A Simple and Unified Approach for Proving Knowledge of Isogenies between Abelian Varieties

Jonathan Komada Eriksen, Riccardo Invernizzi, Jannik Spiessens, and Frederik Vercauteren

KU Leuven, COSIC, Heverlee, Belgium

Abstract. In this paper we introduce a simple and unified approach, based on generic proof systems, to prove knowledge of any isogeny between two principally polarized abelian varieties in any dimension, assuming that the 2^m-torsion is accessible for sufficiently large m. Previous generic proof approaches were only able to prove knowledge of a smooth degree isogeny between elliptic curves, where for each small prime factor ℓ of the degree, bespoke constraints had to be derived, typically from (a variant of) the ℓ-th modular polynomial.

Our approach is much simpler in that it relies on proving knowledge of a 2^n-isogeny between two principally polarized abelian varieties in any dimension. Furthermore, our approach is unified in that the constraints are essentially the same for each dimension, resulting in a simpler and easier-to-optimize algorithm. Our construction has immediate applications to proving knowledge of an isogeny of any degree between two elliptic curves, by using a higher dimensional representation. Indeed, by a result of Robert, any isogeny can be embedded in a 2^n-isogeny by increasing the dimension, and conversely, the knowledge of a 2^n-isogeny between products of varieties implies the knowledge of an isogeny of degree $\leq 2^n$ between a factor of the domain and codomain. Our generic proof does not disclose the degree of the secret isogeny, nor does it rely on knowing the endomorphism ring, thereby solving an open problem posed by Beullens, De Feo, Galbraith, and Petit in 2023.

Two use cases are immediate. First, if one wants to prove knowledge of any isogeny between two supersingular curves over $\mathbb{F}_{p^2}$, e.g. during the generation of an elliptic curve with unknown endomorphism ring. Second, to prove knowledge of a secret isogeny coming from the class group action on oriented supersingular elliptic curves, e.g. CSIDH with curves defined over $\mathbb{F}_p$. Computing such group actions is typically done using qt-Pegasis, which naturally results in a 4-dimensional representation of the isogeny.

Lastly, we propose two tailored zero-knowledge proof systems that improve proving time and proof size without loss of generality and provide the first implementation in dimension 2 and 4 by implementing both proof systems in Rust.

Keywords: Post-quantum cryptography · isogenies · zero-knowledge proofs.

1 Introduction

The task of proving knowledge of a secret isogeny goes back to the very beginning of isogeny based cryptography. Already in the seminal paper by Couveignes [25], an ID-protocol was presented, proving the knowledge of a secret isogeny coming from the class group action on ordinary elliptic curves. Instantiating this protocol proved highly non-trivial, but changing from ordinary elliptic curves to oriented supersingular elliptic curves resulted in a much more efficient construction called CSIDH [18]. The ID-protocol in [25] also applies to CSIDH and is a simple variant of the zero knowledge proof of graph isomorphism by Goldreich, Micali, Wigderson [46]. In particular, it is a sigma protocol with binary challenge space, thus requiring λ repetitions to achieve λ-bit security. The protocol also requires acting by random class group elements, which until recently was impossible and thus required a workaround, for instance using rejection sampling as in SeaSign [36], or via the computation of the class group structure as in CSI-FiSh [12]. Since the introduction of Clapoti [61] and its more efficient variants Pegasis [31] and especially qt-Pegasis [30], it is now possible to act with random class group elements efficiently and the simple ID-protocol applies in its original form.

Proving knowledge of a (non-oriented) isogeny of *secret* degree between supersingular elliptic curves has proven much more difficult. Nevertheless, an important example of such a proof of knowledge is the ID-protocol underlying the (1-dimensional version of the) SQIsign signature scheme [38, 39, 24]. However, SQIsign requires significant setup, and importantly the knowledge of both endomorphism rings of the domain and codomain.

On the other hand, for proving knowledge of an isogeny of *smooth and known degree* between supersingular elliptic curves, there have been many more constructions, both relying on isogeny-based techniques only as well as using generic proof techniques. The first ID-scheme proving knowledge of an isogeny of smooth and known degree, was presented in the original SIDH-paper by De Feo, Jao and Plût [37] (DFJP), where the characteristic of the underlying field was an SIDH prime. However, there was an error in the proof of soundness of the DFJP protocol, pointed out by Ghantous, Pintore, and Veroni [45], and the protocol was later fixed by De Feo, Dobson, Galbraith, and Zobernig [35]. Interestingly, this protocol did not reveal torsion information, despite being "based on SIDH", and is thus still believed to be secure despite SIDH being broken [15, 57, 63] right after the publication of this ID-protocol. The DFJP protocol was then generalized to any characteristic and achieving statistical zero-knowledge in [6]. Similar to the ID-protocol in the class group setting, the protocol in [6] suffers from small challenge space requiring parallel executions.

A recent line of work relies on using generic zero-knowledge proof techniques, such as zk-SNARKs, to prove the existence of an isogeny between two elliptic curves. Inspired by the work of [21], Cong, Lai, and Levin [22] expressed a 2^n-isogeny between two elliptic curves as a sequence of n tuples of j-invariants (j, j_{i+1}) where each tuple satisfies the 2-nd classical modular polynomial $\Phi_2(x, y)$ (which expresses precisely when two elliptic curves are 2-isogeneous). By essen-

tially expressing the equations $\Phi_2(j_i, j_{i+1}) = 0$ for $i = 0, \ldots, n$ (with some added conditions to avoid backtracking) as a system of R1CS constraints, Cong, Lai and Levin then applied generic zk-SNARKs such as Aurora, Ligero and Limbo to derive an isogeny proof of knowledge. It is clear that this approach has two limitations: first, it can only deal with isogenies of degree 2^n and second, the degree of the isogeny is known. By using modular polynomials Φ_ℓ for $\ell > 2$ and more compact versions such as the canonical [47], or Atkin or Weber modular polynomials [48], the degree can be smooth with primes up to $\ell = 71$. However, for each new prime factor ℓ, one needs to derive tailored R1CS constraints to express the equation $\Phi_\ell(j_i, j_{i+1}) = 0$. Furthermore, even when weighted by $\log_2(\ell)$ the constraints for larger ℓ are much less efficient than for $\ell = 2$. Levin and Pedersen [52] replaced the use of the 2-nd modular polynomial by working with radical isogenies, which avoids backtracking and simplifies the resulting R1CS constraint system.

After the aforementioned SIDH-attacks, higher-dimensional representations of isogenies have become ubiquitous in isogeny-based cryptography. For instance, the efficiency of the SQIsign protocol was massively improved by embedding the relevant isogenies in 2^n-isogenies between abelian surfaces, making it a very promising candidate for post-quantum standardization [1, 7, 14] (which also conceptually changed it from proving knowledge of a secret degree isogeny, to an isogeny of known but non-smooth degree). Similarly, (unrestricted) class group action computations, necessary in the ID-protocol by Couveignes [25], were only recently made efficient and scalable by embedding the isogeny in a 2^n-isogeny between 4-dimensional abelian varieties [31, 32]. It is therefore only natural to try to prove knowledge of secret 2^n-isogenies between higher dimensional abelian varieties, which implies a proof of knowledge of an isogeny between elliptic curves that are a factor of the domain and co-domain of these higher dimensional varieties. This approach has two immediate advantages: first, since we are reduced to the case of 2^n-isogenies (although in higher dimension), the resulting constraints will be simple and in fact uniform even in higher dimensions; second, due to the embedding the proof does not reveal nor depend on the (factors of the) degree of the secret isogeny.

Our contributions We provide the first simple and unified framework for proving knowledge of a chain of 2-isogenies in the level-2 theta model between abelian varieties in any dimension, using generic proof techniques. Since any N-isogeny between principally polarized abelian varieties can be embedded in a 2^n-isogeny of higher dimension, for some $2^n > N$ (see Lemma 2), we effectively present the first unified way of proving knowledge of any isogeny, regardless of its degree. Our main contributions can be summarized as follows:

– We show that the radical isogeny framework in the theta model introduced in [51] can be efficiently translated into an R1CS constraint system (Section 3); furthermore, the type of constraints is the same in all dimensions, and is also the same for all types of isogenies (e.g. so-called gluing, splitting, diagonal or generic), resulting in a very simple and unified framework (Section 5);

- Motivated by Lemma 2, we analyze the extension of the radical framework to dimension 4 (in particular, we discuss the additional relations in Appendix A.1), and discuss its security (Section 4);
- We describe a tailor-made multilinear polynomial interactive oracle proof which can be compiled into a zk-SNARK using the BaseFold [69] polynomial commitment scheme without losing any generality (Section 6.1);
- We propose a lattice-based sigma protocol that has significantly smaller proof size and lower prover computation time at the cost of adding an additional hardness assumption (Section 6.2). We note that this proof system forms an independent contribution as a zero-knowledge post-quantum secure proof system for proving small statements over any field with small proof size;
- We solve one of the major open problems in [11], namely whether it is possible to prove knowledge of an arbitrary isogeny between two elliptic curves, without knowledge of the endomorphism ring of the starting curve and without leaking the degree of the secret isogeny; moreover, we discuss how our framework can be adapted to prove all the relations suggested in [11], showing its generality and flexibility (Section 7);
- We show the impact of our work on the recently growing field of higher dimensional isogeny constructions; notably, we give an alternative way of proving a qt-Pegasis isogeny, and we repair the 2-dimensional version of the CGL hash function (among other things, one of the main motivations behind [51]) from the recent attack of [16], by presenting a 2-dimensional trusted setup (Section 7);
- We provide a full implementation of our framework: we show how to augment the most common isogeny computation libraries in dimension 1, 2 and 4 to produce an R1CS transcript, and present a Rust implementation of our construction for proving and verifying such transcript (Section 8). Our implementation is available at

https://github.com/KULeuven-COSIC/unified_zk_iso

Technical overview The goal of this paper is to develop a proof of knowledge for the very general relation

$$\mathcal{R}_{\mathrm{isog}(g)} := \left\{ ((A_0, A_1), \phi) \;\middle|\; \begin{smallmatrix} A_0, A_1 \text{ abelian varieties of dimension } g, \\ \phi : A_0 \to A_1 \text{ is an arbitrary } N\text{-isogeny for some } N \in \mathbb{N} \end{smallmatrix} \right\}.$$

In most practical applications, we will be interested in the above relation for $g = 1$, i.e. the A_i are taken to be elliptic curves, which has usually been referred to simply as $\mathcal{R}_{\mathrm{isog}}$ in earlier works, e.g. [11]. To achieve this, we show it is sufficient to give a proof of knowledge for the relation

$$\mathcal{R}_{\mathrm{isog}(2^n,g)} := \left\{ ((A_0, A_1), \phi) \;\middle|\; \begin{smallmatrix} A_0, A_1 \text{ abelian varieties of dimension } g, \\ \phi : A_0 \to A_1 \text{ is a good } 2^n\text{-isogeny} \end{smallmatrix} \right\},$$

where a good 2^n-isogeny means that its kernel is isomorphic to $(\mathbb{Z}/2^n\mathbb{Z})^g$. Counter-intuitively, but by now a standard fact in isogeny-based cryptography, the second

relation is in fact general enough to encompass all instances of the first one, possibly by increasing 2^n and g: given an (A_0, A_1) in the language of isogeneous principally polarized abelian varieties of dimension g, a witness ϕ for $\mathcal{R}_{\mathrm{isog}(2^n,g)}$ is clearly also a witness for $\mathcal{R}_{\mathrm{isog}(g)}$, however given a witness ψ for $\mathcal{R}_{\mathrm{isog}(g)}$, this can also be turned into a witness Ψ for $(A_0 \times B, A_1 \times B')$ in $\mathcal{R}_{\mathrm{isog}(2^n,8g)}$, for $2^n > \deg \psi$, and some auxiliary varieties B, B', by Robert's embedding lemma (see Lemma 2). The isogeny $\Psi : A_0 \times B \to A_1 \times B'$ now again allows the extraction of a witness for (A_0, A_1) in $\mathcal{R}_{\mathrm{isog}(g)}$ by composing with the projection maps.

To derive a proof of knowledge for the relation $\mathcal{R}_{\mathrm{isog}(2^n,g)}$, we first derive an equivalent system of constraints, e.g. in R1CS, and apply general purpose zero-knowledge proof systems. In order to arithmetize the relation $\mathcal{R}_{\mathrm{isog}(2^n,g)}$, we make use of the level-2 theta model. In the recent work by Kunzweiler, Maino, Moriya, Petit, Pope, Robert, Stopar, and Ti [51], the authors proved that a good 2^n-isogeny can be computed as a composition of 2-isogenies in the theta model, where the theta null-points all satisfy

$$\mathcal{S} \circ \mathcal{H}(\theta^{A_{i+1}}(0_{A_{i+1}})) = \mathcal{H} \circ \mathcal{S}(\theta^{A_i}(0_{A_i}))\,,$$

where $\mathcal{S}$ denotes the coordinate-wise squaring map, and $\mathcal{H}$ denotes a Hadamard-transform. Thus, a witness for $\mathcal{R}_{\mathrm{isog}(2^n,g)}$ can be turned into a witness for the relation

$$\mathcal{R}_{\mathrm{HS}(n,g)} := \left\{ \begin{pmatrix} ((A_0,\theta_0),(A_n,\theta_n)), \\ (T_i)_{i\in[n-1]} \end{pmatrix} \;\middle|\; \begin{array}{c} (A_0,\theta_0),(A_n,\theta_n) \text{ isogenous abelian varieties of} \\ \text{dimension } g \text{ with a level-2 theta structure,} \\ \mathcal{S}\circ\mathcal{H}(T_1)=\mathcal{H}\circ\mathcal{S}(\theta^{A_0}(0_{A_0})), \\ \mathcal{S}\circ\mathcal{H}(\theta^{A_n}(0_{A_n})=\mathcal{H}\circ\mathcal{S}(T_{n-1}), \\ \mathcal{S}\circ\mathcal{H}(T_{i+1})=\mathcal{H}\circ\mathcal{S}(T_i) \quad \forall i\in[n-2] \end{array} \right\}$$

by working with the level-2 theta-model, and taking the T_i's to be the theta null-points in the chain of 2-isogenies. Again applying the results of [51], this relation can also be shown to be equivalent to $\mathcal{R}_{\mathrm{isog}(2^n,g)}$, whenever we restrict to abelian varieties of dimension $g \leq 2$. However, importantly, this is not true for $g > 2$. Thus, in principle, we have to derive additional relations to verify that the witness to $\mathcal{R}_{\mathrm{HS}(n,g)}$ corresponds to an isogeny. For $g = 3$, this was already done in [51]. We discuss the derivation of similar relations for $g = 4$, and show that the resulting extra relations are too complex to be of practical use.

However, we argue that these extra relations can be dropped without compromising security. Although there may a priori exist witnesses to $\mathcal{R}_{\mathrm{HS}(n,g)}$ which do not come from good 2^n-isogenies, we can instead simply define the alternate relation $\tilde{\mathcal{R}}_{\mathrm{HS}(n,g)}$, which is the subset of $\mathcal{R}_{\mathrm{HS}(n,g)}$ corresponding to isogenies. We then provide heuristics that constructing a witness for the relation in $\mathcal{R}_{\mathrm{HS}(n,g)}$ which at the same time is not a witness for $\tilde{\mathcal{R}}_{\mathrm{HS}(n,g)}$, is computationally hard (see Section 4). Thus, assuming the hardness of this problem, $\mathcal{R}_{\mathrm{isog}(2^n,g)}$ and $\mathcal{R}_{\mathrm{HS}(n,g)}$ are still weakly equivalent for any g, and we can limit ourselves to deriving a proof of knowledge for $\mathcal{R}_{\mathrm{HS}(n,g)}$.

