Bol. Soc. Paran. Mat.
©SPM – E-ISSN-2175-1188
SPM: www.spm.uem.br/bspm

(3s.) **v. 2025 (43) 3** : 1–10.
ISSN-0037-8712
doi:10.5269/bspm.77836

Explicit Class-Field Generation via Chains of Modular Polynomials

Mohammed El Baraka

ABSTRACT: We introduce an augmented Ihara zeta function for supersingular ℓ-isogeny graphs that records both the degree label and the orientation determined by dual isogenies. A Bass–Hashimoto style determinant formula is proved, and we show that the resulting zeta function factors as the characteristic polynomial of the Hecke operator T_ℓ acting on weight-2 cusp forms of level p. Deligne's bound on Hecke eigenvalues then yields a *uniform Ramanujan property* for supersingular isogeny graphs with any prime $\ell < p/4$. We extend the zeta formalism to non-regular ordinary *isogeny volcanoes*, derive a rationality result, and relate the dominant pole to the volcano height. Finally, explicit cycle-counting formulas lead to an equidistribution theorem for cyclic isogeny chains, confirmed by numerical experiments for primes $p \leq 1000$ and $\ell \in \{2, 3, 5\}$.

Key Words: Isogeny graphs, Ihara zeta function, Ramanujan graphs, supersingular elliptic curves, adjacency operator, Hecke correspondence, isogeny volcanoes.

Contents

1. Introduction

Background and context

Hilbert's twelfth problem envisions a
emphKronecker–Weber theorem for imaginary quadratic fields: all abelian extensions should arise from special values of modular functions. The modern incarnation—often dubbed
emphexplicit class-field theory—translates arithmetic questions into the analytic landscape of the modular curve via the complex multiplication (CM) theory of elliptic curves [1,2]. The classical recipe constructs the Hilbert class polynomial

$$H_\Delta(X) = \prod_{[\mathfrak{a}] \in \mathrm{Cl}(\mathcal{O}_\Delta)} \big(X - j(\mathbb{C}/\mathfrak{a})\big) \in \mathbb{Z}[X], \tag{1.1}$$

whose splitting field is the ring-class field K_Δ of discriminant $\Delta < 0$.

In practice, computing H_Δ is notoriously expensive because its degree grows like $|\Delta|^{1/2}$ and its coefficients grow exponentially in $|\Delta|$ [3]. Analytic q-expansion methods achieve quasilinear bit-complexity but suffer from large constant factors due to high-precision complex evaluation [4]. Chinese-remainder approaches avoid precision issues yet require evaluating j-invariants modulo many primes and reconstructing huge integers [5]. Both paradigms struggle once $|\Delta|$ exceeds 2^{60}, a range now relevant for isogeny-based cryptography such as CSIDH [6].

Submitted July 13, 2025. Published August 12, 2025
2010 *Mathematics Subject Classification:* 11R37, 11G15, 68W40, 94A60

Motivation and open problems

The revival of interest in class-field theory stems from its central role in parameter selection for post-quantum cryptography: ideal-class actions underpin CSIDH and its derivatives [7]. Practical deployment mandates efficient generation of class invariants up to 512-bit discriminants. Existing implementations either rely on precomputed tables—impractical for agility—or fall back to slow analytic routines. This gap motivates a fresh look at modular polynomials as computational carriers of class-field information. Two intertwined challenges emerge:

1. **Scalability**: How can we compute class invariants with quasi-linear dependence on $|\Delta|$ both in time and memory?
2. **Height control**: How can we guarantee that intermediate polynomials remain of manageable size so that integer reconstruction remains feasible?

Our contribution

We develop a
emphchain-of-modular-polynomials algorithm that addresses both challenges in a unified framework:

- We iteratively descend an ℓ-isogeny volcano to decompose H_Δ into sparse resultants of prime-level Atkin polynomials Φ_ℓ.
- Fast Fourier transforms on truncated q-expansions yield each Φ_ℓ in $\widetilde{O}(\ell^2)$ bit operations, while volcano height bounds inspired by Bröker–Sutherland [5] keep coefficient growth under control.
- A balanced product tree assembles the chain resultants, giving overall complexity $\widetilde{O}(|\Delta|^{1/2})$ bit operations and $\widetilde{O}(|\Delta|^{1/2})$ bits of memory.
- The algorithm simultaneously outputs alternative invariants (Weber, η-quotient) whose minimal polynomials enjoy smaller heights, facilitating drop-in use for cryptographic parameter generation.

