

SQISIGN

Modifications introduced in the third round

September 1, 2026

The changes from the round-2 submission have been relatively minor by SQISIGN standards. The most important ones are:

- (1) **Increased parameter sizes** in reaction to the recently discovered attack [Wes26]. The increase in size leads to a $1.5\text{-}2\times$ slow-down in verification, but is more than offset in key-generation and signing by the improvements listed below.
- (2) **Entirely redesigned quaternions.** The previous iteration relied on standard multiple-precision integer linear algebra algorithms for computations in quaternion algebras. Bounding the size of integers in the computations was notoriously difficult and early attempts at fixed-precision integer arithmetic introduced unbearable slow-downs [KLKL26]. The present iteration introduces a new representation for quaternionic ideals called *inert representation* [Ler26; LS26]. Carefully redesigning algorithms, we are now able to set tight bounds on the growth of integers, replacing multiple-precision with fixed-precision without incurring performance losses.
- (3) **Dropped floating point arithmetic and improved lattice reduction.** New lattice reduction algorithms tailored to the quaternion algebra $B_{p,\infty}$ drop the requirement for floating point types and improve performance. Their outputs enforce canonical forms for quaternion lattices, simplifying reproducibility of test vectors and enabling the possibility of deterministic signatures.
- (4) **Adopted a new ideal-to-isogeny algorithm.** As introduced in [BCE+25] and improved in [DLS26], the new algorithm offers considerable speed-ups in key-generation and signing, reduces complexity and memory footprint, and brings failure cases down to a statistically negligible number.
- (5) **Improved isogeny formulas** introduced in [Dar25; Dup25] simplify the code and give moderate speed-ups.
- (6) **Changed the distribution of responses in the Σ -protocol.** This is the only cryptographically-relevant change. The round-2 version of SQISIGN defined the challenge isogeny φ_{chl} to be a 2-isogeny walk of length $\approx 2\lambda$, then sampled the response isogeny φ_{rsp} as a random isogeny of degree $\approx \sqrt{p}$. This allowed for the possibility that φ_{chl} and φ_{rsp} share a common suffix, significantly complicating the computation and representation of φ_{rsp} . Additionally, soundness of the Σ -protocol could only be ensured by having $2\lambda > \log(p)/2$.

In the present revision, we have modified φ_{chl} to have length λ and φ_{rsp} to have **odd** degree $\approx \sqrt{p^{2\lambda/2}}$, ensuring that φ_{chl} and φ_{rsp} have no common factor. Soundness follows

immediately and the computation of φ_{rsp} is more streamlined, leading to a noticeable reduction in code size and simpler algorithms. Overall, this change leads to an increase in signature size by $\lambda/8$ bytes and a minor speed-up in verification.

This change also affects the provable security of SQISIGN. In both cases, unforgeability reduces to the endomorphism ring problem with *hints* whose degrees are distributed like $\deg \varphi_{\text{rsp}}$. Having changed φ_{rsp} , the hints change accordingly, thus security assumptions for the two versions are technically incomparable. In practice, the difference between arbitrary and odd degree is minimal, and, if anything, increasing the degree of the hints could make the problem harder. There is no reason to believe that the change makes SQISIGN less secure. We discuss this question in depth in the chapter of the specification dedicated to the security.

Additionally, the following improvements have been made to the reference implementation:

- (1) **Dropped the dependency on external libraries.** Formerly SQISIGN relied on GMP for multi-precision integers and DPE for double-precision-exponent floating points. The former has been replaced by a custom-made fixed-precision integer library; the latter is no longer needed, as floating-point numbers have been completely removed.
- (2) **Added support for various platforms.** The code now compiles and runs on x86_64, arm64 and Cortex-M4 platforms, with assembly optimizations for broadwell and AARCH64 architectures.

Overall, the present revision significantly accelerates SQISIGN, at the same time making it simpler to analyse and easier to implement. Though we could not achieve a fully constant-time implementation at this time, the changes to integer arithmetic were a necessary first step and pave the way for it.

Modifications introduced in the second round

The main difference between the *round-1 SQISIGN submission* and the present *round-2 SQISIGN submission* is the implementation of the improvements described in [BDD+24]. The structure of the scheme remains unchanged. The differences amount to the following three points:

