An improved algorithm to solve the isogeny problem in dimension 2 with precomputation

Maria Corte-Real Santos

École Normale Supérieure de Lyon and CNRS
maria.corte_real_santos@ens-lyon.fr

Abstract. We give a new algorithm for the superspecial isogeny problem in dimension 2. After a one-time precomputation depending only on p (which succeeds with overwhelming probability), each instance is solved in expected time and memory $p^{3/4+o(1)}$. This result is unconditional, relying on no heuristics. Including the precomputation, a single instance is solved in time $p^{19/23+o(1)}$, improving on the $p^{1+o(1)}$ of the Costello–Smith algorithm. The improvement holds even with the same memory as (the rigorous version of) their algorithm.
The main ingredient is an algorithm that computes the equations of the Humbert surfaces $\mathcal{H}_{N^2}$ modulo p for all N up to a bound D growing with p. These equations give a membership test for optimally (N, N)-split surfaces, which we evaluate in batches using fast multipoint evaluation. This speeds up the bottleneck of the Costello–Smith algorithm: finding an isogeny from a superspecial principally polarised abelian surface to a product of supersingular elliptic curves.

Keywords: Isogeny-based cryptography, dimension 2, cryptanalysis.

1 Introduction

Let X_1, X_2 be *superspecial principally polarised abelian surfaces* defined over $\overline{\mathbb{F}}_p$. The general dimension-2 *superspecial isogeny problem* asks us to find a (polarised) isogeny

$$\varphi : X_1 \to X_2.$$

The best known algorithm for solving the superspecial isogeny problem is due to Costello and Smith [CS20]. The core idea of their algorithm is that the problem reduces to finding a *splitting*: an isogeny $\varphi_i : X_i \to E_i \times E_i'$, where E_i, E_i' are supersingular elliptic curves. Once these splittings are found, it suffices to solve a dimension-1 isogeny problem, i.e., find isogenies $E_1 \to E_2$ and $E_1' \to E_2'$, which can be done in time $\widetilde{O}(p^{1/2})$. These are then patched together to obtain the path $\pi : E_1 \times E_1' \to E_2 \times E_2'$ connecting φ_1 and φ_2 in order to output the full solution $\varphi := \varphi_2^\dagger \circ \pi \circ \varphi_1$.

Since products of elliptic curves make up a proportion $\Theta(1/p)$ of all superspecial surfaces, a random walk in the superspecial isogeny graph reaches a product after about p steps. The runtime analysis of [CS20] relied on a heuristic on the

expansion of the isogeny graph, so that pseudorandom walks quickly converge to the uniform distribution. This was recently proved in [ATY24], and later improved in [Cor+26]: this algorithm terminates in expected time $(\log p)^{O(1)}p$. The cost of the algorithm is dominated by finding splittings, so it is natural to focus on accelerating this part of the algorithm.

The problem of computing splittings has recently also been studied in its own right. Indeed, Kunzweiler and Shen [KS26a] and Kutas and Shen [KS26b] relate the problem of finding particular splittings, namely those satisfying certain non-triviality conditions, to the one endomorphism and isogeny problems for supersingular elliptic curves.

In the SplitSearcher algorithm [CCF24], concrete improvements were made to the first bottleneck step of the Costello–Smith algorithm. Instead of waiting for a product of elliptic curves, it checks at every step of the walk whether the current surface is (N, N)-split, i.e., (N, N)-isogenous to a product of elliptic curves, for some $N \leq 11$. Equivalently, it checks whether the surface lies on the Humbert surface $\mathcal{H}_{N^2}$ of discriminant N^2. These checks use explicit rational parametrisations of $\mathcal{H}_{N^2}$ due to Kumar [Kum15], which exist only because $\mathcal{H}_{N^2}$ is rational for these N. This likely does not extend to large N: in characteristic 0, Frengley's classification [Fre24] shows that $\mathcal{H}_{N^2}$ is rational only for $N \leq 16$ and $N \in \{18, 20, 24\}$, and is of general type for $N \geq 22$, $N \neq 24$. Since N is bounded, SplitSearcher gains a constant factor, but not an asymptotic one.

Kırımlı and Korpal [KK25; KK26] take a different approach to detecting splittings. They compute the *refined Humbert invariant* q_X of a superspecial p.p. abelian surface X (see Section 3.1), a positive definite quinary quadratic form. Then, (N, N)-splittings can be detected as primitive representations of N^2 by q_X. This detection is efficient once q_X is known. However, their method to compute q_X takes as input the polarisation data of X, given by its Ibukiyama–Katsura–Oort matrix, rather than an explicit model. As the authors point out [KK26, Rem. 2], no algorithm to compute q_X from a model of X (as given in the input of the isogeny problem) is currently known; see [KM25] for a detailed account of what is known.

In this work, we take the alternative approach explored in [CCF24]. Namely, we use an equation describing $\mathcal{H}_{N^2}$ rather than a parametrisation. The image of $\mathcal{H}_{N^2}$ in (weighted) projective space is a hypersurface cut out by a single polynomial, so testing whether a surface X is optimally (N, N)-split amounts to evaluating this polynomial at the invariants of X. Unlike parametrisations, such equations exist for every N. However, they grow quickly with N, so testing one surface at a time is expensive. Indeed, Shaska [Sha26c] observes that evaluating such an equation at a single surface costs $\Theta(N^9)$ while the density gain is only $\Theta(N^3)$, so that raising N does not pay off when surfaces are tested one at a time. We circumvent this issue by batch testing many N and many surfaces at the same time.

Our contributions. We present a new algorithm, called FindSplitting, that finds an $\overline{\mathbb{F}}_p$-isogeny from a superspecial Jacobian defined over $\overline{\mathbb{F}}_p$ to a product

of (necessarily, supersingular) elliptic curves. The insight of this work is that we can test, at the same time and at small amortised cost, whether *many* surfaces are (N, N)-split for *many* values of N. Concretely, we test membership of many surfaces in the union $\bigcup_{N \leq D} \mathcal{H}_{N^2}$ of Humbert surfaces using Kedlaya and Umans fast multipoint evaluation [KU11]. The algorithm FindSplitting has three main steps:

1. **Precomputation (offline).** For each $2 \leq N \leq D$, we interpolate the equation G_N of the characteristic-p locus of (optimally) (N, N)-split Jacobians, using points obtained by gluing random pairs of elliptic curves. To show that this succeeds, we prove that gluing random elliptic curves over $\mathbb{F}_{p^2}$ produces points that are spread out enough on $\mathcal{H}_{N^2}$ (Section 3.7 and Lemma 4.3), using results of Castryck and Hubrechts on the distribution of Frobenius on the N-torsion of elliptic curves over finite fields [CH13].
   We combine these equations into a single polynomial g, which allows us to test membership in the union of the $\mathcal{H}_{N^2}$. This step succeeds with high probability and depends only on p and D, so it can be done once and reused for any number of instances of the problem defined over $\overline{\mathbb{F}}_p$.
2. **Batch testing (online).** We then take many random walks from the input Jacobian and evaluate the polynomial g at all the endpoints at once using the Kedlaya–Umans algorithm for fast multivariate multipoint evaluation [KU11]. By taking enough random walks, the cost of testing each endpoint can be amortised.
3. **Repeat until hit (online).** We repeat Step 2 until g evaluates to 0 (i.e., we have found a surface that is (N, N)-split). This surface and corresponding splitting can then be efficiently recovered.

The key reason this algorithm outperforms Costello–Smith is that the union $\bigcup_{N \leq D} H_{N^2}$ is much larger than the locus of products of elliptic curves: it contains a proportion around D^4/p of all superspecial surfaces, compared to around $1/p$ for products. We prove a lower bound of this order (Proposition 3.23), which, combined with the rapid mixing of the superspecial 2-isogeny graph [Cor+26], shows that we expect to find such a surface after testing the endpoints of about p/D^4 random walks, instead of about p. On the other hand, the cost of precomputing g and the size of the batches grow with D. We bound both using the weight of the characteristic-0 equation of $\mathcal{H}_{N^2}$, after showing that G_N divides its reduction modulo p (Lemma 3.13). Choosing D to balance these costs gives an algorithm that is asymptotically faster than that of Costello and Smith.

One might ask why we test membership in a union of Humbert surfaces rather than in a single $\mathcal{H}_{N^2}$ with N large. A single $N \approx D^{4/3}$ gives the same density and therefore the same online cost, but interpolating its equation is much more expensive. By using many small N, the union reaches the same density at a lower precomputation cost (see Remark 6.1).

Combining this with the dimension-1 step of the Costello–Smith algorithm as detailed in [Cor+26, Thm 6.11], we obtain the main result of this paper.

Theorem 1.1. *Fix $\epsilon > 0$. Let p be a sufficiently large prime, and let $2 \leq D \leq p^{1/16}$ be an integer. There is an algorithm for the superspecial isogeny problem in dimension 2 consisting of two phases:*

1. *a precomputation, depending only on p and D, which runs in expected time $D^{19}p^{\epsilon}$ and memory $D^{12}p^{\epsilon}$, and succeeds (i.e., outputs G_N for all $N \leq D$) except with probability at most $1/p$;*
2. *an online phase which, given the output of the precomputation and superspecial p.p. abelian surfaces X_1, X_2 over $\mathbb{F}_{p^2}$, returns a polarised isogeny $X_1 \to X_2$ with probability 1. If the precomputation succeeded, it runs in expected time $(p/D^4 + D^{12})\, p^{\epsilon}$ using memory $(D^{12} + p^{1/2})\, p^{\epsilon}$.*

The term $(p/D^4 + D^{12})$ in the cost of the online phase exhibits the balance between the probability that each run will succeed (so we have to repeat less often) versus the cost of each run.

We immediately obtain the following consequences. Since ϵ is arbitrary, we write $p^{a+o(1)}$ for a bound that holds with $p^{a+\epsilon}$ for every $\epsilon > 0$.

1. **Optimal choice for a fixed prime and many instances.** Taking $D = p^{1/16}$, after a precomputation in time $p^{19/16+o(1)}$ and memory $p^{3/4+o(1)}$, each instance is solved in expected time $p^{3/4+o(1)}$ and memory $p^{3/4+o(1)}$.
2. **Optimal choice for a single instance.** Taking $D = p^{1/23}$ balances the precomputation against the online phase, and solves the superspecial isogeny problem in total expected time $p^{19/23+o(1)}$ and memory $p^{12/23+o(1)}$.

Time-memory trade-off. As detailed in Theorem 1.1, D also governs the memory, so varying D gives a time–memory trade-off. With memory M, where $p^{1/2} \leq M \leq p^{3/4}$, we show in Corollary 6.2 that each instance is solved in time $p^{1+o(1)}/M^{1/3}$ after a precomputation of time $M^{19/12+o(1)}$. For $M \leq p^{12/23}$, the precomputation is dominated by the online phase, so this also bounds the total cost of a single instance. In particular, with memory $p^{1/2+o(1)}$, the same as the (rigorous) Costello–Smith algorithm of [Cor+26, Thm. 6.11], we solve a single instance in time $p^{5/6+o(1)}$. The restriction $M \geq p^{1/2}$ comes only from the dimension-1 step (see Remark 6.3). A memory-efficient replacement for this step would therefore extend the trade-off down to memory $(\log p)^{O(1)}$, where it recovers the $p^{1+o(1)}$ running time of the Costello–Smith algorithm.

Computing endomorphism rings. We remark that the splitting algorithm of this paper can also be used to speed up the computation of endomorphism rings. Indeed, the algorithm of [Cor+26, Thm. 6.15] computes the endomorphism ring of a superspecial p.p. abelian surface X by computing two splittings from X of coprime degrees. We expect that replacing this step with FindSplitting will give a similar asymptotic improvement; see Section 6 for more details.

Related work. Houben and Pham [HP26] recently gave an algorithm for the one endomorphism problem in dimension 2: given a superspecial p.p. abelian

surface X over $\mathbb{F}_{p^2}$ with $p \equiv 3 \pmod 4$, it computes a non-scalar polarised endomorphism of X in expected time and memory $p^{4/5+o(1)}$, under a heuristic on the smoothness of the multiplier of a short endomorphism. Their approach is a meet-in-the-middle search for short polarised endomorphisms, in the spirit of Wesolowski's algorithm in dimension one [Wes26], and builds on Pham's bound on the minimal multiplier of such an endomorphism [Pha26]. Unlike in dimension one, where the one endomorphism, endomorphism ring and isogeny problems are known to be equivalent [MW25; PW24; Wes22], the corresponding reductions in dimension 2 are not fully established (see, e.g., [Cas+25; Mon26]). Houben and Pham therefore recover the full endomorphism ring under a second heuristic on the distribution of the endomorphisms their algorithm outputs, and then, using the machinery of [PRS26], solve the superspecial isogeny problem in the same complexity and additionally assuming GRH.

By contrast, we directly target the isogeny problem, and our analysis is unconditional. For a single instance (including the precomputation), our algorithm takes time $p^{19/23+o(1)} \approx p^{0.826}$, which is slightly more than that of [HP26], however it uses less memory. Our algorithm also naturally fits the setting of algorithms with preprocessing (see, e.g., [CK18]), where one is willing to invest in a one-off precomputation to lower the cost of each subsequent instance. Since our precomputation depends only on p, it can be reused for every instance over the same prime, as arises whenever a cryptographic scheme fixes its parameters. In this setting, after a precomputation in time $p^{19/16+o(1)}$, each instance is solved in time and memory $p^{3/4+o(1)}$, thus outperforming the algorithm of Houben and Pham (per instance).

Another related work is that of Udovenko [Udo26], who uses a similar approach in dimension one. Their algorithm takes many walks in the isogeny graph and uses batched multipoint evaluation of the classical modular polynomial to find curves E with an N-isogeny $E \to E^{(p)}$, which composed with Frobenius gives a non-scalar endomorphism. This solves the isogeny problem unconditionally in time and memory $p^{2/5+o(1)}$. Faster algorithms are known in dimension one, based on a meet-in-the-middle search rather than multipoint evaluation. Most notably, Wesolowski [Wes26] obtains $p^{1/3+o(1)}$ under a smoothness heuristic.

Udovenko's setting differs from ours in two ways. In dimension one, the modular polynomial Φ_N mod p can be computed in time quasi-linear in its size (see, e.g., [BLS12; Eng09; KR25] for prime N and [Udo26, Prop. 2] for composite N), whereas we compute G_N by interpolation, which is more expensive. Moreover, the structure of the points at which Φ_N is evaluated allows Udovenko to work over $\mathbb{F}_p$, which our invariants do not (see Remark 2.11).

Organisation. Section 2 recalls background on superspecial principally polarised abelian surfaces, the superspecial isogeny problem, and multipoint evaluation. Section 3 defines Humbert surfaces, the equations defining them and their relation to (N, N)-splittings. It also contains two new ingredients: a lower bound on the proportion of superspecial surfaces lying on $\bigcup_{2 \le N \le D} \mathcal{H}_{N^2}$ (Section 3.6),

which governs the success probability of each random walk, and an analysis of gluing elliptic curves over finite fields (Section 3.7), which we use to sample points on $\mathcal{H}_{N^2}$. We then describe each step of our algorithm: Section 4 describes the precomputation, Section 5 describes the online phase, and Section 6 assembles the full algorithm and proves our main results.

Acknowledgements. The author is grateful to Harun Kır and Benjamin Wesolowski for their feedback on an earlier draft of this paper. This work was supported by the European Research Council under grant No. 101116169 (AGATHA CRYPTY).

AI Disclosure. An LLM was used as an editorial aid: given a final draft of the paper, it was used to identify inconsistencies, correct typographical errors, and improve grammar. Beyond these standard (and minor) edits, an LLM was also used to search for relevant literature. Most notably, this brought the work of Castryck and Hubrechts [CH13] to the author's attention. Upon reading it, the author developed the current proof of Lemma 3.26, which was originally stated as a heuristic. The author takes full responsibility for the accuracy, integrity, and final presentation of this work.

2 Preliminaries

From this point onwards, let $p > 7$ be a prime. Throughout Sections 2.1 to 2.3 and 3.1 to 3.3, k denotes an algebraically closed field with $\mathrm{char}(k) \notin \{2, 3, 5\}$ and $N \geq 2$ is an integer not divisible by $\mathrm{char}(k)$. In Sections 2.4 and 3.4, we specialise to $k = \overline{\mathbb{F}}_p$ (the algebraic closure of $\mathbb{F}_p$), which is the case of particular interest to us, and in Section 3.5 to $\mathrm{char}(k) = 0$.

2.1 Principally polarised abelian surfaces

We give a brief account of abelian surfaces and fix notation; general references are [Har77; Mil86; Mum70]. All divisors, Néron–Severi classes, subvarieties and isogenies are taken over k, even when the abelian surfaces involved are defined over a subfield of k.

Polarisations. Let X be an abelian surface (i.e., an abelian variety of dimension 2) over k, and write $X^\vee$ for the *dual* abelian variety.

An abelian variety of dimension 1 is an elliptic curve E, and it comes with a canonical 'principal polarisation', i.e., a canonical isomorphism $E \cong E^\vee$. In dimension 2 this is no longer the case: an abelian surface need not admit a principal polarisation, and when it does, the polarisation need not be unique. We therefore work with abelian surfaces together with a choice of principal polarisation.

The Néron–Severi group. Let $\mathrm{Div}(X)$ denote the group of divisors on X. For $\mathcal{D}_1, \mathcal{D}_2 \in \mathrm{Div}(X)$[1] we write $(\mathcal{D}_1 \cdot \mathcal{D}_2)$ for their intersection number. We say that $\mathcal{D}_1$ and $\mathcal{D}_2$ are *numerically equivalent*, denoted $\mathcal{D}_1 \equiv \mathcal{D}_2$, if $(\mathcal{D}_1 \cdot \mathcal{D}) = (\mathcal{D}_2 \cdot \mathcal{D})$ for all $\mathcal{D} \in \mathrm{Div}(X)$. The *Néron–Severi group*, defined as

$$\mathrm{NS}(X) := \mathrm{Div}(X)/\equiv,$$

is a free $\mathbb{Z}$-module of finite rank $\rho(X)$, the *Picard number*. We write $\mathrm{cl}(\mathcal{D})$ for the class of $\mathcal{D}$ in $\mathrm{NS}(X)$, and $(\xi \cdot \eta)$ for the intersection pairing of $\xi, \eta \in \mathrm{NS}(X)$. For $\theta := \mathrm{cl}(\mathcal{D}) \in \mathrm{NS}(X)$, the map

$$\varphi_\theta \colon X \to X^\vee, \quad x \mapsto [t_x^*\mathcal{D} - \mathcal{D}],$$

where $[\,\cdot\,]$ denotes the linear equivalence class in $\mathrm{Pic}^0(X) \cong X^\vee$, is a homomorphism depending only on the class θ, and $\theta \mapsto \varphi_\theta$ is an injective homomorphism $\mathrm{NS}(X) \to \mathrm{Hom}(X, X^\vee)$.

Principal polarisations. By the Nakai–Moishezon criterion [Har77, Thm. V.1.10], ampleness of a divisor depends only on its class. We call $\theta \in \mathrm{NS}(X)$ *ample* if it is the class of an ample divisor. A *polarisation* on X is a homomorphism $\lambda = \varphi_\theta$ with θ ample. It is an isogeny (i.e., a surjective finite morphism of group varieties), and by the Riemann–Roch theorem $\deg \varphi_\theta = \frac{1}{4}(\theta \cdot \theta)^2$. A polarisation λ is *principal* if it is an isomorphism, i.e., $\deg \lambda = 1$, in which case we call (X, λ) a *principally polarised (p.p.)* abelian surface. Hence, setting

$$\mathcal{P}(X) := \{\theta \in \mathrm{NS}(X) : \theta \text{ ample and } (\theta \cdot \theta) = 2\},$$

the map $\theta \mapsto \varphi_\theta$ is a bijection between $\mathcal{P}(X)$ and the set of principal polarisations on X. We identify a principal polarisation λ with its class $\theta \in \mathcal{P}(X)$, so that $\lambda = \varphi_\theta$. In this section, we mostly write λ, while working with the class θ will be the more convenient point of view in Section 3.

A p.p. abelian surface X always comes with its principal polarisation $\lambda_X = \varphi_{\theta_X}$, which we suppress from the notation when no confusion can arise. Every p.p. abelian surface over k is either the Jacobian $\mathrm{Jac}(C)$ of a genus 2 curve C with its canonical principal polarisation, or a product $E_1 \times E_2$ of elliptic curves with the product polarisation. We often simply write J to denote $\mathrm{Jac}(C)$ when the underlying curve is clear from context.

2.2 Isogenies of p.p. abelian surfaces

A *(polarised) isogeny* between p.p. abelian surfaces (X_1, λ_1) and (X_2, λ_2) is an isogeny $\varphi \colon X_1 \to X_2$ such that

$$\varphi^*\lambda_2 := \varphi^\vee \circ \lambda_2 \circ \varphi = d\lambda_1$$

[1] We use the slightly unconventional notation of $\mathcal{D}$ for divisors to avoid the notation clash with our bound D.

for some integer $d > 0$, where $\varphi^\vee \colon X_2^\vee \to X_1^\vee$ denotes the dual isogeny. Equivalently, in terms of the classes θ_i of λ_i, we have $\varphi^*\theta_2 = d\,\theta_1$ in $\mathrm{NS}(X_1)$. In this case $\deg(\varphi) = d^2$. We call d the *reduced degree* of φ, and φ is called a *d-isogeny*. Reduced degrees are multiplicative: the composite of a d-isogeny and a d'-isogeny is a dd'-isogeny.

Definition 2.1. *Two d-isogenies $\varphi_i \colon (X, \lambda) \to (X_i, \lambda_i)$, $i = 1, 2$, between p.p. abelian surfaces are said to be* equivalent *if there is an isomorphism $\beta \colon X_1 \to X_2$ with $\beta^*\lambda_2 = \lambda_1$ such that $\varphi_2 = \beta\varphi_1$.*

Definition 2.2 (Adjoint isogeny). *Let $\varphi \colon (X_1, \lambda_1) \to (X_2, \lambda_2)$ be a d-isogeny. We define its* adjoint isogeny *by*

$$\varphi^\dagger := \lambda_1^{-1}\varphi^\vee\lambda_2 \colon (X_2, \lambda_2) \to (X_1, \lambda_1).$$

It satisfies $\varphi\varphi^\dagger = [d] = \varphi^\dagger\varphi$.

