma 1 and Proposition 1, we have

$$2\Delta(D_{\mathsf{com}}^{\otimes q}, \pi^{\otimes q}) \le \sqrt{\left(1 + \frac{2}{3p}\right)^q - 1}.$$

Let $\delta = \frac{2}{3p}$. Observe that $(1+\delta)^q \le e^{q\delta}$. For polynomial $q = o(p)$ and our parameter regime with exponential p, we have $0 \le q\delta \le 1$. Applying the inequality $e^x - 1 \le 2x$ for $x \in [0, 1]$, we obtain

$$2\Delta(D_{\mathsf{com}}^{\otimes q}, \pi^{\otimes q}) \le \sqrt{2q\delta} = 2\sqrt{\frac{q}{3p}};$$

the result follows. □

3.2 Effect on SQIsign's Security proof

Since we're interested solely in enhancing the bound present in the key technical result [ABDF+25, Theorem 2] we avoid reproducing the entire proof structure as present in [ABDF+25]. We start with the definition of weak Honet Verifier Zero-Knowledge wHVZK for Σ protocols. It is basically the HVZK property for identification protocols where simulators are equipped with hints oracle.

Definition 3 (Hint-assisted wHVZK advantage)**.** Let Σ be an identification protocol, let $\mathcal{H}$ be a hint distribution, and let $\mathcal{S}$ be a hint-assisted simulator, namely a simulator

with oracle access to hints sampled from $\mathcal{H}$. The q-hint-assisted wHVZK advantage of an adversary $\mathcal{A}$ is

$$\mathsf{Adv}^{\mathsf{hint\text{-}wHVZK}}_{\Sigma,\mathcal{H},\mathcal{S}}(\mathcal{A},q) := \left|\Pr\left[\mathcal{A}(\mathsf{Real}_\Sigma(1^\lambda,q))=1\right] - \Pr\left[\mathcal{A}(\mathsf{HintSim}_{\Sigma,\mathcal{H},\mathcal{S}}(1^\lambda,q))=1\right]\right|.$$

Here $\mathsf{Real}_\Sigma(1^\lambda,q)$ denotes q honestly generated transcripts of Σ for the same statement, and $\mathsf{HintSim}_{\Sigma,\mathcal{H},\mathcal{S}}(1^\lambda,q)$ denotes q transcripts generated by the simulator $\mathcal{S}$ using hints sampled from $\mathcal{H}$.

Next, we use the same security experiments and advantage notions as Aardal et al. We recall only their informal notions in order to fix notation and refer the reader to the relative sections therein, namely section 4 and 5.

- $\mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{SIG}[\Sigma_{\mathsf{SQI}}]}(\mathcal{A})$ denotes the usual existential unforgeability advantage of an adversary $\mathcal{A}$ against the Fiat–Shamir signature scheme $\mathsf{SIG}(\Sigma_{\mathsf{SQI}})$ obtained from the identification protocol Σ_{SQI}. The adversary has access to a random oracle and a signing oracle, and wins by producing a valid signature on a message that was not previously signed.
- $\mathsf{Adv}^{\mathsf{hint\text{-}wHVZK}}_{\Sigma_{\mathsf{SQI}},\mathcal{H}^{\mathsf{sim}},\mathcal{S}_{\mathsf{SQI}}}(\mathcal{A},q)$ denotes the advantage of an adversary $\mathcal{A}$ in distinguishing q real transcripts of the identification protocol from q transcripts generated by the hint-assisted simulator $\mathcal{S}_{\mathsf{SQI}}$ using the simulated hint distribution $\mathcal{H}^{\mathsf{sim}}$.
- $\mathsf{Adv}^{\mathsf{rel}}_{R_{\mathsf{SQI}}}(\mathcal{B})$ denotes the advantage of an algorithm $\mathcal{B}$ in solving the soundness relation associated with the SQIsign identification protocol.
- $\mathsf{Adv}^{\mathsf{hint\text{-}rel}}_{(R_{\mathsf{OneEnd}},\mathcal{H}^{\mathsf{sim}})}(\mathcal{E},s)$ denotes the corresponding hinted-relation advantage for the OneEnd relation with s simulated hints.

We do not modify any of these experiments, the hint distribution, or the simulator. We aim at analyzing the statistical estimates for the commitment distribution that enter the wHVZK and EUF-CMA bounds.