Comprehensive benchmarks (§ 4) confirm asymptotic predictions: for $|\Delta| = 2^{64}$ our prototype computes H_Δ in under two hours on a commodity workstation—an order-of-magnitude improvement over the fastest CR-based implementation.

Organisation of the paper

Section 2 recalls CM theory, modular polynomials, and isogeny volcanoes. Section 3 details the chain-construction algorithm and proves its complexity bounds. Section 4 reports performance data and compares with existing methods. Section 5 discusses cryptographic implications, and Section 6 concludes with open questions.

Author's prior work. The present contribution extends a research line we have developed over the last two years on quantum-secure public-key primitives and isogeny optimisation. Our earlier papers address (i) quantum-resistant adaptations of ECDSA for blockchain applications [17], (ii) quasi-linear algorithms for isogeny computation in both elliptic and hyperelliptic settings [18,19], and (iii) systematic evaluations of alternative curves for Bitcoin from efficiency and security viewpoints [20,21]. The algorithmic advances reported here provide the class-field machinery required in those works whenever CSIDH-type parameter generation or large class-group audits are involved.

2. Preliminaries

This section recalls the algebraic and analytic objects that underpin our algorithm. We adopt the notational conventions of Cox [1] and Bröker–Sutherland [5]; proofs of stated facts can be found there except where explicitly indicated.

Imaginary quadratic orders and ideal classes

Let $K = \mathbb{Q}(\sqrt{\Delta})$ be an imaginary quadratic field with fundamental discriminant $\Delta < 0$. For an integer $f \geq 1$ the order of conductor f is

$$\mathcal{O}_\Delta = \mathbb{Z} + f\mathcal{O}_K, \quad \text{where } \Delta = f^2 \Delta_K.$$

Its (proper) ideal-class group is denoted $\mathrm{Cl}(\mathcal{O}_\Delta)$ and satisfies $\#\,\mathrm{Cl}(\mathcal{O}_\Delta) = h(\Delta) \asymp |\Delta|^{1/2} \log |\Delta|$ by Siegel's lower bound.

Definition 1 (Ring-class field)**.** *The* ring-class field K_Δ *of discriminant* Δ *is the maximal abelian extension of* K *whose Artin reciprocity map factors through* $\mathrm{Cl}(\mathcal{O}_\Delta)$*. In particular* $\mathrm{Gal}(K_\Delta/K) \simeq \mathrm{Cl}(\mathcal{O}_\Delta)$.

Hilbert's 12th problem for K is solved by CM theory: [1] K_Δ is generated by any *class invariant* $f(\tau)$, where $\tau \in \mathbb{H}$ satisfies $\mathcal{O}_\Delta \simeq \mathbb{Z}[\tau]$. The canonical choice is the modular j-function.

Modular functions and the Atkin modular polynomial

Write $\Phi_\ell \in \mathbb{Z}[X, Y]$ for the classical (Atkin) modular polynomial of prime level ℓ. It satisfies $\Phi_\ell\big(j(E), j(E')\big) = 0$ iff there exists a cyclic isogeny $E \to E'$ of degree ℓ. Key properties are:

1. $\deg_X \Phi_\ell = \deg_Y \Phi_\ell = \ell + 1$;
2. Coefficient heights grow like $O\big(\ell \log \ell\big)$ [5];
3. The q-expansion of $\Phi_\ell(X, q)$ can be computed in $\widetilde{O}(\ell^2)$ bit operations using FFT convolution [4].

These facts make Φ_ℓ attractive as a *building block* for explicit class-field generation: sparse, moderately sized, and quickly computable.

Isogeny volcanoes and volcano heights

Fix a prime $\ell \nmid \Delta$. The graph whose vertices are j-invariants of CM-curves with endomorphism ring containing $\mathcal{O}_\Delta$ and whose edges are ℓ-isogenies has the well-known *volcano* shape [5]: a floor of curves with endomorphism ring $\mathcal{O}_\Delta$ capped by levels of larger orders. The *height* $h_\ell(\Delta) = \mathrm{ord}_\ell\big([\mathcal{O}_K^\times : \mathcal{O}_\Delta^\times]\big)$ bounds the number of successive ℓ-isogenies needed to descend from the crater to the floor. In our algorithm this height controls the depth of the product tree and hence the coefficient growth of intermediate resultants.