Finally, we show how to prove knowledge of a witness for the relation $\mathcal{R}_{\mathrm{HS}(n,g)}$. We do this by translating the Hadamard relations into a constraint system

$M_1 z \circ M_1 z = M_2 z$ for public matrices M_1, M_2 and witness z. To prove this constraint system we define multilinear polynomials $\tilde{z}_1, \tilde{z}_2$ such that the previous relation holds if and only if for all $\forall X \in \{0,1\}^m : \tilde{z}_1(X)^2 = \tilde{z}_2(X)$. Then we use two subsequent instantiations of the sumcheck protocol [53] to reduce this statement to proving the evaluation of a multilinear polynomial, for which we use the BaseFold [69] polynomial commitment scheme. We further propose a trade off where we decrease the proving time and proof size at the cost of adding well-known post-quantum lattice-based hardness assumptions. The resulting protocol recursively proves knowledge of valid sumcheck transcripts by committing to them in BDLOP commitments [8] and then using a lattice-based sigma protocol to prove linear relations. Non-linear verification steps in the sumcheck protocol can be masked using Laurent polynomials and then verified in the clear. Verifying the masking itself can again be expressed using linear relations.

Our resulting proof systems are general enough to derive proofs of knowledge for many relations of interest in isogeny-based cryptography. One example of this is the CSIDH-relation

$$\mathcal{R}_{\text{CSIDH}} := \{((E_0, E_1), \phi) \mid \phi : E_0 \to E_1 \text{ is an isogeny defined over } \mathbb{F}_p\}.$$

Compared to the state of the art, running the original ID-protocol by Couveignes, using qt-Pegasis [32], our generic proofs achieve 5.2s prover time with 1.1MiB proof sizes. This shows that our proof system, although very generic, is already quite practical for many applications, with a performance that is comparable to specialized native protocols.

Acknowledgments

This work is supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement ISOCRYPT – No. 101020788), by the Research Council KU Leuven grant C14/ 24/099, as well as by Cybersecurity Research Flanders with reference number VR20192203. Riccardo Invernizzi is funded by Research Foundation Flanders (FWO) under a PhD Fellowship fundamental research (project number 1138925N). Jannik Spiessens is funded by Research Foundation Flanders (FWO) under a PhD Fellowship fundamental research (project number 1139125N).

2 Preliminaries

In this work, we consider isogenies between principally polarized abelian varieties (equipped with a level-2 theta structure) and zero-knowledge proofs of knowledge of such maps.

We first recall some important facts about principally polarized abelian varieties with level-2 theta structure and their relation to the computation of (good) 2^n-isogenies. For a more general introduction to (principally polarized) abelian varieties we refer to Milne [58]. We also recall the necessary background on zero-knowledge proof systems.

2.1 Abelian varieties and isogenies

An abelian variety is a smooth, projective algebraic variety equipped with a commutative group law. Abelian varieties of dimension 1 are elliptic curves. The g-th power of an elliptic curve E^g is an example of a g-dimensional abelian variety. When not clear from the context, we use $A(k)$ to denote the set of k-rational points of A. $A[\ell](k)$ denotes the ℓ-torsion of such set, i.e. the k-rational points of order ℓ. Throughout the rest of this work, p will be a prime, and we will be mostly interested in $k = \mathbb{F}_{p^2}$ or $\mathbb{F}_p$, thus dropping the k from above.

In particular, in dimension 1 we will focus on *supersingular* elliptic curves, i.e. curves such that $E[p] = \{0_E\}$. One important property of such curves is that they always admit a model which is defined over $\mathbb{F}_{p^2}$. One generalization of supersingularity is that of *superspecial* abelian varieties. Superspecial abelian varieties admit a model over $\mathbb{F}_{p^2}$ as well [17]. Products of elliptic curves and all other abelian varieties in this work will be superspecial.

An isogeny is a surjective morphism between abelian varieties, respecting the group law, with finite kernel. *Separable* isogenies can be identified by their kernel, up to post composition with an isomorphism. For every abelian variety A, there is a corresponding dual variety $\hat{A}$. An isogeny $\lambda : A \to \hat{A}$ is called a *polarization* if it is induced by an ample line bundle [58, Section I.11] (over $\overline{k}$ in general, though this distinction is not necessary when k is a finite field [23, Theorem 2.6]); if it is an isomorphism the polarization is *principal*. A pair (A, λ) where λ is a principal polarization of A is called a *principally polarized abelian variety* (PPAV). A principal polarization induces a polarized Weil-pairing $e_N^\lambda : A[N] \times A[N] \to \mu_N$, with μ_N the N-th roots of unity, by precomposing the standard Weil-pairing with the polarization.

An isogeny $\phi : (A, \lambda_A) \to (B, \lambda_B)$ is called compatible with the polarization if $\widehat{\phi} \circ \lambda_B \circ \phi = [N]\lambda_A$ for some N. Such a compatible isogeny is called an N-isogeny. For especially nice N-isogenies, we use the following definition:

Definition 1. *An N-isogeny $\phi : A \to B$ between principally polarized abelian varieties of dimension g is called a good N-isogeny if* $\ker \phi \simeq (\mathbb{Z}/N\mathbb{Z})^g$.

A good N-isogeny is more commonly denoted as an (N, N)-isogeny in dimension 2, an (N, N, N)-isogeny in dimension 3 and so on.

2.2 Good extensions

Let $\phi : A \to B$ and $\psi : B \to C$ be two good N-isogenies, for some prime N, between PPAVs of dimension g. If $g = 1$, the composition $\psi \circ \phi$ can be either a good N^2-isogeny (with kernel $\cong \mathbb{Z}/N^2\mathbb{Z}$) or multiplication by N (with kernel $\cong (\mathbb{Z}/N\mathbb{Z})^2$). In the latter case $\psi = \hat{\phi}$ is the *dual* isogeny of ϕ (recall that in dimension 1, the polarization is canonical). The situation becomes more complex for $g \geq 2$: in general we have

$$\ker(\psi \circ \phi) \cong (\mathbb{Z}/N^2\mathbb{Z})^m \times (\mathbb{Z}/N\mathbb{Z})^{2k}$$

with $m + k = g$. This motivates the following definition, originally due to [17] and later generalized by [51] to arbitrary g.

Definition 2. *Let $\phi : A \to B$ and $\psi : B \to C$ be good N-isogenies between principally polarized abelian varieties of dimension g, and $\ker(\psi \circ \phi)$ as above. Then:*

- *if $k = g$ and $m = 0$, we say that ψ is the* dual extension *of ϕ;*
- *if $k = 0$ and $m = g$, we say that ψ is a* good extension *of ϕ;*
- *otherwise, we say that ψ is a* bad extension *of ϕ.*

Good extensions can also be characterized as isogenies ψ such that $(\ker \psi) \cap \phi(A[N]) = 0$ [17]. For a given good N-isogeny ϕ, there are $N^{g(g+1)/2}$ good extensions [19, Lemma 2]. Composing ϕ with any good extension results in a good N^2-isogeny. In this work, we will be working with radical isogenies. Isogenies built in this way are automatically good extensions [51].

The dual extension of ϕ is the isogeny such that $\ker \psi = \phi(A[N])$, and corresponds to the (polarized) dual of ϕ.

2.3 Embedding isogenies

One of the main outcomes of the SIDH attacks [15, 57, 63] is that one-dimensional isogenies can be *embedded* in higher dimensional isogenies. The main building block that powers this technique is the now famous Kani's Lemma [49]:

Lemma 1. *Given a commutative diagram of isogenies between principally polarized abelian varieties*

$$\begin{array}{ccc} A & \xrightarrow{\phi} & B \\ \downarrow{\scriptstyle \psi} & & \downarrow{\scriptstyle \psi'} \\ A' & \xrightarrow{\phi'} & B' \end{array}$$

where ϕ, ϕ' are both n_1-isogenies, and ψ, ψ' are n_2-isogenies, then

$$\Phi = \begin{pmatrix} \phi & \widetilde{\psi'} \\ -\psi & \widetilde{\phi'} \end{pmatrix} : A \times B' \to B \times A'$$

is an N-isogeny, for $N = n_1 + n_2$.

Proof. See [62, Lemma 2.3] for this formulation. □

Clearly, being able to evaluate Φ is sufficient for being able to evaluate ϕ (and ψ, ϕ' and ψ'), by composing with the appropriate injection/projection maps. Hence, we say that these isogenies are *embedded* in Φ.

The reason why we are interested in this constructively is that the complexity of evaluating an isogeny when given its kernel depends (exponentially) on the largest prime factor of its degree. If n_1 contains a large prime factor but $N =$

$n_1 + n_2$ is smooth, Φ and hence ϕ can be evaluated efficiently at the cost of doubling the dimension.

For efficiency reasons one is generally interested in evaluating a one dimensional isogeny ϕ of degree n_1, and one fixes $N = 2^n$. Without any knowledge of the endomorphism ring of the starting curve A, building an isogeny ψ of the correct degree $2^n - n_1$ (sometimes called *auxiliary isogeny*) is a challenging task on its own. However, we always know $\mathbb{Z} \subset \mathrm{End}(E)$, i.e. we can always use scalar multiplication as an endomorphism. We have the following cases: 1) if n_1 is smooth, ϕ can already be evaluated in dimension 1; 2) if $2^n - n_1 = a^2$ is a square, then taking $\psi = [a]$ the multiplication by a map is a valid choice; 3) if $2^n - n_1 = a^2 + b^2$ is a sum of two squares, then

$$\psi = \begin{pmatrix} a & -b \\ b & a \end{pmatrix}$$

is an endomorphism of $E \times E$ of the correct degree; we must hence work in dimension 4; 4) due to a famous result by Lagrange, every integer can be written as a sum of four squares; thus we can always write an endomorphism of E^4 consisting of integer entries and use it to embed any isogeny in dimension 8.

This leads to the following lemma, due to Robert [62, Section 3].

Lemma 2. *Let $\phi : A \to B$ be an N-isogeny between principally polarized abelian varieties. Then ϕ can be embedded in an N'-isogeny*

$$\Phi : A^4 \times B^4 \to A^4 \times B^4$$

for any $N' > N$.

2.4 Level-2 theta structures

Given a principally polarized abelian variety A of dimension g, a *level-2 symmetric theta structure on A* is in particular a map $\theta^A : A \to \mathbb{P}^{2^g-1}$, which induces a symplectic basis (with respect to the polarized Weil-pairing) of $A[2]$. Such symplectic basis of $A[2] \simeq (\mathbb{Z}/2\mathbb{Z})^{2g}$ is given by $2g$ 2-torsion points $\{P_1, \ldots, P_g, Q_1, \ldots, Q_g\}$ such that $e_2^\lambda(P_i, P_j) = e_2^\lambda(Q_i, Q_j) = 1$ and $e_2^\lambda(P_i, Q_j) = (-1)^{\delta_{i,j}}$ with $\delta_{i,j}$ Kronecker delta for all $1 \leq i, j \leq n$.

We now give the following full definition, which is due to Duparc [41, Section 2.2], which the reader may find simpler than the original definition by Mumford.

First, note that a symplectic basis of $A[2]$ induces an isomorphism

$$\pi : A[2] \simeq (\mathbb{Z}/2\mathbb{Z})^g \times \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$$

(called a symplectic representation) such that

$$e_2^\lambda(P, Q) = (-1)^{\widehat{x_Q}(x_P) - \widehat{x_P}(x_Q)}$$

for all $P, Q \in A[2]$ (here, we denote $\pi(R) = (x_R, \widehat{x_R})$). This allows the following definition.

Definition 3. *Let A be a principally polarized abelian variety of dimension g. A (symmetric) level-2 theta structure is a map $\theta^A : A \to \mathbb{P}^{2^g-1}$, whose coordinates we denote by θ_i^A indexing by $i \in (\mathbb{Z}/2\mathbb{Z})^g$, which satisfies*

$$\theta_i^A(P+Q) = (-1)^{\widehat{x_Q}(i)}\theta_{i+x_Q}^A(P) \tag{1}$$

for all $P \in A$ and $Q \in A[2]$.

Informally, a level-2 theta structure thus corresponds to a map θ from A to the projective space, such that the 2-torsion acts by very specific linear transformations on the points of A. Note that this map is not injective. However, when A is irreducible, it does induce an injection of the kummer variety $A/[\pm 1] \hookrightarrow \mathbb{P}^{2^g-1}$.

An important invariant of a principally polarized abelian variety with a level-2 theta structure is the *theta null-point* $\theta^A(0_A)$. However, note that this is not an isomorphism invariant of A as a principally polarized abelian variety alone (i.e. seen without a theta structure): different choices of theta-structure give different theta-null points for the same A. Note also that a level-2 theta structure induces the symplectic basis of $A[2] = K_1 + K_2$ (which is referred to as *the canonical decomposition*); this basis is easily recovered from the theta null-point by applying Equation (1) with $P = 0_A$ (for an explicit description, see [51, Lemma 1]).

Formulae for the computation of 2^n-isogenies in the theta-2 model in dimension 2 were originally given in [34] (while dimension 4 was done by Dartois [28]). As the theta-structure already induces two choices of maximally isotropic subgroup K_1, K_2 (from $A[2] = K_1 + K_2$), a 2-isogeny $\phi : (A, \theta^A) \to (B, \theta^B)$ will refer to an isogeny with kernel K_2, such that θ_B gives a new kernel, defining a good extension of ϕ. To describe the action of these isogenies on the theta null-points, we recall the notation

$$\mathcal{S} : \mathbb{P}^{2^g-1} \to \mathbb{P}^{2^g-1} : (x_0 : \cdots : x_{2^g}) \to (x_0^2 : \cdots : x_{2^g}^2)$$

denoting coordinate-wise squaring, and

$$\mathcal{H} : \mathbb{P}^{2^g-1} \to \mathbb{P}^{2^g-1} : (x_0 : \cdots : x_{2^g}) \to H_g(x_0, \ldots, x_g)^T$$

where

$$H_g := \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}^{\otimes g}$$

is the g-th Hadamard transform. This gives the following lemma [51, Lemma 3]:

Lemma 3. *Let (A, θ^A) be an abelian variety with a level-2 theta structure, and let $A[2] = K_1 + K_2$ be the canonical decomposition. Let $\phi : (A, \theta^A) \to (B, \theta^B)$ be a 2-isogeny, with kernel K_2, where θ^B is a compatible theta structure on B. Then*

$$\mathcal{S} \circ \mathcal{H}(\theta^B(0_B)) = \mathcal{H} \circ \mathcal{S}(\theta^A(0_A)). \tag{2}$$

Further, there are $g(g+1)/2$ free choices of square-roots corresponding to the different compatible theta-structures on B.

In [51, Lemma 3], the authors give more details on the choices of square-roots, but the description above will be sufficient for us. One important detail is that composing 2-isogenies obtained by Lemma 3 automatically gives good extensions: see [51, Section 3.4].

Since $\mathcal{H}$ is invertible, it is clear that once we fix θ_A, there are $2^g - 1$ different choices of θ_B satisfying the equation (corresponding to the different choices of square roots, projectively). In particular, for dimension $g = 1, 2$, we have $g(g+1)/2 = 2^g - 1$, and thus the number of possible square root choices and compatible theta-structures are the same. Hence, any choice is valid and corresponds to the different possible good extensions.

2.5 Proof systems

Let $\mathcal{R} : X \times W \to \{0,1\}$ be a relation between the input set X and the witness set W, defining the NP-language $\mathcal{L} = \{x \in X \mid \exists w \in W \text{ s. t. } (x,w) \in \mathcal{R}\}$.