Kernels of polarised isogenies. Let $N \geq 2$ be an integer coprime to $\mathrm{char}(k)$. For any abelian surface X we have the N-Weil pairing $X[N] \times X^\vee[N] \to \mu_N$, where μ_N is the group of N-th roots of unity. When X is equipped with a principal polarisation λ, composing with λ gives the N-Weil pairing $e_N \colon X[N] \times X[N] \to \mu_N$. We say that a subgroup $K \subseteq X[N]$ is *isotropic* if $e_N(P, Q) = 1$ for all $P, Q \in K$, and *maximal isotropic* if moreover there is no isotropic subgroup K' with $K \subsetneq K' \subseteq X[N]$. A maximal isotropic subgroup K has order N^2.

By [Kan97, Prop. 1.1], there is a bijection between N-isogenies $\varphi : (X, \lambda) \to (X', \lambda')$, up to equivalence, and maximal isotropic subgroups $K \subseteq X[N]$. Indeed, $\ker\varphi$ is a maximal isotropic subgroup of $X[N]$. Conversely, given such a K, the quotient X/K carries a unique principal polarisation for which the quotient map is an N-isogeny.

Definition 2.3. *Let $\varphi \colon (X, \lambda) \to (X', \lambda')$ be an N-isogeny of p.p. abelian surfaces. We say that φ is an (N, N)-*isogeny *if $\ker\varphi \cong (\mathbb{Z}/N\mathbb{Z})^2$. In this case, $\ker\varphi$ is an (N, N)-*subgroup*.*

For a fixed p.p. abelian surface X, the number of (N, N)-isogenies with domain X, up to equivalence is

$$\kappa_N := N^3 \prod_{\substack{\ell \mid N \\ \ell \text{ prime}}} \frac{1}{\ell^3}(\ell + 1)(\ell^2 + 1). \tag{1}$$

In particular, when N is prime every N-isogeny is an (N, N)-isogeny, and $\kappa_N = (N^2 + 1)(N + 1)$.

Types of polarised isogenies. There are four types of polarised N-isogenies in dimension 2:

1. *Generic:* an isogeny $J \to J'$ between Jacobians of genus 2 curves.

2. *Splitting:* an isogeny $J \to E_1 \times E_2$ from a Jacobian to a product of elliptic curves with the product polarisation.
3. *Gluing:* an isogeny $E_1 \times E_2 \to J$ from a product of elliptic curves to a Jacobian. In this work, gluing isogenies will arise from gluing the two factors E_1, E_2 along an anti-isometry of their N-torsion (see Section 3).
4. *Diagonal:* an isogeny $E_1 \times E_2 \to E_1' \times E_2'$ between products of elliptic curves, typically of the form $\varphi_1 \times \varphi_2$ with $\varphi_i \colon E_i \to E_i'$.

Of particular importance to this work will be (N, N)-splittings.

Definition 2.4. *We say that the Jacobian $J = \mathrm{Jac}(C)$ of a genus 2 curve C is* (N, N)-split *if there exists an (N, N)-isogeny $\varphi \colon J \to E_1 \times E_2$ to a product of elliptic curves with the product polarisation.*

In Section 3, we describe the connection between (N, N)-splittings and points on Humbert surfaces.

2.3 Isomorphism classes

In this work, isomorphisms $\beta \colon X \to X'$ of p.p. abelian surfaces will always be polarised, i.e., $\beta^* \lambda_{X'} = \lambda_X$. Furthermore, $\mathrm{Aut}(X)$ denotes the group of automorphisms of X preserving λ_X. This will matter particularly in our setting, since we will work with *superspecial* abelian surfaces, which are all isomorphic as unpolarised abelian varieties (see Section 2.4).

We write $\mathcal{A}_2$ for the coarse moduli space of p.p. abelian surfaces. Since k is algebraically closed, $\mathcal{A}_2(k)$ is the set of isomorphism classes of p.p. abelian surfaces over k. Similarly, $\mathcal{M}_2$ denotes the coarse moduli space of genus 2 curves. By the Torelli theorem, $C \mapsto \mathrm{Jac}(C)$ identifies $\mathcal{M}_2$ with an open subvariety of $\mathcal{A}_2$, whose complement is the locus of products of elliptic curves. We therefore regard $\mathcal{M}_2(k) \subset \mathcal{A}_2(k)$.

Igusa–Clebsch invariants. Igusa–Clebsch invariants will provide us with isomorphism invariants for Jacobians of genus-2 curves. From the invariant theory of the binary sextic, to any genus 2 curve C/k we may associate its Igusa–Clebsch invariants $I_2(C)$, $I_4(C)$, $I_6(C)$, $I_{10}(C)$, constructed as follows. Suppose C is given by a Weierstrass equation

$$C \colon y^2 = (x - a_0) \cdots (x - a_5),$$

where $a_0, \ldots, a_5 \in k$. We define

$$\begin{aligned}
I_2(C) &:= \sum_{15} (01)^2 (23)^2 (45)^2, \\
I_4(C) &:= \sum_{10} (01)^2 (12)^2 (20)^2 (34)^2 (45)^2 (53)^2, \\
I_6(C) &:= \sum_{60} (01)^2 (12)^2 (20)^2 (34)^2 (45)^2 (53)^2 (03)^2 (14)^2 (25)^2, \\
I_{10}(C) &:= \prod_{i<j} (a_i - a_j)^2,
\end{aligned}$$

where, for any permutation $\sigma \in S_6$, we let (ij) denote the difference $a_{\sigma(i)} - a_{\sigma(j)}$. Here the sums are taken over all distinct expressions in the a_i as σ ranges over S_6; the subscripts denote the number of expressions in each sum. Note that $(I_2(C) : I_4(C) : I_6(C) : I_{10}(C)) \in \mathbb{P}(2,4,6,10)$, and that $I_{10}(C) \neq 0$. Moreover, since $\mathrm{char}(k) \neq 2,3,5$, the isomorphism class of C is uniquely determined by its Igusa–Clebsch invariants [Igu60; Mes91]. We call the resulting morphism

$$\mathrm{IC}\colon \mathcal{M}_2 \longrightarrow \mathbb{P}(2,4,6,10), \qquad C \longmapsto \big(I_2(C) : I_4(C) : I_6(C) : I_{10}(C)\big)$$

the *Igusa–Clebsch map*; it is an isomorphism onto the open subvariety $\{I_{10} \neq 0\}$ of $\mathbb{P}(2,4,6,10)$. For a Jacobian $\mathrm{Jac}(C)$ defined over $\mathbb{F}_q$, the Igusa–Clebsch invariants of C are $\mathbb{F}_q$-rational, and have a representative $(I_2, I_4, I_6, I_{10}) \in \mathbb{F}_q^4$. We always work with such a representative.

Normalised invariants. Choose $u \in k$ with $u^5 = I_{10}(C)^{-1}$. We define the *normalised invariants* of C to be the coordinates of the point

$$(i_1, i_2, i_3) := \big(uI_2(C),\, u^2 I_4(C),\, u^3 I_6(C)\big) \in \mathbb{A}^3(k). \tag{2}$$

This point depends on the choice of u only up to the action $(i_1, i_2, i_3) \mapsto (\zeta i_1, \zeta^2 i_2, \zeta^3 i_3)$ of $\zeta \in \mu_5$. It is a representative of the weighted-projective point $(I_2 : I_4 : I_6 : I_{10})$ with last coordinate 1, obtained by rescaling with t, where $t^2 = u$. Hence, for every weighted-homogeneous form G, we have $G(I_2, I_4, I_6, I_{10}) = 0$ if and only if $G(i_1, i_2, i_3, 1) = 0$, independently of the choice of u.

Remark 2.5. When $\mathrm{Jac}(C)$ is defined over $\mathbb{F}_{p^2}$ (e.g., when it is superspecial and $k = \overline{\mathbb{F}}_p$; see Section 2.4), u lies in $\mathbb{F}_{p^2}$ or $\mathbb{F}_{p^{10}}$, and computing it costs $(\log p)^{O(1)}$. Evaluating a weighted-homogeneous form G at the normalised invariants may therefore require computations over $\mathbb{F}_{p^{10}}$. This can be avoided by partitioning the points according to the class of I_{10} in $\mathbb{F}_{p^2}^\times / (\mathbb{F}_{p^2}^\times)^5$: for each class, fix a representative c, rescale each point by some $u \in \mathbb{F}_{p^2}^\times$ via $(uI_2, u^2 I_4, u^3 I_6, u^5 I_{10})$ so that $I_{10} = c$, and evaluate $G(x_1, x_2, x_3, c)$. All computations then remain over $\mathbb{F}_{p^2}$. We do not do this, for simplicity, as it does not affect the asymptotic complexity.

As discussed above, the Igusa–Clebsch invariants of a genus 2 curve define a point in weighted projective space. Since we will construct forms by interpolating at such points, we work with weighted-homogeneous polynomials.

Definition 2.6. *For a field $\mathbb{F}$ and $w \in \mathbb{Z}$, define $V_w(\mathbb{F}) \subseteq \mathbb{F}[I_2, I_4, I_6, I_{10}]$ to be the* space of weighted-homogeneous polynomials of weight w*, where I_{2j} has weight $2j$. We say a non-zero $f \in V_w$ has* weight w*. When the field $\mathbb{F}$ is clear from context, we simply write V_w.*

Setting $I_{10} = 1$ identifies V_w with the space of polynomials in x_1, x_2, x_3 (of weights $2, 4, 6$) all of whose monomials have weighted degree $w - 10k$ for some $k \geq 0$.

2.4 The isogeny problem

We specialise to the case $k = \overline{\mathbb{F}}_p$ to define the isogeny problem.

Superspeciality. For cryptographic applications the most natural generalisation of the set of supersingular elliptic curves to dimension 2 is the set of superspecial p.p. abelian surfaces.

Definition 2.7. *We say that X is* superspecial *if Frobenius acts as zero on $H^1(X, \mathcal{O}_X)$. For $X = \mathrm{Jac}(C)$, this means that the Hasse–Witt matrix $M \in \overline{\mathbb{F}}_p^{2\times 2}$ of C vanishes identically.*

By results of Oort, Deligne, Ogus and Shioda (see [LO98, §1.6]) we have an alternative characterisation for superspeciality: X is superspecial if and only if it is isomorphic, as an unpolarised abelian variety, to a product of supersingular elliptic curves. Furthermore, all superspecial abelian surfaces are isomorphic as unpolarised abelian varieties. In particular, if $X \cong E_1 \times E_2$ (as p.p. abelian surfaces), then X is superspecial if and only if E_1 and E_2 are supersingular.

Every superspecial p.p. abelian surface X is isomorphic (as a p.p. abelian surface) to one defined over $\mathbb{F}_{p^2}$ [IK94, Thm. 1]. In fact, we can choose a model for X on which the p^2-power Frobenius acts as $[\pm p]$ [CDS20]. We therefore assume that our superspecial p.p. abelian surfaces are given in such a model.

The isogeny problem. The dimension 2 superspecial isogeny problem may be stated precisely as follows.

Problem 2.8. Given a pair of superspecial p.p. abelian surfaces X_1, X_2 defined over $\mathbb{F}_{p^2}$, find a polarised isogeny $X_1 \to X_2$ (over $\overline{\mathbb{F}}_p$).

The superspecial isogeny graph. The superspecial isogeny graph allows us to reframe Problem 2.8 as a path-finding problem. Let $\mathrm{SS}_g(p)$ denote the set of isomorphism classes (as p.p. abelian varieties over $\overline{\mathbb{F}}_p$) of superspecial p.p. abelian varieties of dimension g (for $g = 1, 2$).

Fix a prime number $\ell \neq p$, and consider the ℓ-isogeny graph $\mathcal{G}(p; \ell)$ whose vertex set is $\mathrm{SS}_2(p)$ and whose edges from X to X' are the ℓ-isogenies $X \to X'$, up to equivalence as in Definition 2.1. We call $\mathcal{G}(p; \ell)$ the *superspecial ℓ-isogeny graph*. Note that this is a directed multigraph (see, e.g., [Cor+26, §3]).

From any $X \in \mathrm{SS}_2(p)$ there are κ_ℓ outgoing ℓ-isogenies, so the graph is κ_ℓ-(out-)regular. Unlike the case in dimension 1, the superspecial ℓ-isogeny graph is not Ramanujan, i.e., it does not have optimal expansion [JZ26, §9.1]. Despite this, recent work shows that $\mathcal{G}(p; \ell)$ still has good enough expansion for cryptographic purposes: random walks in the graph converge rapidly to a stationary distribution [Cor+26]. This result will play a role in Section 3.6.

The Costello–Smith algorithm. As outlined in the introduction, the algorithm of Costello and Smith [CS20] solves Problem 2.8 in two steps:

1. *Finding splittings:* for $i = 1, 2$, find a polarised isogeny $\varphi_i \colon X_i \to E_i \times E_i'$ to a product of supersingular elliptic curves, by taking random walks in $\mathcal{G}(p; 2)$ from X_i;
2. *Solving in dimension* 1*:* find isogenies $E_1 \to E_2$ and $E_1' \to E_2'$, combine them into a polarised isogeny $\pi \colon E_1 \times E_1' \to E_2 \times E_2'$ (e.g., by following the proof of [Cor+26, Thm. 6.11]), and output $\varphi_2^\dagger \circ \pi \circ \varphi_1$.

Since products of elliptic curves form a proportion $O(1/p)$ of $\mathrm{SS}_2(p)$, the first step dominates. Using the rapid mixing of $\mathcal{G}(p; 2)$, it takes time $p(\log p)^{O(1)}$ [Cor+26, Thm. 6.9], while the second step takes time $\widetilde{O}(p^{1/2})$; see [Cor+26, Thm. 6.11] for a rigorous treatment of the full algorithm. Our contribution is a faster algorithm for the first step (Algorithm 4), whilst the second step is used unchanged.

2.5 Efficient evaluation of multivariate polynomials

We now detail the result of Kedlaya and Umans [KU11] that will be crucial to the algorithms in this paper to perform efficient multipoint evaluation of multivariate polynomials. Though the statements given in [KU11] are more general, we will restrict our attention to the setting relevant for this work (i.e., working over a finite field $\mathbb{F}_q$ with 3 variables). Note that, as the number of variables is small and fixed, we do not find any improvement using the algorithm due to Bhargava, Ghosh, Guo, Kumar, and Umans [Bha+24].

Let $f \in \mathbb{F}_q[x_1, x_2, x_3]$, with individual degrees less than d in each variable. The *box size* of f is d^3: it is the number of coefficients in the dense representation of f, where we count one coefficient for every exponent vector in the cube $\{0, \dots, d-1\}^3$ (including zero coefficients). A polynomial whose degrees in x_1, x_2, x_3 are at most d_1, d_2, d_3 is treated as having individual degrees less than $d = \max_i d_i + 1$, by padding with zero coefficients.

Theorem 2.9 ([KU11, Cor. 4.5]). *Let $\mathbb{F}_q$ be a finite field, and fix $\delta > 0$. For sufficiently large d (depending on δ), given $f \in \mathbb{F}_q[x_1, x_2, x_3]$ with box size n, and points $\alpha_1, \dots, \alpha_m \in \mathbb{F}_q^3$, one can compute $f(\alpha_1), \dots, f(\alpha_m)$ using $(n + m)^{1+\delta}(\log q)^{1+o(1)}$ bit operations.*

By padding with zeros if necessary, we may assume that d is large enough for Theorem 2.9 to apply, at the cost of a constant factor depending only on δ.

Remark 2.10. We remark that due to our reliance on fast multipoint evaluation, our bounds are asymptotic. The input sizes at which these algorithms outperform classical methods are very large [HL20], so the concrete cost of our algorithm remains to be determined.

Remark 2.11. Over a prime field, multipoint evaluation has better bounds [Gho+23, Thm. 3.2], which Udovenko uses in dimension one [Udo26]. There, Φ_N is evaluated at conjugate pairs $(j, j^p) = (x + \beta y, x - \beta y)$, so writing $j = x + \beta y$ turns

two variables over $\mathbb{F}_{p^2}$ into two variables over $\mathbb{F}_p$. Our invariants have no such structure: they lie in $\mathbb{F}_{p^2}$ at best (see Remark 2.5), and writing each of them in coordinates over $\mathbb{F}_p$ turns our three variables into six, each with the same degree bound. If d bounds the degree of G_N in each variable, the number of coefficients grows from d^3 to d^6, so working over $\mathbb{F}_p$ would make interpolation far more expensive.

Notation and conventions. Throughout, we fix $\delta > 0$ and integers $N \leq D < p$. In particular, $p \nmid N$. All cost bounds hold for p sufficiently large in terms of δ, and the implied constants may depend on δ.

Our analysis of the precomputation requires $D \leq p^{1/9-\epsilon}$ for some $\epsilon > 0$ (Propositions 4.5 and 4.8), and our density estimate requires $200D^7 \leq p$ (Proposition 3.23). As the cost of the online phase is minimised at $D = p^{1/16}$ and increases for larger D, we assume $D \leq p^{1/16}$ throughout, which implies both conditions for p sufficiently large (in particular, for cryptographically sized primes).

Throughout, log denotes the natural logarithm. All logarithms appear either inside $O(\cdot)$, in exponents of the form $(\log p)^{O(1)}$, or in ratios (as in Proposition 3.16), so the choice of base does not affect any statement. We interpret $\log\log N$ as $\max(1, \log\log N)$, so that the corresponding bounds also hold for small N.

3 Humbert Surfaces

In this section we introduce Humbert surfaces, which are central to our algorithm. Fix an integer $N \geq 2$ coprime to p. General references for this section are [Kan94; Kan97; Kan16; Kan18]; the content of Sections 3.6 and 3.7 is new to this work.

3.1 The refined Humbert invariant

We first recall the refined Humbert invariant of a p.p. abelian surface, introduced by Kani [Kan94], following the presentation of [Kan16; Kan18]. Let X be a p.p. abelian surface with principal polarisation $\theta \in \mathcal{P}(X)$. As θ is a principal polarisation, the class θ is primitive in $\mathrm{NS}(X)$, and the *polarised Néron–Severi group* $\mathrm{NS}(X,\theta) := \mathrm{NS}(X)/\mathbb{Z}\theta$ is free of rank $\rho(X) - 1$.

Definition 3.1. *The* refined Humbert invariant *of X is the quadratic form*

$$q_X \colon \mathrm{NS}(X,\theta) \to \mathbb{Z}, \quad \xi \mapsto (\xi \cdot \theta)^2 - 2(\xi \cdot \xi).$$

It is well-defined and positive definite [Kan94, pg. 208]. An isomorphism $X \to X'$ of p.p. abelian surfaces induces an isometry $(\mathrm{NS}(X',\theta'), q_{X'}) \cong (\mathrm{NS}(X,\theta), q_X)$, so q_X depends only on the class of X in $\mathcal{A}_2(k)$. The refined Humbert invariant detects whether a p.p. abelian surface is a Jacobian.

Proposition 3.2 (Prop. 6, [Kan16]). *For a p.p. abelian surface, we have $X \notin \mathcal{M}_2(k)$ if and only if q_X represents 1.*

3.2 Humbert surfaces and elliptic subgroups

We want to detect not only products of elliptic curves, but also surfaces that are close to products, in the sense that they are (N, N)-split for a particularly small N. This leads to the study of Humbert surfaces.

Definition 3.3. *For an integer $N \geq 1$, the* Humbert surface $\mathcal{H}_{N^2}$ *is the locus*

$$\mathcal{H}_{N^2}(k) = \{X \in \mathcal{A}_2(k) : q_X \text{ primitively represents } N^2\}.$$

Kani's definition makes sense over an algebraically closed field of any characteristic. Over $\mathbb{C}$, the integers primitively represented by q_X are exactly the classical Humbert invariants of X [Kan94, Rem. 5.6], so $\mathcal{H}_{N^2}$ recovers the classical Humbert surface of discriminant $\Delta = N^2$ [Hum99; Hum00; Hum01] (see [Gee88] for a modern treatment).

Elliptic subgroups. Proposition 3.2 shows that $\mathcal{H}_1$ is precisely the locus of products of elliptic curves with the product polarisation, i.e., $\mathcal{H}_1 = \mathcal{A}_2 \setminus \mathcal{M}_2$. For $N \geq 2$, $\mathcal{H}_{N^2}$ admits the following geometric description.

Definition 3.4. *An* elliptic subgroup *of X is an abelian subvariety $E \subset X$ of dimension 1. Its* degree *is $\deg_\theta(E) = (E \cdot \theta)$. Equivalently, given the inclusion $\iota_E \colon E \hookrightarrow X$, we have that $\iota_E^\vee \circ \lambda \circ \iota_E = [\deg_\theta(E)]$ on $E \cong E^\vee$.*

We write $\mathcal{E}_N(X)$ for the set of elliptic subgroups of X of degree N. As $\#\mathcal{E}_N(X)$ is finite [Kan94, Cor. 2.10], we define $\varepsilon_N(X) := \#\mathcal{E}_N(X)$. The following result due to Kani allows us to relate elliptic subgroups of degree N to points on $\mathcal{H}_{N^2}$.

Theorem 3.5 (Cor. 3.9, [Kan94]). *A p.p. abelian surface (X, θ) has an elliptic subgroup $E \subseteq X$ of degree N if and only if $(X, \theta) \in \mathcal{H}_{N^2}$.*

Now that we have established the connection between elliptic subgroups and $\mathcal{H}_{N^2}$, we work towards understanding how they connect to (N, N)-splittings.

The gluing construction. The first ingredient is the following construction due to Frey and Kani [FK91]. Let E, E' be elliptic curves over k, with N-Weil pairings denoted e_N and e'_N, respectively.