We now explain how the preceding product bound modifies the statistical term in Lemma 4.6 of Aardal et al. and consequently the subsequent theorems.

In the real experiment, the q commitments are distributed as $D_{\mathsf{com}}^{\otimes q}$, whereas in the simulated experiment they are distributed as $\pi^{\otimes q}$. The original proof bounds the resulting mismatch by the linear hybrid estimate

$$2\Delta(D_{\mathsf{com}}^{\otimes q},\pi^{\otimes q}) \le 2q\Delta(D_{\mathsf{com}},\pi).$$

Using Corollary 2 we have a refined bound for the advantage of any wHVZK adversary in [ABDF+25, Lemma 4.6]. We specifically refer the readers to experiment 3 and Algorithm 4 in [ABDF+25] for the hint distribution $\mathcal{H}^{\mathsf{sim}}$ and the hint-assisted simulator $\mathcal{S}_{\mathsf{SQI}}$ used in SQIsign.

Proposition 2 (Refined wHVZK statistical term)**.** *Let $\mathcal{S}_{\mathsf{SQI}}$ be the hint-assisted simulator receiving hints according to the $\mathcal{H}^{\mathsf{sim}}$ distribution. For any $q = \mathsf{poly}(\lambda)$ and any PPT adversary $\mathcal{A}$ against the hinte-assisted wHVZK of Σ_{SQI} with q hints, there exists a PPT adversart $\mathcal{B}$ such that*

$$\mathsf{Adv}^{\mathsf{hint\text{-}wHVZK}}_{\Sigma_{\mathsf{SQI}},\mathcal{H}^{\mathsf{sim}},S_{\mathsf{SQI}}}(\mathcal{A},q) \le \mathsf{Adv}^{\mathsf{rel}}_{R_{\mathsf{SQI}}}(\mathcal{B}) + 2\sqrt{\frac{q}{3p}}.$$

Proof. The proof is identical to the proof of Lemma 4.6 in Aardal et al. up to the point where the commitment marginals are compared in equation (2) therein. At that point, the real and simulated commitment distributions on q transcripts are

$$D_{\mathsf{com}}^{\otimes q} \qquad \text{and} \qquad \pi^{\otimes q},$$

and the statistical contribution is bounded by $2\Delta(D_{\mathsf{com}}^{\otimes q},\pi^{\otimes q})$. By applying Corollary 2 to these product distributions, the claimed bound follows. □

Optimality of the statistical bound We briefly mention why we believe the term $\sqrt{q/p}$ is optimal for Proposition 2. Lemma 4.6 of Aardal et al. or equivalently Proposition 2 is a computational wHVZK statement, since the adversary is restricted to be PPT. However, the proof bounds the adversary's advantage by an information-theoretic distance between the real and simulated commitment marginals. The term $\sqrt{q/p}$ should be viewed as a limitation of this statistical simulator comparison. Indeed, an unbounded adversary can distinguish with an advantage achieving the statistical distance $\Delta(D_{\mathsf{com}}^{\otimes q}, \pi^{\otimes q})$. Since a single instance $\Delta(D_{\mathsf{com}}, \pi)$ is of order $\asymp 1/\sqrt{p}$, the q-fold L_1-distance $\Delta(D_{\mathsf{com}}^{\otimes q}, \pi^{\otimes q})$ is heuristically of order $\sqrt{q/p}$. This does not necessarily imply the existence of a PPT distinguisher with the same advantage. Rather, removing this term requires exploiting computational structure or avoiding full product-distribution comparisons.

We note, however, that this limitation does not indicate a fundamental flaw in SQIsign. Rather, it highlights that the generic Fiat–Shamir-with-hints proof framework introduces an inherent proof artifact. Consequently, achieving full λ-bit security level requires a more tailored security analysis that exploits the specific structure of SQIsign rather than relying on the generic framework.

To state our technical theorem, we also give an informal definition of the following advantages:

- $\mathsf{Adv}^{\mathsf{hint\text{-}OneEnd}_p}(\mathcal{B}, q)$ denotes the success probability of $\mathcal{B}$ in the q-hint-OneEnd$_p$ problem: given a stationary supersingular curve E and q uniform hints for E, output a non-scalar endomorphism of E in efficient representation.
- $\mathsf{Adv}^{\mathsf{hint\text{-}dist}}(\mathcal{D}, q)$ denotes the advantage of an adversary $\mathcal{D}$ in distinguishing the real hint distribution from the pushed-hint distribution. More precisely, $\mathcal{D}$ is given a tuple $(E, h_1, \ldots, h_q)$ sampled either from $\mathsf{RealHints}(1^\lambda, q)$ or from $\mathsf{PushedHints}(1^\lambda, q)$ (where the hints are pushed forward by another isogeny), each with probability $1/2$, and must decide from which distribution the tuple was sampled.