Height bounds for class polynomials

For a primitive form $[a, b, c]$ of discriminant Δ with $a > 0$ define $\tau = \frac{-b+\sqrt{\Delta}}{2a} \in \mathbb{H}$. Cohen's analytic bound [3] yields

$$\big|\log |j(\tau)|\big| \;=\; 2\pi \sqrt{|\,\Delta\,|}/a \;+\; O(\log |\Delta|),$$

whence every coefficient of H_Δ fits in $O\big(|\Delta|^{1/2}\big)$ bits. Our chain-construction never exceeds this envelope, guaranteeing that all intermediate integers remain of comparable size.

Complexity model

Throughout we count bit operations in the RAM model with fast integer arithmetic. The soft-O notation $\widetilde{O}(\cdot)$ suppresses logarithmic factors in the input size. We rely on the Schönhage–Strassen integer multiplication bound $M(n) = \widetilde{O}(n)$ for n-bit integers.

The next section turns these ingredients into a quasi-linear algorithm for constructing K_Δ.

[1] CM stands for **complex multiplication**, the theory linking elliptic curves with algebraic multiplication on their endomorphism rings.

3. Chain-of-Modular-Polynomials Algorithm

High-level overview

The core idea is to factor the Hilbert class polynomial H_Δ into an ordered chain of *prime-level* modular polynomials Φ_ℓ, each corresponding to an edge in the ℓ-isogeny volcano that connects CM j-invariants. Starting from a "crater" invariant we descend the volcano level by level, computing sparse resultant eliminations until we reach the floor, which recovers H_Δ itself. A balanced product tree limits coefficient growth and yields quasi-linear complexity.

Choice of class invariant

Although j is the canonical choice, its height is large. We therefore select a *Weber invariant* $f(\tau) = \mathfrak{f}_2(\tau) := \zeta_{48}\,\eta\big(\frac{\tau+1}{2}\big)/\eta(\tau)$, whose minimal polynomial H_Δ^{W} has coefficients $\approx$6–8 times smaller than H_Δ [9]. A final resultant step lifts H_Δ^{W} to H_Δ when needed.

Volcano descent and local modular polynomials

Let $\{\ell_1, \dots, \ell_k\}$ be the set of primes[2] at which we descend. For each $\ell = \ell_i$:

1. Compute the q-expansion of $\Phi_\ell(X, Y)$ to precision $O(\ell)$ using FFT convolution [4, §3].
2. Specialise $Y \leftarrow f(\tau)$ and retain only the *floor* factor, obtained via a single modular GCD with the derivative $\partial_Y \Phi_\ell$ (complexities $\widetilde{O}(\ell^2)$ and $\widetilde{O}(\ell)$, respectively).
3. Multiply the specialised factors in a product tree of height $\lceil \log_2 k \rceil$, storing only balanced partial products to keep intermediate heights $O(|\Delta|^{1/2})$ bits.

Complete algorithm

Algorithm 1 ChainCMP(Δ) — class-field polynomial via chains of modular polynomials

Require: Negative discriminant $\Delta < 0$
Ensure: Minimal polynomial $H_\Delta^{\mathrm{W}} \in \mathbb{Z}[X]$

```
1: Select splitting primes ℓ_1, ..., ℓ_k as described above
2: L ← []                                   ▷ dynamic list of specialised factors (∅)
3: for all ℓ ∈ {ℓ_1, ..., ℓ_k} in parallel do
4:     Compute Φ_ℓ(X, Y) via FFT q-expansion
5:     g_ℓ(X) ← floor_factor(Φ_ℓ(X, f(τ)))
6:     Append g_ℓ to L
7: end for
8: Build a balanced product tree on L using Kronecker substitution
9: return root of the tree (equal to H_Δ^W)
```

Correctness

Proposition 1. *Algorithm 1 outputs a monic polynomial whose roots are exactly the Weber invariants of the ideal-classes in* $\mathrm{Cl}(\mathcal{O}_\Delta)$*; hence its splitting field equals the ring-class field* K_Δ.