Lemma 4 (Schwartz-Zippel Lemma). *Let $f \in \mathbb{F}[X_1, \ldots, X_m]$ be a non-zero polynomial of total degree d over a field $\mathbb{F}$. Let S be any finite subset of $\mathbb{F}$, and let $r_1, \ldots, r_m$ be m field elements selected independently and uniformly from the set S, then*

$$\Pr[f(r_1, \ldots, r_m) = 0] \leq \frac{d}{|S|}.$$

Definition 4 (Multilinear extensions). *For every function $f : \{0,1\}^m \to \mathbb{F}$, there is a unique multilinear polynomial $\tilde{f} \in \mathbb{F}[X_1, \ldots, X_m]$ such that $\forall b \in \{0,1\}^m : \tilde{f}(b) = f(b)$. We call $\tilde{f}$ the multilinear extension of f and $\tilde{f}$ can be expressed as*

$$\tilde{f}(X) = \sum_{b \in \{0,1\}^m} f(b) \cdot \tilde{\mathsf{eq}}(b, X)$$

where $\mathsf{eq} : \{0,1\}^m \times \{0,1\}^m \to \{0,1\} : (x,y) \mapsto \mathbb{1}[x = y]$. We can compute

$$\tilde{\mathsf{eq}}(x,y) = \prod_{i=1}^{m} \big(x_i y_i + (1 - x_i)(1 - y_i)\big).$$

Definition 5 (Sumcheck protocol [53]). *For a multivariate polynomial $f \in \mathbb{F}[X_1, \ldots, X_m]$ with m variables with degree at most ℓ in each variable and some $S \in \mathbb{F}$, the sumcheck protocol is a public-coin interactive protocol between a prover P and a verifier V that reduces the claim $S \stackrel{?}{=} \sum_{b \in \{0,1\}^m} f(b)$ to the claim that $f(r) \stackrel{?}{=} e$ for some uniformly random $e \in \mathbb{F}$ and $r \in \mathbb{F}^m$. The protocol has $\ell \cdot m/|\mathbb{F}|$ soundness error and $\mathcal{O}(\ell \cdot m)$ communication cost.*

3 Path of theta-null points

In this section, we discuss the relationship between knowing a path of level-2 theta null-points satisfying Equation (2), and knowing a 2^n-isogeny between abelian varieties of dimension g. Most of the material in this section follows from [51], but will be essential to our generic proofs. Of particular importance is Definition 6, and Problems 1 and2, which formalize the issue that a path of theta null-points is only equivalent to an isogeny-path when $g \leq 2$.

3.1 Encoding an isogeny-path as a set of solutions

Consider a good 2^n-isogeny expressed as a composition of n good 2-isogenies

$$A_0 \xrightarrow{\phi_0} A_1 \xrightarrow{\phi_1} \dots \xrightarrow{\phi_{n-1}} A_n$$

between abelian varieties A_i of dimension g. We note that in practical applications, e.g. when the 2^n-isogeny corresponds to a 4 dimensional representation of a class group action, both A_0 and A_n typically are products of abelian varieties of lower dimension. Since the connecting isogenies however are still good, these gluing (at the start) and splitting isogenies (at the end) are part of the above diagram and do not require separate treatment. This is in stark contrast to the actual computation of these isogenies, which requires tailored functions for each type of isogeny.

By Lemma 3, there exists level-2 theta structures θ^{A_i} on A_i for all i, such that the theta-null points all satisfy

$$\mathcal{S} \circ \mathcal{H}(\theta^{A_{i+1}}(0_{A_{i+1}})) = \mathcal{H} \circ \mathcal{S}(\theta^{A_i}(0_{A_i})). \tag{3}$$

The whole path can therefore be expressed as a concatenation of $(n+1)$ theta null points, each consisting of 2^g finite field elements. More in detail, the path is given by the set of $2^g \times (n+1)$ values X_i^j for $0 \leq i \leq 2^g - 1$, $0 \leq j \leq n$ such that $X^j = (X_0^j : \dots : X_{2^g-1}^j) = \theta^{A_j}(0_{A_j})$. The first and last varieties of the chain A_0 and A_n, and consequently X_0, X_n, are public, so the variables X_i^j satisfy the following relations

$$\begin{cases} X^0 = \theta^{A_0}(0_{A_0}), \\ \mathcal{S} \circ \mathcal{H}(X^{j+1}) = \mathcal{H} \circ \mathcal{S}(X^j), \quad 0 \leq j \leq n-1, \\ X^n = \theta^{A_n}(0_{A_n}). \end{cases} \tag{4}$$

Note that since theta null points are projective, the above equalities also have to be interpreted as such, i.e. they allow scaling each theta null point by a different non-zero scalar. By normalizing the first theta null point $\theta^{A_0}(0_{A_0})$ and imposing that the above equalities hold exactly (and not just projectively), we will impose a specific representative for each following theta null point $\theta^{A_i}(0_{A_i})$.

Given a good 2^n-isogeny, it is almost trivial to recover the corresponding theta null points $\theta^{A_i}(0_{A_i})$ and thus the full path X_i^j: they are induced by any symplectic basis $A_i[2] = K_1 + K_2$, where K_2 corresponds to the kernel of the isogeny ϕ_i. More details are given in Section 8.1.

Remark 1. There is a subtlety in going from an efficient representation of a 2^n-isogeny $\phi : A_0 \to A_n$, to a chain of 2-isogenies $\phi_{n-1} \circ \cdots \circ \phi_0 = \phi$, as this is not an easy task in general. In order to factor the isogeny, we need access to the kernel $K \subset A_0[2^n]$. This can be obtained from the representation of ϕ, provided the 2^n-torsion is accessible, or even $2^{\lfloor \frac{n}{2} \rfloor}$-torsion, by working with the duals, and doing a meet-in-the-middle approach. This is why we will need to require that the 2^m-torsion is accessible for sufficiently large m in general.

3.2 Recovering an isogeny-path from a solution set

We now discuss how to recover a 2^n-isogeny from a set of solutions to Equation (4). As already mentioned, we are forced to split our discussion into two cases. Indeed, from Lemma 1, we see that if we fix A_i there are in total $2^g - 1$ solutions to Equation (3), of which only $g(g+1)/2$ correspond to valid theta structures of abelian varieties A_{i+1}. For $g = 1, 2$ these two quantities coincide, but for $g > 2$ this is no longer the case. We therefore make the following case distinction.

Dimension $g \leq 2$. The simplest case is dimension $g \leq 2$. Assume we are given the theta null-points $\theta^{A_0}(0_{A_0}), \theta^{A_n}(0_{A_n})$, and we wish to recover the 2^n-isogeny between them, encoded by the set of solutions $T_1, \dots T_{n-1} \in \mathbb{P}^{2^g-1}$, all satisfying Equation (4), where the systems of equations holds affinely. As the solutions to

$$\mathcal{S} \circ \mathcal{H}(x_0 : \cdots : x_{2^g-1}) = \mathcal{H} \circ \mathcal{S}(\theta^{A_0}(0_{A_0})) \tag{5}$$

correspond exactly to the different good extensions from (A_0, θ^{A_0}) (see Section 2.4), we know that any valid solution-set must correspond to a theta-null point of a 2-isogenous variety, inducing a good extension. By induction, we can thus extract a good 2^n-isogeny from the solution set.

Dimension $g > 2$. The same argument fails as soon as $g > 2$. As mentioned before, there are $2^g - 1$ solutions to Equation (5) but only $g(g+1)/2$ correspond to valid isogenies. Thus, there are solutions to the above equations, which do not correspond to paths of 2-isogenies. To make this distinction clear, we introduce the following terminology.

Definition 6. *Let $(A_0, \theta^{A_0}), (A_n, \theta^{A_n})$ be two isogenous abelian varieties of dimension g with a given level-2 theta-structure. Let $T_1, \dots, T_{n-1} \in \mathbb{P}^{2^g-1}$ be points satisfying Equation (4). We call the tuple $(\theta^{A_0}(0_{A_0}), T_1, \dots, T_{n-1}, \theta^{A_n}(0_{A_n}))$ a pseudo 2^n-isogeny path.*

Computing a pseudo-isogeny path is formalized in the following problem.

Problem 1 (Pseudo-isogeny path problem). Given two isogenous abelian varieties of dimension g with a given level-2 theta structure (A_0, θ^{A_0}) and (A_n, θ^{A_n}), find a pseudo 2^n-isogeny path between them for some n.

The previous problem can be seen as a natural generalization of the isogeny-problem in dimension $g > 2$. However, it seems unlikely that the above problem is significantly easier than the isogeny-problem itself, especially considering they are equivalent in dimension $g \leq 2$. Since we already assume the hardness of the isogeny-path problem, we are mainly interested in the hardness of computing a pseudo-isogeny path that does not correspond to an actual isogeny path.

Problem 2 (Non-isogeny path problem). Find points $T_0, T_1, \ldots, T_n \in \mathbb{P}^{2^g-1}$, such that $T_0 = \theta^{A_0}(0_{A_0})$ and $T_n = \theta^{A_n}(0_{A_n})$, where (A_0, θ^{A_0}) and (A_n, θ^{A_n}) are isogenous abelian varieties of dimension g with a given level-2 theta structure, but at least one pair T_i, T_{i+1} does not occur as the theta-null points of 2-isogenous abelian varieties of dimension g.

We analyze the hardness of Problem 1 and Problem 2 in Section 4, though note already now that for $g \leq 2$, Problem 2 is not only a hard problem, but in-fact an impossible problem.

In Appendix A we show that, at least in theory, additional constraints can be derived to assure for dimension $g > 2$ that a path really corresponds to a 2^n-isogeny. However, especially for $g \geq 4$, these constraints quickly become unwieldy and it would simply not be practical to explicitly include them. In the next section, we will argue that these extra constraints can in fact be left out without compromising security.

4 Equivalence of isogeny relations

The general relation that we want to give a proof of knowledge of is

$$\mathcal{R}_{\text{isog}(g)} := \left\{ ((A_0, A_1), \phi) \;\middle|\; \begin{smallmatrix} A_0, A_1 \text{ abelian varieties of dimension } g, \\ \phi : A_0 \to A_1 \text{ is an arbitrary } N\text{-isogeny for some } N \in \mathbb{N} \end{smallmatrix} \right\}.$$

In Section 3, we discussed the following two relations

$$\mathcal{R}_{\text{isog}(2^n,g)} := \left\{ ((A_0, A_1), \phi) \;\middle|\; \begin{smallmatrix} A_0, A_1 \text{ abelian varieties of dimension } g, \\ \phi : A_0 \to A_1 \text{ is a good } 2^n\text{-isogeny} \end{smallmatrix} \right\}.$$

and

$$\mathcal{R}_{\text{HS}(n,g)} := \left\{ \begin{pmatrix} ((A_0,\theta_0),(A_n,\theta_n)), \\ (T_i)_{i\in[n-1]} \end{pmatrix} \;\middle|\; \begin{smallmatrix} (A_0,\theta_0),(A_1,\theta_1) \text{ isogenous abelian varieties of} \\ \text{dimension } g \text{ with a level-2 theta structure,} \\ \mathcal{S}\circ\mathcal{H}(T_1) = \mathcal{H}\circ\mathcal{S}(\theta^{A_0}(0_{A_0})), \\ \mathcal{S}\circ\mathcal{H}(\theta^{A_n}(0_{A_n})) = \mathcal{H}\circ\mathcal{S}(T_{n-1}), \\ \mathcal{S}\circ\mathcal{H}(T_{i+1}) = \mathcal{H}\circ\mathcal{S}(T_i) \quad \forall i \in [n-2] \end{smallmatrix} \right\},$$

that are related to the difference of Problem 1 and Problem 2.

In Section 5 we will show how to give a zero-knowledge proof for $\mathcal{R}_{\text{HS}(n,g)}$. The goal for this section is to prove that (subrelations of) these relations are (weakly) equivalent, i.e. that there exists a polynomial time algorithm turning a witness for one into a witness for the other ones.

Before doing so, we must be more precise about what we mean by witness. We assume to be working over a fixed finite field k (generally $\mathbb{F}_p$ or $\mathbb{F}_{p^2}$ for some

prime p). Then a witness $(T_i)_{i\in[n-1]}$ to $\mathcal{R}_{\mathrm{HS}(n,g)}$ is a set of $2^g \times (n-1)$ values in k satisfying Equation (4). A witness for $\mathcal{R}_{\mathrm{isog}(g)}$ (resp. $\mathcal{R}_{\mathrm{isog}(2^n,g)}$) is an isogeny (resp. of degree 2^n) defined over k, that can be described and evaluated on points of $E_0(k)$ (resp. $A_0(k)$) in polynomial time. This goes by the name of *efficient isogeny representation* (see [64] for a survey).

4.1 Relation between $\mathcal{R}_{\mathbf{isog}(g)}$ and $\mathcal{R}_{\mathbf{isog}(2^n,g')}$

Our first goal will be to prove that $\mathcal{R}_{\mathrm{isog}(2^n,g')}$ is general enough to express elements of $\mathcal{R}_{\mathrm{isog}(g)}$. This is a direct consequence of Lemma 2.

We start with the following lemma.

Lemma 5. *There exists an algorithm, polynomial in* $\log n$*, that converts a witness* ϕ *for the statement* (A_0, A_1) *in* $\mathcal{R}_{isog(g)}$ *to a witness* Φ *for the statement* $(A_0^4 \times A_1^4, A_0^4 \times A_1^4)$ *in the relation* $\mathcal{R}_{isog(2^n,8g)}$ *for any* n *such that* $2^n > \deg(\phi)$*.*

Proof. This is just a rephrasing of Lemma 2. □

However, there is also a sense in which the corresponding relations in $\mathcal{R}_{\mathrm{isog}(2^n,g)}$ give back relations in $\mathcal{R}_{\mathrm{isog}(g)}$. To make this precise, we first define the subset $\mathcal{R}_{\mathrm{isog}(g)\leq B} \subseteq \mathcal{R}_{\mathrm{isog}(g)}$ as

$$\mathcal{R}_{\mathrm{isog}(g)\leq B} := \left\{((A_0, A_1), \phi) \in \mathcal{R}_{\mathrm{isog}(g)} \;\middle|\; \phi \text{ is an } N\text{-isogeny, for } N \leq B\right\}.$$

We must also consider a subset of $\mathcal{R}_{\mathrm{isog}(2^n,g)}$, which acts as witnesses for the same statements as in $\mathcal{R}_{\mathrm{isog}(g)}$. To do this, we define the relation

$$\mathcal{S}_{\mathrm{isog}(2^n,g,d)} := \left\{((A_0, A_1), \phi) \;\middle|\; \begin{array}{c} A_0, A_1 \text{ abelian varieties of dimension } g, \\ ((A_0\times B_0, A_1\times B_1), \Phi) \in \mathcal{R}_{\mathrm{isog}(2^n,dg)} \text{ for some ppav's} \\ B_0, B_1 \text{ and some isogeny } \Phi \end{array}\right\}$$

which injects into $\mathcal{R}_{\mathrm{isog}(2^n,dg)}$ by fixing a choice of representatives (B_0, B_1) for each (A_0, A_1). Notice also that any witness for (A_0, A_1) in $\mathcal{S}_{\mathrm{isog}(2^n,g,d)}$ is also a witness for the same statement in the relation $\mathcal{R}_{\mathrm{isog}(g)}$ (and even $\mathcal{R}_{\mathrm{isog}(g)\leq 2^n}$), since Φ encodes an efficient representation of an isogeny between A_0 and A_1 as one of its components.

As a consequence we obtain the following corollary.

Lemma 6. *The relations* $\mathcal{R}_{isog(g)\leq 2^n}$ *and* $\mathcal{S}_{isog(2^n,g,8)}$ *are equivalent.*

Proof. One direction is given by Lemma 5, and the other by a special case of the discussion above. □

4.2 Equivalence of $\mathcal{R}_{\mathbf{isog}(2^n,g)}$ and $\mathcal{R}_{\mathbf{HS}(n,g)}$ for $g = 1, 2$

A polynomial time algorithm turning a witness for the statement (A_0, A_1) in $\mathcal{R}_{\mathrm{isog}(2^n,g)}$ into a witness for (A_0, A_1) in $\mathcal{R}_{\mathrm{HS}(n,g)}$ (provided the $2^{\lfloor \frac{n}{2} \rfloor}$-torsion is accessible, see Remark 1) was described in Section 3.1. However, as we saw in Section 3.2, going the other way is more subtle. A witness $(T_i)_{i\in[n-1]}$ for

$\mathcal{R}_{\mathrm{HS}(n,g)}$ composed of a path of theta null points of abelian varieties can be used to evaluate the corresponding isogeny $\Phi : A_0 \to A_1$. Indeed at each step the theta null points of the codomain contain all the information needed to evaluate Φ (see [65, Chap. 8, Sec. 12]; more concretely, [33, Alg. 6] in dimension 2 and [27, Alg. 1] in dimension 4). In dimension 1 and 2 Lemma 3 guarantees that $(T_i)_{i\in[n-1]}$ will always encode an isogeny path. This is enough to prove the following partial result (note that the requirement on the torsion is only needed in one direction, in special cases).