Definition 3.6. *An* anti-isometry *is a group isomorphism $\alpha \colon E[N] \to E'[N]$ such that $e'_N(\alpha(P), \alpha(Q)) = e_N(P, Q)^{-1}$ for all $P, Q \in E[N]$. Write $\mathcal{I}_N(E, E')$ for the set of all such anti-isometries.*

A count for the number of anti-isometries is given in [Kan97, pg. 114]:

$$\#\mathcal{I}_N(E, E') = \#\mathrm{SL}_2(\mathbb{Z}/N\mathbb{Z}) = N^3 \prod_{\ell \mid N} \left(1 - \ell^{-2}\right). \tag{3}$$

Given $\alpha \in \mathcal{I}_N(E, E')$, the graph $\Gamma_\alpha = \{(P, \alpha(P)) : P \in E[N]\}$ is a maximal isotropic subgroup of $(E \times E')[N]$. The quotient $X_\alpha := (E \times E')/\Gamma_\alpha$ is a p.p. abelian surface.

Let $E \subset X$ be an elliptic subgroup of degree N. Following [Kan19, Prop. 10], there is a unique elliptic subgroup $E^\perp \subset X$ such that $E^\perp$ has degree N, $E \cap E^\perp = E[N] = E^\perp[N]$, and $\alpha_E := -\mathrm{id}\colon E[N] \to E^\perp[N]$ is an anti-isometry. We call $E^\perp$ the *complement* of E. Noting that the addition map

$$\mu_E\colon E \times E^\perp \to X,\ (x, y) \mapsto x + y,$$

is an isogeny of p.p. abelian surfaces with kernel Γ_{α_E}, we have $X \cong X_{\alpha_E}$.

Let Z_N denote the set of isomorphism classes of triples (E, E', α) over k, where (E_1, E_1', α_1) and (E_2, E_2', α_2) are isomorphic if there are isomorphisms $u\colon E_1 \to E_2$ and $u' : E_1' \to E_2'$ with $u' \circ \alpha_1 = \alpha_2 \circ u$.[2] We obtain the *gluing map*:

$$\mathrm{gl}_N\colon Z_N \longrightarrow \mathcal{A}_2(k), \qquad (E_1, E_2, \alpha) \longmapsto X_\alpha. \tag{4}$$

The image of gl_N is exactly $\mathcal{H}_{N^2}(k)$: each X_α has an elliptic subgroup of degree N, namely the image of $E \times \{0\}$, so $X_\alpha \in \mathcal{H}_{N^2}(k)$ by Theorem 3.5. Conversely, every $X \in \mathcal{H}_{N^2}(k)$ has an elliptic subgroup E of degree N, and $X \cong X_{\alpha_E} = \mathrm{gl}_N(E, E^\perp, \alpha_E)$.

Dictionary with optimal splittings. Let C/k be a genus 2 curve with Jacobian $J = \mathrm{Jac}(C)$.

Definition 3.7. *A finite morphism $f\colon C \to E$ from a genus 2 curve to an elliptic curve is* minimal *if it does not factor through an isogeny $E' \to E$ of degree greater than 1.*[3] *An (N, N)-splitting $\varphi\colon J \to E_1 \times E_2$ is* optimal *if the cover $C \to E_1$, obtained by composing an Abel–Jacobi map $C \hookrightarrow J$ with φ and the projection to E_1, is minimal. We say that J is* optimally (N, N)-split *if it admits an optimal (N, N)-splitting.*

The following sets are in natural bijection:

1. elliptic subgroups $E \subset J$ of degree N;
2. minimal covers $f\colon C \to E$ of degree N, up to isomorphism of the genus-1 curve E, via $f \mapsto f^*(E)$ [Kan94, Cor. 4.3];
3. optimal (N, N)-splittings $\varphi\colon J \to E_1 \times E_2$, up to equivalence (in the sense of Definition 2.1), together with an ordering of the factors, via $E \mapsto \mu_E^\dagger$.

Swapping the factors in item 3 corresponds to replacing E by $E^\perp$ in item 1. Since the kernel of $\mu_E^\dagger$ is $E[N] = E^\perp[N]$, and a (principally polarised) product of elliptic curves determines its two factors, the number of (N, N)-subgroups giving an optimal splitting is $\varepsilon_N(J)/2$.

[2] This is the analogue over k of the surface $Z_{N,-1}$ of Kani and Schanz [KS98].
[3] Such covers are called *optimal* in [CCF24].

3.3 Irreducibility and equations of Humbert surfaces

Since $\mathrm{char}(k) \nmid N$, $\mathcal{H}_{N^2}$ is a closed irreducible surface in $\mathcal{A}_2$ [Kan94, Cor. 3.10]. We now look at its image under the Igusa–Clebsch map IC from Section 2.3.

Lemma 3.8. *Let $\mathcal{H}^{\mathrm{IC}}_{N^2} \subseteq \mathbb{P}(2,4,6,10)$ be the Zariski closure of $\mathrm{IC}(\mathcal{H}_{N^2} \cap \mathcal{M}_2)$ over k. Then $\mathcal{H}^{\mathrm{IC}}_{N^2}$ is an irreducible surface, whose intersection with $\{I_{10} \neq 0\}$ is $\mathrm{IC}(\mathcal{H}_{N^2} \cap \mathcal{M}_2)$. Its homogeneous ideal in $k[I_2, I_4, I_6, I_{10}]$ is principal, generated by an irreducible weighted-homogeneous polynomial $\mathcal{G}$ (unique up to a scalar).*

Proof. Since $\mathcal{M}_2 = \mathcal{A}_2 \setminus \mathcal{H}_1$ is open in $\mathcal{A}_2$ and $\mathcal{H}_{N^2} \not\subseteq \mathcal{H}_1$ (indeed, $\mathrm{gl}_N(E_1, E_2, \alpha)$ is a Jacobian whenever E_1 and E_2 are not isogenous [Kan97, Cor. 2.4]), the intersection $\mathcal{H}_{N^2} \cap \mathcal{M}_2$ is a non-empty open, hence dense, subset of the irreducible surface $\mathcal{H}_{N^2}$. In particular, it is an irreducible surface, and it is closed in $\mathcal{M}_2$, as $\mathcal{H}_{N^2}$ is closed in $\mathcal{A}_2$.

Since IC is an isomorphism from $\mathcal{M}_2$ onto $\{I_{10} \neq 0\}$, the image $\mathrm{IC}(\mathcal{H}_{N^2} \cap \mathcal{M}_2)$ is an irreducible surface which is closed in $\{I_{10} \neq 0\}$. The closure $\mathcal{H}^{\mathrm{IC}}_{N^2}$ is therefore an irreducible surface whose intersection with $\{I_{10} \neq 0\}$ is $\mathrm{IC}(\mathcal{H}_{N^2} \cap \mathcal{M}_2)$.

As $\mathcal{H}^{\mathrm{IC}}_{N^2}$ is an irreducible surface in $\mathbb{P}(2,4,6,10)$, its homogeneous ideal is a homogeneous prime ideal of height 1 in the UFD $k[I_2, I_4, I_6, I_{10}]$. Therefore, it is principal and generated by a single irreducible weighted-homogeneous polynomial $\mathcal{G}$, unique up to a scalar in $k^\times$. □

Remark 3.9. Let $k = \overline{\mathbb{F}}_p$ (resp. $\overline{\mathbb{Q}}$) and $k_0 = \mathbb{F}_p$ (resp. $\mathbb{Q}$). Definition 3.3 is intrinsic, so $\mathcal{H}_{N^2}$ is stable under $\mathrm{Gal}(k/k_0)$, and the Igusa–Clebsch invariants commute with this action. Hence, $\mathcal{H}^{\mathrm{IC}}_{N^2}$ is Galois-stable, so $\mathcal{G}^\sigma = c_\sigma \mathcal{G}$ with $c_\sigma \in k^\times$ for every $\sigma \in \mathrm{Gal}(k/k_0)$. Normalising one non-zero coefficient of $\mathcal{G}$ to 1 forces $c_\sigma = 1$ for all σ, so $\mathcal{G} \in k_0[I_2, I_4, I_6, I_{10}]$.

Theorem 3.10. *Let $\mathcal{G}$ be as in Lemma 3.8, let C be a genus-2 curve defined over k with Jacobian J, and let (i_1, i_2, i_3) be its normalised invariants as in Eq. (2). The following are equivalent:*

1. *$\mathcal{G}(I_2(C), I_4(C), I_6(C), I_{10}(C)) = 0$;*
2. *$\mathcal{G}(i_1, i_2, i_3, 1) = 0$;*
3. *$J \in \mathcal{H}_{N^2}(k)$;*
4. *J is optimally (N, N)-split.*

Proof. The equivalence of (1) and (2) holds because $\mathcal{G}$ is weighted-homogeneous; see Section 2.3. Since $I_{10}(C) \neq 0$, Lemma 3.8 shows that (1) holds if and only if the Igusa–Clebsch invariants of C lie in the image of $\mathcal{H}_{N^2} \cap \mathcal{M}_2$, which by injectivity of the Igusa–Clebsch map is equivalent to (3). Finally, (3) and (4) are equivalent by Theorem 3.5 and the dictionary in Section 3.2. □

3.4 The equation modulo p

We now take $k = \overline{\mathbb{F}}_p$. We write $\mathcal{H}_{N^2,p} := \mathcal{H}^{\mathrm{IC}}_{N^2}(\overline{\mathbb{F}}_p)$, and let $G_N \in \mathbb{F}_p[I_2, I_4, I_6, I_{10}]$ be the generator of its homogeneous ideal given by Lemma 3.8, normalised so

that its first non-zero coefficient (in lexicographic order) is 1. We let d'_N denote the weight of G_N. Since G_N generates the homogeneous ideal of $\mathcal{H}_{N^2,p}$, for every even w the forms of weight w vanishing on $\mathcal{H}_{N^2,p}$ are exactly $G_N \cdot V_{w-d'_N}$, where we set $V_w = 0$ for $w < 0$. We also write

$$g_N(x_1, x_2, x_3) := G_N(x_1, x_2, x_3, 1) \in \mathbb{F}_p[x_1, x_2, x_3]$$

for its dehomogenisation on the chart $I_{10} = 1$. By Theorem 3.10, g_N gives the membership test underlying our algorithm:

Corollary 3.11. *Let C be a genus-2 curve defined over $\overline{\mathbb{F}}_p$ with normalised invariants (i_1, i_2, i_3). Then $g_N(i_1, i_2, i_3) = 0$ if and only if $\mathrm{Jac}(C)$ is optimally (N, N)-split.*

The polynomial G_N is not known a priori. In Section 4 we compute it by interpolation using the gluing map gl_N to sample points on $\mathcal{H}_{N^2,p}$. By relating it to the characteristic-0 equation F_N introduced in the next section, we bound its weight d'_N, thus bounding the cost of interpolation.

3.5 The characteristic-zero equation

The precomputation algorithm of Section 4 computes G_N directly in characteristic p. To bound the size of this computation, we compare it with the classical equation of the Humbert surface in characteristic 0, which we now recall. In this section, we take $k = \overline{\mathbb{Q}}$. By Lemma 3.8, the ideal of $\mathcal{H}^{\mathrm{IC}}_{N^2}$ is generated by an irreducible polynomial with coefficients in $\mathbb{Q}$. Clearing denominators, we obtain a primitive polynomial $F_N \in \mathbb{Z}[I_2, I_4, I_6, I_{10}]$, unique up to sign. We let d_N denote its weight, and write $\bar{F}_N \in \mathbb{F}_p[I_2, I_4, I_6, I_{10}]$ for its reduction modulo p.

Genus-2 curves with split Jacobians over k, and their moduli, have appeared extensively in the literature. Indeed, for $N = 2, 3, 4$, generic families were known classically from work of Legendre, Jacobi, Hermite, Goursat, Burkhardt, Brioschi and Bolza (see the introduction of [Kum15] for details). In modern literature, Kuhn [Kuh88] revisited the case $N = 3$, and Shaska [Sha01a] gave a general method for computing the locus of curves with a degree-N elliptic subcover, which was carried out for $N = 2, 3, 4, 5$ in joint works with Magaard, Völklein, Wijesiri, Wolf and Woodland [MSV09; Sha01b; Sha04; SV04; Sha+08]; see also [Brö+15; Dju19; GS01] for $N = 3$.

Bruin and Doerksen [BD11] instead used the description of splittings via anti-isometries of N-torsion, and computed F_4 by interpolation. Kumar [Kum15] gave explicit rational models of the surfaces $\mathcal{H}_{N^2}$ for all $N \leq 11$, from which F_N can be interpolated for $N \leq 5$ [CCF24]. By Frengley's classification [Fre24], such rational models cannot exist for most larger N. Most recently, Shaska [Sha26b] determined the weight d_N for every N and recomputed F_5 by linear algebra in spaces of Siegel modular forms.

The polynomials F_N grow quickly: $F_2, \ldots, F_5$ have weights $30, 80, 180, 480$ and $34, 318, 2699, 43410$ monomials, respectively, with coefficients of average

bitlength around 17, 64, 197 and 617 [CCF24, Table 2]. The direct computation of F_7 is estimated in [Sha26b] to require about two tebibytes of memory per prime. For this reason, our algorithm never computes F_N.

What we use is the weight d_N, which is known in closed form. Let $J_2(N) = N^2 \prod_{\ell \mid N}(1 - \ell^{-2})$ be Jordan's totient function. Then, by [Sha26b, Thm. 5.5] and [Sha26a, Prop. 5.5], we have

$$d_N = \begin{cases} 30 & \text{if } N = 2, \\ 5(N-1)\, J_2(N) & \text{if } N \geq 3, \end{cases}$$

For instance, $d_3 = 80$, $d_4 = 180$, $d_5 = 480$ and $d_7 = 1440$. In particular, since $J_2(N) < N^2$, we have $d_N < 5N^3$ for all $N \geq 2$. Furthermore, we have that

$$\dim V_{d_N} = O(d_N^3) = O(N^9). \tag{5}$$

This dimension governs the cost of our precomputation: it bounds the number of coefficients of G_N, and hence the number of samples and the size of the linear system we solve to interpolate G_N in Section 4.

Remark 3.12. In [Sha26b], Shaska works with Igusa's invariants J_2, J_4, J_6, J_{10} rather than the Igusa–Clebsch invariants. The two are related by

$$I_2 = 8J_2, \quad I_4 = 4J_2^2 - 96J_4, \quad I_6 = 8J_2^3 - 160J_2J_4 - 576J_6, \quad I_{10} = 4096J_{10}.$$

This change of variables preserves weights and is invertible over $\mathbb{Z}[1/6]$. Hence, the two primitive equations have the same weight, and, since $p > 7$, their reductions modulo p define the same subvariety of $\mathbb{P}(2, 4, 6, 10)$.

Relationship between F_N and G_N. To summarise, we work with two polynomials. The characteristic-0 equation $F_N \in \mathbb{Z}[I_2, I_4, I_6, I_{10}]$ has known weight d_N, but it is large and our algorithm never computes it. The characteristic-p equation $G_N \in \mathbb{F}_p[I_2, I_4, I_6, I_{10}]$ is what our algorithm computes and evaluates; its weight d'_N is not known a priori. We now show that the two are related by $G_N \mid \bar{F}_N$, so $d'_N \leq d_N$. Furthermore, as discussed in Remark 4.6, we expect $\bar{F}_N = c\, G_N$, although we do not need this.

Lemma 3.13. *We have $\bar{F}_N \neq 0$, and $\bar{F}_N$ vanishes on $\mathcal{H}_{N^2,p}$. Consequently, $G_N \mid \bar{F}_N$ and $d'_N \leq d_N$.*

Proof. Since F_N is primitive, $\bar{F}_N \neq 0$. By Lemma 3.8, it suffices to show that $\bar{F}_N(I_2(C), I_4(C), I_6(C), I_{10}(C)) = 0$ for every genus-2 curve $C/\overline{\mathbb{F}}_p$ with $J = \mathrm{Jac}(C) \in \mathcal{H}_{N^2}$. By Section 3.2, $J \cong X_\alpha$ for some $(E_1, E_2, \alpha) \in Z_N$, and we may assume all of these are defined over a finite field $\mathbb{F}$. Let $L/\mathbb{Q}_p$ be a finite extension whose residue field contains $\mathbb{F}$, with ring of integers $\mathcal{O}_L$ and maximal ideal $\mathfrak{m}_L$, and fix an algebraic closure $\bar{L}$ of L; we enlarge L inside $\bar{L}$ when necessary.

The proof strategy goes as follows. Recall that the polynomial F_N is defined in characteristic 0, so to show that $\bar{F}_N$ vanishes at J, we construct a split Jacobian in characteristic 0 whose reduction modulo $\mathfrak{m}_L$ is J, and then reduce

the identity $F_N = 0$ at this lift. Since $p \nmid N$, all the torsion and isogenies involved have order and degree prime to p, and we use the following standard facts: (i) if E/L is an elliptic curve with good reduction, reduction modulo $\mathfrak{m}_L$ is injective on the prime-to-p torsion of $E(L)$ [Sil09, Prop. VII.3.1]; (ii) if X has good reduction and $X \to X'$ is a surjective homomorphism of abelian varieties over L, then X' has good reduction [ST68, §1, Cor. 2]; and (iii) homomorphisms between abelian varieties with good reduction reduce to homomorphisms between their reductions, compatibly with composition and degree.[4] In particular, isomorphisms and polarisations reduce to isomorphisms and polarisations.

We first lift the gluing data. Choose elliptic curves $\tilde{E}_i/\mathcal{O}_L$ lifting E_i so that $\tilde{E}_i$ has good reduction. Enlarge L so that all points of $\tilde{E}_i[N]$ are defined over L. By (i), reduction modulo $\mathfrak{m}_L$ induces a group isomorphism $\tilde{E}_i[N] \to E_i[N]$, and likewise $\mu_N(L) \to \mu_N(\overline{\mathbb{F}}_p)$, where $\mu_N \subseteq L$ since the Weil pairing on $\tilde{E}_1[N]$ is surjective. Moreover, the reduction is compatible with the Weil pairing. Hence, there is a unique isomorphism $\tilde{\alpha}\colon \tilde{E}_1[N] \to \tilde{E}_2[N]$ reducing to α, and it is an anti-isometry.

The quotient $\tilde{X} := (\tilde{E}_1 \times \tilde{E}_2)/\Gamma_{\tilde{\alpha}}$ is a principally polarised abelian surface over L. Being isogenous to $\tilde{E}_1 \times \tilde{E}_2$, it has good reduction by (ii). By (iii), the quotient isogeny reduces to an isogeny $E_1 \times E_2 \to \tilde{X}_{\mathrm{red}}$ of degree N^2, whose kernel contains the reduction Γ_α of $\Gamma_{\tilde{\alpha}}$, which has order N^2. Hence $\tilde{X}_{\mathrm{red}}$ is $(E_1 \times E_2)/\Gamma_\alpha = X_\alpha \cong J$, and by (iii) this isomorphism holds as principally polarised surfaces.

The surface $\tilde{X}$ has an elliptic subgroup of degree N, namely the image of $\tilde{E}_1 \times \{0\}$, so $\tilde{X} \in \mathcal{H}_{N^2}(\bar{L})$ by Theorem 3.5 with $k = \bar{L}$. Moreover, $\tilde{X}$ is not a product of elliptic curves. Indeed, if $\tilde{X} \cong E' \times E''$ over $\bar{L}$ with the product polarisation, then after enlarging L both E' and E'' have good reduction by (ii), and by (iii) reducing would give $J \cong \bar{E}' \times \bar{E}''$ as principally polarised surfaces, contradicting that J is a Jacobian. Hence, after enlarging L, we have $\tilde{X} \cong \mathrm{Jac}(\tilde{C})$ for a genus-2 curve $\tilde{C}/L$, and $F_N(I(\tilde{C})) = 0$ by the definition of F_N.[5]

Since $\mathrm{Jac}(\tilde{C})$ has good reduction and its reduction $J = \mathrm{Jac}(C)$ is the Jacobian of a smooth curve, the curve $\tilde{C}$ itself has good reduction, with reduction isomorphic to C (e.g., see the proof of [Sha26b, Prop. 7.11]). Concretely, after enlarging L and a change of coordinates, $\tilde{C}$ has a model $y^2 = \tilde{f}(x)$ with $\tilde{f} \in \mathcal{O}_L[x]$ of degree 6 whose reduction modulo $\mathfrak{m}_L$ is a sextic model of C. Its Igusa–Clebsch invariants then lie in $\mathcal{O}_L$, with I_{10} a unit since the reduction is smooth, and reduce to those of C. Reducing $F_N(I(\tilde{C})) = 0$ modulo $\mathfrak{m}_L$ gives $\bar{F}_N(I(C)) = 0$. □

Remark 3.14. Shaska [Sha26c, Thm. 2] showed that, for p outside a finite set called B_N, the reduction modulo p of the characteristic-0 Humbert form vanishes at the invariants of $\mathrm{Jac}(C)$, for C a genus-2 curve over $\overline{\mathbb{F}}_p$, if and only if $\mathrm{Jac}(C)$ is optimally (N, N)-split. Here B_N consists of the primes dividing $2N$ or an auxiliary integer depending on N, together with the primes for which $\bar{F}_N$ is

[4] This is the Néron mapping property; see [BLR12, §1.2].

[5] Here we are using the fact that F_N has rational coefficients and that membership in $\mathcal{H}_{N^2}$ is algebraic, so we may transport the statement along an embedding $\bar{L} \hookrightarrow \mathbb{C}$.

reducible or not generically reduced [Sha26c, §5.2]. The implication from vanishing to splitting is [Sha26b, Prop. 7.11(1)]. For the converse, [Sha26c, §5.1] lifts the kernel of the dual splitting to characteristic 0. This is also the strategy taken by Lemma 3.13, where we lift the gluing data. Lemma 3.13 proves only that optimally (N, N)-split Jacobians are zeros of $\bar{F}_N$, but it does so for every $p > 7$ not dividing N. This suffices for our purposes, since F_N enters our algorithm only through its weight. Indeed, our membership test will use G_N, which detects optimal (N, N)-splittings exactly by Corollary 3.11.