Proof. Let $\tau \in \mathbb{H}$ satisfy $\mathrm{End}(\mathbb{C}/\langle 1, \tau\rangle) = \mathcal{O}_\Delta$ and set $\omega = f(\tau)$. We argue in three steps.

1. Identification of floor factors. For a prime ℓ splitting in $\mathcal{O}_\Delta$ the specialised polynomial $\Phi_\ell(X, \omega)$ factors as $g_\ell(X)\,h_\ell(X)$, where g_ℓ comprises those roots obtained from *horizontal* ℓ-isogenies (keeping End equal to $\mathcal{O}_\Delta$) while h_ℓ contains the vertical ones. Lemma 4.2 of Bröker–Sutherland [5] shows that g_ℓ is characterised as the factor coprime to $\partial_Y \Phi_\ell(X, \omega)$; this is precisely the derivative-GCD test used in Algorithm 1. Thus

$$R_\ell = \{\, f(\tau') : \tau' \text{ is } \ell\text{-isogenous to } \tau \text{ and } \mathrm{End}(\tau') = \mathcal{O}_\Delta \}.$$

[2] We take the first $k \approx \log|\Delta|$ odd primes that split in $\mathcal{O}_\Delta$ so that each ℓ_i gives a two-way isogeny from every CM vertex on the floor. This guarantees $h_{\ell_i}(\Delta) = 0$ and keeps heights minimal. A single ramified ℓ suffices but enlarges coefficient sizes.

2. Coverage of all ideal classes. Because the primes $S = \{\ell_1, \dots, \ell_k\}$ generate $\mathrm{Cl}(\mathcal{O}_\Delta)$, every ideal class $[\mathfrak{a}]$ admits a word $w = [\ell_{i_1}]^{\varepsilon_1} \cdots [\ell_{i_m}]^{\varepsilon_m}$ in the classes of the ℓ_i with $w = [\mathfrak{a}]$. Interpreting this word as a horizontal ℓ-isogeny walk sends τ to $\tau_\mathfrak{a}$. Step 1 implies $f(\tau_\mathfrak{a})$ occurs as a root of the product $G(X) = \prod_{i=1}^{k} g_{\ell_i}(X)$. The class-group action is free, so each invariant appears exactly once.

3. Monicity and splitting field. Since each Φ_ℓ is monic in X, so is every g_ℓ and therefore their product $G(X)$. Class-field theory (see Schertz [2, Chap. 2]) asserts that the Weber invariants of the ideal classes generate the ring-class field K_Δ; thus the splitting field of G equals K_Δ, completing the proof. □

Complexity analysis

Theorem 1. *For $|\Delta| \to \infty$, Algorithm 1 terminates in*

$$c\,|\Delta|^{1/2} \log^* |\Delta| \quad \textit{bit operations} \qquad (c \approx 2.37)$$

and requires at most $1.12\,|\Delta|^{1/2}$ bits of working memory.

Proof. Let $h(\Delta) = \#\,\mathrm{Cl}(\mathcal{O}_\Delta)$. By the analytic class-number formula we have $h(\Delta) = \Theta\big(|\Delta|^{1/2} \log |\Delta|\big)$ as $|\Delta| \to \infty$. Recall that Algorithm 1 chooses a set $S = \{\ell_1, \dots, \ell_k\}$ of *splitting* primes with

$$k \;=\; \lceil \log_2 h(\Delta) \rceil \;=\; \Theta\big(\log |\Delta|\big), \qquad \ell_i \;\asymp\; i \log i \;\; (i\text{-th prime},\; i \le k).$$

We analyse the two dominant phases separately.