Lemma 7. *Let $g \leq 2$. The relations $\mathcal{R}_{isog(2^n,g)}$ and $\mathcal{R}_{HS(n,g)}$ are equivalent, for statements where the $2^{\lfloor \frac{n}{2} \rfloor}$-torsion is accessible.*

In order to say something about the relation to $\mathcal{R}_{\mathrm{isog}(g)}$, we again define the analogous

$$\mathcal{S}_{\mathrm{HS}(n,g,d)} := \left\{ \begin{pmatrix} (A_0,A_1), \\ (T_i)_{i\in[n-1]} \end{pmatrix} \;\middle|\; \begin{array}{c} A_0,A_1 \text{ abelian varieties of dimension } g, \\ (((A_0\times B_0,\theta_0),(A_1\times B_1,\theta_1)),(T_i)_{i\in[n-1]})\in\mathcal{R}_{\mathrm{HS}(n,dg)}, \\ \text{for some ppav's } B_0,B_1, \text{ and level-2 theta structures } \theta_0,\theta_1 \end{array} \right\}.$$

Lemma 8. *There exists a polynomial time algorithm that converts a witness ϕ for the statement (E_0, E_1) in $\mathcal{S}_{isog(2^n,1,2)}$ to a witness Φ for the same statement in the relation $\mathcal{R}_{isog(1)\leq 2^n}$.*

Proof. The equivalence from Lemma 7 extends to an equivalence between $\mathcal{S}_{\mathrm{isog}(2^n,1,2)}$ and $\mathcal{S}_{\mathrm{HS}(n,1,2)}$ in the obvious way (and does not require the condition on the $2^{\lfloor \frac{n}{2} \rfloor}$-torsion in this direction). The result now follows from Lemma 6. □

4.3 Weak equivalence of $\mathcal{R}_{\mathbf{isog}(2^n,g)}$ and $\mathcal{R}_{\mathbf{HS}(n,g)}$ for $g > 2$

In dimension $g \geq 3$ the lack of direct correspondence between solutions to Equation (4) and isogenies prevents Lemma 7 from generalizing directly. A solution often adopted for instance in lattice based cryptography (see e.g. [56, Def. 2.7]) is to define an alternative relation $\tilde{\mathcal{R}}_{\mathrm{HS}(n,g)}$ where the transcript $\mathcal{T}$ is forced to correspond to an isogeny. We then say that $\mathcal{R}_{\mathrm{isog}(2^n,g)}$ and $\mathcal{R}_{\mathrm{HS}(n,g)}$ are *weakly equivalent* if $\mathcal{R}_{\mathrm{isog}(2^n,g)}$ is equivalent to $\tilde{\mathcal{R}}_{\mathrm{HS}(n,g)}$ (which is now implied by Section 4.2) and for a given statement (A_0, A_1) finding a witness in $\mathcal{R}_{\mathrm{HS}(n,g)} \setminus \tilde{\mathcal{R}}_{\mathrm{HS}(n,g)}$ is hard regardless of the knowledge of a witness in $\tilde{\mathcal{R}}_{\mathrm{HS}(n,g)}$. In fact, we will argue something even stronger: that finding any element at all in $\mathcal{R}_{\mathrm{HS}(n,g)} \setminus \tilde{\mathcal{R}}_{\mathrm{HS}(n,g)}$ is hard, i.e. that Problem 2 is a hard problem.

Either way, the assumption that Problem 1 is hard is exactly saying that $\mathcal{R}_{\mathrm{HS}(n,g)}$ is a hard relation. Problem 2 results in the following lemma.

Lemma 9. *Assuming the hardness of Problem 2, the relations $\mathcal{R}_{isog(2^n,g)}$ and $\mathcal{R}_{HS(n,g)}$ (resp. $\mathcal{S}_{isog(2^n,g,d)}$ and $\mathcal{S}_{HS(n,g,d)}$) are weakly equivalent for any n, g (resp. any n, g, d), for statements where the $2^{\lfloor \frac{n}{2} \rfloor}$-torsion is accessible.*

Proof. Clear from previous discussion. □

We then finally obtain the following theorem.

Theorem 1. *Assuming the hardness of Problem 2, the relations $\mathcal{R}_{isog(g)\leq 2^n}$ and $\mathcal{S}_{HS(n,g,8)}$ are weakly equivalent, for statements where the $2^{\lfloor \frac{n}{2} \rfloor}$-torsion is accessible.*

Proof. Combine Lemma 6 and Lemma 9. □

We again stress that the requirement on the $2^{\lfloor \frac{n}{2} \rfloor}$-torsion is only necessary for one direction of the equivalence above.

4.4 Cryptanalysis of Problem 2

As we have seen, if we want $\mathcal{R}_{\mathrm{HS}(n,g)}$ to be a hard relation, we rely on the hardness of Problem 1. Further, if we wish for proving statements in $\mathcal{R}_{\mathrm{HS}(n,g)}$ to be sufficient for proving statements in $\mathcal{R}_{\mathrm{isog}(g)}$, we rely on the hardness of Problem 2.

To the best of our knowledge, neither problem has been studied before. We will in this section limit our discussion to Problem 2, since the hardness of Problem 2 already implies that the usual isogeny-path problem reduces to Problem 1.

If we fix a certain theta null point X^0 in Equation (4) we see that the possible solutions are given by

$$X^1 = \frac{1}{2^g}\mathcal{H} \circ \mathcal{T}_s \circ \mathcal{H} \circ \mathcal{S}(X^0) \tag{6}$$

where

$$\mathcal{T}_s : \mathbb{P}^{2^g-1} \to \mathbb{P}^{2^g-1} : (x_0 : \cdots : x_{2^g-1}) \quad \to (x_0 : s_1\sqrt{x_0x_1} \cdots : s_{2^g-1}\sqrt{x_0x_{2^g-1}}),$$

is the map defined in [51, Cor. 5] with the difference that we now allow all $2^g - 1$ possible sign choices. If we at any point take an "invalid" choice of square roots, we are likely to step out of the moduli space defined by the theta-null points of abelian varieties. Heuristically, if we then model the resulting X^1 as a random element of $\mathbb{P}^{2^g-1}$, there is a priori no reason why the map $\mathcal{T}_s$ should produce new elements in $\mathbb{P}^{2^g-1}(k)$. This happens if and only if all the square-roots are rational, which happens with probability $1/2^{2^g-1}$ when working over finite fields, thus most "invalid" path choices end immediately. We illustrate this in Figure 1 in Appendix B, where we explore the neighbourhood of a theta-null point coming from a superspecial abelian threefold.

Even ignoring the fact that most counterfeit paths end immediately, the ability to produce random values would cryptographically not be enough to break Problem 2. Under the heuristic that the pseudo-path null points acts as random points of $\mathbb{P}^{2^g-1}$, the moduli space defined by the theta null-points of abelian varieties of dimension $g > 2$ always has co-dimension ≥ 1, and thus we roughly expect such a random, bruteforcing strategy to have complexity $O(p)$ in the absolute best case of $g = 3$ (and typically, much worse). To test this analysis, we also ran a depth-first search for solutions to Problem 2 in low characteristic

for superspecial abelian threefolds, which strongly suggests that the complexity of such a bruteforce search for solutions to Problem 2 does grow very quickly with the characteristic p. See Table 1 for the results.

	$p = 3$	$p = 7$	$p = 11$	$p = 19$
Avg. #vertices visited	328	9994	72442	852349
Median #vertices visited	230	1361	16554	514930

Table 1: Complexity of depth-first search against Problem 2 in dim. $g = 3$ over 100 runs.

Further, we also ran several breath-first searchers around singular theta-null points (such as those with extra zeroes, or those coming from reducible varieties) under the hypothesis that such theta-null points could be a source of solutions to Problem 2, but these searches did not yield any solutions in higher characteristic. In practice, we did this by traversing the full graph of pseudo-paths starting from E_0^3, for $E_0 : y^2 = x^3 + x$, and building the subgraph consisting of only isogeny-paths along the way (i.e. the edges and vertices in black in Figure 1). The code used for our tests can be found in our repository.

Remark 2. Given the novelty of Problem 2, we can of course not exclude that there is extra structure in these pseudo-isogeny graphs which can be used to leverage an attack. However, at least for superspecial abelian varieties, there is a vague relation to another very difficult problem; namely the problem of producing "non-backdoored" superspecial abelian varieties. In dimension 1, this problem, usually referred to as hashing into the supersingular isogeny graph, has been notoriously difficult, and has still eluded a solution (see [13] for several failed attempts). Since solutions to Problem 2 do not correspond to isogenies, it is unclear if such a path leaks anything meaningful about the endomorphism ring of the final abelian variety, and thus a solution to Problem 2 could potentially give a solution to the higher-dimensional analogue of hashing into the supersingular isogeny graph, though more insight into Problem 2 is clearly required to assert this claim.

5 Constraint systems for proofs of knowledge of isogenies

In the previous sections, we have shown that proving knowledge of a length n path of 2-isogenies between two principally polarized abelian varieties A_0, A_n of dimension g is computationally equivalent to proving knowledge of a vector $v \in (k^{2^g})^{n+1}$ such that $v[0] = \theta_{A_1}, v[n] = \theta_{A_n}$, and

$$\forall i \in [n-1] \, : \, \mathcal{S} \circ \mathcal{H}(v[i+1]) = \mathcal{H} \circ \mathcal{S}(v[i]). \tag{7}$$

We now transform the relation in Equation (7) to an equivalent relation that is more conducive to the construction of a zero-knowledge proof of knowledge. Define the Hadamard matrix $H \in k^{2^g \times 2^g}$ such that $\mathcal{H} : k^{2^g} \to k^{2^g} : x \mapsto Hx$. We prove knowledge of a vector $z \in k^{2^{g+1}(n+1)}$ such that $z[0 \dots 2^g - 1] = \theta_{A_0} \wedge z[2^g n \dots 2^g(n+1) - 1] = \theta_{A_n}$ and

$$\left(\left[\begin{array}{cccc|c} & & & & \\ & I & & & 0 \\ & & & & \\ \hline 0 & H & 0 & 0 & \\ 0 & 0 & \ddots & 0 & 0 \\ 0 & 0 & 0 & H & \end{array}\right] z \right)^2 = \left[\begin{array}{cccc|c} & & & & \\ & 0 & & & I \\ & & & & \\ \hline & & & & H \ 0 \ 0 \ 0 \\ & 0 & & & 0 \ \ddots \ 0 \ 0 \\ & & & & 0 \ 0 \ H \ 0 \end{array}\right] z, \quad (8)$$

where the square of a vector denotes the Hadamard product with itself. Clearly, $z = (v, v \circ v)$, with v the vector containing the theta-null points of the isogeny path, satisfies this relation. In the following we will denote the matrix on the left by M_1 and the one on the right by M_2, so the above constraint system can be written as $(M_1 z) \circ (M_1 z) = M_2 z$.

Optimizing constraint systems. Previous research on proving knowledge of isogeny relations often follows the strategy: 1) construct an alternative relation that is equivalent to the isogeny relation, 2) express this alternative relation in a constraint system for which there exist post-quantum secure proof systems, 3) gauge the performance of the resulting proof system based on the size of the resulting constraint system. Usually, the relations are expressed in the R1CS constaint system which is defined by three sparse matrices $A, B, C \in k^{m \times n}$ with $\mathcal{O}(m) = \mathcal{O}(n)$ non-zero entries such that a transcript $z \in k^n$ satisfies this constraint system if and only if $(Az) \circ (Bz) = Cz$.

As a result, previous research has focused mostly on minimizing the dimensions and the number of non-zero entries when expressing the isogeny relation in an R1CS constraint system. Research on proof systems has conclusively shown that, even though this cannot give any asymptotic improvement in performance, practical performance can be greatly improved by expressing the relation in specialized constraint systems such as Plonk, CCS and GKR. Let us illustrate this with a simple example. The relation in Equation (8) can be expressed in R1CS by setting $A = B = M_1$ and $C = M_2$. However, it should be clear that the performance of any proof system for R1CS can be improved by specializing it to R1CS instances where $A = B$. Additionally, many of the most efficient proof systems will actually pad the number of rows and columns to a power of two and thus have limited performance difference within the same power of two.

For these reasons, we will also report actual concrete implementation results in the following sections, whereas most previous works relied on extrapolated estimates.

Representing $\mathbb{F}_{p^2}$ as $\mathbb{F}_p \times \mathbb{F}_p$. In the most important use case of dimension 4, namely the CSIDH class group action, we will have $k = \mathbb{F}_p$ for some prime p. In dimension 1 and 2 however, most use cases will have $k = \mathbb{F}_{p^2}$. Due to many proof system implementations only natively supporting prime fields, previous work has discussed constraint system sizes resulting from both expression over the native field $\mathbb{F}_{p^2}$, as well as expression over $\mathbb{F}_p$ by emulating $\mathbb{F}_{p^2}$. We will also discuss both approaches using the same techniques for representing $\mathbb{F}_{p^2}$ as $\mathbb{F}_p \times \mathbb{F}_p$ and report implementation results comparing both.

Note that any element in $\mathbb{F}_{p^2}$ can be written as $a + b\alpha$ where $\alpha^2 = d$ is a non-square residue. To express the square $c + d\alpha := (a + b\alpha)^2$ using the R1CS constraint system, we append $\mathbb{F}_p$ elements $u := (a + b)(a + bd)$ and $v := ab$ to the transcript. Expressing the relation between a, b, u, v can be achieved using only two R1CS constraints. The $\mathbb{F}_p$ elements representing the square can now be expressed "for free" as $c = u - (d + 1)v$ and $d = 2v$.

Recall the matrices $M_1, M_2 \in k^{(2^{g+1}(n+1)-2^g)\times 2^{g+1}(n+1)}$ such that Equation (8) is equivalent to $(M_1 z) \circ (M_1 z) = M_2 z$. We use this constraint system to prove isogeny relations in dimension $g = 4$ for $k = \mathbb{F}_p$ and in dimensions $g = 1, 2$ for $k = \mathbb{F}_{p^2}$. Let us briefly describe our approach for proving isogeny relations in dimension $g = 1$ using proof systems over $k = \mathbb{F}_p$. We write one step of Equation (7) for $v[0] = [x_1 + y_1\alpha, x_2 + y_2\alpha]$ and $v[1] = [\mathsf{x}_1 + \mathsf{y}_1\alpha, \mathsf{x}_2 + \mathsf{y}_2\alpha]$ as

$$\begin{aligned} x_1^2 + y_1^2 d \pm (x_2^2 + y_2^2 d) &= \mathsf{x}_1^2 + \mathsf{y}_1^2 d \pm 2(\mathsf{x}_1\mathsf{x}_2 + \mathsf{y}_1\mathsf{y}_2 d) + \mathsf{x}_2^2 + \mathsf{y}_2^2 d \\ 2x_1y_1 \pm 2x_2y_2 &= 2\mathsf{x}_1\mathsf{y}_1 \pm 2(\mathsf{x}_1\mathsf{y}_2 + \mathsf{x}_2\mathsf{y}_1) + 2\mathsf{x}_2\mathsf{y}_2. \end{aligned}$$

which we can rewrite (using the u and v notation introduced above) as

$$\begin{aligned} u_1 - (d+1)v_1 \pm u_2 - (d+1)v_2 &= \mathsf{u}_1 - (d+1)\mathsf{v}_1 \pm 2(\mathsf{x}_1\mathsf{x}_2 + \mathsf{y}_1\mathsf{y}_2 d) + \mathsf{u}_2 - (d+1)\mathsf{v}_2 \\ v_1 \pm v_2 &= \mathsf{v}_1 \pm 2(\mathsf{x}_1\mathsf{y}_2 + \mathsf{x}_2\mathsf{y}_1) + \mathsf{v}_2. \end{aligned}$$

Notice that we can not express these four equations as only four rows of an R1CS instance. Instead, we add the cross-terms x_1x_2, x_1y_2 to the transcript (which each require one R1CS row to express), and then express the previous equations as four R1CS constraints. We follow a similar approach for dimension $g = 2$. In this case, the square of the Hadamard results in six cross-terms and similarly we add two finite field elements per cross-term. We refer to Table 2 for total sizes of the resulting R1CS constraint systems.