Theorem 2. *For any PPT adversary $\mathcal{A}$ against* EUF-CMA *of* $\mathsf{SIG}(\Sigma_{\mathsf{SQI}})$ *there exists an expected PPT algorithm $\mathcal{E}$ with*

$$\mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{SIG}[\Sigma_{\mathsf{SQI}}]}(\mathcal{A}) \leq (q+1)\left(2\mathsf{Adv}^{\mathsf{hint\text{-}rel}}_{(R_{\mathsf{OneEnd}}, \mathcal{H}^{\mathsf{sim}})}(\mathcal{E}, s) + 2\sqrt{\frac{s}{3p}} + 2^{-e_{\mathsf{chl}}}\right) + \frac{16(q+s+1)s}{p},$$

where s and q are the number of queries to the signing and random oracles respectively and $e_{\mathsf{chl}} = \lambda$ is the security parameter.

Proof. The proof proceeds analogously to that of [ABDF+25, Theorem 2], except we replace their weak min-entropy bound with our refined bound from Corollary 1, and substitute their Lemma 4.6 with our refined Proposition 2. □

With Theorem 2 in hand we are able to reduce the EUF-CMA security of SQIsign to the one endomorphism problem OneEnd with hints.

Theorem 3. *For any PPT adversary $\mathcal{A}$ against* EUF-CMA *of* $\mathsf{SIG}(\Sigma_{\mathsf{SQI}})$*, there exist expected PPT algorithms $\mathcal{B}$ and $\mathcal{D}$ such that*

$$\mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{SIG}[\Sigma_{\mathsf{SQI}}]}(\mathcal{A}) \leq (q+1)\left(2\mathsf{Adv}^{\mathsf{hint\text{-}OneEnd}_p}(\mathcal{B}, s) + 2\mathsf{Adv}^{\mathsf{hint\text{-}dist}}(\mathcal{D}, s) + \frac{2\sqrt{s}+1}{\sqrt{p}} + 2^{-e_{\mathsf{chl}}}\right) + \frac{16(q+s+1)s}{p},$$

where s and q are the numbers of signing-oracle and random-oracle queries, respectively, and $e_{\mathsf{chl}} = \lambda$ is the security parameter.

Proof. By Theorem 2, there exists an expected PPT algorithm $\mathcal{E}$ such that

$$\mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{SIG}[\Sigma_{\mathsf{SQI}}]}(\mathcal{A}) \leq (q+1)\left(2\mathsf{Adv}^{\mathsf{hint\text{-}rel}}_{(R_{\mathsf{OneEnd}},\mathcal{H}^{\mathsf{sim}})}(\mathcal{E},s) + 2\sqrt{\frac{s}{3p}} + 2^{-e_{\mathsf{chl}}}\right) + \frac{16(q+s+1)s}{p}.$$

As in Aardal et al., we replace the simulated hint distribution $\mathcal{H}^{\mathsf{sim}}$ by the uniform hint distribution $\mathcal{H}^{\mathsf{unif}}$. This gives algorithms $\mathcal{B}$ and $\mathcal{D}$ such that

$$\mathsf{Adv}^{\mathsf{hint\text{-}rel}}_{(R_{\mathsf{OneEnd}},\mathcal{H}^{\mathsf{sim}})}(\mathcal{E},s) \leq \mathsf{Adv}^{\mathsf{hint\text{-}OneEnd}}_{\mathsf{OneEnd}_p}(\mathcal{B},s) + \mathsf{Adv}^{\mathsf{hint\text{-}dist}}(\mathcal{D},s) + \Delta(E_{\mathsf{pk}},\pi).$$