3.6 Density in the superspecial locus

For a p.p. abelian surface X, define the *mass of* X to be $\mu(X) := 1/\#\operatorname{Aut}(X)$. For a set S we define $\mu(S) := \sum_{X \in S} \mu(X)$.

Example 3.15. We have

$$\mu(\mathrm{SS}_1(p)) = \frac{p-1}{24}, \text{ and } \mu(\mathrm{SS}_2(p)) = \frac{(p-1)(p^2+1)}{5760}.$$

This is a classical result; see Hashimoto–Ibukiyama [HI80, Prop. 9] and Equation (5.3) of Ekedahl [Eke87].

Using Lemma 2.10 and Thm. 4.1(3) from [Cor+26], we have the following proposition.

Proposition 3.16. *Let $V = \mathrm{SS}_2(p)$ be the vertex set of $\mathcal{G}(p; 2)$ and $S \subseteq V$. Fix some $X_0 \in V$ and set*

$$n_S := \left\lceil \log\left(2 \,\#\operatorname{Aut}(X_0)^{1/2}\, \mu(V)\, \mu(S)^{-1/2}\right) \cdot (\log(9/8))^{-1} \right\rceil.$$

For every $n \geq n_S$, taking a random walk in $\mathcal{G}(p; 2)$ of length n starting at X_0, the probability of landing in S is at least $\frac{1}{2}\frac{\mu(S)}{\mu(V)}$.

Remark 3.17. The constant $\log(9/8)$ in Proposition 3.16 comes from the unconditional bound $8/9$ on the normalised non-trivial eigenvalues of $\mathcal{G}(p; 2)$ [Cor+26, Thm. 4.1(3)]. Assuming the twisted weighted fundamental lemma, this bound improves to $4/5$ [Cor+26, Thm. 4.1(4) and Rem. 5.35]. One can then replace $\log(9/8)$ by $\log(5/4)$ in the definition of n_S, and the walks become about 1.9 times shorter. This only affects the constants, as in both cases $n_S = O(\log p)$.

Let $\mathcal{H}^{\mathrm{ss}}_{N^2} := \mathcal{H}_{N^2}(\overline{\mathbb{F}}_p) \cap \mathrm{SS}_2(p)$. For the product locus $\mathcal{H}^{\mathrm{ss}}_1$, by [Cor+26, Lemma 6.5] we have

$$\frac{\mu(\mathcal{H}^{\mathrm{ss}}_1)}{\mu(\mathrm{SS}_2(p))} \approx \frac{5}{p}. \tag{6}$$

Set $\mathcal{H}^{\mathrm{ss}} := \bigcup_{2 \leq N \leq D} \mathcal{H}^{\mathrm{ss}}_{N^2}$. We are interested in (a lower bound on) the density of $\mathcal{H}^{\mathrm{ss}}$ in $\mathrm{SS}_2(p)$. This will allow us to determine (a lower bound on) the probability that a random surface in $\mathrm{SS}_2(p)$ lies on $\mathcal{H}^{\mathrm{ss}}$ via Proposition 3.16.

Recall from Theorem 3.5 that $X \in \mathcal{H}_{N^2}$ if and only if $\varepsilon_N(X) = \#\mathcal{E}_N(X) > 0$. Moreover, elliptic subgroups come in complementary pairs $\{E, E^\perp\}$ (c.f. Section 3.2). Let

$$\omega(X) := \frac{1}{2} \sum_{2 \le N \le D} \varepsilon_N(X)$$

be the number of complementary pairs $\{E, E^\perp\}$ of degree N with $2 \le N \le D$. With this, we have that $\mathcal{H}^{\mathrm{ss}} = \{X \in \mathrm{SS}_2(p) \colon \omega(X) \ge 1\}$, using which we can rewrite

$$\mu(\mathcal{H}^{\mathrm{ss}}) = \sum_{X \in \mathrm{SS}_2(p)} \mathbb{1}_{\omega(X) \ge 1} \mu(X).$$

Noting that $\mathbb{1}_{m \ge 1} \ge m - \binom{m}{2}$ for every integer $m \ge 0$, we obtain the bound

$$\mu(\mathcal{H}^{\mathrm{ss}}) \ge \sum_{X \in \mathrm{SS}_2(p)} \omega(X)\mu(X) - \sum_{X \in \mathrm{SS}_2(p)} \binom{\omega(X)}{2} \mu(X). \tag{7}$$

In the next lemmas we provide a lower bound for the first sum, and an upper bound for the second sum; this provides a lower bound on $\mu(\mathcal{H}^{\mathrm{ss}})$ as desired.

Both sums are sums over pairs (X, E) with $E \in \mathcal{E}_N(X)$. Two such pairs (X, E) and (X', E') are isomorphic if there is an isomorphism $\beta \colon X \to X'$ with $\beta(E) = E'$, and $\mathrm{Aut}(X, E)$ denotes the stabiliser of E in $\mathrm{Aut}(X)$. Recall also that to an anti-isometry $\alpha \in \mathcal{I}_N(E_1, E_2)$ we attach the p.p. abelian surface X_α, and write E_α for the image of $E_1 \times \{0\}$ in X_α, an elliptic subgroup of degree N.

Lemma 3.18. *Let $2 \le N < p$. For every function Φ on isomorphism classes of pairs (X, E) with supersingular $E \in \mathcal{E}_N(X)$ we have*

$$\sum_{X \in \mathrm{SS}_2(p)} \mu(X) \sum_{E \in \mathcal{E}_N(X)} \Phi(X, E) = \sum_{E_1, E_2 \in \mathrm{SS}_1(p)} \frac{1}{\#\mathrm{Aut}(E_1) \#\mathrm{Aut}(E_2)} \sum_{\alpha \in \mathcal{I}_N(E_1, E_2)} \Phi(X_\alpha, E_\alpha),$$

where the outer sum on the right runs over ordered pairs.

Proof. We use the following consequence of the orbit–stabiliser theorem: if a finite group G acts on a finite set Y and Ψ is constant on G-orbits, then

$$\frac{1}{\#G} \sum_{y \in Y} \Psi(y) = \sum_{[y] \in Y/G} \frac{\Psi(y)}{\#\mathrm{Stab}_G(y)}. \tag{8}$$

We show that both sides of the lemma equal $\sum_{(X,E)} \Phi(X, E)/\#\mathrm{Aut}(X, E)$, where the sum runs over isomorphism classes of pairs with X superspecial.

For the left-hand side, fix $X \in \mathrm{SS}_2(p)$ and let $G = \mathrm{Aut}(X)$ act on $Y = \mathcal{E}_N(X)$ by $\beta \cdot E = \beta(E)$. Two pairs (X, E) and (X, E') are isomorphic exactly when $E' = \beta(E)$ for some $\beta \in \mathrm{Aut}(X)$, so the orbits are the isomorphism

classes of pairs (X, E), and the stabiliser of E is $\mathrm{Aut}(X, E)$. Applying (8) with $\Psi = \Phi(X, \cdot)$, and using $\mu(X) = 1/\#\operatorname{Aut}(X)$, gives $\mu(X)\sum_{E\in\mathcal{E}_N(X)}\Phi(X,E) = \sum_{[E]}\Phi(X,E)/\#\operatorname{Aut}(X,E)$.

For the right-hand side, fix $E_1, E_2 \in \mathrm{SS}_1(p)$ and let $G = \mathrm{Aut}(E_1)\times\mathrm{Aut}(E_2)$ act on $Y = \mathcal{I}_N(E_1,E_2)$ by $(u,v)\star\alpha = v\circ\alpha\circ u^{-1}$. By definition, the orbits are the isomorphism classes of triples (E_1, E_2, α), and the stabiliser of α is $\mathrm{Aut}(E_1, E_2, \alpha)$. Applying (8) with $\Psi(\alpha) = \Phi(X_\alpha, E_\alpha)$ turns the inner sum, divided by $\#\operatorname{Aut}(E_1)\#\operatorname{Aut}(E_2)$, into $\sum_{[\alpha]}\Phi(X_\alpha,E_\alpha)/\#\operatorname{Aut}(E_1,E_2,\alpha)$.

Finally, by Section 3.2, $(E_1, E_2, \alpha) \mapsto (X_\alpha, E_\alpha)$ is a bijection between isomorphism classes of triples and of pairs, and $\mathrm{Aut}(E_1, E_2, \alpha) \cong \mathrm{Aut}(X_\alpha, E_\alpha)$: an automorphism $u \times v$ of $E_1 \times E_2$ descends to X_α exactly when it preserves Γ_α, i.e., when $v\alpha = \alpha u$. The bijection restricts to one between triples with E_1, E_2 supersingular and pairs with X superspecial, since $\mu_E\colon E\times E^\perp \to X$ has degree N^2 prime to p. □

Note that setting $\Phi = 1$ and using Eq. (3) immediately gives us the first sum:

$$\sum_{X\in\mathrm{SS}_2(p)}\omega(X)\mu(X) = \frac{1}{2}\mu(\mathrm{SS}_1(p))^2\sum_{2\le N\le D}\#\operatorname{SL}_2(\mathbb{Z}/N\mathbb{Z}) \ge \frac{1}{8\zeta(2)}\,\mu(\mathrm{SS}_1(p))^2 D^4, \tag{9}$$

where we use that $\sum_{2\le N\le D} N^3 \ge \frac{D^4}{4}$ for $D \ge 2$.

Remark 3.19. Dividing Eq. (9) by $\mu(\mathrm{SS}_2(p))$ shows that the expected number of complementary pairs of degree N, with respect to the stationary measure of $\mathcal{G}(p;2)$, is $5\,\#\operatorname{SL}_2(\mathbb{Z}/N\mathbb{Z})(p-1)/(p^2+1)$. Since $\#\operatorname{SL}_2(\mathbb{Z}/N\mathbb{Z}) = NJ_2(N)$, where J_2 is Jordan's totient, this recovers [Sha26a, Thm. 2.1(1)], which is proved there for $p \equiv 3 \pmod 4$ using Siegel's theorem on genus averages. Our argument via Lemma 3.18 is the elementary proof sketched in [Sha26a, Rem. 5.4].

To bound the second sum, we need to understand surfaces with two different splittings, i.e., pairs $E \in \mathcal{E}_N(X)$ and $F \in \mathcal{E}_{N'}(X)$ with $F \notin \{E, E^\perp\}$. We show that a second splitting forces a short isogeny between the elliptic curves of the first one, which makes such pairs (relatively) rare.

Definition 3.20. *For elliptic curves E, E' and an integer $m \ge 1$, we write*

$$\operatorname{Hom}_m(E,E') := \{f \in \operatorname{Hom}(E,E') : \deg f = m\}$$

for the set of isogenies $E \to E'$ of degree m.

Lemma 3.21. *Let X be a p.p. abelian surface, and let $N, N' \ge 2$ coprime to p. For $E \in \mathcal{E}_N(X)$, we have*

$$\#\{F \in \mathcal{E}_{N'}(X) : F \notin \{E, E^\perp\}\} \le \sum_{d\mid NN'}\sum_{a=1}^{d-1}\#\operatorname{Hom}_{a(d-a)}(E, E^\perp).$$

Proof. Let θ be the class of the principal polarisation of X, so that the degree of an elliptic subgroup F is $(F \cdot \theta)$. For distinct elliptic subgroups $F \neq F'$ of X, the intersection number $(F \cdot F')$ equals the order of the finite group scheme $F \cap F'$ [Har77, Prop. V.1.4]. In particular, $(F \cdot F') \geq 1$ (as they always intersect at 0).

We have $E \cap E^\perp = E[N] = E^\perp[N]$ and $\mathrm{cl}(E) + \mathrm{cl}(E^\perp) = N\theta$ [Kan19, Prop. 10(a), Rem. 11(a)]. Write ι_E and $\iota_{E^\perp}$ for the inclusions of E and $E^\perp$ in X. By [Kan19, Lem. 19], the homomorphism $\varphi_1 := \iota_E^\vee \circ \lambda \colon X \to E^\vee \cong E$ is surjective with kernel $E^\perp$. Since the conditions defining the complement are symmetric in E and $E^\perp$, we have $(E^\perp)^\perp = E$, so the same lemma shows that $\varphi_2 := \iota_{E^\perp}^\vee \circ \lambda \colon X \to E^\perp$ is surjective with kernel E. Set $\varphi := (\varphi_1, \varphi_2) \colon X \to E \times E^\perp$. Its kernel is $E \cap E^\perp = E[N]$, which is finite, so φ is an isogeny.

Let $F \in \mathcal{E}_{N'}(X)$ with $F \notin \{E, E^\perp\}$, and put $s = (F \cdot E^\perp)$ and $t = (F \cdot E)$. Then $s, t \geq 1$ and $s + t = (F \cdot N\theta) = NN'$. The restrictions $\varphi_1|_F \colon F \to E$ and $\varphi_2|_F \colon F \to E^\perp$ have kernels $F \cap E^\perp$ and $F \cap E$, so they are isogenies of degrees s and t.

Let $C = \varphi(F) \subseteq E \times E^\perp$, an elliptic subgroup, and let $e = \deg(\varphi|_F \colon F \to C)$. Write $\pi_1 \colon C \to E$ and $\pi_2 \colon C \to E^\perp$ for the restrictions of the projections. Since $\varphi_i|_F = \pi_i \circ \varphi|_F$, we have $\deg \pi_1 = s/e$ and $\deg \pi_2 = t/e$, which are positive integers. Writing $a := \deg \pi_1$ and $d := \deg \pi_1 + \deg \pi_2 = NN'/e$, we have that $d \mid NN'$ and $1 \leq a \leq d - 1$.

Define $f := \pi_2 \circ \hat{\pi}_1 \in \mathrm{Hom}(E, E^\perp)$. It is non-zero of degree $a(d-a)$. We show that (a, f) determines F.

First we recover C from (a, f). Write $\iota \colon C \hookrightarrow E \times E^\perp$ for the inclusion, so that $\iota = (\pi_1, \pi_2)$. Since $\pi_1 \circ \hat{\pi}_1 = [a]$, we have

$$\iota \circ \hat{\pi}_1 = (\pi_1 \circ \hat{\pi}_1,\ \pi_2 \circ \hat{\pi}_1) = ([a], f) \colon E \longrightarrow E \times E^\perp.$$

As $\hat{\pi}_1 \colon E \to C$ is surjective, C is the image of the map $([a], f)$, and hence is determined by (a, f).

Now we recover F from C. Since $C = \varphi(F)$, we have $\varphi^{-1}(C) = F + \ker \varphi$. As $\ker \varphi = E[N]$ is finite, this is a finite union of translates of F, and F is the connected component of $\varphi^{-1}(C)$ containing 0. Hence, F is determined by C.

Therefore, the map $F \mapsto (a, f)$ is injective, and the bound follows by counting, for each $d \mid NN'$, the possible pairs (a, f) with $1 \leq a \leq d - 1$ and $\deg f = a(d - a)$. □

We now combine the result above with Lemma 3.18 to obtain a bound for the second sum.

Lemma 3.22. *We have*

$$\sum_{X \in \mathrm{SS}_2(p)} \binom{\omega(X)}{2} \mu(X) \leq \frac{5\zeta(3)^2}{96} \mu(\mathrm{SS}_1(p)) D^{11}.$$

Proof. For $2 \le N' \le D$, write $\Phi_{N'}(X,E) := \#\{F \in \mathcal{E}_{N'}(X) : F \notin \{E, E^\perp\}\}$, and for supersingular E_1, E_2 put

$$B_{NN'}(E_1,E_2) := \sum_{d \mid NN'} \sum_{a=1}^{d-1} \#\operatorname{Hom}_{a(d-a)}(E_1,E_2).$$

Let $X \in \mathrm{SS}_2(p)$. The set $\bigcup_{2\le N\le D} \mathcal{E}_N(X)$ has $2\omega(X)$ elements, each E in it has a complement $E^\perp \neq E$ of the same degree, so counting ordered pairs (E,F) with $F \notin \{E, E^\perp\}$ gives

$$8\binom{\omega(X)}{2} = 2\omega(X)\big(2\omega(X)-2\big) = \sum_{2\le N,N'\le D} \sum_{E \in \mathcal{E}_N(X)} \Phi_{N'}(X,E).$$

Multiplying by $\mu(X)$ and summing over X, we obtain

$$\begin{aligned} 8\sum_X \binom{\omega(X)}{2}\mu(X) &= \sum_{2\le N,N'\le D} \sum_{E_1,E_2 \in \mathrm{SS}_1(p)} \frac{1}{\#\operatorname{Aut}(E_1)\#\operatorname{Aut}(E_2)} \sum_{\alpha \in \mathcal{I}_N(E_1,E_2)} \Phi_{N'}(X_\alpha, E_\alpha) \\ &\le \sum_{2\le N,N'\le D} N^3 \sum_{E_1,E_2 \in \mathrm{SS}_1(p)} \frac{B_{NN'}(E_1,E_2)}{\#\operatorname{Aut}(E_1)\#\operatorname{Aut}(E_2)} \\ &\le \mu(\mathrm{SS}_1(p)) \sum_{2\le N,N'\le D} N^3 \sum_{d\mid NN'} \sum_{a=1}^{d-1} \sigma_1\big(a(d-a)\big). \end{aligned}$$

Here, the first equality follows from Lemma 3.18 applied for each N, N' with $\Phi = \Phi_{N'}$. For the first inequality, Lemma 3.21 gives $\Phi_{N'}(X_\alpha, E_\alpha) \le B_{NN'}(E_1,E_2)$ for every α, since the quotient map identifies E_1 and E_2 with E_α and $E_\alpha^\perp$, and there are $\#\mathcal{I}_N(E_1,E_2) = \#\,\mathrm{SL}_2(\mathbb{Z}/N\mathbb{Z}) \le N^3$ anti-isometries. For the second inequality, we use that for every $m \ge 1$ and supersingular E_1,

$$\sum_{E_2 \in \mathrm{SS}_1(p)} \frac{\#\operatorname{Hom}_m(E_1,E_2)}{\#\operatorname{Aut}(E_2)} = \#\{\text{subgroup schemes of } E_1 \text{ of order } m\} \le \sigma_1(m),$$

where $\sigma_1(n) = \sum_{d\mid n} d$, and then sum over E_1 with weight $1/\#\operatorname{Aut}(E_1)$.

It remains to bound the inner sum. Since $\sigma_1(mn) \le \sigma_1(m)\sigma_1(n)$, Glaisher's identity [Gla85] (see also [Kan97, Eq. (4.16)]) gives, for every $d \ge 2$,

$$\sum_{a=1}^{d-1} \sigma_1\big(a(d-a)\big) \le \sum_{a=1}^{d-1} \sigma_1(a)\sigma_1(d-a) = \frac{5\sigma_3(d) + (1-6d)\sigma_1(d)}{12} \le \frac{5}{12}\sigma_3(d).$$

Using $\sigma_3(d) \le \zeta(3) d^3$ and summing over $d \mid NN'$ bounds the inner sum by $\frac{5\zeta(3)^2}{12} N^3 (N')^3$. Finally, $\sum_{N\le D} N^6 \le D^7$ and $\sum_{N'\le D} (N')^3 \le D^4$ give the statement. □

We reach the main result of this section.

Proposition 3.23. *Let $D \geq 2$ with $200 \cdot D^7 \leq p$, and let $V = \mathrm{SS}_2(p)$. Then*

$$\frac{\mu(\mathcal{H}^{ss})}{\mu(V)} \geq \frac{D^4}{4p}.$$

In particular, $\mu(\mathcal{H}^{ss})/\mu(V) = \Omega(D^4/p)$ for $D \leq (p/200)^{1/7}$.

Proof. Putting Eq. (9) and Lemma 3.22 into Eq. (7) gives

$$\frac{\mu(\mathcal{H}^{\mathrm{ss}})}{\mu(V)} \geq \frac{5(p-1)}{p^2+1} \sum_{2 \leq N \leq D} \#\,\mathrm{SL}_2(\mathbb{Z}/N\mathbb{Z}) - \frac{19D^{11}}{p^2+1}.$$

For the main term, the sum is at least $\frac{D^4}{4\zeta(2)} = \frac{3D^4}{2\pi^2}$. Moreover, $\frac{p-1}{p^2+1} \geq \frac{1}{2p}$ for $p \geq 3$, so the main term is at least $\frac{15D^4}{4\pi^2 p}$. For the error term, $200D^7 \leq p$ gives $\frac{19D^{11}}{p^2+1} \leq \frac{19D^4}{200p}$. Therefore,

$$\frac{\mu(\mathcal{H}^{\mathrm{ss}})}{\mu(V)} \geq \left(\frac{15}{4\pi^2} - \frac{19}{200}\right)\frac{D^4}{p} \geq \frac{D^4}{4p},$$

as required. □

We use this to obtain a lower bound on the probability that a random walk of length n in $\mathcal{G}(p;2)$ starting at X_0 hits a point in $\mathcal{H}^{\mathrm{ss}} \cup H_1^{\mathrm{ss}}$. To this end, let n be as in Proposition 3.16 with $\mu(S)$ replaced by $\mu(\mathcal{H}_1^{ss}) = \mu(\mathrm{SS}_1(p))^2/2$ [Cor+26, Eq. (24)]. Since $\mu(\mathcal{H}_1^{ss}) \leq \mu(\mathcal{H}_1^{ss} \cup \mathcal{H}^{\mathrm{ss}})$ and n_S is decreasing in $\mu(S)$, we have $n \geq n_S$ both for $S = \mathcal{H}_1^{ss}$ and for $S = \mathcal{H}_1^{ss} \cup \mathcal{H}^{\mathrm{ss}}$. We obtain the following result.