6 Proof systems for proving knowledge of isogenies

In this section we will discuss different proof systems that prove the constraint systems described in Section 5. Since we argue that proof systems for proving knowledge of isogeny relations should ideally be compared based on measured runtimes instead of sizes of R1CS constraint systems, we provide concrete runtimes in Section 8.

	$g = 1, k = 256$			$g = 2, k = 108$			$g = 4, k = 498$		
	m	n	nz	m	n	nz	m	n	nz
$\mathbb{F}_p$	2556	2560	12772	4096	4104	28560	16384	15936	270400
$\mathbb{F}_{p^2}$	1024	1024	3064	1024	864	4288		n/a	

Table 2: The number of rows m, number of columns n and the number of non-zero elements of the R1CS constraint system resulting from expressing Equation (7) in R1CS over either $\mathbb{F}_p$ or $\mathbb{F}_{p^2}$ for dimensions $g = 1, 2, 4$ where k is the path length corresponding to the applications described in Section 7. In the cases that do not require emulation, we do not count non-zero elements of matrix B since it is exactly equal to A.

Furthermore, we want to point out that all succinct non-interactive proof systems must rely on non-falsifiable computational assumptions [44]. Since no succinct isogeny-based proof system is currently known in the literature, proof systems for proving knowledge of isogeny relations must make additional assumptions besides the isogeny path problem. Ideally, these proof systems only additionally assume the existence of (quantum) random oracles, which is already common in isogeny-based cryptography. For this reason, previous work refers to hash-based succinct proof systems such as Aurora [9] and Ligero [3] whose soundness relies on the proximity gaps phenomenon for Reed-Solomon codes. However, the most efficient instantiations of hash-based proof systems (in terms of proof size and verification time) rely on parameter sets for the proximity gaps that were only conjectured to be secure. Recent work has now disproven this conjecture [40, 26]. Instantiating these proof systems with provably secure parameters will increase proof size and verification time by at least a factor of 2-3x.

For this reason, we will describe both a hash-based proof system and a lattice-based proof system that is able to achieve smaller proof size and verification time at the cost of making additional assumptions. In Section 6.1, we describe a multilinear proof system that uses the BaseFold [69] hash-based polynomial commitment scheme as a building block. We choose BaseFold since it is completely field-agnostic. Note that many efficient proof systems will require that $\mathbb{F}_p$ contains 2^k-roots of unity with $k \gg 0$, which is incompatible with the $p = 3 \bmod 4$ requirement common for many isogeny use cases. In Section 6.2, we design a lattice-based proof system that achieves a lower proving time and a smaller proof size.

6.1 Multilinear PIOP for Equation (8)

The section describes a Polynomial Interactive Oracle Proof (PIOP) [10] for the relation implicit in Equation (8). This is an adaptation of generic techniques for multilinear PIOPs [67] to our constraint system. As discussed in Section 5, emulating $\mathbb{F}_{p^2}$ for dimensions $g = 1, 2$ will instead require PIOPs for the full R1CS constraint system.

Let us define the transcript size $N := 2^{g+1}(n+1)$ and $m := \log_2(N)$. For ease of notation we assume M_1, M_2 were padded to become square matrices. Our proof of knowledge starts by proving knowledge of vectors $z_1, z_2 \in k^N$ such that $z_1 \circ z_1 = z_2$. We start by restating this relation using multilinear polynomials

$$\begin{aligned} z_1^2 = z_2 &\Leftrightarrow \forall X \in \{0,1\}^m \,:\, \tilde{z}_1(X)^2 = \tilde{z}_2(X) \\ &\Leftrightarrow F_{rc}(X) := \sum_{b \in \{0,1\}^m} \tilde{\mathsf{eq}}(X,b)(\tilde{z}_1(b)^2 - \tilde{z}_2(b)) = 0\,. \end{aligned}$$

For some random $\tau \in k^m$, the Schwartz-Zippel lemma (Lemma 4) states that $F_{rc}(\tau) = 0$ implies $F_{rc} = 0$ except with probability $1/|k|$. The multilinear sumcheck protocol reduces this sum to the claim that

$$h_m(r_m) = \tilde{\mathsf{eq}}(\tau, r)(\tilde{z}_1(r)^2 - \tilde{z}_2(r)) \tag{9}$$

where $h_m \in k^{\leq 3}[Y]$ is the last polynomial sent in the sumcheck protocol and $r \in k^m$ is the random evaluation point. The verifier checks the claim $h_m(r_m) \stackrel{?}{=} \tilde{\mathsf{eq}}(\tau, r)(a_1^2 - a_2(r))$ where $a_1, a_2 \in k$ were provided by the prover and the evaluations $\tilde{\mathsf{eq}}(\tau, r), h_m(r_m)$ can be computed by the verifier themself. Now the verifier only has to check the validity of the claim that $a_1 \stackrel{?}{=} \tilde{z}_1(r)$ and $a_2 \stackrel{?}{=} \tilde{z}_2(r)$.

Again, we restate this claim using multilinear polynomials so that we can leverage the sumcheck protocol. For some $i \in \{1,2\}$ and $\tilde{z}$ a multilinear extension of z, we can see that

$$v_i = \tilde{z}_i(r) \Leftrightarrow a_i = \sum_{b \in \{0,1\}^m} \tilde{M}_i(r,b)\tilde{z}(b)$$

where $\tilde{M}_i$ is a multilinear extension of the M_i matrix such that the first n variables represent the row index and the last n variables represent the column index. Again, we can leverage the Schwartz-Zippel lemma (Lemma 4) to check this claim for all i by instead verifying whether $r_1^* a_1 + r_2^* a_2 \stackrel{?}{=} r_1^* \tilde{z}_1(r) + r_2^* \tilde{z}_2(r)$. The sumcheck protocol will reduce this claim to

$$h'_m(r'_m) \stackrel{?}{=} (r_1^* \tilde{M}_1(r,r') + r_2^* \tilde{M}_2(r,r')) \cdot \tilde{z}(r').$$

Notice that $\tilde{M}_1(r,r'), \tilde{M}_2(r,r')$ can be efficiently computed by the verifier themself, but $\tilde{z}(r')$ has to be provided by the prover in a verifiable way. We refer to Figure 2 in Appendix C for a formal description of the multilinear PIOP described in this section.

This PIOP can be compiled into a zero-knowledge Succinct Non-interactive ARgument of Knowledge (zkSNARK) by applying standard techniques [10, 67]. As previously mentioned, our implementation will prove the final evaluation of $\tilde{z}$ using the BaseFold Polynomial Commitment Scheme (PCS). Adding zero-knowledge to this PIOP and making it non-interactive will not be discussed since it has negligible impact on performance.

6.2 Lattice-based sigma protocol with linear and quadratic relations

In this section, we propose a new type of proof system based on lattice hardness assumptions instead of the proximity gaps conjecture. Importantly, this proof system should have the same post-quantum secure, field-agnostic and zero-knowledge properties. Relative to the proof system described in Section 6.1, we achieve both lower prover computation costs and lower proof size. Note that, even though it is asymptotically succinct, the previous proof system has proof size 7.6MiB in dimension 4 for a transcript that has size 0.96MiB. This can be ascribed to the fact that Basefold is not practically succinct for small instances. The proof system below will achieve a 1.2MiB proof size while being almost twice as fast. In particular, we propose a lattice-based sigma protocol that proves Equation (7) by recursively proving knowledge of a verifiable proof for the PIOP described in Section 6.1. Firstly, we describe the lattice-based sigma protocol and next we describe how it can prove Equation (7).

The protocol from Figure 3 in Appendix C proves knowledge of an opening $c \in k^{m_2}$ for a BDLOP [8]-style commitment Com constructed as

$$\begin{bmatrix} t_1 \\ t_2 \end{bmatrix} = \begin{bmatrix} B_1 \\ B_2 \end{bmatrix} \cdot s + \begin{bmatrix} 0 \\ c \end{bmatrix} \tag{10}$$

for some public parameters $B_1, B_2 \in k^{(m_1+m_2)\times r}$ and some secret $s \in k^r$ such that $\|s\| \leq \mathsf{B}_s$. We define the challenge distribution $\Gamma \coloneqq \{\gamma \mid \gamma \leq \mathsf{B}_\gamma\} \subset k$ and the discrete Gaussian distribution $\mathcal{D}_\sigma$ over k with standard deviation σ. Notice that this commitment scheme is binding if $\mathsf{SIS}_{k,m_1,r,2\mathsf{B}_s}$ is hard. Similarly, the hiding property depends on the hardness of $\mathsf{LWE}_{k,r-m_1-m_2,m_1+m_2,\mathsf{B}_s}$. This sigma protocol is able to prove v linear relations expressed as $Rc = u$ using the matrix $R \in k^{v\times m_2}$ and $u \in k^v$.

The sigma protocol from Figure 3 derives its simulatability from both the hiding property of the commitment scheme combined with the rejection sampling which ensures that z is statistically indistinguishable from a random value sampled from $\mathcal{D}$. Let us discuss the knowledge soundness property in more detail. Notice that an extractor can perform rewinding to obtain two accepting transcripts $(w, v, \gamma^{(1)}, z^{(1)}), (w, v, \gamma^{(2)}, z^{(2)}))$ such that

$$B_1(z^{(1)} - z^{(2)}) = t_1(\gamma^{(1)} - \gamma^{(2)}) \tag{11}$$

and thereby effectively extracting an opening

$$\bar{s} \coloneqq \bar{z}/\bar{\gamma} \coloneqq (z^{(1)} - z^{(2)})/(\gamma^{(1)} - \gamma^{(2)}).$$

It should be clear that the opening also satisfies $R\bar{c} = u$ for $\bar{c} \coloneqq t_2 - B_2\bar{s}$. Next, we show that $\bar{s}$ (and therefore also $\bar{c}$) is unique given the hardness of $\mathsf{SIS}_{k,m_1,r,8\mathsf{B}_\gamma\sigma\sqrt{2r}}$. Any extracted second opening $\bar{s}'$ has to also satisfy Equation (11) and thus

$$B_1(\bar{s} - \bar{s}') = 0 \Leftrightarrow B_1(\bar{s}\bar{\gamma}\bar{\gamma}' - \bar{s}'\bar{\gamma}\bar{\gamma}') = 0 \Leftrightarrow B_1(\bar{z}\bar{\gamma}' - \bar{z}'\bar{\gamma}) = 0 \tag{12}$$

for $\|\bar{z}\bar{\gamma}'\|, \|\bar{z}'\bar{\gamma}\| \leq 4\mathsf{B}_\gamma\sigma\sqrt{2r}$. Lastly, we show this protocol is complete except with negligible probability. Clearly, the two linear relations checked by the verifier will always satisfy for honest provers. Following [55, Lemma 2.2], we can state that the inequality will hold with overwhelming probability as long as $r \geq 640$. Following [55, Lemma 2.14], the rejection sampling subroutine returns after an expected $M \approx 3$ repetitions if we set $\sigma = 13\mathsf{B}_\gamma\mathsf{B}_s$. The proof size of this protocol is approximately $|t| + |z| + |\gamma|$ in the non-interactive setting. In that case, the verifier uses the linear relations to compute w, v and then uses them to audit the sampling of γ through Fiat-Shamir with aborts [54].

Notice that almost all verification steps in the protocol from Section 6.1 are linear and can therefore be expressed as a linear relation $Rc = u$ where c contains the proof transcript. The only exception is Equation (9), which we can not prove using the protocol from Figure 3. We can restate this as proving knowledge of h_m, a_1, a_1', a_2 such that $a_1' = a_2 + h_m(r_m)/\widetilde{\mathsf{eq}}(\tau, r)$ and $a_1' = a_1^2$. To achieve this we define the Laurent polynomials

$$p(Y) \coloneqq a_1(Y + Y^{-1}) + a_1'Y^2 + a_3Y^3$$
$$q(Y) \coloneqq p(Y) \cdot (p(Y) - 2Y^{-2})$$

for some random $a_3 \in k$. Now notice that $q(0) = 0$ if and only if $a_1' = a_1^2$. First, the prover commits to a_1, a_1', a_3, a_2 and to $q(Y) = q_1Y^{-3} + q_2Y^{-2} + q_3Y^{-1} + q_4Y + q_5Y^2 + q_6Y^3 + q_7Y^4 + q_8Y^5 + q_9Y^6$ such that $q(0) = 0$ by definition. Then, the verifier responds with a randomly sampled $y \leftarrow k$. The prover responds with $p(y), q(y)$, which the verifier can use to check $q(y) \stackrel{?}{=} p(y)(p(y) - 2y^{-2})$. Since only one evaluation of $p(Y)$ will be known to the verifier, the addition of a_3 is sufficient for zero-knowledge.

Proof composition. In the interest of minimizing the proof size, we propose performing one recursive step. In other words, we recursively prove knowledge of a valid proof for the proof system decribed in Figure 2 using the proof system described in Figure 3. Specifically, the prover will simulate execution of the multilinear PIOP to generate a transcript c, which will become the message committed to in Equation (10). Then the prover executes the "outer" proof system, i.e. the one described in Figure 3.

Notice that this proof system is capable of efficiently representing all verification operations of the "inner" proof system except the Fiat-Shamir transform required to make it non-interactive. We can avoid proving the Fiat-Shamir transform in our outer proof system by first transforming the inner proof system using the commit-and-prove transform before proof composition. Concretely, every prover message in Figure 2 is replaced by a commitment to that message. The verifier does not perform any operation on the messages and only sends back random challenges. As final message, the prover opens all those commitments. Then, the verifier can verify those openings and perform the deferred verification operations.

To summarize, we transform the protocol from Figure 2 to a commit-and-prove protocol that uses the lattice-based commitment scheme from Equation (10).

Next, we use the Fiat-Shamir transform to make the proof system non-interactive. Lastly, we prove all verification operations on commitment openings in the resulting protocol recursively using the proof system from Figure 3. We present a formal description of the interactive variant (before the Fiat-Shamir transform) of the resulting protocol in Figure 4 in Appendix C.

7 Applications

In this section we discuss the implications, theoretical as well as practical, of our construction.

7.1 Proofs of isogeny relations

Our main theoretical contribution is providing a zero-knowledge proof of knowledge for the relation

$$\mathcal{R}_{\mathrm{isog}(g)} := \left\{ ((A_0, A_1), \phi) \;\middle|\; \substack{A_0, A_1 \text{ abelian varieties of dimension } g, \\ \phi : A_0 \to A_1 \text{ is an arbitrary } N\text{-isogeny for some } N \in \mathbb{N}} \right\}.$$

This fully solves (a generalization of) one of the major open problems stated in [11], namely how to prove $\mathcal{R}_{\mathrm{isog}}$ without knowledge of endomorphism rings and without leaking the degree of ϕ. As already observed there, tailored protocols like SQIsign rely on nontrivial information about the curves involved (e.g. their endomorphism rings) while generic proofs in dimension 1 are constrained to $\deg(\phi)$ being smooth (on top of leaking $\deg(\phi)$ itself).