Since the statistical distance satisfies

$$\Delta(E_{\mathsf{pk}},\pi) \leq \frac{1}{2\sqrt{p}},$$

we get

$$2\mathsf{Adv}^{\mathsf{hint\text{-}rel}}_{(R_{\mathsf{OneEnd}},\mathcal{H}^{\mathsf{sim}})}(\mathcal{E},s) \leq 2\mathsf{Adv}^{\mathsf{hint\text{-}OneEnd}}_{\mathsf{OneEnd}_p}(\mathcal{B},s) + 2\mathsf{Adv}^{\mathsf{hint\text{-}dist}}(\mathcal{D},s) + \frac{1}{\sqrt{p}}.$$

Substituting this into the previous bound proves the theorem. □

Effect on the security parameters. We now discuss the provable security of SQIsign in light of Theorem 3. As noted in the introduction, the quadratic dependence on the number of signatures in the bound of Aardal et al. is problematic: after normalization by $\sqrt{p} \approx 2^{\lambda}$, it yields only about half of the desired bit-security, since $s^2/\sqrt{p} \leq s/2^{\lambda/2}$. Our refined min-entropy bound replaces this contribution by

$$\frac{s^2+qs}{p} = \mathcal{O}\left(\frac{q^2}{p}\right),$$

in the relevant regime $s \leq q$. In particular, when $p \approx 2^{2\lambda}$ and $q \leq 2^{\lambda}$, this term is at most $\mathcal{O}(q/2^{\lambda})$.

The remaining bottleneck is the term $q\sqrt{s}/\sqrt{p}$. Under the same assumption $s \leq q$, we obtain

$$\frac{q\sqrt{s}}{\sqrt{p}} \leq \frac{q^{3/2}}{2^{\lambda}}.$$

Consequently, the reduction only supports the range $q \lesssim 2^{2\lambda/3}$, preserving roughly two thirds of the target bit-security. Note that this $(1/\sqrt{p})$-term comes from the statistical distance $\Delta(D_{\mathsf{com}},\pi)$ where the commitment isogeny has degree $N_{\mathrm{mix}} \approx 2^{4\lambda}$. Thus, within the existing proof strategy, the most plausible way to recover full λ-bit security is to increase the commitment degree to $N_{\mathrm{mix}} \approx 2^{5\lambda}$, rather than $2^{8\lambda}$ as suggested by [ABDF+25]. This follows directly from the fact that such N_{mix} gives a bound of order $2^{-3\lambda}$ in Proposition 1 and hence in the subsequent estimates. Moreover, it increases the cost of commitment generation by roughly 25%, which is a relatively modest overhead.

It would be also interesting to see if the $(q+1)$ term in the statement of Theorem 3 can be improved. However, this factor arises from the generic black-box Fiat–Shamir proof, and therefore seems unlikely to be removed without a more tailored analysis. Nevertheless, compared with [ABDF+25, Theorem 3], this gives a nontrivial improvement in the multi-signature regime. For instance, when $q = 2^{64}$ and $\lambda = 128$, the dominant term is approximately 2^{-32}. Thus the reduction still yields 32 bits of security, provided that $\mathsf{Adv}^{\mathsf{hint\text{-}OneEnd}_p}$ and $\mathsf{Adv}^{\mathsf{hint\text{-}dist}}$ remain sufficiently small.[3]

[3] As noted in [AAA+25, Section 10.2], the classical advantage $\mathsf{Adv}^{\mathsf{hint\text{-}OneEnd}_p}$ of OneEnd is believed to be $\mathcal{O}(q/2^{\lambda})$, and $\mathsf{Adv}^{\mathsf{hint\text{-}dist}}$ is argued to be computationally comparable.

Unfortunately we don't have EUF-CMA security analysis of SQIsign in the plain quantum random oracle model QROM. However, we have such results in the QROM and algebraic isogeny model (QROM+AIM). In particular, Aardal et. al [ABR26, Theorem 3] show the following bound for the EUF-CMA security of SQIsign

$$\begin{aligned}\mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{SIG}[\Sigma_{\mathsf{SQI}}]}(\mathcal{A}) \leq 4(q+1)^2 \cdot \mathsf{Adv}^{\mathsf{hint\text{-}OneEnd}_p}(\mathcal{B}) + \mathsf{Adv}^{\mathsf{hint\text{-}dist}}(\mathcal{D}) \\ + \frac{4(q+1)^2}{2^{e_{\mathrm{ch1}}}} + \frac{2(q+1)^2+s}{\sqrt{p}} + \frac{2s\sqrt{s+q+1}}{\sqrt[4]{p}}.\end{aligned} \tag{2}$$