Corollary 3.24. *Let $D \geq 2$ with $200D^7 \leq p$, and let $S := \mathcal{H}_1^{ss} \cup \mathcal{H}^{ss}$. With n as above, let X be the codomain of a random walk of length n in $\mathcal{G}(p;2)$ starting at X_0. Then $\Pr[X \in S] \geq D^4/(8p)$, and $\Pr[X \in \mathcal{H}_1^{ss}] \geq \frac{1}{2}\mu(\mathcal{H}_1^{ss})/\mu(V) > 0$.*

Remark 3.25. In comparison to Lemma 3.22 and Proposition 3.23, Shaska [Sha26a] obtains sharper results for a fixed set of N. For every finite set S of N and a sufficiently large prime $p \equiv 3 \pmod 4$, he shows that the second term in Eq. (7) is exactly of the form $C_S/(p^2+1)$, with a constant C_S depending only on S, and deduces the proportion of optimally (N,N)-split surfaces with $N \in S$ [Sha26a, Cor. 6.8]. The constants C_S are determined explicitly only for $N \leq 7$, and their growth with S is not controlled. Our bound in Lemma 3.22 is of the same shape, but with a constant that is not sharp. On the other hand, it holds for every prime p, with an error term explicit in D. This is what our algorithm requires, since D grows with p.

3.7 Gluing over finite fields

For elliptic curves E_1, E_2 over $\mathbb{F}_q$, write $\mathcal{I}_N(E_1,E_2)(\mathbb{F}_q) \subseteq \mathcal{I}_N(E_1,E_2)$ for the set of anti-isometries $E_1[N] \to E_2[N]$ that commute with the q-power Frobenius. Equivalently, Γ_α is defined over $\mathbb{F}_q$. Let $Z_N(\mathbb{F}_q)$ denote the set of $\mathbb{F}_q$-isomorphism classes of triples (E_1, E_2, α) with E_1, E_2 defined over $\mathbb{F}_q$ and $\alpha \in$

$\mathcal{I}_N(E_1,E_2)(\mathbb{F}_q)$, where isomorphisms are as in Section 3.2 but defined over $\mathbb{F}_q$, and write $\mathrm{Aut}(E_1,E_2,\alpha)(\mathbb{F}_q)$ for the group of automorphisms over $\mathbb{F}_q$. We define the mass of $Z_N(\mathbb{F}_q)$ as

$$\mu(Z_N(\mathbb{F}_q)) := \sum_{(E_1,E_2,\alpha)\in Z_N(\mathbb{F}_q)} \frac{1}{\#\,\mathrm{Aut}(E_1,E_2,\alpha)(\mathbb{F}_q)}.$$

Lemma 3.26. *Let E_1, E_2 be independent random elliptic curves over $\mathbb{F}_q$, where an $\mathbb{F}_q$-isomorphism class E is drawn with probability $1/(q\#\,\mathrm{Aut}(E))$, and let $\nu(E_1,E_2) := \#\mathcal{I}_N(E_1,E_2)(\mathbb{F}_q)$. Then*

$$\mu(Z_N(\mathbb{F}_q)) = \mathbb{E}[\nu]\,q^2, \quad \textit{and} \quad \big|\mathbb{E}[\nu]-1\big| \le c\,\frac{N^6}{q},$$

for some absolute constant c. In particular, if $q=p^2$ and $N \le \frac{1}{3}p^{1/3}$, then $\mathbb{E}[\nu] \ge 1/2$ for p large enough.

Proof. The measure in the statement is a probability measure. Indeed, the $\mathbb{F}_q$-isomorphism classes of curves with a given j-invariant have total mass 1 [How93, Prop. 2.1].

We first compute the mass of $Z_N(\mathbb{F}_q)$, arguing as in the proof of Lemma 3.18. For representatives E_1, E_2 of two $\mathbb{F}_q$-isomorphism classes, the group $\mathrm{Aut}(E_1)\times\mathrm{Aut}(E_2)$ acts on $\mathcal{I}_N(E_1,E_2)(\mathbb{F}_q)$ by $(u_1,u_2)\cdot\alpha = u_2\circ\alpha\circ u_1^{-1}$, since such automorphisms preserve the Weil pairing and commute with Frobenius. Its orbits are the isomorphism classes of triples with underlying curves E_1, E_2, and the stabiliser of α is $\mathrm{Aut}(E_1,E_2,\alpha)$. Applying (8) with $\Psi = 1$ and summing over all pairs of isomorphism classes gives

$$\mu(Z_N(\mathbb{F}_q)) = q^2 \sum_{E_1}\sum_{E_2} \frac{\nu(E_1,E_2)}{q\,\#\,\mathrm{Aut}(E_1)\cdot q\,\#\,\mathrm{Aut}(E_2)} = q^2\,\mathbb{E}[\nu].$$

It remains to estimate $\mathbb{E}[\nu]$, which we reduce to a question about matrices. Fix a primitive N-th root of unity ζ, and let $\mathcal{S}_q := \{\mathbf{M}\in \mathrm{GL}_2(\mathbb{Z}/N\mathbb{Z}) : \det\mathbf{M} = q\}$, so that $\#\mathcal{S}_q = \#\mathrm{SL}_2(\mathbb{Z}/N\mathbb{Z}) \le N^3$. For $E/\mathbb{F}_q$, let $\mathcal{C}_E \subseteq \mathcal{S}_q$ be the $\mathrm{SL}_2(\mathbb{Z}/N\mathbb{Z})$-conjugacy class of the matrix of the q-power Frobenius on $E[N]$ in a basis (P,Q) with $e_N(P,Q)=\zeta$. This is well defined, as two such bases differ by an element of $\mathrm{SL}_2(\mathbb{Z}/N\mathbb{Z})$. For $\Pi_1,\Pi_2\in\mathcal{S}_q$, let

$$\nu(\Pi_1,\Pi_2) := \#\{\mathbf{M}\in\mathrm{GL}_2(\mathbb{Z}/N\mathbb{Z}) : \det\mathbf{M}=-1,\ \mathbf{M}\Pi_1 = \Pi_2\mathbf{M}\}.$$

In such bases, anti-isometries $E_1[N]\to E_2[N]$ are exactly the matrices of determinant -1, so $\nu(E_1,E_2) = \nu(\Pi_{E_1},\Pi_{E_2}) \le N^3$. Moreover, $\nu(\Pi_1,\Pi_2)$ only depends on the classes of Π_1,Π_2 (conjugating by $\mathbf{S}_i\in\mathrm{SL}_2(\mathbb{Z}/N\mathbb{Z})$ corresponds to $\mathbf{M}\mapsto\mathbf{S}_2\mathbf{M}\mathbf{S}_1^{-1}$), so we write $\nu(\mathcal{C}_1,\mathcal{C}_2)$. Each $\mathbf{M}$ of determinant -1 contributes to exactly one $\Pi_2 = \mathbf{M}\Pi_1\mathbf{M}^{-1}\in\mathcal{S}_q$, so for every $\Pi_1\in\mathcal{S}_q$,

$$\frac{1}{\#\mathcal{S}_q}\sum_{\Pi_2\in\mathcal{S}_q}\nu(\Pi_1,\Pi_2) = \frac{\#\{\mathbf{M}:\det\mathbf{M}=-1\}}{\#\mathcal{S}_q} = 1, \tag{10}$$

and symmetrically in Π_2. Thus if Frobenius matrices were uniformly distributed in $\mathcal{S}_q$, we would have $\mathbb{E}[\nu] = 1$ exactly.

Frobenius matrices are close to uniform. For a class $\mathcal{C} \subseteq \mathcal{S}_q$, let $P(\mathcal{C})$ be the probability that $j(E) \notin \{0, 1728\}$ and $\mathcal{C}_E = \mathcal{C}$, and set $\varepsilon(\mathcal{C}) := P(\mathcal{C}) - \#\mathcal{C}/\#\mathcal{S}_q$. Then

$$|\varepsilon(\mathcal{C})| \leq \frac{c_N \, \#\mathcal{C}}{\sqrt{q}}, \qquad \text{with } c_N = 5 \text{ for } N > 2, \tag{11}$$

and c_2 an absolute constant. For $N > 2$, this follows from the Chebotarev density theorem [CH13, Thm. 4], applied to the image of $\mathcal{C}$ in $G := \mathrm{GL}_2^H(\mathbb{Z}/N\mathbb{Z})/\{\pm 1\}$ with $H = \langle q \rangle$. Here we use that: (i) the $\mathrm{GL}_2^H(\mathbb{Z}/N\mathbb{Z})$-class of $\Pi \in \mathcal{S}_q$ equals its $\mathrm{SL}_2(\mathbb{Z}/N\mathbb{Z})$-class, since if $\det \mathbf{M} = q^k$ then $\mathbf{M}\Pi\mathbf{M}^{-1} = \mathbf{S}\Pi\mathbf{S}^{-1}$ with $\mathbf{S} = \mathbf{M}\Pi^{-k} \in \mathrm{SL}_2(\mathbb{Z}/N\mathbb{Z})$; (ii) that quadratic twists exchange $\mathcal{C}$ and $-\mathcal{C}$ [CH13, Lemmas 5–6]; (iii) that $\#\operatorname{Aut}(E) = 2$ for $j(E) \notin \{0, 1728\}$; and, (iv) that the genus of the relevant modular curve is $g_L = 1 + (N-6)m/(12N) \leq m$ with $m = \#\mathcal{S}_q/2$ [CH13, Eq. (1)]. For $N = 2$, we can apply a similar argument at level 4 to obtain (11); see the end of [CH13, §2].

We now put everything together. Pairs (E_1, E_2) with $j(E_1)$ or $j(E_2)$ in $\{0, 1728\}$ have total mass at most $4/q$, so they contribute at most $4N^3/q$ to $\mathbb{E}[\nu]$. By (10), the remaining pairs contribute

$$\sum_{\mathcal{C}_1, \mathcal{C}_2} P(\mathcal{C}_1) P(\mathcal{C}_2) \, \nu(\mathcal{C}_1, \mathcal{C}_2) = 1 + 2 \sum_{\mathcal{C}} \varepsilon(\mathcal{C}) + \sum_{\mathcal{C}_1, \mathcal{C}_2} \varepsilon(\mathcal{C}_1) \varepsilon(\mathcal{C}_2) \, \nu(\mathcal{C}_1, \mathcal{C}_2).$$

The second term equals $2\big(\sum_{\mathcal{C}} P(\mathcal{C}) - 1\big) \in [-4/q, 0]$. By Eqs. (10) and (11), the last term is at most $\frac{c_N^2}{q} \sum_{\Pi_1, \Pi_2} \nu(\Pi_1, \Pi_2) = \frac{c_N^2 (\#\mathcal{S}_q)^2}{q} \leq \frac{c_N^2 N^6}{q}$ in absolute value. Hence $|\mathbb{E}[\nu] - 1| \leq (4N^3 + 4 + c_N^2 N^6)/q$, which is at most $33N^6/q$ for $N > 2$ and $O(1/q)$ for $N = 2$. In particular, if $q = p^2$ and $N \leq p^{1/3}/3$, then $|\mathbb{E}[\nu] - 1| < 1/2$ for p large enough. □

4 Precomputation: building the membership test

We showed in Corollary 3.11 that testing whether a genus 2 Jacobian J is optimally (N, N)-split amounts to testing whether g_N vanishes at the normalised Igusa–Clebsch invariants of J. The first step of our algorithm is therefore to compute, for each $2 \leq N \leq D$, the equation G_N of $\mathcal{H}_{N^2,p}$, and to combine these into a single membership test for $\bigcup_{2 \leq N \leq D} \mathcal{H}_{N^2}$. Our algorithm will succeed with high probability.

We follow a similar approach to that of Bruin and Doerksen. They compute the equation of the Humbert surface $\mathcal{H}_{16}$ by interpolating the image of an explicit rational parametrisation modulo many primes, followed by rational reconstruction [BD11, Proof of Thm. 1.2] (see also [CCF24, §4.4] for $N \leq 5$, and Shaska [Sha26c, §4.4] for $N = 6$, which use Kumar's rational parametrisation [Kum15]). We will similarly use interpolation, but instead sample points by gluing pairs of elliptic curves along their N-torsion, which requires no parametrisation. This therefore applies to every N, including those for which $\mathcal{H}_{N^2}$ is not

rational. Moreover, we only need the equation modulo the prime p of interest, so we neither reconstruct it over $\mathbb{Q}$ nor verify it; any false positive is instead detected using the algorithm ComputeSplitting (see Section 5).

We remark that the computation described in this section is independent of the particular instance of the problem we are solving, and can be performed once for a fixed prime p and reused for any number of instances. In this way, it can be viewed as an offline precomputation step.

Sampling points. Let $q = p^2$. We sample points of $\mathcal{H}_{N^2,p}$ defined over $\mathbb{F}_q$ by gluing pairs of random elliptic curves over $\mathbb{F}_q$. We require $\mathbb{F}_q$-rationality so that the linear algebra in Algorithm 2 takes place over $\mathbb{F}_q$. In theory, we could instead set $q = p^k$ for some small $k \geq 2$. Indeed, points defined over larger extensions would also lie on $\mathcal{H}_{N^2,p}$. We chose $k = 2$ for simplicity.

We sample a point using the construction in Section 3.2, taking care of this rationality condition. It works as follows.

1. For $i = 1, 2$, sample $(a_i, b_i) \in \mathbb{F}_q^2$ uniformly at random with $4a_i^3 + 27b_i^2 \neq 0$, and let $E_i : y^2 = x^3 + a_i x + b_i$.
2. For $i = 1, 2$, compute a basis (P_i, Q_i) of $E_i[N]$ with $e_N(P_1, Q_1) = e_N(P_2, Q_2)$, and the matrix $\Pi_i \in \mathrm{GL}_2(\mathbb{Z}/N\mathbb{Z})$ of the q-power Frobenius on $E_i[N]$ in this basis. The points of $E_i[N]$ are defined over an extension of $\mathbb{F}_q$ of degree at most $O(N^2)$ (i.e., the order of Π_i).
3. Sample a uniformly random $\mathbf{M} \in \mathrm{GL}_2(\mathbb{Z}/N\mathbb{Z})$ with $\det \mathbf{M} = -1$. Such matrices correspond exactly to the anti-isometries $\alpha : E_1[N] \to E_2[N]$ in the chosen bases. If $\mathbf{M}\Pi_1 \neq \Pi_2\mathbf{M}$, go back to Step 1.
4. Let $\Gamma_\alpha = \{(P, \alpha(P)) : P \in E_1[N]\} \subseteq (E_1 \times E_2)[N]$. Since α commutes with Frobenius, Γ_α is defined over $\mathbb{F}_q$. We then compute the quotient $X := (E_1 \times E_2)/\Gamma_\alpha$. To do so, we decompose the (N, N)-isogeny into a chain of (ℓ, ℓ)-isogenies for primes $\ell \mid N$ with kernels also defined over $\mathbb{F}_q$. We compute each (ℓ, ℓ)-isogeny using [LR23] for odd ℓ and [Dar+24] for $\ell = 2$.
5. Compute the Igusa–Clebsch invariants I_2, I_4, I_6, I_{10} of X, rescaling them if necessary to a representative in $\mathbb{F}_q^4$. If $I_{10} \neq 0$ (so X is a Jacobian), output the invariants; otherwise, go back to Step 1.

On input N, s and p, the algorithm SamplePoints runs the procedure above repeatedly to produce s samples.

Lemma 4.1. *Let $\epsilon > 0$, and suppose that $N \leq p^{1/3-\epsilon}$. For p sufficiently large in terms of ϵ, on input N, s, p, the algorithm* SamplePoints *outputs s independent points of $\mathcal{H}_{N^2,p}(\mathbb{F}_q)$ in expected time $s \cdot N^{7+o(1)}(\log p)^{O(1)}$. Each output arises from the isomorphism class of a triple in $Z_N(\mathbb{F}_q)$, and every class is output with probability at most $2/\mu(Z_N(\mathbb{F}_q))$.*

Proof. Correctness follows from Section 3.2 and Lemma 3.8: an anti-isometry commutes with the q-power Frobenius exactly when Γ_α is defined over $\mathbb{F}_q$, in which case X_α is defined over $\mathbb{F}_q$. If X_α is a Jacobian, then $X_\alpha \in \mathcal{H}_{N^2} \cap \mathcal{M}_2$, so its Igusa–Clebsch invariants lie on $\mathcal{H}_{N^2,p}$.

Given the pair and the bases of Step 2, the anti-isometries correspond bijectively to the matrices of determinant -1, and $\mathbf{M}$ is uniform among the $\#\mathrm{SL}_2(\mathbb{Z}/N\mathbb{Z})$ such matrices. The matrices $\mathbf{M}$ that give a fixed class (E_1, E_2, α) form an orbit of $\mathrm{Aut}(E_1)\times\mathrm{Aut}(E_2)$, of size $\#\,\mathrm{Aut}(E_1)\#\,\mathrm{Aut}(E_2)/\#\,\mathrm{Aut}(E_1, E_2, \alpha)$. Hence a pass through Steps 1–3 produces this class with probability

$$\frac{1}{\big(q^2\,\#\mathrm{SL}_2(\mathbb{Z}/N\mathbb{Z})\,\#\,\mathrm{Aut}(E_1, E_2, \alpha)\big)}.$$

Summing over all classes, we accept with probability $\mu(Z_N(\mathbb{F}_q))/\big(q^2\,\#\mathrm{SL}_2(\mathbb{Z}/N\mathbb{Z})\big)$, and conditionally on acceptance, the class (E_1, E_2, α) is produced with probability

$$\frac{1/\#\,\mathrm{Aut}(E_1, E_2, \alpha)}{\mu(Z_N(\mathbb{F}_q))} \le \frac{1}{\mu(Z_N(\mathbb{F}_q))}.$$

It remains to account for Step 5, which rejects the triples whose quotient is a product of elliptic curves. Since, conditionally on acceptance at Step 3, triples are distributed according to their mass, Step 5 rejects with probability equal to the proportion of the mass of $Z_N(\mathbb{F}_q)$ carried by such triples. We show that this proportion is at most $1/2$. It then follows that every class is output with probability at most $2/\mu(Z_N(\mathbb{F}_q))$, and successive outputs are independent, as they use fresh randomness.

The quotient is a product only if there is an isogeny $E_1 \to E_2$ of degree $m(N-m)$ for some $0 < m < N$ [Kan97, Cor. 2.4]. The number of such pairs $(j_1, j_2) \in \mathbb{F}_q^2$ is at most $q\sum_{0<m<N}\sigma_1\big(m(N-m)\big) = O(N^3 q)$ by Glaisher's identity. Each pair contributes mass at most $\nu(E_1, E_2) \le N^3$ to $Z_N(\mathbb{F}_q)$, so the triples with a product quotient have mass $O(N^6 q)$. As $\mu(Z_N(\mathbb{F}_q)) \ge q^2/2$, they form a proportion $O(N^6/p^2) = O(p^{-6\epsilon}) \le 1/2$ for p large enough, as required.

We now compute the cost. In Step 2, we compute a basis of $E_i[N]$ by standard methods: factoring the N-division polynomial over $\mathbb{F}_q$ gives an extension $\mathbb{F}_{q^r}$ with $r \le 2N^2$ over which $E_i[N]$ is defined, and a basis is then obtained from random points of $E_i(\mathbb{F}_{q^r})$ by cofactor multiplication and the Weil pairing. We then compute Π_i by applying Frobenius to this basis and computing discrete logarithms of Weil pairings in μ_N. The factorisation takes $\widetilde{O}(N^4 \log q)$ operations in $\mathbb{F}_q$, and the remaining steps take $O(r \log q)$ operations in $\mathbb{F}_{q^r}$, each costing $N^{2+o(1)}(\log p)^{1+o(1)}$. Thus, a pass through Steps 1–3 costs $N^{4+o(1)}(\log p)^{O(1)}$.

It is accepted with probability $\mathbb{E}[\nu]/\#\mathrm{SL}_2(\mathbb{Z}/N\mathbb{Z}) \ge 1/(2N^3)$. After acceptance, each (ℓ, ℓ)-isogeny in Step 4 costs $\tilde{O}(\ell^2)$. As the isogeny chain has $O(\log N)$ steps, the whole (N, N)-isogeny costs $\tilde{O}(N^2)$ operations in the field of definition of the points of Γ_α, an extension of $\mathbb{F}_q$ of degree $O(N^2)$ (see Step 2). The total cost of Step 4 is therefore $N^{4+o(1)}(\log p)^{O(1)}$. Step 5 succeeds with probability at least $1/2$. Hence, each output costs an expected $N^{7+o(1)}(\log p)^{O(1)}$. □

Remark 4.2. The only property of the isogeny algorithms of [Dar+24; LR23] that we use (in Lemma 4.1 and Proposition 5.2) is that an (N, N)-isogeny can be computed in N^c operations for some constant $c \le 5$; any such algorithm leaves all our bounds unchanged. In particular, our results do not depend on the finer details of these algorithms.

To bound the failure probability of Algorithm 2, we need the samples output by SamplePoints to be 'generic enough'. This is the content of the following lemma.

Lemma 4.3. *Let $\epsilon > 0$ and suppose that $N \leq p^{1/3-\epsilon}$. Let $f \in \mathbb{F}_q[I_2, I_4, I_6, I_{10}]$ be a weighted-homogeneous polynomial of weight $w \leq d_N$ that does not vanish identically on $\mathcal{H}_{N^2,p}$, i.e., $G_N \nmid f$. A point produced by* SamplePoints *lies on $V(f)$ with probability $O(N^9 \log\log N/q)$.*

Proof. Let $C = V(f) \cap \mathcal{H}_{N^2,p}$. As $\mathcal{H}_{N^2,p}$ is an irreducible surface by Lemma 3.8, and f does not vanish identically on $\mathcal{H}_{N^2,p}$, C is a curve (or empty).