Another important relation mentioned in [11] is for CSIDH isogenies

$$\mathcal{R}_{\mathrm{CSIDH}} := \{((E_0, E_1), \phi) \mid \phi : E_0 \to E_1 \text{ is an isogeny defined over } \mathbb{F}_p\}.$$

Thanks to qt-Pegasis, such an isogeny ϕ can always be embedded into a 2^n-isogeny in dimension 4, i.e. $\mathcal{R}_{\mathrm{CSIDH}}$ can be expressed in $\mathcal{R}_{\mathrm{isog}(2^n,4)}$. Our framework thus applies directly, providing a completely novel proof of knowledge for $\mathcal{R}_{\mathrm{CSIDH}}$ compared to the graph isomorphism protocol due to Couveignes. This shows once more the great generality of our approach. The fact that, as discussed in Section 8, our proving time is comparable to the ID-protocol by Couveignes using the recent heavily optimized implementation of qt-Pegasis [29] (although proof sizes are 1 to 2 orders of magnitude larger, see Section 8) furthermore shows that this flexibility comes at a moderate cost.

7.2 Additional isogeny relations

Given the generality of our framework, one natural question is how easy it is to extend it to prove more specific isogeny relations. We leave a detailed answer to this question for future work, but we sketch below a possible approach in two examples.

Proving the evaluation of points Sometimes it is important to prove, together with the knowledge of an isogeny, the (scaled) evaluation of points through that isogeny (see e.g. [4, Sec 5.2]). Evaluation of points in the theta model resembles very closely codomain computation (compare, for instance, Algorithm 5 and 6 in [34]). We thus expect relations similar to Equation (4) to hold for theta coordinates of points.

Proving the degree of an isogeny Another relevant property of an isogeny that one might be interested to prove is its degree. This is for instance the case of PRISM [5], and is captured by the relation $\mathcal{R}_{\text{deg}}$ in [11]. When applying Lemma 2 we effectively hide the degree of the embedded isogeny (which is in general a desirable zero-knowledge property). To add it to the proof we can perform the verification mechanism of PRISM in zero-knowledge. This requires evaluating points (which we discussed above) and computing a pairing. Pairing computation in zero knowledge is a well known problem [42, 60]. We thus expect that the wide literature on the topic can be adapted to the setting of the level 2 theta model.

Evaluating points and degrees are some of the most important properties of an isogeny that one might be interested in proving. For instance, they immediately imply a proof for the relation $\mathcal{R}_{\text{M-SIDH}}$ (the only one from [11] not explicitly covered above).

7.3 Applications to higher dimensional isogenies

As a more concrete application of our construction we already mentioned an alternative proof of knowledge of a CSIDH isogeny built with qt-Pegasis. Another interesting application is a trusted setup for the 2-dimensional CGL hash function [17, 51]. In a similar fashion to how KLPT [50] showed that the original CGL hash function [20] is insecure if instantiated from a curve of known endomorphism ring, the recent 2-dimensional variant of KLPT [16] showed that the same is true for the 2-dimensional variant of CGL. The trusted setup for CGL proposed in [6] obviates this problem in dimension 1, and is currently one of the main applications of generic proofs of knowledge of isogenies. Thanks to our construction, this multiparty trusted setup can be replicated directly in dimension 2. We provide a proof of concept implementation proving a random walk from [51]. To the best of our knowledge, this is the first instantiation of the 2 dimensional CGL hash function which is secure against [16]. Notice that, since we are working in dimension 2, we do not need to rely on the hardness of Problem 2.

Finally, we also mention that recent work by Robert [66] uses 2^n-isogenies in dimension 4, to construct a non interactive key exchange named $\otimes$-MIKE. Our framework can directly be applied to proving such an isogeny. Although this has no direct applications for $\otimes$-MIKE itself, it is likely to be a useful feature in future protocols built on $\otimes$-MIKE.

8 Implementation

We implement our constuction in the following cases:

- in dimension 1 and 2, we prove knowledge of a random walk; as discussed in Section 7.3, in dimension 2 this can be used to securely instantiate the dimension 2 CGL hash function preventing the attack from [16];
- in dimension 4, we prove the computation of a qt-Pegasis isogeny [32]; this provides an alternative proof of the $\mathcal{R}_{\text{CSIDH}}$ relation discussed in Section 7.1.

We could not implement our construction in dimension 8 since, to the best of our knowledge, 8-dimensional isogenies in the level 2 theta model have not been implemented yet. However, we stress that once such an implementation exists, adapting it to also produce transcripts to our framework will not be more difficult than in dimension 2 or 4. Our implementation can be found at

https://github.com/KULeuven-COSIC/unified_zk_iso

8.1 Building the transcript

We implement the generation of a transcript satisfying Equation (4) directly on top of some of the most common isogeny computation libraries, to facilitate the usability of our code. In particular, we make the following choices:

- in dimension 1, we build on the implementation of [34]; we work over $\mathbb{F}_{p^2}$ where $p = 5 \cdot 2^{248} - 1$ is the SQIsign level I prime, and perform a random walk of length 256 from $E_0 : y^2 = x^3 + x$;
- in dimension 2, we prove a random walk as an evaluation of the CGL hash function from [51] (originally due to [68]); we adopt their parameters, i.e. $\mathbb{F}_{p^2}$ where $p = 2^{127} - 1$, and perform a random walk of length 108 from $E_0 \times E_0$;
- in dimension 4, we integrate the transcript generation into the qt-Pegasis code; we thus work over $\mathbb{F}_p$ where $p = 27 \cdot 2^{500} - 1$ and perform 498 steps.

The procedure is the same in all cases: for each isogeny step we can use the theta coordinates X^j in the transcript, and then update the coordinates of X^{j+1} so that Equation (3) holds affinely. Since it must hold projectively, it is enough to compute for instance the coefficient c_j such that

$$\sum_{i=0}^{2^g-1} \left(X_i^j\right)^2 = c_j \left(\sum_{i=0}^{2^g-1} X_i^{j+1}\right)^2,$$

and then divide X^{j+1} by $\sqrt{c_j}$. This costs a single square root per step. Alternatively, one could also keep track of the coefficients c_j and modify the proof system accordingly. Since all the above mentioned libraries work projectively, it is always possible to scale X^j by a given factor. Notice that this procedure works for all types of isogenies (e.g. also for gluing and splitting isogenies).

8.2 Proof systems

In Table 3, we report measured results from implementing the Basefold-based proof system described in Section 6.1 in Rust and running it on an Apple Mac mini M4 Pro with 64GB RAM.

		$g = 1, k = 256$	$g = 2, k = 108$	$g = 4, k = 498$
$\mathbb{F}_p$	Prover time	1288ms	966ms	10961ms
	Proof size	4826 KiB	6320KiB	7833KiB
	Verif. time	74ms	45ms	228ms
$\mathbb{F}_{p^2}$	Prover time	1846ms	1074ms	n/a
	Proof size	3344KiB	3767KiB	
	Verif. time	286ms	179ms	

Table 3: The prover time, proof size and verification time over either $\mathbb{F}_p$ or $\mathbb{F}_{p^2}$ for dimensions $g = 1, 2, 4$ where k is the path length corresponding to the applications described in Section 7.

We would like to note that using similar extrapolation techniques as in [48], we would arrive at smaller and faster proofs, e.g. 8.5s prover time and 1.08MiB proof size for our dimension 4 case using the Aurora proof system. However, as explained before, the parameters used in these implementations are not provably secure. We refer to Table 4 for our instantiations of the lattice parameters described in Section 6.2 such that the SIS and LWE instances mentioned are secure for $\lambda = 128, g = 4$ using the lattice estimator by Albrecht et al. [2].

m_1	m_2	r	B_s	B_γ	p
2^6	16496	$2^{14} + 2^{12} + m_1 + m_2$	2^{28}	2^{128}	$27 \cdot 2^{500} - 1$

Table 4: Lattice-based sigma protocol parameters described in Section 6.2.

Measured results from our implementation are reported in Table 5. We would like to note that even though the proof system described in Section 6.2 is based on the LWE problem for simplicity, we instead implemented a version of this proof system extended to $k[X]/(X^{2^d}+1)$, i.e. the RLWE version, using standard techniques. This results in increased verifier performance (compare to $d = 0$), at the cost of slightly worse prover performance.

In dimension 1 our R1CS instance is roughly a factor 2 bigger than the current state of the art. This is due to the fact that we use 2 projective coordinates instead of a single affine one. The reason for this choice is to show the generality

d	Prover time (s)	Proof size (KiB)	Verification time (s)
0	5.2	1147	105.5
6	6.0	1156	37.4
8	6.5	1200	10.8
11	8.6	1648	3.8

Table 5: Results for proof system described in Section 6.2 for $g = 4$ and $k = 498$.

of our approach; our dimension 1 instance can indeed be seen as an introduction to dimension 2 and 4.

In dimension 4 we can compare with the tailored proof of CSI-FiSh instantiated with the recent optimized implementation of qt-Pegasis [29]. The cost of performing 128 actions is estimated around 3 seconds, and the proof size is 8.3KiB. Even though our approach is fully general, our proving time is less than 2 times slower. On the other hand, proofs are 1.2MiB. We stress that both proving time and proof sizes (especially in the lattice setting) are largely inflated by working with a longer chain over a larger field. Results for applications not threatened by Kuperberg's algorithm should rather be extrapolated by our dimension 1 and 2 implementation.

In dimension 2, our application is fully novel. The timing and sizes provided in Table 3 can be used to estimate the efficiency of the trusted setup for the CGL hash function described in Section 7.3.

References

[1] M. A. Aardal, G. Adj, D. F. Aranha, A. Basso, I. A. Canales Martínez, J. Chávez-Saab, M. C. Santos, P. Dartois, L. De Feo, M. Duparc, J. K. Eriksen, T. B. Fouotsa, D. L. G. Filho, B. Hess, D. Kohel, A. Leroux, P. Longa, L. Maino, M. Meyer, K. Nakagawa, H. Onuki, L. Panny, S. Patranabis, C. Petit, G. Pope, K. Reijnders, D. Robert, F. Rodríguez Henríquez, S. Schaeffler, and B. Wesolowski. *SQIsign*. Tech. rep. available at https://csrc.nist.gov/Projects/pqc-dig-sig/round-2-additional-signatures. National Institute of Standards and Technology, 2024.

[2] M. Albrecht, R. Player, and S. Scott. "On the concrete hardness of Learning with Errors". In: *Journal of Mathematical Cryptology* 9 (Oct. 2015). DOI: 10.1515/jmc-2015-0016.

[3] S. Ames, C. Hazay, Y. Ishai, and M. Venkitasubramaniam. "Ligero: Lightweight Sublinear Arguments Without a Trusted Setup". In: *ACM CCS 2017*. Ed. by B. M. Thuraisingham, D. Evans, T. Malkin, and D. Xu. Dallas, TX, USA: ACM Press, 2017, pp. 2087–2104. DOI: 10.1145/3133956.3134104.

[4] A. Basso. "A Post-Quantum Round-Optimal Oblivious PRF from Isogenies". In: *SAC 2023*. Ed. by C. Carlet, K. Mandal, and V. Rijmen. Vol. 14201. LNCS. Fredericton, Canada: Springer, Cham, Switzerland, 2024, pp. 147–168. DOI: 10.1007/978-3-031-53368-6_8.

[5] A. Basso, G. Borin, W. Castryck, M. Corte-Real Santos, R. Invernizzi, A. Leroux, L. Maino, F. Vercauteren, and B. Wesolowski. "PRISM: Simple and Compact Identification and Signatures from Large Prime Degree Isogenies". In: *PKC 2025, Part III*. Ed. by T. Jager and J. Pan. Vol. 15676. LNCS. Røros, Norway: Springer, Cham, Switzerland, 2025, pp. 300–332. DOI: 10.1007/978-3-031-91826-1_10.

[6] A. Basso, G. Codogni, D. Connolly, L. De Feo, T. B. Fouotsa, G. M. Lido, T. Morrison, L. Panny, S. Patranabis, and B. Wesolowski. "Supersingular Curves You Can Trust". In: *EUROCRYPT 2023, Part II*. Ed. by C. Hazay and M. Stam. Vol. 14005. LNCS. Lyon, France: Springer, Cham, Switzerland, 2023, pp. 405–437. DOI: 10.1007/978-3-031-30617-4_14.

[7] A. Basso, P. Dartois, L. De Feo, A. Leroux, L. Maino, G. Pope, D. Robert, and B. Wesolowski. "SQIsign2D-West - The Fast, the Small, and the Safer". In: *ASIACRYPT 2024, Part III*. Ed. by K.-M. Chung and Y. Sasaki. Vol. 15486. LNCS. Kolkata, India: Springer, Singapore, Singapore, 2024, pp. 339–370. DOI: 10.1007/978-981-96-0891-1_11.

[8] C. Baum, I. Damgård, V. Lyubashevsky, S. Oechsner, and C. Peikert. "More Efficient Commitments from Structured Lattice Assumptions". In: *SCN 18*. Ed. by D. Catalano and R. De Prisco. Vol. 11035. LNCS. Amalfi, Italy: Springer, Cham, Switzerland, 2018, pp. 368–385. DOI: 10.1007/978-3-319-98113-0_20.

[9] E. Ben-Sasson, A. Chiesa, M. Riabzev, N. Spooner, M. Virza, and N. P. Ward. "Aurora: Transparent Succinct Arguments for R1CS". In: *EUROCRYPT 2019, Part I*. Ed. by Y. Ishai and V. Rijmen. Vol. 11476. LNCS.

Darmstadt, Germany: Springer, Cham, Switzerland, 2019, pp. 103–128. DOI: 10.1007/978-3-030-17653-2_4.

[10] E. Ben-Sasson, A. Chiesa, and N. Spooner. "Interactive Oracle Proofs". In: *TCC 2016-B, Part II*. Ed. by M. Hirt and A. D. Smith. Vol. 9986. LNCS. Beijing, China: Springer Berlin Heidelberg, Germany, 2016, pp. 31–60. DOI: 10.1007/978-3-662-53644-5_2.

[11] W. Beullens, L. De Feo, S. D. Galbraith, and C. Petit. "Proving knowledge of isogenies: a survey". In: *DCC* 91.11 (2023), pp. 3425–3456. DOI: 10.1007/s10623-023-01243-3.

[12] W. Beullens, T. Kleinjung, and F. Vercauteren. "CSI-FiSh: Efficient Isogeny Based Signatures Through Class Group Computations". In: *ASIACRYPT 2019, Part I*. Ed. by S. D. Galbraith and S. Moriai. Vol. 11921. LNCS. Kobe, Japan: Springer, Cham, Switzerland, 2019, pp. 227–247. DOI: 10.1007/978-3-030-34578-5_9.

[13] J. Booher, R. Bowden, J. Doliskani, T. B. Fouotsa, S. D. Galbraith, S. Kunzweiler, S. Merz, C. Petit, B. Smith, K. E. Stange, Y. B. Ti, C. Vincent, J. F. Voloch, C. Weitkämper, and L. Zobernig. "Failing to Hash Into Supersingular Isogeny Graphs". In: *Comput. J.* 67.8 (2024), pp. 2702–2719. DOI: 10.1093/COMJNL/BXAE038. URL: https://doi.org/10.1093/comjnl/bxae038.

[14] G. Borin, M. Corte-Real Santos, J. K. Eriksen, R. Invernizzi, M. Mula, S. Schaeffler, and F. Vercauteren. *Qlapoti: Simple and Efficient Translation of Quaternion Ideals to Isogenies*. Cryptology ePrint Archive, Report 2025/1604. 2025. URL: https://eprint.iacr.org/2025/1604.

[15] W. Castryck and T. Decru. "An Efficient Key Recovery Attack on SIDH". In: *EUROCRYPT 2023, Part V*. Ed. by C. Hazay and M. Stam. Vol. 14008. LNCS. Lyon, France: Springer, Cham, Switzerland, 2023, pp. 423–447. DOI: 10.1007/978-3-031-30589-4_15.