Our refined min-entropy result thus replaces the term $\frac{s\sqrt{s+q+1}}{\sqrt[4]{p}}$ by $\frac{s\sqrt{s+q+1}}{\sqrt{p}}$. However due to the term $(q+1)^2$ in the above bound, which enters through generic bounds on Fiat-Shamir signatures in [DFM20, Theorem 2] and [ABR26, Lemma C.3], we are only able to obatin half of the desired bit-security. Assuming $s \ll q$, we obtain a bound of $O(q^2/\sqrt{p}) = O(q^2/2^\lambda)$ which only supports the regime $q \lesssim 2^{\lambda/2}$. Nevertheless our result still improves on the upper bound (2) which only supports the case $q \lesssim 2^{\lambda/3}$.

Finally, we note that the recent cryptanalytic algorithm for the isogeny path problem [Wes26] will likely affect security parameters. However, due to its exponential space complexity, further assessment is needed to determine concrete parameter choices to achieve λ-bit security.

Acknowledgment. The authors acknowledge support of the NSERC Alliance Consortia Quantum Grant "Accelerating the transition to quantum-resistant cryptography" (ALLRP 578463–2022), and the NSERC Discovery Grant program.

References

[AAA+25] Marius A. Aardal, Gora Adj, Diego F. Aranha, Andrea Basso, Isaac Andrés Canales Martínez, Jorge Chávez-Saab, Maria Corte-Real Santos, Pierrick Dartois, Luca De Feo, Max Duparc, Jonathan Komada Eriksen, Tako Boris Fouotsa, Décio Luiz Gazzoni Filho, Basil Hess, David Kohel, Antonin Leroux, Patrick Longa, Luciano Maino, Michael Meyer, Kohei Nakagawa, Hiroshi Onuki, Lorenz Panny, Sikhar Patranabis, Christophe Petit, Giacomo Pope, Krijn Reijnders, Damien Robert, Francisco Rodríguez-Henríquez, Sina Schaeffler, and Benjamin Wesolowski. SQIsign. Technical report, National Institute of Standards and Technology, 2025. URL: https://sqisign.org.

[ABDF+25] M. A. Aardal, A. Basso, L. De Feo, S. Patranabis, and B. Wesolowski. A complete security proof of SQIsign. In Y. Tauman Kalai and S. F. Kamara, editors, *Advances in Cryptology – CRYPTO 2025*, volume 16005 of *Lecture Notes in Computer Science*. Springer, Cham, 2025. doi:10.1007/978-3-032-01887-8_7.

[ABR26] M. A. Aardal, A. Basso, and D. Riepel. The algebraic isogeny model: A general model with applications to SQIsign and key exchanges. In J. Daemen and E. Thomé, editors, *Advances in Cryptology – EUROCRYPT 2026*, volume 16544 of *Lecture Notes in Computer Science*. Springer, Cham, 2026. doi:10.1007/978-3-032-25327-9_21.

[BCC+23] Andrea Basso, Giulio Codogni, Deirdre Connolly, Luca De Feo, Tako Boris Fouotsa, Guido Maria Lido, Travis Morrison, Lorenz Panny, Sikhar Patranabis, and Benjamin Wesolowski. Supersingular curves you can trust. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023*, pages 405–437, Cham, 2023. Springer Nature Switzerland.

[Del80] Pierre Deligne. La conjecture de weil. II. *Publications Mathématiques de l'IHÉS*, 52:137–252, 1980.

[DFM20] Jelle Don, Serge Fehr, and Christian Majenz. The measure-and-reprogram technique 2.0: Multi-round fiat-shamir and more. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology – CRYPTO 2020, Part III*, volume 12172 of *Lecture Notes in Computer Science*, pages 602–631, Cham, August 2020. Springer. doi:10.1007/978-3-030-56877-1_21.

[Mur03] M. Ram Murty. Ramanujan graphs. *Journal of the Ramanujan Mathematical Society*, 18(1):33–52, 2003.

[Nat16] National Institute of Standards and Technology. Post-quantum cryptography standardization: Security (evaluation criteria). https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/evaluation-criteria/security-(evaluation-criteria), 2016. Accessed: 2026-09-23.

[Wes26] Benjamin Wesolowski. The supersingular isogeny problem in time and memory $p^{1/3+o(1)}$. Cryptology ePrint Archive, Paper 2026/1486, 2026. URL: https://eprint.iacr.org/2026/1486.