Let $K \subseteq \mathbb{A}^4$ be the zero set of f and G_N, viewed as polynomials in the four variables I_2, I_4, I_6, I_{10}. K is the affine cone over C, and has dimension 2. As f and G_N have total degrees at most $w/2$ and $d'_N/2$, respectively, by [LR15, Cor. 2.2] we have

$$\#K(\mathbb{F}_q) \leq \frac{w}{2} \cdot \frac{d'_N}{2} \cdot q^2.$$

Every sample lying on C is given by a point $(I_2, I_4, I_6, I_{10}) \in K(\mathbb{F}_q)$, and the action $\lambda \cdot (I_2, I_4, I_6, I_{10}) = (\lambda^2 I_2, \lambda^4 I_4, \lambda^6 I_6, \lambda^{10} I_{10})$ of $\mathbb{F}_q^\times$ gives at least $(q-1)/10$ distinct such representatives of the same point of C. Indeed, if λ stabilises (I_2, I_4, I_6, I_{10}), then $\lambda^{10} I_{10} = I_{10}$, so $\lambda^{10} = 1$, as $I_{10} \neq 0$. Hence, the stabiliser lies in $\mu_{10}(\mathbb{F}_q)$. Moreover, every point of $C(\mathbb{F}_q)$ has a representative in $\mathbb{F}_q^4$. Thus, we have

$$\#C(\mathbb{F}_q) \leq \frac{10}{q-1} \cdot \frac{w \cdot d'_N}{4} \cdot q^2 = O(N^6 q).$$

We now bound the number of classes of $Z_N(\mathbb{F}_q)$ mapping to a given point of $C(\mathbb{F}_q)$, corresponding to a Jacobian J. By Section 3.2, the $\overline{\mathbb{F}}_p$-isomorphism classes of triples (E_1, E_2, α) with $X_\alpha \cong J$ correspond to the $\mathrm{Aut}(J)$-orbits on $\mathcal{E}_N(J)$, so there are at most $\varepsilon_N(J)$ of them. Moreover, $\varepsilon_N(J)$ is twice the number of (N,N)-subgroups of $J[N]$ giving an optimal splitting, so $\varepsilon_N(J) \leq 2\kappa_N = O(N^3 \log\log N)$ by Eq. (1). Each $\overline{\mathbb{F}}_p$-isomorphism class of triples splits into at most $\#\mathrm{Aut}(E_1, E_2, \alpha) \leq 24^2$ classes over $\mathbb{F}_q$ (its twists). So $O(N^3 \log\log N)$ classes of $Z_N(\mathbb{F}_q)$ map to each point of $C(\mathbb{F}_q)$, and so a total of $O(N^9 q \log\log N)$ classes of $Z_N(\mathbb{F}_q)$ map into C.

By Lemma 4.1, each is output with probability at most $2/\mu(Z_N(\mathbb{F}_q))$. Using Lemma 3.26, a sample lies on $V(f)$ with probability $O(N^9 \log\log N/q)$. □

Remark 4.4. We could restrict this sampling to work only with supersingular curves, in which case rationality statements are more straightforward, and would result in a more practical sampling algorithm. However, understanding the distribution of such samples is more challenging, and we are focused mainly on asymptotic cost statements.

Computing the polynomial. After running SamplePoints to obtain a number of samples, we use these to run interpolation. Since d'_N is not known in advance, we proceed in two passes. We first interpolate at weight d_N, which is an upper

Algorithm 1 Interpolate

Input: An integer $N \geq 2$, an even weight $w \geq 0$, and a prime p.
Output: The dimension r of the space of forms in V_w vanishing at the samples, a nonzero such form f, and the samples S.

$v_w \leftarrow \dim V_w$
Enumerate the monomials $m_1, \ldots, m_{v_w}$ of V_w
$S = (P_1, \ldots, P_{v_w}) \leftarrow \mathsf{SamplePoints}(N, v_w, p)$ ▷ $P_i = (I_2, I_4, I_6, I_{10}) \in \mathbb{F}_{p^2}^4$
$\mathbf{M} \leftarrow (m_j(P_i))_{i,j}$ ▷ accessed only as a black box, see Proposition 4.7
repeat ▷ fresh randomness in each iteration
    $\rho_1, \rho_2 \leftarrow$ ranks of $\mathbf{M}$ from two independent runs ▷ [KS91, Thm. 3]
    $\rho \leftarrow \max(\rho_1, \rho_2)$
    $r \leftarrow v_w - \rho$
    $c \leftarrow$ random element of $\ker \mathbf{M}$ ▷ [KS91, Thms. 2 and 4]
until $c \neq 0$ and $Mc = 0$
return $\left(r,\ f := \sum_j c_j m_j,\ S\right)$

bound for d'_N by Lemma 3.13. With high probability, the forms of weight d_N vanishing at the samples are exactly $G_N \cdot V_{d_N - d'_N}$, so their dimension is $\dim V_{d_N - d'_N}$. Since $\dim V_0 = \dim V_2 = 1$ and $\dim V_e < \dim V_{e+2}$ for even $e \geq 2$, this dimension determines $e = d_N - d'_N$, except that dimension 1 allows $e \in \{0, 2\}$. If the dimension is at least 2, we interpolate again at weight $d'_N = d_N - e$, where the kernel is $\langle G_N \rangle$. If the dimension is 1, the kernel is spanned by G_N or by $I_2 G_N$, and we need to distinguish the two cases. In particular, if $\bar{F}_N = c \cdot G_N$ for some non-zero constant c (see Remark 4.6), a single pass suffices.

The algorithm only ever works in characteristic p. Indeed, the samples lie on $\mathcal{H}_{N^2,p}$, so its output is a multiple of G_N, not necessarily of $\bar{F}_N = F_N \bmod p$. The characteristic-zero polynomial F_N enters only through its weight d_N via Lemma 3.13.

Proposition 4.5. *Let $\epsilon > 0$, and suppose that $N \leq p^{1/9-\epsilon}$. For p sufficiently large in terms of ϵ,* $\mathsf{HumbertSurfacePolynomial}(N, p)$ *returns $h_N = G_N$, except with probability $O\big(N^{36}(\log\log N)^2/p^4\big)$. In this case, $h_N \in \mathbb{F}_p[I_2, I_4, I_6, I_{10}]$ and $V(h_N) = \mathcal{H}_{N^2,p}$.*

Proof. Interpolate can fail in two ways: the samples may not cut out the right space of forms, or the Monte Carlo rank computation may return a wrong value. We bound the probability of each, and then show that if neither occurs, the algorithm returns G_N.

We discuss the first source of failure. For an even weight w with $d'_N \leq w \leq d_N$, let $W_w := G_N \cdot V_{w-d'_N}$ be the space of forms of weight w vanishing on $\mathcal{H}_{N^2,p}$, and let K_w be the kernel computed in $\mathsf{Interpolate}(N, w, p)$. We claim that

$$\Pr[K_w \neq W_w] = O\big(N^{36}(\log\log N)^2/p^4\big).$$

Every sample lies on $\mathcal{H}_{N^2,p}$ by Lemma 4.1, so $K_w \supseteq W_w$, and $\dim W_w \geq 1$. Let $K_{w,t}$ denote the kernel after the first t samples. Then $K_{w,t} \supseteq W_w$ for all t, and

Algorithm 2 HumbertSurfacePolynomial

Input: An integer $N \geq 2$ and a prime p.
Output: A form $h_N \in \mathbb{F}_p[I_2, I_4, I_6, I_{10}]$, or $\perp$.

$d_N \leftarrow \deg_w F_N$
$(r, f, S) \leftarrow \mathsf{Interpolate}(N, d_N, p)$
if $r = 1$ **then** ▷ $f = G_N$ or $f = I_2 G_N$
  if $I_2 \mid f$ **and** f/I_2 vanishes at every point of S **then**
    $f \leftarrow f/I_2$
else
  $e \leftarrow$ the even integer $e \geq 4$ with $\dim V_e = r$ ▷ $e = d_N - d'_N$
  if none exists, **return** $\perp$
  $(r', f, S') \leftarrow \mathsf{Interpolate}(N, d_N - e, p)$ ▷ fresh samples
  if $r' \neq 1$ **then return** $\perp$
$h_N \leftarrow$ scalar multiple of f whose first non-zero coefficient (in lexicographic order) is 1
return h_N

once $K_{w,t} = W_w$ it remains so. Let $v_w = \dim V_w$. If $K_w \neq W_w$, then fewer than $v_w - \dim W_w$ of the v_w samples decreased the dimension, so at least $\dim W_w + 1$ did not, and all of these occurred while $K_{w,t} \neq W_w$. At each such step, pick $f \in K_{w,t} \setminus W_w$. This is a form of weight $w \leq d_N$ that does not vanish identically on $\mathcal{H}_{N^2,p}$, so, conditionally on the past samples, it vanishes at the next sample with probability at most $\eta = O(N^9 \log\log N/p^2)$ by Lemma 4.3 with $q = p^2$. Moreover, multiplication by $I_2^{(d_N-w)/2}$ is an injective linear map $V_w \hookrightarrow V_{d_N}$, so $v_w \leq \dim V_{d_N} = O(N^9)$ by Eq. (5). Hence

$$\begin{aligned}\Pr[K_w \neq W_w] &\leq \binom{v_w}{\dim W_w + 1} \eta^{\dim W_w + 1} \leq (v_w \eta)^{\dim W_w + 1} \\ &\leq (v_w \eta)^2 = O\big(N^{36} (\log\log N)^2 / p^4\big),\end{aligned}$$

as $\dim W_w \geq 1$ and $v_w \eta = O\big(N^{18} \log\log N/p^2\big) = O\big(p^{-18\epsilon} \log\log p\big) < 1$ for p large enough.

We now look at the second source of failure. The rank computation is one-sided: each run returns a value at most $\operatorname{rank} \mathbf{M}$ [KS91, §2 and §4], so $\rho = \max(\rho_1, \rho_2)$ is correct as soon as one of the two runs is. Hence a single iteration of the loop returns a wrong rank with probability at most $\eta' = O(v_w^4/p^4)$. The returned r comes from the final iteration. Let T be the number of iterations; whether iteration k is reached depends only on the randomness of earlier iterations, so

$$\Pr[r \text{ is wrong}] \leq \sum_{k \geq 1} \Pr[T \geq k]\, \eta' = \mathbb{E}[T]\, \eta' \leq 2\eta' = O(N^{36}/p^4),$$

using $\mathbb{E}[T] \leq 2$ (see the proof of Proposition 4.7).

Suppose now that in each pass the rank is correct and $K_w = W_w$, at $w = d_N$ and, if a second pass runs, at $w = d'_N$. As the second pass uses

fresh samples, a union bound shows that this happens except with probability $O\big(N^{36}(\log\log N)^2/p^4\big)$.

Set $e := d_N - d'_N$, which is even and non-negative by Lemma 3.13. Since $K_{d_N} = W_{d_N}$, we have $r = \dim V_e$. Suppose first that $r = 1$, so that $e \in \{0, 2\}$ and $f = c\,G_N$ or $f = c\,I_2 G_N$ for some $c \neq 0$. If $f = c\,I_2 G_N$, then $f/I_2 = c\,G_N$ vanishes at every sample, since the samples lie on $\mathcal{H}_{N^2,p}$, so the algorithm replaces f by $c\,G_N$. If $f = c\,G_N$ and $I_2 \nmid G_N$, the test fails and f is unchanged. If $f = c\,G_N$ and $I_2 \mid G_N$, then $G_N = c' I_2$ with $c' \neq 0$, since G_N is irreducible. So f/I_2 is a nonzero constant, which vanishes at no sample, and f is again unchanged. In all cases, $f = c\,G_N$ after Line 5.

Suppose now that $r \geq 2$. Since $\dim V_2 = 1$ and $\dim V_e < \dim V_{e+2}$ for even $e \geq 2$, the integer $e = d_N - d'_N$ is the unique even $e \geq 4$ with $\dim V_e = r$, and the algorithm finds it. The second pass therefore runs at weight d'_N. Since $K_{d'_N} = W_{d'_N} = G_N \cdot V_0 = \langle G_N \rangle$, we get $r' = 1$ and $f = c\,G_N$ for some $c \neq 0$.

Finally, G_N can be chosen in $\mathbb{F}_p[I_2, I_4, I_6, I_{10}]$ by Remark 3.9. The scalar multiple of $f = c\,G_N$ whose first nonzero coefficient is 1 is therefore unique, and equals G_N normalised in this way. In particular, it has coefficients in $\mathbb{F}_p$, even though f was computed over $\mathbb{F}_{p^2}$. That $V(h_N) = \mathcal{H}_{N^2,p}$ is Lemma 3.8. □

Remark 4.6. We expect $\bar{F}_N = c \cdot G_N$ for some $c \in \mathbb{F}_p^\times$, in which case $d'_N = d_N$ and Algorithm 2 terminates after a single pass. Since $G_N \mid \bar{F}_N$ by Lemma 3.13, this equality holds if and only if $\bar{F}_N$ is irreducible over $\overline{\mathbb{F}}_p$. It has been verified experimentally for all $N \leq 5$ (see, e.g., the software attached to [CCF24]), and for $N = 6$ at the prime used in [Sha26c, §4.4]. In [Sha26c, Thm. 2], primes at which $\bar{F}_N$ is reducible are excluded by assumption, as part of the set B_N. We do not need this equality: our algorithm computes G_N directly by interpolation, so if $\bar{F}_N$ is reducible, this only costs a second pass at weight $d'_N < d_N$, and does not affect the correctness of the membership test.

Proposition 4.7. *Let $\epsilon > 0$, and $N \leq p^{1/9-\epsilon}$. For p sufficiently large in terms of ϵ,* HumbertSurfacePolynomial(N, p) *runs in expected time $N^{18+9\delta}(\log p)^{O(1)}$ and uses memory $N^{9+9\delta}(\log p)^{O(1)}$.*

Proof. The algorithm makes at most two calls to Interpolate, at weights $w \leq d_N$. It therefore suffices to bound the cost of one call, with $v_w := \dim V_w \leq \dim V_{d_N} = O(N^9)$ by Eq. (5).

By Lemma 4.1, sampling costs $v_w N^{7+o(1)}(\log p)^{O(1)} = N^{16+o(1)}(\log p)^{O(1)}$. For each sample P_i we also compute u_i with $u_i^5 = I_{10}(P_i)^{-1}$ and the normalised invariants $(i_1^{(i)}, i_2^{(i)}, i_3^{(i)})$ of Eq. (2), at total cost $v_w(\log p)^{O(1)}$.

We access the matrix $\mathbf{M} = (m_j(P_i))_{i,j}$ only through products $v \mapsto \mathbf{M}v$, and compute these as follows. The normalised invariants of P_i are the rescaling of P_i by some t_i with $t_i^2 = u_i$. Each m_j is a monomial of even weight w, so $m_j(P_i) = u_i^{-w/2}\, m_j(i_1^{(i)}, i_2^{(i)}, i_3^{(i)}, 1)$. Hence, writing $f_v := \sum_j v_j m_j$,

$$(\mathbf{M}v)_i = u_i^{-w/2}\, f_v\big(i_1^{(i)}, i_2^{(i)}, i_3^{(i)}, 1\big).$$

The polynomial $f_v(x_1, x_2, x_3, 1)$ has individual degrees at most $w/2$, $w/4$ and $w/6$, so its box size is $O(w^3) = O(N^9)$, as $w \le d_N = O(N^3)$. Thus, one product with $\mathbf{M}$ is one multipoint evaluation at $v_w = O(N^9)$ points of $\mathbb{F}_{p^{10}}^3$, followed by v_w multiplications, at cost $N^{9+9\delta}(\log p)^{O(1)}$ by Theorem 2.9. Note that $\mathbf{M}$ itself has entries in $\mathbb{F}_{p^2}$; only the evaluation points lie in $\mathbb{F}_{p^{10}}$.

We take all random field elements in [KS91] from $\mathbb{F}_{p^2}$. Each iteration of the loop in Lines 5–10 exits the loop except with probability $O(v_w^2/p^2) = O(p^{-18\epsilon})$,[6] which is at most $1/2$ for p large enough. The expected number of iterations is at most 2, and the total cost of one call is

$$O(v_w) \cdot N^{9+9\delta}(\log p)^{O(1)} + \tilde{O}(v_w^2)(\log p)^{O(1)} = N^{18+9\delta}(\log p)^{O(1)},$$

which dominates the cost of sampling. The memory is dominated by multipoint evaluation, which uses $v_w^{1+\delta}(\log p)^{O(1)} = N^{9+9\delta}(\log p)^{O(1)}$ memory, as the black-box linear algebra stores $O(1)$ vectors of length v_w at a time [KS91, §2]. □

4.1 The precomputation algorithm

On input a prime p and an integer $2 \le D \le p^{1/16}$, the algorithm Precomputation runs HumbertSurfacePolynomial(N, p) for each $2 \le N \le D$, rerunning it whenever it returns $\perp$, to obtain forms $h_N \in \mathbb{F}_p[I_2, I_4, I_6, I_{10}]$. It outputs `precomp` $= \big((h_N)_N, h\big)$, where

$$h := \prod_{2 \le N \le D} h_N(x_1, x_2, x_3, 1) \in \mathbb{F}_p[x_1, x_2, x_3].$$

If every call succeeds, then $h_N = G_N$ for all N, so that $h = g := \prod_N g_N$.

Proposition 4.8. *Let $\epsilon > 0$, and let $2 \le D \le p^{1/9-\epsilon}$. For p sufficiently large in terms of ϵ,* Precomputation *runs in expected time $D^{19+9\delta}(\log p)^{O(1)}$ and uses memory $D^{12}(\log p)^{O(1)}$, and the polynomial h it outputs has box size $O(D^{12})$.*

Moreover, except with probability $O\big(D^{37}(\log\log D)^2/p^4\big)$, we have $h_N = G_N$ for all $2 \le N \le D$. In this case, for every genus 2 curve C defined over $\overline{\mathbb{F}}_p$ with normalised invariants (i_1, i_2, i_3), we have $h(i_1, i_2, i_3) = 0$ if and only if $\mathrm{Jac}(C)$ is optimally (N, N)-split for some $2 \le N \le D$.

Remark 4.9. The failure probability is less than 1 only for $D < p^{4/37}$. Under our standing assumption $D \le p^{1/16}$, it is $O\big(p^{-27/16}(\log\log p)^2\big)$.

Proof. Every $2 \le N \le D$ satisfies $N \le p^{1/9-\epsilon}$, so Propositions 4.5 and 4.7 apply to each call HumbertSurfacePolynomial(N, p), with implied constants independent of N.

[6] An iteration exits the loop if $\rho = \operatorname{rank} \mathbf{M}$, the leading $\rho \times \rho$ block of the preconditioned matrix $\mathbf{UML}$ is nonsingular, and the resulting c is nonzero. The first two conditions fail with probability $O(v_w^2/p^2)$ by [KS91, Thms. 2 and 3], and under them [KS91, Thm. 4] returns a uniformly random element of $\ker \mathbf{M}$, which is nonzero except with probability at most p^{-2} since $\ker \mathbf{M} \supseteq W_w \ne 0$.

We first show correctness. By the proof of Proposition 4.5, a call returns G_N except with probability $\gamma_N = O\big(N^{36}(\log\log N)^2/p^4\big)$. For p large enough we have $\gamma_N \le 1/2$, so the expected number of calls for each N is at most 2, and the final output h_N differs from G_N with probability at most $\sum_{k\ge1}\gamma_N^k \le 2\gamma_N$. By a union bound, $h_N \neq G_N$ for some N with probability

$$\sum_{2\le N\le D} O\big(N^{36}(\log\log N)^2/p^4\big) = O\big(D^{37}(\log\log D)^2/p^4\big).$$

If $h_N = G_N$ for all N, then $h = \prod_N g_N$, which vanishes at (i_1, i_2, i_3) if and only if some g_N does. The last claim therefore follows from Corollary 3.11.

We now compute the cost of the algorithm. Since the expected number of calls for each N is at most 2, Proposition 4.7 gives an expected time of

$$\sum_{N\le D} N^{18+9\delta}(\log p)^{O(1)} = D^{19+9\delta}(\log p)^{O(1)}$$

for all calls. Each $h_N(x_1, x_2, x_3, 1)$ has individual degrees at most $d_N/2$, so the individual degrees of h are at most $\sum_N d_N/2 = O(D^4)$, and its box size is $O(D^{12})$. Computing h with a product tree and Kronecker substitution costs $D^{12}(\log D)^{O(1)}(\log p)^{O(1)}$, which is dominated by the cost above.

In terms of memory, storing h uses $D^{12}(\log p)^{O(1)}$ memory, and storing the h_N uses

$$\sum_N O(N^9)(\log p)^{O(1)} = O(D^{10})(\log p)^{O(1)}.$$

Each call uses memory $N^{9+9\delta}(\log p)^{O(1)}$ by Proposition 4.7, and memory is reused across calls. Assuming $\delta \le 1/3$, this is at most $D^{12}(\log p)^{O(1)}$. □

Remark 4.10. One can further test each h_N at a few fresh outputs of SamplePoints. If h_N does not vanish identically on $\mathcal{H}_{N^2,p}$, each test detects this except with probability $O(N^9 \log\log N/p^2)$ by Lemma 4.3. This reduces the failure probability of the precomputation, but it cannot detect outputs that are proper multiples of G_N.

5 The online phase: solving each instance

In this section, we fix a superspecial Jacobian J of a genus-2 curve defined over $\mathbb{F}_{p^2}$, together with the output $\texttt{precomp} = ((h_N)_N, h)$ of Precomputation from Section 4.1. Our goal is to find a polarised isogeny $\varphi\colon J \to E_1 \times E_2$, where E_1, E_2 are supersingular elliptic curves. Our algorithm has two main steps.

Step 1: walking to random surfaces. We take B independent random walks in $\mathcal{G}(p;2)$ starting at J, each of length n as in Corollary 3.24 with $X_0 = J$. Each step of the walk chooses one of the 15 possible outgoing 2-isogenies uniformly at random, so that the length n ensures each endpoint lands in $S = \mathcal{H}^{\mathrm{ss}} \cup H_1^{ss}$ with

probability at least $D^4/(8p)$. We obtain polarised 2^n-isogenies $\varphi_j \colon J \to X_j$ for $j = 1, \ldots, B$, where each $X_j \in \mathrm{SS}_2(p)$. Rather than storing each φ_j, it suffices to store the seed $\mathbf{s}_j$ of the j-th walk, which costs $O(\log p)$ memory, and to recompute φ_j when needed.

Step 2: batched membership test. If some $X_j \cong E_1 \times E_2$, we terminate, outputting φ_j, E_1 and E_2. Otherwise, each X_j is a Jacobian, and we compute its normalised Igusa–Clebsch invariants $(i_1^{(j)}, i_2^{(j)}, i_3^{(j)})$ as in (2), and evaluate h at all B points at once using multipoint evaluation. By Proposition 4.8, if the precomputation succeeded, $h(i_1^{(j)}, i_2^{(j)}, i_3^{(j)}) = 0$ if and only if X_j is optimally (N, N)-split for some $2 \leq N \leq D$.