[16] W. Castryck, T. Decru, P. Kutas, A. Laval, C. Petit, and Y. B. Ti. "KLPT^2: Algebraic Pathfinding in Dimension Two and Applications". In: *CRYPTO 2025, Part I*. Ed. by Y. T. Kalai and S. F. Kamara. Vol. 16000. LNCS. Santa Barbara, CA, USA: Springer, Cham, Switzerland, 2025, pp. 167–200. DOI: 10.1007/978-3-032-01855-7_6.

[17] W. Castryck, T. Decru, and B. Smith. "Hash functions from superspecial genus-2 curves using Richelot isogenies". In: *J. Math. Cryptol.* 14.1 (2020), pp. 268–292. DOI: 10.1515/JMC-2019-0021. URL: https://doi.org/10.1515/jmc-2019-0021.

[18] W. Castryck, T. Lange, C. Martindale, L. Panny, and J. Renes. "CSIDH: An Efficient Post-Quantum Commutative Group Action". In: *ASIACRYPT 2018, Part III*. Ed. by T. Peyrin and S. Galbraith. Vol. 11274. LNCS. Brisbane, Queensland, Australia: Springer, Cham, Switzerland, 2018, pp. 395–427. DOI: 10.1007/978-3-030-03332-3_15.

[19] Castryck, Wouter and Decru, Thomas. "Multiradical isogenies". eng. In: *Arithmetic, gometry, cryptography, and coding theory 2021: 18th International Conference Arithmetic, Geometry, Cryptography, and Coding The-*

ory May 31–June 4, 2021 Centre International de Rencontres Mathématiques, Marseille, France. Ed. by Anni, Samuele and Karemaker, Valentijn and Lorenzo García, Elisa. Vol. 779. Contemporary mathematics. Marseille, France: American Mathematical Society, 2022, 57–89. ISBN: 9781470467944. URL: {http://doi.org/10.1090/conm/779}.

[20] D. X. Charles, K. E. Lauter, and E. Z. Goren. "Cryptographic Hash Functions from Expander Graphs". In: *Journal of Cryptology* 22.1 (Jan. 2009), pp. 93–113. DOI: 10.1007/s00145-007-9002-x.

[21] J. Chávez-Saab, F. Rodríguez-Henríquez, and M. Tibouchi. "Verifiable Isogeny Walks: Towards an Isogeny-Based Postquantum VDF". In: *SAC 2021: 28th*. Ed. by R. AlTawy and A. Hülsing. Vol. 13203. LNCS. Virtual Event: Springer, Cham, Switzerland, 2022, pp. 441–460. DOI: 10.1007/978-3-030-99277-4_21.

[22] K. Cong, Y.-F. Lai, and S. Levin. "Efficient Isogeny Proofs Using Generic Techniques". In: *ACNS 2023, Part II*. Ed. by M. Tibouchi and X. Wang. Vol. 13906. LNCS. Kyoto, Japan: Springer, Cham, Switzerland, 2023, pp. 248–275. DOI: 10.1007/978-3-031-33491-7_10.

[23] B. Conrad. *Polarizations*. Online lecture notes. VIGRE 2004, Stanford University. 2004. URL: https://math.stanford.edu/~conrad/vigregroup/vigre04/polarization.pdf.

[24] M. Corte-Real Santos, J. K. Eriksen, M. Meyer, and K. Reijnders. "AprèsSQI: Extra Fast Verification for SQIsign Using Extension-Field Signing". In: *EUROCRYPT 2024, Part I*. Ed. by M. Joye and G. Leander. Vol. 14651. LNCS. Zurich, Switzerland: Springer, Cham, Switzerland, 2024, pp. 63–93. DOI: 10.1007/978-3-031-58716-0_3.

[25] J.-M. Couveignes. *Hard Homogeneous Spaces*. Cryptology ePrint Archive, Report 2006/291. 2006. URL: https://eprint.iacr.org/2006/291.

[26] E. Crites and A. Stewart. *On Reed–Solomon Proximity Gaps Conjectures*. Cryptology ePrint Archive, Paper 2025/2046. 2025. URL: https://eprint.iacr.org/2025/2046.

[27] P. Dartois. *Fast computation of 2-isogenies in dimension 4 and cryptographic applications*. Cryptology ePrint Archive, Report 2024/1180. 2024. URL: https://eprint.iacr.org/2024/1180.

[28] P. Dartois. "Fast computation of higher dimensional isogenies for cryptographic applications". PhD thesis. Université de Bordeaux, 2025.

[29] P. Dartois and M. Duparc. *Chasing Rabbits Through Hypercubes: Better algorithms for higher dimensional 2-isogeny computations*. Cryptology ePrint Archive, Paper 2026/114. 2026. URL: https://eprint.iacr.org/2026/114.

[30] P. Dartois, J. K. Eriksen, T. B. Fouotsa, A. H. L. Merdy, R. Invernizzi, D. Robert, R. Rueger, F. Vercauteren, and B. Wesolowski. *PEGASIS: Practical Effective Class Group Action using 4-Dimensional Isogenies*. Cryptology ePrint Archive, Report 2025/401. 2025. URL: https://eprint.iacr.org/2025/401.

[31] P. Dartois, J. K. Eriksen, T. B. Fouotsa, A. H. L. Merdy, R. Invernizzi, D. Robert, R. Rueger, F. Vercauteren, and B. Wesolowski. "PEGASIS: Practical Effective Class Group Action using 4-Dimensional Isogenies". In: *CRYPTO 2025, Part I*. Ed. by Y. T. Kalai and S. F. Kamara. Vol. 16000. LNCS. Santa Barbara, CA, USA: Springer, Cham, Switzerland, 2025, pp. 67–99. DOI: 10.1007/978-3-032-01855-7_3.

[32] P. Dartois, J. K. Eriksen, R. Invernizzi, and F. Vercauteren. *qt-Pegasis: Simpler and Faster Effective Class Group Actions*. Cryptology ePrint Archive, Report 2025/1859. 2025. URL: https://eprint.iacr.org/2025/1859.

[33] P. Dartois, L. Maino, G. Pope, and D. Robert. *An Algorithmic Approach to* $(2, 2)$*-isogenies in the Theta Model and Applications to Isogeny-based Cryptography*. Cryptology ePrint Archive, Report 2023/1747. 2023. URL: https://eprint.iacr.org/2023/1747.

[34] P. Dartois, L. Maino, G. Pope, and D. Robert. "An Algorithmic Approach to (2, 2)-Isogenies in the Theta Model and Applications to Isogeny-Based Cryptography". In: *ASIACRYPT 2024, Part III*. Ed. by K.-M. Chung and Y. Sasaki. Vol. 15486. LNCS. Kolkata, India: Springer, Singapore, Singapore, 2024, pp. 304–338. DOI: 10.1007/978-981-96-0891-1_10.

[35] L. De Feo, S. Dobson, S. D. Galbraith, and L. Zobernig. "SIDH Proof of Knowledge". In: *ASIACRYPT 2022, Part II*. Ed. by S. Agrawal and D. Lin. Vol. 13792. LNCS. Taipei, Taiwan: Springer, Cham, Switzerland, 2022, pp. 310–339. DOI: 10.1007/978-3-031-22966-4_11.

[36] L. De Feo and S. D. Galbraith. "SeaSign: Compact Isogeny Signatures from Class Group Actions". In: *EUROCRYPT 2019, Part III*. Ed. by Y. Ishai and V. Rijmen. Vol. 11478. LNCS. Darmstadt, Germany: Springer, Cham, Switzerland, 2019, pp. 759–789. DOI: 10.1007/978-3-030-17659-4_26.

[37] L. De Feo, D. Jao, and J. Plût. *Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies*. Cryptology ePrint Archive, Report 2011/506. 2011. URL: https://eprint.iacr.org/2011/506.

[38] L. De Feo, D. Kohel, A. Leroux, C. Petit, and B. Wesolowski. "SQISign: Compact Post-quantum Signatures from Quaternions and Isogenies". In: *ASIACRYPT 2020, Part I*. Ed. by S. Moriai and H. Wang. Vol. 12491. LNCS. Daejeon, South Korea: Springer, Cham, Switzerland, 2020, pp. 64–93. DOI: 10.1007/978-3-030-64837-4_3.

[39] L. De Feo, A. Leroux, P. Longa, and B. Wesolowski. "New Algorithms for the Deuring Correspondence - Towards Practical and Secure SQISign Signatures". In: *EUROCRYPT 2023, Part V*. Ed. by C. Hazay and M. Stam. Vol. 14008. LNCS. Lyon, France: Springer, Cham, Switzerland, 2023, pp. 659–690. DOI: 10.1007/978-3-031-30589-4_23.

[40] B. E. Diamond and A. Gruen. *On the Distribution of the Distances of Random Words*. Cryptology ePrint Archive, Report 2025/2010. 2025. URL: https://eprint.iacr.org/2025/2010.

[41] M. Duparc. *Superglue: Fast formulae for (2,2)-gluing isogenies*. Cryptology ePrint Archive, Report 2025/736. 2025. URL: https://eprint.iacr.org/2025/736.

[42] Y. El Housni. "Pairings in Rank-1 Constraint Systems". In: *ACNS 2023, Part I*. Ed. by M. Tibouchi and X. Wang. Vol. 13905. LNCS. Kyoto, Japan: Springer, Cham, Switzerland, 2023, pp. 339–362. DOI: 10.1007/978-3-031-33488-7_13.

[43] H. M. Farkas, S. Grushevsky, and R. Salvati Manni. "An explicit solution to the weak Schottky problem". In: *Algebraic Geometry* 8.3 (2021), pp. 358–373. DOI: 10.14231/AG-2021-009.

[44] C. Gentry and D. Wichs. "Separating succinct non-interactive arguments from all falsifiable assumptions". In: *43rd ACM STOC*. Ed. by L. Fortnow and S. P. Vadhan. San Jose, CA, USA: ACM Press, 2011, pp. 99–108. DOI: 10.1145/1993636.1993651.

[45] W. Ghantous, F. Pintore, and M. Veroni. *Collisions in Supersingular Isogeny Graphs and the SIDH-based Identification Protocol*. Cryptology ePrint Archive, Report 2021/1051. 2021. URL: https://eprint.iacr.org/2021/1051.

[46] O. Goldreich, S. Micali, and A. Wigderson. "Proofs That Yield Nothing But Their Validity Or All Languages in NP Have Zero-Knowledge Proof Systems". In: *Journal of the ACM* 38.3 (July 1991), pp. 691–729. DOI: 10.1145/116825.116852.

[47] T. den Hollander, S. Kleine, M. Mula, D. Slamanig, and S. A. Spindler. "More Efficient Isogeny Proofs of Knowledge via Canonical Modular Polynomials". In: *CRYPTO 2025, Part I*. Ed. by Y. T. Kalai and S. F. Kamara. Vol. 16000. LNCS. Santa Barbara, CA, USA: Springer, Cham, Switzerland, 2025, pp. 131–166. DOI: 10.1007/978-3-032-01855-7_5.

[48] T. den Hollander, M. Mula, D. Slamanig, and S. A. Spindler. *On the Use of Atkin and Weber Modular Polynomials in Isogeny Proofs of Knowledge*. Cryptology ePrint Archive, Paper 2026/193. 2026. URL: https://eprint.iacr.org/2026/193.

[49] E. Kani. "The number of curves of genus two with elliptic differentials". In: *Journal für die reine und angewandte Mathematik* 1997.485 (1997), pp. 93–122. DOI: 10.1515/crll.1997.485.93.

[50] D. Kohel, K. Lauter, C. Petit, and J.-P. Tignol. "On the quaternion-isogeny path problem". In: *LMS Journal of Computation and Mathematics* 17.A (2014), pp. 418–432.

[51] S. Kunzweiler, L. Maino, T. Moriya, C. Petit, G. Pope, D. Robert, M. Stopar, and Y. B. Ti. "Radical 2-Isogenies and Cryptographic Hash Functions in Dimensions 1, 2 and 3". In: *PKC 2025, Part III*. Ed. by T. Jager and J. Pan. Vol. 15676. LNCS. Røros, Norway: Springer, Cham, Switzerland, 2025, pp. 265–299. DOI: 10.1007/978-3-031-91826-1_9.

[52] S. Levin and R. Pedersen. *Faster Proofs and VRFs from Isogenies*. Cryptology ePrint Archive, Report 2024/1626. 2024. URL: https://eprint.iacr.org/2024/1626.

[53] C. Lund, L. Fortnow, H. Karloff, and N. Nisan. "Algebraic methods for interactive proof systems". In: *J. ACM* 39.4 (Oct. 1992), 859–868. ISSN: 0004-5411. DOI: 10.1145/146585.146605. URL: https://doi.org/10.1145/146585.146605.

[54] V. Lyubashevsky. "Fiat-Shamir with Aborts: Applications to Lattice and Factoring-Based Signatures". In: *ASIACRYPT 2009*. Ed. by M. Matsui. Vol. 5912. LNCS. Tokyo, Japan: Springer Berlin Heidelberg, Germany, 2009, pp. 598–616. DOI: 10.1007/978-3-642-10366-7_35.

[55] V. Lyubashevsky, N. K. Nguyen, and M. Plançon. "Lattice-Based Zero-Knowledge Proofs and Applications: Shorter, Simpler, and More General". In: *CRYPTO 2022, Part II*. Ed. by Y. Dodis and T. Shrimpton. Vol. 13508. LNCS. Santa Barbara, CA, USA: Springer, Cham, Switzerland, 2022, pp. 71–101. DOI: 10.1007/978-3-031-15979-4_3.

[56] V. Lyubashevsky, N. K. Nguyen, and G. Seiler. "Shorter Lattice-Based Zero-Knowledge Proofs via One-Time Commitments". In: *PKC 2021, Part I*. Ed. by J. Garay. Vol. 12710. LNCS. Virtual Event: Springer, Cham, Switzerland, 2021, pp. 215–241. DOI: 10.1007/978-3-030-75245-3_9.

[57] L. Maino, C. Martindale, L. Panny, G. Pope, and B. Wesolowski. "A Direct Key Recovery Attack on SIDH". In: *EUROCRYPT 2023, Part V*. Ed. by C. Hazay and M. Stam. Vol. 14008. LNCS. Lyon, France: Springer, Cham, Switzerland, 2023, pp. 448–471. DOI: 10.1007/978-3-031-30589-4_16.

[58] J. S. Milne. *Abelian Varieties (v2.00)*. Available at www.jmilne.org/math/. 2008.

[59] D. Mumford. "On the equations defining abelian varieties I". In: *Inventiones Mathematicae* 1.4 (1966), pp. 287–354. DOI: 10.1007/BF01389737.

[60] A. Novakovic and L. Eagen. *On Proving Pairings*. Cryptology ePrint Archive, Report 2024/640. 2024. URL: https://eprint.iacr.org/2024/640.

[61] A. Page and D. Robert. *Introducing Clapoti(s): Evaluating the isogeny class group action in polynomial time*. Cryptology ePrint Archive, Report 2023/1766. 2023. URL: https://eprint.iacr.org/2023/1766.

[62] D. Robert. *Evaluating isogenies in polylogarithmic time*. Cryptology ePrint Archive, Report 2022/1068. 2022. URL: https://eprint.iacr.org/2022/1068.

[63] D. Robert. "Breaking SIDH in Polynomial Time". In: *EUROCRYPT 2023, Part V*. Ed. by C. Hazay and M. Stam. Vol. 14008. LNCS. Lyon, France: Springer, Cham, Switzerland, 2023, pp. 472–503. DOI: 10.1007/978-3-031-30589-4_17.

[64] D. Robert. *On the efficient representation of isogenies (a survey)*. Cryptology ePrint Archive, Report 2024/1071. 2024. URL: https://eprint.iacr.org/2024/1071.

[65] D. Robert. *Some notes on algorithms for abelian varieties*. Cryptology ePrint Archive, Report 2024/406. 2024. URL: https://eprint.iacr.org/2024/406.

[66] D. Robert. *The module action for isogeny based cryptography*. Cryptology ePrint Archive, Report 2024/1556. 2024. URL: https://eprint.iacr.org/2024/1556.