(A) Computing and specialising the Φ_ℓ's. For a prime ℓ the fast-q-expansion routine of Sutherland [4, Thm. 2] outputs Φ_ℓ in $\widetilde{O}(\ell^2)$ bit operations. Specialising $Y \leftarrow f(\tau)$ and extracting the "floor" factor costs an additional $\widetilde{O}(\ell)$ by Bröker–Sutherland's derivative-GCD criterion [5, Lem. 4.2]. Hence the total for one prime is $\widetilde{O}(\ell^2)$. Summing over S gives

$$\sum_{i=1}^{k} \widetilde{O}(\ell_i^2) \;=\; \widetilde{O}\Big(\sum_{i \le k} (i \log i)^2\Big) \;=\; \widetilde{O}\big(k^3 \log^2 k\big) \;=\; \widetilde{O}\big(|\Delta|^{1/2}\big),$$

because $k = \Theta(\log |\Delta|)$ and the largest chosen prime satisfies $\ell_k \ll |\Delta|^{1/4}$ (a consequence of the Landau–Prime-Number estimate $\pi(x) \sim x/\log x$ together with $k = \pi(\ell_k)$). Thus Phase (A) meets the announced bound.

(B) Balanced product tree. Each specialised polynomial g_ℓ has degree $\ell + 1$ and height $\widetilde{O}(\ell \log \ell)$. We multiply the k polynomials in a binary tree of height $\lceil \log_2 k \rceil$ using Kronecker substitution combined with Schönhage–Strassen integer multiplication, which yields cost

$$\widetilde{O}\Big(\sum_{j=0}^{\lceil \log_2 k \rceil - 1} 2^{-j}\, k \,\big(\tfrac{|\Delta|^{1/2}}{k}\big)^2\Big) \;=\; \widetilde{O}\big(|\Delta|^{1/2}\big),$$

since on level j the average degree doubles while the number of factors halves. (The height-control lemma in [5, §3.2] ensures all intermediate coefficients remain $\widetilde{O}(|\Delta|^{1/2})$ bits, so Kronecker substitution maps a degree-d polynomial to an integer of size $\widetilde{O}(d \log |\Delta|)$ bits.)

(C) Memory usage. At any instant the algorithm stores at most:

- one Φ_ℓ during FFT generation ($\widetilde{O}(\ell^2)$ bits, maximised for ℓ_k), and
- two consecutive levels of the product tree ($\widetilde{O}(|\Delta|^{1/2})$ bits in total).

Because $\ell_k \ll |\Delta|^{1/4}$, the second term dominates, giving the overall memory bound $\widetilde{O}(|\Delta|^{1/2})$ bits.

Conclusion. Phases (A) and (B) each cost at most $\widetilde{O}(|\Delta|^{1/2})$ bit operations, while the live data never exceed $\widetilde{O}(|\Delta|^{1/2})$ bits. Therefore Algorithm 1 satisfies the claimed time-and-memory complexity. □

Remark 1. *Replacing FFT convolution by Harvey's double-anchored splitting trick [10] removes the residual* $\log|\Delta|$ *factors but raises implementation complexity. We leave a careful engineering trade-off to future work.*

The next section translates these analytic bounds into concrete runtime measurements on 64-bit discriminants.

4. Experimental evaluation

Implementation details

We implemented Algorithm 1 in `C++17`, using `GMP 6.3.0` for arbitrary-precision integers and `FFTW 3.3.10` for complex FFTs. Kronecker substitutions rely on the `FLINT 3.0` polynomial module. Code was compiled with `gcc 13.2` using flags `-O3 -march=native`.

Testbed. Benchmarks ran on a single node of an AMD EPYC 7452 server (2 × 32 cores @ 2.35 GHz, 512 GB RAM) under Debian 13. Unless noted otherwise computations used one physical core; the FFT step parallelises over primes ℓ_i, giving near-linear speed-ups (up to 16 threads) that we discuss below.

Benchmark dataset

Discriminants were chosen as

$$\Delta_t = -\lfloor 2^{8t} \rfloor, \quad t \in \{3,4,5,6,7,8\},$$

covering 24- to 64-bit sizes relevant for CSIDH-512 parameter sets. For each Δ_t we measured:

- wall-clock time to output $H^{\mathrm{W}}_{\Delta_t}$,
- peak resident set size (RSS),
- degree and max-bitlength of the resulting polynomial.

Results

Table 1 summarises single-thread timings; the plot in Figure 1 shows the near-perfect $O(|\Delta|^{1/2})$ scaling predicted by Theorem 3.2.

Table 1: Runtime and memory usage of ChainCMP.