- (1) **Uniform keys.** The key generation procedure now selects a uniformly random supersingular elliptic curve. This improves the theoretical security guarantee, because the underlying computational problem (the endomorphism ring problem) benefits from a worst-case to average-case self-reduction for the uniform distribution: key recovery is now provably as hard as the hardest instance of the endomorphism ring problem.
- (2) **Response sampled from a well-understood distribution.** The response phase requires sampling a “response isogeny” $\varphi_{\text{rsp}} : E_{\text{com}} \rightarrow E_{\text{chl}}$ from a collection of possible responses. In the previous version of SQISIGN, this response was sampled in an *ad hoc* manner, which was hard to analyse (a concern for the *zero-knowledge* property of the scheme), and forced the degree of φ_{rsp} to be very large (causing the scheme to be slow). The response is now sampled from a natural, well-understood distribution: the uniform distribution on the finite set of isogenies $E_{\text{com}} \rightarrow E_{\text{chl}}$ of bounded degree. This improves the theoretical security guarantee by removing *ad hoc* assumptions from the zero-knowledge property.
- (3) **Response represented by interpolation data.** The response isogeny is now represented by interpolation data: the images of a few points through the isogeny. The previous version of SQISIGN used another representation (the isogeny path representation) which only

works for special isogenies. The interpolation method allows one to represent any isogeny. This allows one to represent isogenies sampled from the aforementioned uniform distribution. Combined with the aforementioned improvement on the degree of φ_{rsp} , this method results in a significant speedup of SQISIGN. This interpolation method requires the computation of isogenies between abelian surfaces (a two-dimensional analogue of elliptic curves).

- (4) **Much better performance, in all metrics.** As a result of the changes outlined above, together with an improved implementation, the round-2 version of SQISign drastically outperforms the round-1 version: for security level I, the optimised implementation of signing is now nearly $20\times$ faster, at 103.0 Mcycles, and verification is more than $6\times$ faster, at 5.1 Mcycles. At higher security levels, the improvements are even larger. Beyond the running time improvements, signatures in the round-2 version are also smaller, by about 14%.

Concretely, the material on the KLPT algorithm has been removed (Section 2.5.2. *The KLPT algorithm and generalisations* in the previous version), and material on the computation of isogenies in dimension 2 has been added. Modifications throughout the document reflect this new approach.

Bibliography

- [BCE+25] Giacomo Borin, Maria Corte-Real Santos, Jonathan Komada Eriksen, Riccardo Invernizzi, Marzio Mula, Sina Schaeffler, and Frederik Vercauteren. “Qlapoti: Simple and Efficient Translation of Quaternion Ideals to Isogenies”. In: *ASIACRYPT 2025, Part IV*. Ed. by Goichiro Hanaoka and Bo-Yin Yang. Vol. 16248. LNCS. Springer, Singapore, Dec. 2025, pp. 174–205. doi: [10.1007/978-981-95-5113-2_6](https://doi.org/10.1007/978-981-95-5113-2_6) (cit. on p. 1).
- [BDD+24] Andrea Basso, Pierrick Dartois, Luca De Feo, Antonin Leroux, Luciano Maino, Giacomo Pope, Damien Robert, and Benjamin Wesolowski. “SQIsign2D-West - The Fast, the Small, and the Safer”. In: *ASIACRYPT 2024, Part III*. Ed. by Kai-Min Chung and Yu Sasaki. Vol. 15486. LNCS. Springer, Singapore, Dec. 2024, pp. 339–370. doi: [10.1007/978-981-96-0891-1_11](https://doi.org/10.1007/978-981-96-0891-1_11) (cit. on p. 2).
- [Dar25] Pierrick Dartois. “Fast computation of higher dimensional isogenies for cryptographic applications”. PhD thesis. Université de Bordeaux, France, July 2025. URL: <https://theses.fr/s364682> (cit. on p. 1).
- [DLS26] Max Duparc, Antonin Leroux, and Sina Schaeffler. *Qlapoty: Improved analysis and efficiency for quaternionic ideal to isogeny transformation*. Cryptology ePrint Archive, Paper 2026/1700. 2026. URL: <https://eprint.iacr.org/2026/1700> (cit. on p. 1).
- [Dup25] Max Duparc. “Superglue: Fast Formulae for (2,2)-Gluings Isogenies”. In: *ASIACRYPT 2025, Part IV*. Ed. by Goichiro Hanaoka and Bo-Yin Yang. Vol. 16248. LNCS. Springer, Singapore, Dec. 2025, pp. 372–400. doi: [10.1007/978-981-95-5113-2_12](https://doi.org/10.1007/978-981-95-5113-2_12) (cit. on p. 1).
- [KLKL26] Won Kim, Jeonghwan Lee, Hyeonhak Kim, and Changmin Lee. “SQIsign with Fixed-Precision Integer Arithmetic”. In: *PKC 2026, Part III*. Ed. by Shi Bai and Edoardo Persichetti. Vol. 16553. LNCS. Springer, Cham, May 2026, pp. 3–30. doi: [10.1007/978-3-032-26737-5_1](https://doi.org/10.1007/978-3-032-26737-5_1) (cit. on p. 1).
- [Ler26] Antonin Leroux. “Efficient quaternion algorithms for the deuring correspondence, and application to the evaluation of modular polynomials”. In: *ANTS XVII (2026)* (cit. on p. 1).
- [LS26] Antonin Leroux and Sina Schaeffler. *New algorithms for quaternion ideals in SQIsign*. to appear. 2026 (cit. on p. 1).
- [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> (cit. on p. 1).