[67] S. Setty. "Spartan: Efficient and General-Purpose zkSNARKs Without Trusted Setup". In: *CRYPTO 2020, Part III*. Ed. by D. Micciancio and T. Ristenpart. Vol. 12172. LNCS. Santa Barbara, CA, USA: Springer, Cham, Switzerland, 2020, pp. 704–737. DOI: 10.1007/978-3-030-56877-1_25.

[68] K. Takashima. "Efficient Algorithms for Isogeny Sequences and Their Cryptographic Applications". In: *Mathematical Modelling for Next-Generation Cryptography: CREST Crypto-Math Project*. Ed. by T. Takagi, M. Wakayama, K. Tanaka, N. Kunihiro, K. Kimoto, and D. H. Duong. Mathematics for Industry. Springer Singapore, 2017, pp. 97–114. DOI: 10.1007/978-981-10-5065-7_6. URL: https://doi.org/10.1007/978-981-10-5065-7_6.

[69] H. Zeilberger, B. Chen, and B. Fisch. "BaseFold: Efficient Field-Agnostic Polynomial Commitment Schemes from Foldable Codes". In: *CRYPTO 2024, Part X*. Ed. by L. Reyzin and D. Stebila. Vol. 14929. LNCS. Santa Barbara, CA, USA: Springer, Cham, Switzerland, 2024, pp. 138–169. DOI: 10.1007/978-3-031-68403-6_5.

A Additional relations

A.1 Additional relations for $g > 2$

In principle, it is possible to prove that the solutions to Equation (4) actually come from theta null-points of principally polarized abelian varieties even for $g > 2$, by providing additional relations. The missing ingredient is that we have to express that each of the g-tuples T_i in the solution actually corresponds to the theta-null point of a principally polarized abelian variety.

These extra relations can be derived explicitly from the Riemann relations [59] between level-4 theta null points (see [51, Proposition 7] for this proof technique). From these Riemann relations, we will then derive $(2^g - 1) - (g(g+1)/2)$ extra conditions that define valid level-2 theta null points.

Dimension $g = 3$. An example of this in dimension $g = 3$ is the following, which is taken directly from [51]. For $g = 3$ we have $2^g - 1 = 7$ and $g(g+1)/2 = 6$ and we thus need a single additional equation, that can be derived directly from the Riemann relations.

Theorem 2. *Let $T := (a_{000}, \ldots, a_{111}) \in \mathbb{A}^8$ (indices are in binary). Defining*

$$(\beta_{00}, \ldots, \beta_{11}, \delta_{00}, \ldots, \delta_{11}) = \mathcal{H} \circ \mathcal{S}(T)$$

and for $i \in (\mathbb{Z}/2\mathbb{Z})^2$, $\gamma_i := 2\sum_{j \in (\mathbb{Z}/2\mathbb{Z})^2} (-1)^{\langle i,j \rangle} a_{0j} a_{1j}$,

$$R_1 = \prod_i \beta_i, \quad R_2 = \prod_i \gamma_i, \quad R_3 = \prod_i \delta_i$$

Then $(a_{000} : \cdots : a_{111}) \in \mathbb{P}^7$ is the level-2 theta null point of an abelian variety of dimension 3 if and only if

$$R_1^2 + R_2^2 + R_3^2 - 2(R_1R_2 + R_1R_3 + R_2R_3) = 0. \tag{13}$$

Further, this relation is enough to completely determine the last sign in Lemma 3.

Proof. See [51, Proposition 7] and [51, Theorem 8]. □

The above theorem shows that for $g = 3$ it suffices to add one extra non-linear condition given by Equation (13) on T for it to be a valid level-2 theta null point. Note that even in this case the degree of this extra relation is already 8.

Dimension $g = 4$. A similar approach can be taken for dimension $g = 4$. Let $T := (a_{0000}, \ldots, a_{1111}) \in \mathbb{A}^{16}$, then in this case we will need to derive $5 = (2^g - 1) - (g(g+1)/2)$ extra relations amongst the coordinates of T for it to represent a valid level-2 theta null point. We will not derive these equations fully, but only sketch the complexity involved. Following a similar technique as [51, Proposition 7], these extra relations will be derived from the Riemann relations between level-4 theta null coordinates. Compared to a level-2 theta structure, a level-4 theta structure defines a map $\theta^A : A \to \mathbb{P}^{2^{2g}-1}$ and we can thus index these level-4 coordinates by elements of $(\mathbb{Z}/2\mathbb{Z})^{2g}$.

The Riemann relations (see for instance [43, Equation 2.1] for an explicit expression) define identities between level-4 theta null coordinates of abelian varieties. To illustrate the nature of these relations, we provide an example.

Example 1. Since we have $2^{2g} = 256$ level-4 theta coordinates for $g = 4$, we will index them with integers $0 \leq i < 256$. One of the many Riemann relations that needs to hold is then given by:

$$\begin{aligned} 3 \cdot [0, 64, 128, 192] =& [1, 65, 129, 193] + [2, 66, 130, 194] + [3, 67, 131, 195] + \\ & [16, 80, 144, 208] + [18, 82, 146, 210] + [32, 96, 160, 224] + \\ & [33, 97, 161, 225] + [48, 112, 176, 240] + [51, 115, 179, 243] \end{aligned}$$

In the above, each vector with 4 indices corresponds to a product of the corresponding level-4 theta null coordinates, e.g. $[0, 64, 128, 192] = \theta_0 \cdot \theta_{64} \cdot \theta_{128} \cdot \theta_{192}$. The above relation between level-4 theta null coordinates can then be turned into a relation between level-2 theta null coordinates by using the fact that the squares of the level-4 theta null coordinates can be written as a combination of the level-2 theta null coordinates. If we denote each of the 10 terms in the above equation as r_i for $i = 0, \ldots, 9$, then the above equation now reads $3r_0 - \sum_{i=1}^{9} r_i = 0$. To derive an equation involving only the squares of the r_i (since these are expressions in the level-2 theta null points), we need to consider the product over all possible sign twists

$$\prod_{\delta_i \in \{0,1\}} (3r_0 - \sum_{i=1}^{9} (-1)^{\delta_i} r_i) = 0\,.$$

Note that this equation defines a relation of degree 2^8 between the squares of the r_i. Since each r_i is a product of four level-4 theta null coordinates, and each square of a level-4 theta null coordinate can can be written as a linear combination of a product of two level-2 theta null coordinates, we conclude that the above equation defines a relation of degree 2048 between the level-2 theta null coordinates. □

As the above example illustrates, the extra condition coming from the Riemann relations is given by a polynomial of degree 2^8 in 10 variables, namely the squares r_i^2, where each of the r_i^2 is an expression of degree 8 in the level-2 theta null coordinates. Although these equations are very explicit, their complexity

is totally prohibitive (the homogeneous polynomial of degree 2^8 in 10 variables would involve $> 2^{53}$ terms) to verify explicitly in every step of the isogeny chain. It might be possible to derive different equations of lower complexity, but it is rather unlikely these will be practical.

B Random walk graph

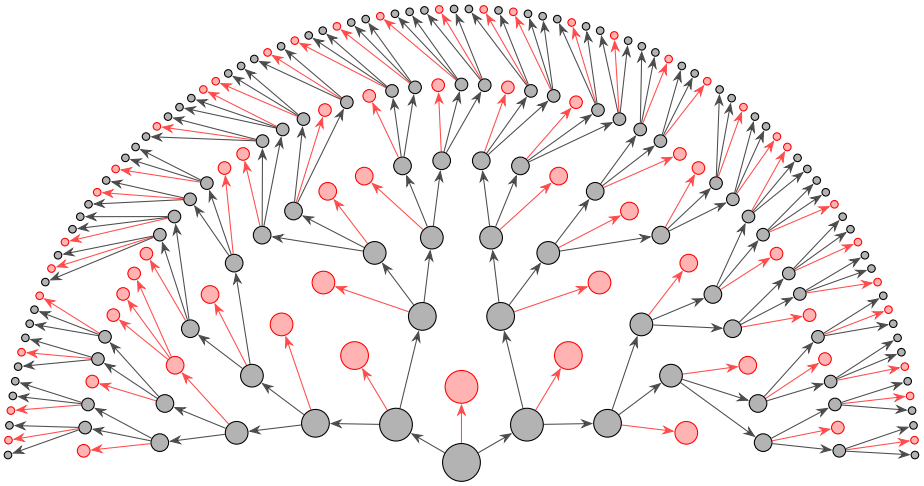

Fig. 1: Random walks starting from a superspecial theta null point of an abelian threefold over $\mathbb{F}_{p^2}$ for $p = 2^{61} - 1$. Red vertices correspond to tuples that do not correspond to abelian varieties. Note that from each vertex, we only choose 3 outgoing paths, if possible. This is for illustrative purposes (in total, there are $2^{2^g-1} = 2^7$ valid out-going paths, if there are any)

C Proof of Knowledge protocols

PoK of $v \in (k^{2^g})^n$ that satisfies Equation (7)

Prover		**Verifier**
$z = (v, v \circ v)$	$\xrightarrow{z}$	
		$\tau \leftarrow k^m$
	$\xleftarrow{\tau}$	
Sumcheck for $0 = \sum_{b \in \{0,1\}^m} \tilde{\mathsf{eq}}(\tau, b)(\tilde{z}_1(b)^2 - \tilde{z}_2(b))$	$\xleftrightarrow[r_1, \ldots, r_m]{h_1(Y), \ldots, h_m(Y)}$	$h_{i-1}(r_{i-1}) \stackrel{?}{=} h_i(0) + h_i(1)$
		$r_i \leftarrow k^m$
$a_1 = \tilde{z}_1(r), a_2 = \tilde{z}_2(r)$		
	$\xrightarrow{a_1, a_2}$	
		$r_1^*, r_2^* \leftarrow k$
	$\xleftarrow{r_1^*, r_2^*}$	
Sumcheck for $r_1^* a_1 + r_2^* a_2 = \sum_{b \in \{0,1\}^m} (\tilde{M}_1(r, b) + \tilde{M}_2(r, b))\tilde{z}(b)$	$\xleftrightarrow[r'_1, \ldots, r'_m]{h'_1(Y), \ldots, h'_m(Y)}$	$h_{i-1}(r_{i-1}) \stackrel{?}{=} h_i(0) + h_i(1)$
		$r'_i \leftarrow k^m$
$a_3 = \tilde{z}(r)$		
	$\xrightarrow{a_3}$	$h'_m(r'_m) \stackrel{?}{=} (\tilde{M}_1(r, r') + \tilde{M}_2(r, r'))a_3$
		$a_3 \stackrel{?}{=} \tilde{z}(r)$

Fig. 2: Multilinear PIOP for Equation (8)

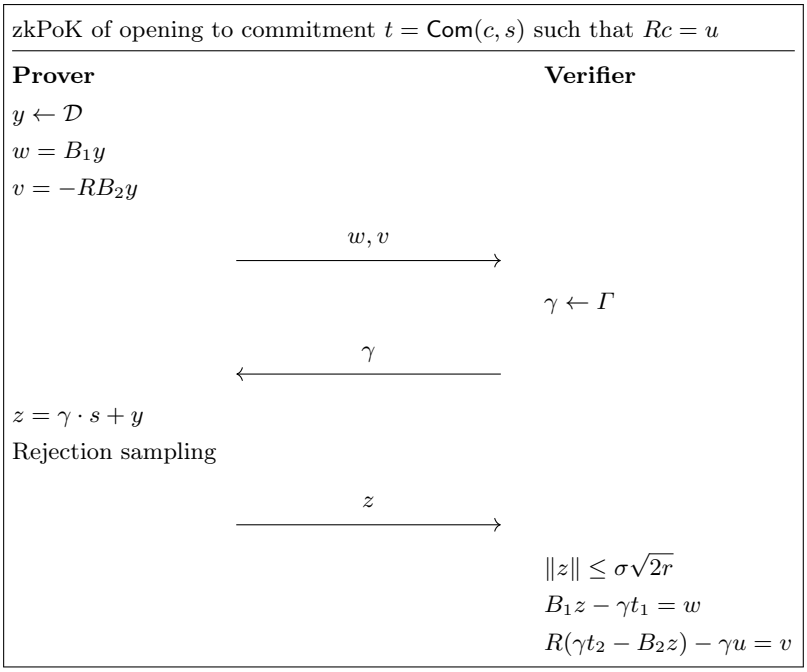

Fig. 3: Lattice-based zero-knowledge proof of knowledge with linear relation

zkPoK of $v \in (k^{2^g})^n$ that satisfies Equation (7)

Prover		**Verifier**
$t_1 = B_1 s, t_{2,1} = B_{2,1} s + z$	$\xrightarrow{t_1, t_{2,1}}$	
	$\xleftarrow{\tau}$	$\tau \leftarrow k^n$
..... Sumcheck for $\sum_{b \in \{0,1\}^n} \tilde{\mathsf{eq}}(\tau, b)(\tilde{z}_1(b)^2 - \tilde{z}_2(b)) = 0$		
		$h_0(r_0) \coloneqq 0$
$t_{2,i} = B_{2,i} s + (h_i(-1), h_i(0), h_i(1), h_i(2))$	$\xrightarrow{t_{2,i}}$	$r_i \leftarrow k^n$
	$\xleftarrow{r_i}$	Append $R, u : h_i(0) + h_i(1) \stackrel{?}{=} h_{i-1}(r_{i-1})$
.....		
$a_3 \leftarrow k, t_{2,n+1} = B_{2,n+1} s + (\tilde{z}_1(r), \tilde{z}_1(r)^2, a_3, \tilde{z}_2(r))$	$\xrightarrow{t_{2,n+1}}$	Append $R, u : a_1' \stackrel{?}{=} a_2 + h_n(r_n)/\tilde{\mathsf{eq}}(\tau, r)$
$t_{2,n+2} = B_{2,n+2} s + (q_1, \ldots, q_9)$	$\xrightarrow{t_{2,n+2}}$	
	$\xleftarrow{y}$	$y \leftarrow k$
	$\xrightarrow{p(y), q(y)}$	$q(y) \stackrel{?}{=} p(y)(p(y) - 2y^{-2})$
		Append $R, u : p(y) \stackrel{?}{=} p\big\vert_y, q(y) \stackrel{?}{=} q\big\vert_y$
	$\xleftarrow{r_1^*, r_2^*}$	$r_1^*, r_2^* \leftarrow k$
..... Sumcheck for $r_1^* \tilde{z}_1(r) + r_2^* \tilde{z}_2(r) = \sum_{b \in \{0,1\}^n} (\tilde{M}_1(r, b) + \tilde{M}_2(r, b)) \tilde{z}(b)$		
		$h_0'(r_0) \coloneqq r_1^* \tilde{z}_1(r) + r_2^* \tilde{z}_2(r)$
$t_{2,n+2+i} = B_{2,n+2+i} s + (h_i'(-1), h_i'(0), h_i'(1))$	$\xrightarrow{t_{2,n+2+i}}$	$r_i' \leftarrow k^n$
	$\xleftarrow{r_i'}$	Append $R, u : h_i'(0) + h_i'(1) \stackrel{?}{=} h_{i-1}'(r_{i-1}')$
.....		
$t_{2,2n+3} = B_{2,2n+3} s + \tilde{z}(r')$	$\xrightarrow{t_{2,2n+3}}$	Append $R, u : \tilde{z}(r') \stackrel{?}{=} \tilde{z}\big\vert_r'$
$y \leftarrow \mathcal{D}, w = B_1 y, v = -R B_2 y$	$\xrightarrow{w, v}$	
	$\xleftarrow{\gamma}$	$\gamma \leftarrow \Gamma$
$z = \gamma \cdot s + y$, Rejection sampling	$\xrightarrow{z}$	$\|z\| \leq \sigma\sqrt{2r}$
		$B_1 z - \gamma t_1 = w$
		$R(\gamma t_2 - B_2 z) - \gamma u = v$

Fig. 4: Zero-knowledge proof system for Equation 7 with small proof size