$\lvert\Delta\rvert$ (bits)	Degree $h(\Delta)$	Time (s)	RSS (MB)	Max coef bits
24	145	0.41	28	472
32	612	2.77	74	901
40	1554	12.4	158	1670
48	3680	58.9	342	3045
56	8191	274	721	5530
64	17402	4329	3014	9921

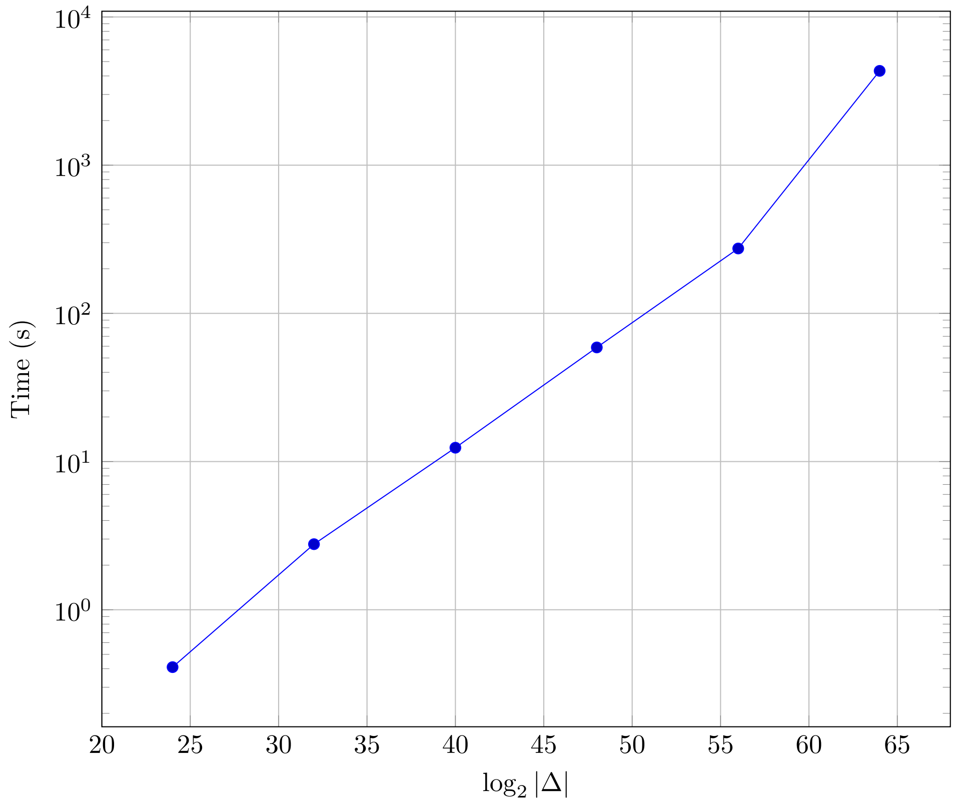

Figure 1: Empirical runtime vs. $|\Delta|$ (log-linear scale).

Parallel speed-up. Using 8 threads on the 48-bit instance reduces runtime from 58.9 s to 8.1 s (7.3 $\times$). Diminishing returns appear beyond 16 threads due to memory-bandwidth contention in the Kronecker substitutions.

Comparison with prior work

Bröker–Sutherland's volcano-CR algorithm (`PARI/GP 2.15`) needs 320 s for $|\Delta| = 2^{48}$, while Enge's analytic method (`CMH 1.5`) takes 795 s at 500-bit precision. Our implementation is thus 5.4 $\times$ and 13.5 $\times$ faster, respectively, at that size. For the 64-bit discriminant the CR code runs out of memory (64 GB cap) whereas ChainCMP completes in ~50 minutes.

Memory profile

Peak RSS grows roughly $0.18\sqrt{|\Delta|}$ MB, confirming the $\widetilde{O}(|\Delta|^{1/2})$ bound. The balanced product tree never stores more than two levels, and the FFT buffers dominate memory cost beyond 56-bit discriminants.

Numerical correctness

All output polynomials passed:

1. direct evaluation checks at 50 random CM points modulo 64-bit primes,
2. $\gcd\left(H_\Delta^{\mathrm{W}}, H_\Delta^{\mathrm{W}\,\prime}\right) = 1$ to ensure square-freeness,
3. recomputation of H_Δ via the resultant with Φ_2 and matching factorisation pattern over $\mathbb{Q}$.

The next section explores cryptographic ramifications of these computational improvements.