If h vanishes at none of these points, we repeat the process with a new batch of random walks from J. Otherwise, for each j with $h(i_1^{(j)}, i_2^{(j)}, i_3^{(j)}) = 0$, we try to recover the corresponding (N, N)-splitting $\varphi' \colon X_j \to E_1 \times E_2$, and if we are successful, we output $\varphi = \varphi' \circ \varphi_j \colon J \to E_1 \times E_2$.

The algorithm. Putting this together, we obtain the algorithm FindSplitting, given in Algorithm 4. In FindSplitting, once a zero of h is found, the algorithm uses the subroutine ComputeSplitting (Algorithm 3) to extract the corresponding (N, N)-splitting. The purpose of this subroutine is twofold: to certify that a splitting exists and, if so, to output the corresponding splitting.

Algorithm 3 ComputeSplitting

Input: A superspecial Jacobian J of a genus-2 curve defined over $\mathbb{F}_{p^2}$, and `precomp` from Section 4.1.
Output: Supersingular elliptic curves E_1, E_2 and an (N, N)-isogeny $\varphi \colon J \to E_1 \times E_2$ for some $2 \leq N \leq D$, or $\perp$.

Extract $(h_N)_{2 \leq N \leq D}$ from `precomp`
Compute the normalised invariants (i_1, i_2, i_3) of J ▷ Eq. (2)
for $N = 2, \ldots, D$ **do**
    if $h_N(i_1, i_2, i_3, 1) = 0$ **then**
        for each (N, N)-subgroup $K \leq J[N]$ **do**
            Compute the (N, N)-isogeny $\varphi \colon J \to J/K$ ▷ Step 4, SamplePoints
            if $J/K \cong E_1 \times E_2$ **then**
                return E_1, E_2, φ
return $\perp$

Proposition 5.1. *Let J be a superspecial Jacobian of a genus-2 curve defined over $\mathbb{F}_{p^2}$, and let* `precomp` *be the output of* Precomputation. *Every output of* FindSplitting$(J,$ `precomp`$)$ *consists of supersingular elliptic curves E_1, E_2 and a polarised isogeny $\varphi \colon J \to E_1 \times E_2$. Furthermore,* FindSplitting *terminates with probability* 1*, for any output* `precomp` *of* Precomputation.

Algorithm 4 FindSplitting

Input: A superspecial Jacobian J of a genus-2 curve defined over $\mathbb{F}_{p^2}$, and the precomputed data `precomp` from Section 4.

Output: E_1, E_2 supersingular elliptic curves, and a polarised isogeny $\varphi : J \to E_1 \times E_2$.

1: Extract h from `precomp`
2: $B \leftarrow \max(\text{box size of } h,\ D^{12})$
3: Set n to be as in Corollary 3.24.
4: **while** `true` **do**
5: $\quad\mathcal{P} \leftarrow \emptyset$ ▷ points to be evaluated in this batch
6: $\quad$**for** $j = 1, \dots, B$ **do** ▷ B random walks of length n
7: $\quad\quad$Draw a fresh seed $\mathtt{s}_j$ ▷ e.g., StartingSeed in [CCF24]
8: $\quad\quad$Compute a 2^n-isogeny $\varphi_j : J \to X_j$ from $\mathtt{s}_j$
9: $\quad\quad$**if** $X_j \cong E_1 \times E_2$ **then** ▷ X_j is split
10: $\quad\quad\quad$**return** E_1, E_2, φ_j
11: $\quad\quad$Compute invariants $(i_1^{(j)}, i_2^{(j)}, i_3^{(j)})$ of X_j ▷ X_j a Jacobian, use [CCF24, Remark 5]
12: $\quad\quad$Add $\big(\mathtt{s}_j, (i_1^{(j)}, i_2^{(j)}, i_3^{(j)})\big)$ to $\mathcal{P}$
13: $\quad$Evaluate h at all points of $\mathcal{P}$ ▷ multipoint evaluation, Theorem 2.9
14: $\quad$**for** each $\big(\mathtt{s}_j, (i_1^{(j)}, i_2^{(j)}, i_3^{(j)})\big) \in \mathcal{P}$ with $h(i_1^{(j)}, i_2^{(j)}, i_3^{(j)}) = 0$ **do**
15: $\quad\quad$Recompute the walk $\varphi_j : J \to X_j$ from $\mathtt{s}_j$
16: $\quad\quad$output $\leftarrow$ ComputeSplitting$(X_j, \mathtt{precomp})$
17: $\quad\quad$**if** output $\neq \perp$ **then** ▷ certifies the splitting
18: $\quad\quad\quad E_1, E_2, \varphi_{\text{split}} \leftarrow$ output
19: $\quad\quad\quad$**return** $E_1, E_2, \varphi_{\text{split}} \circ \varphi_j$

Proof. Every output of ComputeSplitting is an (N, N)-isogeny to a product of elliptic curves, as this is checked in Line 7 of Algorithm 3, and the elliptic curves are supersingular since J is superspecial. Hence every output of FindSplitting is the composite of a polarised 2^n-isogeny with either the identity or an (N, N)-isogeny, and is therefore a polarised isogeny to a product of supersingular elliptic curves. This holds whatever the output of Precomputation: false positives of the membership test only cause ComputeSplitting to return $\perp$.

For termination, by Corollary 3.24 each walk ends on a product of elliptic curves with probability at least $\frac{1}{2}\mu(\mathcal{H}_1^{\text{ss}})/\mu(V) > 0$, in which case the algorithm returns in Line 10. Each iteration of the while loop runs in finite time and uses fresh seeds, so the algorithm terminates with probability 1. □

Next, we analyse the cost of this algorithm, assuming that Precomputation succeeded. This assumption is needed only here: it ensures that h vanishes at every Jacobian in $\mathcal{H}^{\text{ss}}$, so that each walk succeeds whenever its endpoint lies in $\mathcal{H}_1^{\text{ss}} \cup \mathcal{H}^{\text{ss}}$, which happens with probability at least $D^4/(8p)$ by Corollary 3.24. This bounds the expected number of iterations of the while loop.

Proposition 5.2. *Assume* Precomputation *correctly outputs* `precomp`. *Then, Algorithm 4 runs in expected time* $\big(p/D^4 + D^{12}\big)D^{12\delta}(\log p)^{O(1)}$ *with memory* $D^{12+12\delta}(\log p)^{O(1)}$.

Proof. Let $B := \max(b_h, D^{12})$, where b_h denotes the box size of h. Since $b_h = O(D^{12})$ by Proposition 4.8, we have $B = \Theta(D^{12})$.

We first determine the cost of each iteration of the **while** loop (Lines 4–19). Using, for example, [CCF24, Alg. 4] to compute the 2^n-isogenies from the seeds, each walk costs $(\log p)^{O(1)}$, so Lines 6–12 cost $B(\log p)^{O(1)}$. By Theorem 2.9, Line 13 costs $(2B)^{1+\delta}(\log p)^{O(1)} = D^{12+12\delta}(\log p)^{O(1)}$.

We now determine the cost of ComputeSplitting. Writing b_N for the box size of h_N, evaluating the h_N at one point costs $\sum_{N \le D} b_N = O(D^{10})$ operations in $\mathbb{F}_{p^{10}}$.

We first determine the cost of computing (N, N)-subgroups of $J[N]$. We note that for a given N, we have $J[N] \subseteq J(\mathbb{F}_{p^{2r}})$ with $r < N$.[7] Furthermore, $J(\mathbb{F}_{p^{2r}}) = J[M] \cong (\mathbb{Z}/M\mathbb{Z})^4$, where $M = |(\pm p)^r - 1|$ and $N \mid M$. Hence, we can find generators for $J[N]$ by sampling uniformly random points $P \in J(\mathbb{F}_{p^{2r}})$ (using, e.g., [Bru13, Alg. 3.7][8]), and computing $[M/N]P$.

As $\mathrm{GL}_4(\mathbb{Z}/N\mathbb{Z})$ acts simply transitively on the ordered bases of $J[N]$, four points $Q_1, \dots, Q_4$ sampled in this way form a basis of $J[N]$ with probability

$$\frac{\#\mathrm{GL}_4(\mathbb{Z}/N\mathbb{Z})}{N^{16}} = \prod_{\ell \mid N} \prod_{i=1}^{4} (1 - \ell^{-i}) \ge \frac{\varphi(N)}{N\,\zeta(2)^3} = \Omega(1/\log\log N).$$

We test this using the N-Weil pairing e_N: fixing a primitive N-th root of unity ζ, let $\mathbf{W}$ be a 4×4 matrix over $\mathbb{Z}/N\mathbb{Z}$ such that $\mathbf{W}_{ij} = \log_\zeta(e_N(Q_i, Q_j))$. Then, by non-degeneracy of e_N, $\mathbf{W}$ is invertible modulo N if and only if the Q_i form a basis.

Sampling a point and multiplying it by M/N costs $O(r \log p)$ operations in $\mathbb{F}_{p^{2r}}$, and an expected $O(\log\log N)$ samples suffice. The pairings and discrete logarithms in μ_N cost $O(N)$ further operations. As each operation in $\mathbb{F}_{p^{2r}}$ costs $N^{1+o(1)}(\log p)^{1+o(1)}$, we obtain a basis of $J[N]$ in expected time $N^{2+o(1)}(\log p)^{O(1)}$.

There are $O(N^3 \log\log N)$ maximal isotropic subgroups $K \simeq (\mathbb{Z}/N\mathbb{Z})^2$ of $J[N]$. For each such K, we compute $J \to J/K$ as a chain of (ℓ, ℓ)-isogenies as in Step 4 of SamplePoints. Testing whether J/K is a product costs $O(1)$ operations. Hence a single N costs $N^{6+o(1)}(\log p)^{O(1)}$, and even trying every $N \le D$ costs $D^{7+o(1)}(\log p)^{O(1)}$. In total, ComputeSplitting costs $O(D^{10})(\log p)^{O(1)}$, which is dominated by Line 13 of FindSplitting.

We now bound the expected number of iterations. The B walks are independent. By Corollary 3.24, which applies since $200 \cdot D^7 \le p$ (this follows from $D \le p^{1/16}$ as soon as $p \ge 200^{16/9}$), each X_j lies in $\mathcal{H}^{\mathrm{ss}} \cup \mathcal{H}_1^{\mathrm{ss}}$ with probability $\gamma \ge D^4/(8p)$. If X_j is a product, the algorithm returns in Line 10. Otherwise,

[7] All isogenies taken are $\mathbb{F}_{p^2}$-rational, so the endpoint of each walk has an $\mathbb{F}_{p^2}$-model on which the p^2-Frobenius acts as $[\pm p]$. In particular, it acts on $J[N]$ as the scalar $\pm p$, so every subgroup of $J[N]$ is stable under it and thus defined over $\mathbb{F}_{p^2}$. Moreover, $\pm p$ has order at most $\varphi(N) < N$ in $(\mathbb{Z}/N\mathbb{Z})^\times$, so $J[N] \subseteq J(\mathbb{F}_{p^{2r}})$ for some $r < N$.

[8] Bruin's algorithm needs the zeta function, but for us this is immediate as Frobenius acts as $[\pm p]$.

$X_j \in \mathcal{H}_{N^2}$ for some $2 \le N \le D$, so h vanishes at its normalised invariants by Proposition 4.8, and ComputeSplitting returns a splitting. Hence, each iteration returns a splitting with probability at least

$$1 - (1-\gamma)^B \ge 1 - e^{-B\gamma} \ge \tfrac{1}{2}\min(1, B\gamma) = \Omega\big(\min(1, D^{16}/p)\big),$$

and so the expected number of iterations is $O(1 + p/D^{16})$. Multiplying by the cost of one iteration, the expected running time is

$$O\big(1 + p/D^{16}\big) \cdot D^{12+12\delta}(\log p)^{O(1)} = \big(D^{12} + p/D^4\big) D^{12\delta}(\log p)^{O(1)}.$$

In terms of memory cost: the algorithm stores h, whose box size is $O(D^{12})$, the polynomials h_N, whose box sizes sum to $O(D^{10})$, and one batch of points. Within one iteration, the algorithm stores the seed and normalised invariants of each of the B walks, requiring $B(\log p)^{O(1)} = D^{12}(\log p)^{O(1)}$ memory as the invariants lie in $\mathbb{F}_{p^{10}}$. The multipoint evaluation in Line 13 uses at most its running time in memory, namely $D^{12+12\delta}(\log p)^{O(1)}$. A walk is recomputed from its seed in memory $(\log p)^{O(1)}$, and ComputeSplitting uses $O(D^3 \log p)$, as it stores a basis of $J[N]$ over $\mathbb{F}_{p^{2r}}$, enumerates the subgroups one at a time, and computes one (N, N)-isogeny at a time. Memory is reused across iterations. In total, the memory is $D^{12+12\delta}(\log p)^{O(1)}$. □

Remark 5.3. In practice, we would detect if each intermediate step of our walks in $\mathcal{G}(p; 2)$ lands on a product, and terminate if so. For simplicity, we do not include this in the analysis; it would only decrease the expected time.

6 An algorithm for the isogeny problem

We now solve the superspecial isogeny problem (Problem 2.8) following the algorithm of [Cor+26, Thm. 6.11], a rigorous version of the Costello–Smith algorithm [CS20]. The only change is the first step, namely finding splittings: we replace [Cor+26, Thm. 6.9] by FindSplitting (Algorithm 4). Given superspecial p.p. abelian surfaces X_1, X_2 defined over $\mathbb{F}_{p^2}$, our algorithm runs as follows. Fix an integer $2 \le D \le p^{1/16}$.

1. Run Precomputation(p, D) to obtain $\texttt{precomp} = ((h_N)_N, h)$.
2. For $i = 1, 2$: if X_i is a product of elliptic curves, let φ_i be the identity. Otherwise, call FindSplitting$(X_i, \texttt{precomp})$ to obtain a polarised isogeny $\varphi_i \colon X_i \to E_i \times E_i'$.
3. Run [PW23, Alg. 7] to find isogenies $\psi \colon E_1 \to E_2$ and $\psi' \colon E_1' \to E_2'$ whose degrees are powers of 2 with even exponent, and form the product isogeny $\pi \colon E_1 \times E_1' \to E_2 \times E_2'$.
4. Return $\varphi_2^\dagger \circ \pi \circ \varphi_1 \colon X_1 \to X_2$.

With this algorithm, and the results of the previous sections, we can immediately prove Theorem 1.1.

Proof (of Theorem 1.1). Given $\epsilon > 0$, we will apply Propositions 4.8, 5.1 and 5.2 with $\delta := \epsilon$. Since $D \leq p^{1/16}$, the factors $D^{O(\delta)}$ and $(\log p)^{O(1)}$ are then at most p^{ϵ} for p large enough, after rescaling ϵ. The first part of the theorem then follows directly from applying Proposition 4.8 to Step 1. Since $D \leq p^{1/16}$, its failure probability is $O(p^{-27/16}(\log\log p)^2) \leq 1/p$ for p large enough (see Remark 4.9).

Step 2. For $i = 1, 2$, if X_i is a product of elliptic curves we take φ_i to be the identity; otherwise, by Proposition 5.1, $\mathsf{FindSplitting}(X_i, \texttt{precomp})$ terminates with probability 1 and returns supersingular elliptic curves E_i, E_i' and a polarised isogeny $\varphi_i \colon X_i \to E_i \times E_i'$. If the precomputation succeeded, by Proposition 5.2, each call runs in expected time $(p/D^4 + D^{12})\, p^{\epsilon}$ and uses memory $D^{12}p^{\epsilon}$.

Step 3. Correctness of Step 3 follows from [Cor+26, Thm. 6.11]: the algorithm of [PW23, Alg. 7] returns isogenies $\psi \colon E_1 \to E_2$ and $\psi' \colon E_1' \to E_2'$ whose degrees are 2^{2a} and 2^{2b} for some integers $a, b \geq 0$. Composing the one of smaller degree with $[2^{|a-b|}]$, we may assume that $\deg\psi = \deg\psi'$, so that $\pi = \psi \times \psi'$ is a polarised isogeny $E_1 \times E_1' \to E_2 \times E_2'$. By [PW24, Prop. 8.7], this step terminates in expected time $p^{1/2}(\log p)^{O(1)}$. Since $D \leq p^{1/16}$, the cost of Step 3 is absorbed into that of Step 2. Each run stores a hash table of $p^{1/2}$ walks (or their seeds), so it uses memory $p^{1/2}(\log p)^{O(1)}$.

Step 4. The output $\varphi_2^{\dagger} \circ \pi \circ \varphi_1 \colon X_1 \to X_2$ is a composite of polarised isogenies, hence is itself a polarised isogeny. Computing this composition has cost negligible compared to Step 2.

Combining Steps 2–4 gives us the second part of the theorem. □

Consequences. We now discuss immediate consequences of our main theorem.

Solving many instances. Using Theorem 1.1 to solve T instances of the supersingular isogeny problem over a field of fixed characteristic p gives us the following complexities. If $T \leq p^{7/16}$, taking $D = (Tp)^{1/23}$ provides an algorithm that runs in time $(Tp)^{19/23+o(1)}$, or $p^{19/23+o(1)}T^{-4/23}$ per instance. For $T \geq p^{7/16}$ the per-instance cost is $p^{3/4+o(1)}$.

Solving one instance. For a single instance, our algorithm takes time $p^{19/23+o(1)}$ and memory $p^{12/23+o(1)}$, obtained by taking $D = p^{1/23}$ in Theorem 1.1 to balance the precomputation (which costs $D^{19+o(1)}$) against the online phase (which costs p/D^4). The bottleneck is therefore the precomputation: interpolating a single h_N costs $N^{18+o(1)}$. Any improvement in this step would bring the single-instance cost closer to $p^{3/4+o(1)}$. For example, computing all h_N with $N \leq D$ in time $D^{12+o(1)}$ would give a single-instance cost of $p^{3/4+o(1)}$ in both time and memory, matching the cost per instance after precomputation.

The exponent 3/4. The online cost $(p/D^4 + D^{12})p^{o(1)}$ of Theorem 1.1 reflects a simple relation between density and degree. By Remark 3.19 and Section 3.5, the density of optimally (N, N)-split surfaces, about $5\#\mathrm{SL}_2(\mathbb{Z}/N\mathbb{Z})/p$, and the weight $d_N = 5(1 - 1/N)\#\mathrm{SL}_2(\mathbb{Z}/N\mathbb{Z})$ of the equation of $\mathcal{H}_{N^2}$ agree up to the factor p for $N \geq 3$. Hence, the total weight $W := \sum_{2 \leq N \leq D} d_N \approx D^4$ of h

roughly governs both quantities that determine the online cost: h vanishes on a proportion of about W/p of superspecial surfaces (Proposition 3.23), and has box size about W^3 (Proposition 4.8). Balancing p/W against W^3 gives $W = p^{1/4}$, i.e., $D = p^{1/16}$, and the online cost $p^{3/4+o(1)}$. Since this argument only uses the relation between weight and density, it suggests that the exponent $3/4$ is intrinsic to testing membership with a single polynomial in three variables using fast multipoint evaluation, and does not depend on which values of N are used to build h. For instance, a single value of N gives the same online exponent; the choice of N only affects the cost of the precomputation, as the following remark shows.

Remark 6.1. Restricting the membership test to a single N gives the same online exponent. The arguments of Section 3.6 applied to the single value N show that a proportion $\Omega(N^3/p)$ of superspecial surfaces are optimally (N, N)-split, for $N \leq c\,p^{1/6}$ with c an absolute constant, and h_N has box size $O(N^9)$. With batches of size N^9, the online phase therefore costs $(p/N^3 + N^9)\,p^{o(1)}$. Taking $N = p^{1/12}$ gives time and memory $p^{3/4+o(1)}$, as in the first consequence of Theorem 1.1.

The two differ only in the precomputation. To reach density of order D^4/p, a single N must be of size about $D^{4/3}$, and interpolating its equation costs $D^{24+o(1)}$. In contrast, the union reaches the same density using the equations for all $N \leq D$, at total cost $D^{19+o(1)}$. As a result, the single-instance cost would rise from $p^{19/23+o(1)} \approx p^{0.826}$ to $p^{6/7+o(1)} \approx p^{0.857}$ (taking $N = p^{1/21}$).

Computing endomorphism rings in dimension 2. The splitting algorithm of this paper can also be used to speed up the computation of endomorphism rings. The algorithm of [Cor+26, Thm. 6.15] computes $\mathrm{End}(X)$ for a superspecial p.p. abelian surface X from two splittings $\varphi\colon X \to E_1 \times E_2$ and $\varphi'\colon X \to E_1' \times E_2'$ of coprime reduced degrees. Conjugating the endomorphism rings $\mathrm{End}(E_1 \times E_2)$ and $\mathrm{End}(E_1' \times E_2')$ by φ and φ' gives two subrings of $\mathrm{End}(X)$ of coprime index, which together generate $\mathrm{End}(X)$. The cost of this algorithm is dominated by finding the two splittings [Cor+26, Thm. 6.9].

If we compute φ and φ' instead using two calls to FindSplitting, we expect a similar asymptotic improvement. Some care must be taken, as we must ensure that their degrees remain coprime. At a high level, this can be done by taking the second batch of walks in $\mathcal{G}(p; 3)$, and restricting the membership test to integers N coprime to 3 in the first call and coprime to $\deg(\varphi)$ in the second. We expect this to yield an algorithm computing $\mathrm{End}(X)$ with the same complexity as Theorem 1.1. We leave a formal analysis of this for future work.

Time-memory trade-off. Unlike the Costello–Smith algorithm, which uses essentially no memory in their algorithm to find a splitting, our algorithm stores the polynomials h, and $\{h_N\}$, and one batch of points. This allows a smooth trade-off between time and memory for the splitting algorithm FindSplitting.