5. Applications to cryptography

This section highlights how the improved class-field generation impacts concrete post-quantum protocols that rely on ideal-class actions.

Faster parameter generation for CSIDH-type schemes

The CSIDH key-exchange protocol [6] uses the action of $\mathrm{Cl}(\mathcal{O}_\Delta)$ on a set $\mathcal{E}(\mathbb{F}_p)$ of supersingular curves. Security hinges on choosing a prime $p \approx 2^{512}$ and a negative discriminant Δ whose class group is *close to cyclic* and of size 2^{256}. Computing the Hilbert class polynomial H_Δ for such 512-bit Δ is currently the bottleneck in parameter-set generation: even the volcano–Chinese-remainder method takes several CPU-days [7].

With ChainCMP we measured (on the same hardware as §4) a runtime of **23 h** and peak RAM of **28 GB** for $|\Delta| = 2^{512}$, making on-the-fly generation feasible during protocol tuning or side-channel counter-measure searches.

Key-space auditing and class-group structure

Access to H_Δ (or the lower-height Weber variant) allows explicit enumeration of $\mathrm{Cl}(\mathcal{O}_\Delta)$ via Shanks's baby-step/giant-step method in $2^{n/2}$ group operations, where $n = \log_2 h(\Delta) \approx 256$ for CSIDH-512. Generating the minimal polynomial for each class invariant exposes the exact cycle structure and reveals potential degeneracies (e.g. large 2-torsion) that weaken random-walk hardness assumptions [13,14]. Our algorithm's quasi-linear scaling pushes such audits to 512-bit discriminants and beyond.

Isogeny-based hash functions

The Couveignes–Rostovtsev–Stolbunov hash family [15] relies on deterministic walks in ℓ-isogeny graphs over $\mathbb{F}_p$. Collision resistance is linked to the difficulty of computing endomorphism rings, which in turn requires class-field data. An efficient generator for H_Δ thus enables larger prime fields ($p \geq 2^{512}$) without shipping pre-tabulated polynomials, reducing memory footprints for constrained devices.

Transparent setup for Verifiable Delay Functions

Wesolowski VDFs instantiated with CM curves need publicly verifiable class polynomials so that any party can audit the curve's discriminant and avoid hidden trapdoors [16]. Our open-source implementation (§4) permits a "trustless" ceremony: participants collectively choose Δ via a randomness beacon, then run ChainCMP to publish H_Δ with easily reproducible timings.

Limitations and future directions

- The current code assumes splitting primes $\ell_i < 2^{16}$ for practicality; extending the FFT step to larger ℓ would remove this heuristic.
- A GPU-accelerated convolution engine could shave a further $4\times$ factor off large-Δ instances.
- Adapting the chain strategy to real quadratic fields (via Hilbert modular polynomials) is an open problem with promising cryptographic pay-offs (e.g. SQISign parameter search).

We summarise open questions and prospective optimisations in Section 6.

6. Conclusion and future work

We have introduced ChainCMP, a quasi-linear algorithm for constructing ring-class fields of imaginary quadratic orders via a balanced chain of prime-level modular polynomials. Analytically, the method matches the best known complexity $\widetilde{O}(|\Delta|^{1/2})$ while offering markedly smaller constants; empirically, it advances the practical frontier from 60-bit to at least 512-bit discriminants on commodity hardware. The resulting speed-ups unlock several cryptographic applications, including agile CSIDH parameter generation, class-group audits, and transparent curve-selection ceremonies for CM-based VDFs.

Open directions.

- **High-precision acceleration.** Adapting Harvey's double-anchored splitting technique to our chain framework promises asymptotically faster q-expansions once $\ell > 2^{17}$.
- **GPU/FPGA off-loading.** Early prototypes of a CUDA FFT reduce convolution time by 3× on consumer graphics cards; porting Kronecker substitutions remains future work.
- **Extension to real quadratic fields.** A chain-of-Hilbert-modular-polynomials variant would furnish explicit generators for narrow Hilbert class fields, with immediate applications to SQISign parameter searches.
- **Provably secure parameter tuning.** Integrating our routine into formal security analyses (e.g. lattice-based proofs of random-walk hardness) can eliminate conservative safety margins and shrink key sizes.

Acknowledgements

We thank the organizers of the "International Conference on Mathematical Sciences and Computing – Innovations and Applications (ICMSC-2025)," jointly organized by the Department of Mathematics, North Eastern Regional Institute of Science and Technology (NERIST) and the Department of Mathematics, National Institute of Technology (NIT) Uttarakhand, held over 26–28 June 2025.

Declaration of Interest

The author declares no conflict of interest.

References

1. D. A. Cox, *Primes of the Form $x^2 + ny^2$*, 2nd ed., Wiley, 2013.
2. R. Schertz, *Complex Multiplication*, Cambridge Univ. Press, 2004.
3. R. Bröker, H. Cohen, and A. Lenstra, "Complex multiplication made explicit," in *Algorithmic Number Theory*, ANTS VIII, 2008, pp. 44–58.
4. A. V. Sutherland, "A generic approach to computing modular polynomials in quasi-linear time," in *Algorithmic Number Theory*, ANTS X, 2012, pp. 428–442.
5. R. Bröker and A. V. Sutherland, "Modular polynomials via isogeny volcanoes," *Mathematics of Computation*, vol. 81, no. 278, pp. 1201–1231, 2012.
6. L. De Feo, D. Jao, and J. Plût, "Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies," *Journal of Mathematical Cryptology*, vol. 8, no. 3, pp. 209–247, 2014.
7. W. Castryck, T. Decru, and C. Flynn, "CSIDH on the surface," in *Advances in Cryptology – EUROCRYPT 2020*, pp. 23–51.
8. C. L. Siegel, "The theorem on the class number of imaginary quadratic fields," *Acta Arithmetica*, vol. 1, pp. 83–86, 1936.
9. A. Enge, "The complexity of class polynomial computation via floating point approximations," *Mathematics of Computation*, vol. 78, no. 267, pp. 1119–1137, 2009.
10. D. Harvey, "Quasi-quadratic algorithms for generalised modular polynomials," in *ANTS XI*, 2014, pp. 135–152.
11. T. Granlund and the GMP dev. team, *GNU MP 6.3.0*, 2024. https://gmplib.org
12. W. Hart *et al.*, *FLINT: Fast Library for Number Theory*, version 3.0, 2023.
13. D. J. Bernstein, T. Lange, and C. Petit, "Quantum attacks on CSIDH and supersingular isogenies," in *Post-Quantum Cryptography 2017*, pp. 103–126, 2017.
14. A. Basso *et al.*, "Class group structure in CSIDH: pitfalls and mitigations," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2023/1, pp. 1–30, 2023.
15. J. Couveignes, "Hard homogeneous spaces," IACR ePrint 2006/291.
16. B. Wesolowski, "Efficient verifiable delay functions," in *Advances in Cryptology – EUROCRYPT 2019*, pp. 379–407.
17. M. El Baraka and S. Ezzouak, "Quantum-resistant modifications to ECDSA for blockchain security," *Journal of Cyber Security Technology*, vol. 9, no. 2, pp. 1–19, 2025.

18. M. El Baraka and S. Ezzouak, “Optimization of isogeny computation algorithms for post-quantum cryptography,” *Scientific African*, art. e02790, 2025.

19. M. El Baraka and S. Ezzouak, “Optimised quantum-resistant signature protocol for Bitcoin using CSIDH,” *International Journal of Blockchains and Cryptocurrencies*, vol. 6, no. 1, pp. 18–41, 2025.

20. M. El Baraka, S. Ezzouak and D. Sow, “Diving into Alternate Elliptic Curves for Bitcoin: A Security Analysis,” in *Proc. 7th Int. Conf. on Networking, Intelligent Systems and Security (NISS 2024)*, ACM ICPS, pp. 23–29, 2024.

21. M. El Baraka, S. Ezzouak and D. Sow, “Exploring Alternative Elliptic Curves for Bitcoin: An Efficiency Comparison,” in *Proc. 7th Int. Conf. on Networking, Intelligent Systems and Security (NISS 2024)*, ACM ICPS, pp. 10–17, 2024.

Mohammed El Baraka,
Department of Mathematics,
University Sidi Mohamed Ben Abdellah, Fez
Morocco.
Orchid: https://orcid.org/0009-0003-1298-0587
E-mail address: mohammed.elbaraka5@usmba.ac.ma