Corollary 6.2. *Fix $\epsilon > 0$. For every M with $p^{1/2} \leq M \leq p^{3/4}$, the algorithm of Theorem 1.1 with $D = \lfloor M^{1/12} \rfloor$ uses memory Mp^{ϵ} and, after a precomputation in time $M^{19/12}p^{\epsilon}$, solves each instance in expected time $p^{1+\epsilon}/M^{1/3}$.*

Proof. Since $M \geq p^{1/2}$, we have $D \geq 2$ for p large enough, and $M^{1/12}/2 \leq D \leq M^{1/12}$. By Theorem 1.1, the online phase takes expected time

$$(p/D^4 + D^{12})\, p^\epsilon \leq (16\, p/M^{1/3} + M)\, p^\epsilon \leq 17\, p^{1+\epsilon}/M^{1/3},$$

using $M \leq p/M^{1/3}$, which holds as $M \leq p^{3/4}$. Its memory is $(D^{12} + p^{1/2})\, p^\epsilon \leq 2Mp^\epsilon$, as $M \geq p^{1/2}$, and the precomputation takes time $D^{19}p^\epsilon \leq M^{19/12}p^\epsilon$. Rescaling ϵ absorbs the constants. □

Remark 6.3. The elliptic step of [Cor+26, Thm. 6.11] uses memory $p^{1/2+o(1)}$, which is why Corollary 6.2 is stated for $M \geq p^{1/2}$. We chose this variant to allow for simple reconstruction of the product isogeny. In future work, it would be interesting to find a memory-efficient variant of the dimension-1 part of the algorithm, e.g., by replacing the table-based collision search of [PW23, Alg. 7] by Pollard-rho-style cycle finding. In particular, if the dimension-1 step is replaced by one using time $p^{1/2}(\log p)^{O(1)}$ and memory $(\log p)^{O(1)}$, Corollary 6.2 holds for all $(\log p)^{O(1)} \leq M \leq p^{3/4}$.[9]

References

[ATY24] Yusuke Aikawa, Ryokichi Tanaka, and Takuya Yamauchi. "Isogeny graphs on superspecial abelian varieties: eigenvalues and connection to Bruhat–Tits buildings". In: *Canadian Journal of Mathematics* 76.6 (2024), pp. 1891–1916. DOI: 10.4153/S0008414X23000676 (cit. on p. 2).

[Bha+24] Vishwas Bhargava, Sumanta Ghosh, Zeyu Guo, Mrinal Kumar, and Chris Umans. "Fast multivariate multipoint evaluation over all finite fields". In: *Journal of the ACM* 71.3 (2024), pp. 1–32 (cit. on p. 12).

[BLR12] Siegfried Bosch, Werner Lütkebohmert, and Michel Raynaud. *Néron models*. Springer Science & Business Media, 2012 (cit. on p. 19).

[Brö+15] R. Bröker, E. W. Howe, K. E. Lauter, and P. Stevenhagen. "Genus-2 curves and Jacobians with a given number of points". In: *LMS J. Comput. Math.* 18.1 (2015), pp. 170–197 (cit. on p. 17).

[BLS12] Reinier Bröker, Kristin Lauter, and Andrew Sutherland. "Modular polynomials via isogeny volcanoes". In: *Mathematics of Computation* 81.278 (2012), pp. 1201–1231 (cit. on p. 5).

[BD11] N. Bruin and K. Doerksen. "The arithmetic of genus two curves with $(4, 4)$-split Jacobians". In: *Canad. J. Math.* 63.5 (2011), pp. 992–1024 (cit. on pp. 17, 27).

[Bru13] Peter Bruin. "Computing in Picard Groups of Projective Curves over Finite Fields". In: *Mathematics of Computation* 82.283 (2013), pp. 1711–1756. ISSN: 00255718, 10886842 (cit. on p. 38).

[9] For $M < 2^{12}$, where $D \leq 1$, we simply take $D = 1$, i.e., the original Costello–Smith algorithm.

[CDS20] W. Castryck, T. Decru, and B. Smith. "Hash functions from superspecial genus-2 curves using Richelot isogenies". In: *Journal of Math. Crypt.* 14.1 (2020), pp. 268–292 (cit. on p. 11).

[Cas+25] Wouter Castryck, Thomas Decru, Péter Kutas, Abel Laval, Christophe Petit, and Yan Bo Ti. "KLPT^2: Algebraic Pathfinding in Dimension Two and Applications". In: *CRYPTO 2025, Part I*. Ed. by Yael Tauman Kalai and Seny F. Kamara. Vol. 16000. LNCS. Springer, Cham, Aug. 2025, pp. 167–200. DOI: 10.1007/978-3-032-01855-7_6 (cit. on p. 5).

[CH13] Wouter Castryck and Hendrik Hubrechts. "The distribution of the number of points modulo an integer on elliptic curves over finite fields". In: *The Ramanujan Journal* 30.2 (2013), pp. 223–242 (cit. on pp. 3, 6, 27).

[CK18] Henry Corrigan-Gibbs and Dmitry Kogan. "The Discrete-Logarithm Problem with Preprocessing". In: *EUROCRYPT 2018, Part II*. Ed. by Jesper Buus Nielsen and Vincent Rijmen. Vol. 10821. LNCS. Springer, Cham, Apr. 2018, pp. 415–447. DOI: 10.1007/978-3-319-78375-8_14 (cit. on p. 5).

[CCF24] Maria Corte-Real Santos, Craig Costello, and Sam Frengley. "An Algorithm for Efficient Detection of (N, N)-Splittings and Its Application to the Isogeny Problem in Dimension 2". In: *PKC 2024, Part II*. Ed. by Qiang Tang and Vanessa Teague. Vol. 14603. LNCS. Springer, Cham, Apr. 2024, pp. 157–189. DOI: 10.1007/978-3-031-57725-3_6 (cit. on pp. 2, 15, 17, 18, 27, 33, 37, 38).

[Cor+26] Maria Corte-Real Santos, Harun Kır, Olivier Taïbi, and Benjamin Wesolowski. *Rapid mixing in superspecial isogeny graphs*. Cryptology ePrint Archive, Paper 2026/2429. 2026. URL: https://eprint.iacr.org/2026/2429 (cit. on pp. 2–4, 11, 12, 20, 25, 39–42).

[CS20] Craig Costello and Benjamin Smith. "The Supersingular Isogeny Problem in Genus 2 and Beyond". In: *PQCrypto 2020*. Ed. by Jintai Ding and Jean-Pierre Tillich. Springer, Cham, 2020, pp. 151–168. DOI: 10.1007/978-3-030-44223-1_9 (cit. on pp. 1, 12, 39).

[Dar+24] Pierrick Dartois, Luciano Maino, Giacomo Pope, and Damien 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 Kai-Min Chung and Yu Sasaki. Vol. 15486. LNCS. Springer, Singapore, Dec. 2024, pp. 304–338. DOI: 10.1007/978-981-96-0891-1_10 (cit. on pp. 28, 29).

[Dju19] Martin Djukanović. *Families of (3,3)-split Jacobians*. 2019. arXiv: 1811.10075 [math.AG]. URL: https://arxiv.org/abs/1811.10075 (cit. on p. 17).

[Eke87] Torsten Ekedahl. "On supersingular curves and abelian varieties". In: *Mathematica Scandinavica* 60 (1987), pp. 151–178 (cit. on p. 20).

[Eng09] Andreas Enge. “Computing modular polynomials in quasi-linear time”. In: *Mathematics of Computation* 78.267 (2009), pp. 1809–1824 (cit. on p. 5).

[Fre24] Sam Frengley. *On the geometry of the Humbert surface of square discriminant*. 2024. arXiv: 2408.09830 [math.AG]. URL: https://arxiv.org/abs/2408.09830 (cit. on pp. 2, 17).

[FK91] Gerhard Frey and Ernst Kani. “Curves of genus 2 covering elliptic curves and an arithmetical application”. In: *Arithmetic algebraic geometry*. Springer, 1991, pp. 153–176 (cit. on p. 14).

[GS01] P. Gaudry and É. Schost. “On the invariants of the quotients of the Jacobian of a curve of genus 2”. In: *Applied algebra, algebraic algorithms and error-correcting codes (Melbourne, 2001)*. Vol. 2227. Lecture Notes in Comput. Sci. Springer, Berlin, 2001, pp. 373–386 (cit. on p. 17).

[Gee88] Gerard van der Geer. “Humbert Surfaces”. In: *Hilbert Modular Surfaces*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1988, pp. 205–221. ISBN: 978-3-642-61553-5. DOI: 10.1007/978-3-642-61553-5_11. URL: https://doi.org/10.1007/978-3-642-61553-5_11 (cit. on p. 14).

[Gho+23] Sumanta Ghosh, Prahladh Harsha, Simao Herdade, Mrinal Kumar, and Ramprasad Saptharishi. “Fast numerical multivariate multipoint evaluation”. In: *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2023, pp. 1426–1439 (cit. on p. 12).

[Gla85] James Whitbread Lee Glaisher. “On the square of the series in which the coefficients are the sums of the divisors of the exponents”. In: *Mess. Math* 14 (1885), pp. 156–163 (cit. on p. 24).

[Har77] Robin Hartshorne. *Algebraic Geometry*. Vol. 52. Graduate Texts in Mathematics. New York: Springer-Verlag, 1977. ISBN: 978-0-387-90244-9 (cit. on pp. 6, 7, 23).

[HI80] Ki-ichiro Hashimoto and Tomoyoshi Ibukiyama. “On class numbers of positive definite binary quaternion hermitian forms”. In: *J. Fac. Sci. Univ. Tokyo* 27 (1980), pp. 549–601 (cit. on p. 20).

[HL20] Joris van der Hoeven and Grégoire Lecerf. “Fast multivariate multipoint evaluation revisited”. In: *Journal of Complexity* 56 (2020), p. 101405 (cit. on p. 12).

[HP26] Marc Houben and Lam L. Pham. *Computing endomorphisms of superspecial principally polarized abelian surfaces in time and memory $p^{4/5+o(1)}$*. Cryptology ePrint Archive, Paper 2026/2530. 2026. URL: https://eprint.iacr.org/2026/2530 (cit. on pp. 4, 5).

[How93] Everett W. Howe. “On the group orders of elliptic curves over finite fields”. In: *Compositio mathematica* 85.2 (1993), pp. 229–247 (cit. on p. 26).

[Hum99] G. Humbert. “Sur les fonctions abéliennes singulières. I”. In: *J. Math. Pures Appl. (5)* 5 (1899), pp. 233–350 (cit. on p. 14).

[Hum00] G. Humbert. “Sur les fonctions abéliennes singulières. II”. In: *J. Math. Pures Appl. (5)* 6 (1900), pp. 279–386 (cit. on p. 14).

[Hum01] G. Humbert. “Sur les fonctions abéliennes singulières. III”. In: *J. Math. Pures Appl. (5)* 7 (1901), pp. 97–123 (cit. on p. 14).

[IK94] T. Ibukiyama and T. Katsura. “On the field of definition of superspecial polarized abelian varieties and type numbers”. In: *Compositio Math.* 91.1 (1994), pp. 37–46 (cit. on p. 11).

[Igu60] J. Igusa. “Arithmetic variety of moduli for genus two”. In: *Ann. of Math. (2)* 72 (1960), pp. 612–649 (cit. on p. 10).

[JZ26] Bruce W. Jordan and Yevgeny Zaytman. “Isogeny graphs of superspecial abelian varieties and Brandt matrices”. In: *Journal of the Mathematical Society of Japan* (2026). Advance publication, pp. 1–40. DOI: 10.2969/jmsj/91629162. URL: https://projecteuclid.org/journals/journal-of-the-mathematical-society-of-japan/advance-publication/Isogeny-graphs-of-superspecial-abelian-varieties-and-Brandt-matrices/10.2969/jmsj/91629162.short (cit. on p. 11).

[KS91] Erich Kaltofen and David Saunders. “On Wiedemann’s method of solving sparse linear systems”. In: *International Symposium on Applied Algebra, Algebraic Algorithms, and Error-Correcting Codes*. Springer. 1991, pp. 29–38 (cit. on pp. 31, 32, 34).

[Kan94] Ernst Kani. “Elliptic curves on abelian surfaces”. In: *Manuscripta Mathematica* 84.1 (1994), pp. 199–223 (cit. on pp. 13–16).

[Kan97] Ernst Kani. “The number of curves of genus two with elliptic differentials”. In: *Journal für die reine und angewandte Mathematik* 485 (1997), pp. 93–122 (cit. on pp. 8, 13, 14, 16, 24, 29).

[Kan16] Ernst Kani. “The moduli spaces of Jacobians isomorphic to a product of two elliptic curves”. In: *Collectanea mathematica* 67.1 (2016), pp. 21–54 (cit. on p. 13).

[Kan18] Ernst Kani. “Elliptic subcovers of a curve of genus 2 II. The refined Humbert invariant”. In: *Journal of Number Theory* 193 (2018), pp. 302–335 (cit. on p. 13).

[Kan19] Ernst Kani. “Elliptic subcovers of a curve of genus 2. I. The isogeny defect”. In: *Annales mathématiques du Québec* 43.2 (2019), pp. 281–303 (cit. on pp. 15, 23).

[KS98] Ernst Kani and Wolfgang Schanz. “Modular diagonal quotient surfaces”. In: *Mathematische Zeitschrift* 227.2 (1998), p. 337 (cit. on p. 15).

[KU11] Kiran S Kedlaya and Christopher Umans. “Fast polynomial factorization and modular composition”. In: *SIAM Journal on Computing* 40.6 (2011), pp. 1767–1802 (cit. on pp. 3, 12).

[KK25] Eda Kırımlı and Gaurish Korpal. *Refined Humbert Invariants in Supersingular Isogeny Degree Analysis*. Cryptology ePrint Archive, Report 2025/1605. 2025. URL: https://eprint.iacr.org/2025/1605 (cit. on p. 2).

[KK26] Eda Kırımlı and Gaurish Korpal. *Detecting split surfaces, RM surfaces, and minimum walks on isogeny graphs*. Cryptology ePrint Archive, Paper 2026/2088. 2026. URL: https://eprint.iacr.org/2026/2088 (cit. on p. 2).

[KM25] Eda Kırımlı and Chloe Martindale. *The computational refined Humbert invariant problem is equivalent to the computational isogeny problem*. Cryptology ePrint Archive, Report 2025/1295. 2025. URL: https://eprint.iacr.org/2025/1295 (cit. on p. 2).

[Kuh88] R. M. Kuhn. "Curves of genus 2 with split Jacobian". In: *Trans. Amer. Math. Soc.* 307.1 (1988), pp. 41–49 (cit. on p. 17).

[Kum15] A. Kumar. "Hilbert modular surfaces for square discriminants and elliptic subfields of genus 2 function fields". In: *Research in the Mathematical Sciences* 2.1 (2015), pp. 1–46 (cit. on pp. 2, 17, 27).

[KR25] Sabrina Kunzweiler and Damien Robert. "Computing modular polynomials by deformation". In: *Research in Number Theory* 11.1 (2025), p. 10 (cit. on p. 5).

[KS26a] Sabrina Kunzweiler and Min-Yi Shen. "Endomorphisms via Splittings". In: *Post-Quantum Cryptography*. Ed. by Magali Bardet and Ruben Niederhagen. Cham: Springer Nature Switzerland, 2026, pp. 3–38. ISBN: 978-3-032-22698-3 (cit. on p. 2).

[KS26b] Péter Kutas and Min-Yi Shen. *A Note on Splittings and Endomorphisms*. Cryptology ePrint Archive, Paper 2026/1198. 2026. URL: https://eprint.iacr.org/2026/1198 (cit. on p. 2).

[LR15] Gilles Lachaud and Robert Rolland. "On the number of points of algebraic sets over finite fields". In: *Journal of Pure and Applied Algebra* 219.11 (2015), pp. 5117–5136 (cit. on p. 30).

[LO98] Ke-Zheng Li and Frans Oort. *Moduli of supersingular abelian varieties*. English. Vol. 1680. Lect. Notes Math. Berlin: Springer, 1998. ISBN: 3-540-63923-3. DOI: 10.1007/BFb0095931 (cit. on p. 11).

[LR23] David Lubicz and Damien Robert. "Fast change of level and applications to isogenies". In: *Research in Number Theory* 9.1 (2023), p. 7 (cit. on pp. 28, 29).

[MSV09] K. Magaard, T. Shaska, and H. Völklein. "Genus 2 curves that admit a degree 5 map to an elliptic curve". In: *Forum Math.* 21.3 (2009), pp. 547–566 (cit. on p. 17).

[MW25] Arthur Herlédan Le Merdy and Benjamin Wesolowski. "Unconditional Foundations for Supersingular Isogeny-Based Cryptography". In: *TCC 2025, Part III*. Ed. by Benny Applebaum and Huijia (Rachel) Lin. Vol. 16270. LNCS. Springer, Cham, Dec. 2025, pp. 266–297. DOI: 10.1007/978-3-032-12296-4_9 (cit. on p. 5).

[Mes91] Jean-François Mestre. "Construction de courbes de genre 2 à partir de leurs modules". In: *Effective methods in algebraic geometry*. Springer. 1991, pp. 313–334 (cit. on p. 10).

[Mil86] J. S. Milne. “Abelian varieties”. In: *Arithmetic geometry (Storrs, Conn., 1984)*. Springer, New York, 1986, pp. 103–150. ISBN: 0-387-96311-1 (cit. on p. 6).

[Mon26] Mickaël Montessinos. *Equivalent computational problems for superspecial abelian surfaces*. Cryptology ePrint Archive, Paper 2026/120. 2026. URL: https://eprint.iacr.org/2026/120 (cit. on p. 5).

[Mum70] D. Mumford. *Abelian varieties*. 2nd edn. Oxford University Press, London, 1970, pp. viii+242 (cit. on p. 6).

[PRS26] Aurel Page, Damien Robert, and Julien Soumier. “On the Conversion of Module Representations for Higher Dimensional Supersingular Isogenies”. In: *Advances in Cryptology – CRYPTO 2026*. Ed. by Nadia Heninger and Mike Rosulek. Cham: Springer Nature Switzerland, 2026, pp. 415–446. ISBN: 978-3-032-35398-6 (cit. on p. 5).

[PW23] Aurel Page and Benjamin Wesolowski. *The supersingular Endomorphism Ring and One Endomorphism problems are equivalent*. Cryptology ePrint Archive, Report 2023/1399. 2023. URL: https://eprint.iacr.org/2023/1399 (cit. on pp. 39, 40, 42).

[PW24] Aurel Page and Benjamin Wesolowski. “The Supersingular Endomorphism Ring and One Endomorphism Problems are Equivalent”. In: *EUROCRYPT 2024, Part VI*. Ed. by Marc Joye and Gregor Leander. Vol. 14656. LNCS. Springer, Cham, May 2024, pp. 388–417. DOI: 10.1007/978-3-031-58751-1_14 (cit. on pp. 5, 40).

[Pha26] Lam L. Pham. “A uniform bound for the minimal degree of an isogeny between principally polarized superspecial abelian surfaces”. In: (2026). arXiv: 2610.12080 [math.NT]. URL: https://arxiv.org/abs/2610.12080 (cit. on p. 5).

[ST68] Jean-Pierre Serre and John Tate. “Good reduction of abelian varieties”. In: *Annals of Mathematics* 88.3 (1968), pp. 492–517 (cit. on p. 19).

[Sha01a] T. Shaska. “Curves of genus 2 with (N, N) decomposable Jacobians”. In: *Journal of Symbolic Computation* 31.5 (2001), pp. 603–617 (cit. on p. 17).

[Sha01b] T. Shaska. *Curves of genus two covering elliptic curves*. University of Florida, 2001 (cit. on p. 17).

[Sha04] T. Shaska. “Genus 2 fields with degree 3 elliptic subfields”. In: *Forum Math.* 16.2 (2004), pp. 263–280 (cit. on p. 17).

[SV04] T. Shaska and H. Völklein. “Elliptic subfields and automorphisms of genus 2 function fields”. In: *Algebra, arithmetic and geometry with applications*. Springer, 2004, pp. 703–723 (cit. on p. 17).

[Sha+08] T. Shaska, G. S. Wijesiri, S. Wolf, and L. Woodland. “Degree 4 coverings of elliptic curves by genus 2 curves”. In: *Albanian J. Math.* 2.4 (2008), pp. 307–318 (cit. on p. 17).

[Sha26a] Tanush Shaska. *Superspecial points on Humbert surfaces and the probability of (N,N)-splitting*. Cryptology ePrint Archive, Paper 2026/2522.

2026. URL: https://eprint.iacr.org/2026/2522 (cit. on pp. 18, 22, 25).

[Sha26b] Tony Shaska. "Degrees of Genus-Two Split-Jacobian Loci and Humbert-Form Reconstruction". In: *arXiv preprint arXiv:2608.16426* (2026) (cit. on pp. 17–20).

[Sha26c] Tony Shaska. *The Humbert Form of Discriminant 36, and What It Says About Splitting Detection*. Cryptology ePrint Archive, Paper 2026/2327. 2026. URL: https://eprint.iacr.org/2026/2327 (cit. on pp. 2, 19, 20, 27, 33).

[Sil09] Joseph H Silverman. *The arithmetic of elliptic curves*. Vol. 106. 2. Springer, 2009 (cit. on p. 19).

[Udo26] Aleksei Udovenko. *Solving the supersingular isogeny problem in time $p^{2/5+o(1)}$ using bivariate multipoint evaluation*. Cryptology ePrint Archive, Paper 2026/1575. 2026. URL: https://eprint.iacr.org/2026/1575 (cit. on pp. 5, 12).

[Wes22] Benjamin Wesolowski. "The supersingular isogeny path and endomorphism ring problems are equivalent". In: *62nd FOCS*. IEEE Computer Society Press, Feb. 2022, pp. 1100–1111. DOI: 10.1109/FOCS52979.2021.00109 (cit. on p. 5).

[Wes26] Benjamin Wesolowski. *The supersingular isogeny problem in time and memory $p^{1/3+o(1)}$*. Cryptology ePrint Archive, Paper 2026/1486. 2026. URL: https://eprint.iacr.org/2026/1486 (cit. on p. 5).