

HAWK- n Key Recovery Reduces to SVP in Dimension $n/2 + 1$

Zygimantas Straznickas and Stephen A. Weis

Anthropic
{zygi,sweis}@anthropic.com

Abstract. HAWK is a lattice signature scheme that is currently a third-round candidate in NIST’s post-quantum signature competition. We give an unconditional, deterministic polynomial-time reduction from HAWK- n key recovery over $K_n = \mathbb{Q}(\zeta_{2^\ell})$ to $\text{poly}(n)$ calls to an exact Shortest Vector Problem (SVP) oracle in dimension $n/2 + 1$, where $n = 2^{\ell-1}$ is the ring degree. The reduction uses a nontrivial automorphism of the key lattice, supplied by the Galois involution $\tau : \zeta \mapsto -\zeta$ and recoverable as a shortest vector of a public rank- n lattice isometric, up to scaling, to $\mathbb{Z}^{n/2+1} \oplus \sqrt{2}\mathbb{Z}^{n/2-1}$. Ducas’s block reduction [Duc23; BN24] on this near-hypercubic class finds the automorphism, and the descent of van Gent and Pulles [GP25] recovers the key from it. In the gate-count model of [AGPS20], the attack lowers the key-recovery cost of HAWK-512 from 2^{150} to 2^{108} and of HAWK-1024 from 2^{288} to 2^{182} . We demonstrate this with a practical implementation that recovers a HAWK-256 secret key end-to-end in a few hours on a single server. The construction does not transfer to Falcon. Conductors $m \in \{p^k, 2p^k\}$ (p an odd prime), i.e. the $m > 4$ with cyclic $(\mathbb{Z}/m)^\times$, evade the attack.

Keywords: HAWK · Lattice Isomorphism Problem · Module-LIP · Cyclotomic fields · Key recovery · SVP

1 Introduction

NIST’s post-quantum signature competition [NIST22] attracted several proposals based on hard lattice problems outside the LWE/SIS family. Among these, HAWK [HAWK25], originally proposed in [DPPvW22], bases its security on the module Lattice Isomorphism Problem (module-LIP) over power-of-two cyclotomic rings, and has fast, floating-point-free signing and compact keys. It was recently selected for the third round of evaluation [Ala26], the only lattice-based candidate to advance. The secret key is a short basis $B \in \text{SL}_2(R_n)$ of the standard lattice R_n^2 over $K_n = \mathbb{Q}(\zeta)$, where $\zeta := \zeta_{2^\ell}$ is a primitive 2^ℓ -th root of unity ($\ell \geq 3$) and R_n is its ring of integers. The public key is its Gram matrix $Q = B^*B$ ($*$ is the conjugate-transpose), Hermitian with respect to complex conjugation. Key recovery entails finding any B' with $B'^*B' = Q$ and reduces to the search module-LIP problem (smLIP) for the isometry class of R_n^2 .

The best previously known key-recovery attack is direct Block–Korkine–Zolotarev (BKZ) lattice reduction on the rank- $2n$ key lattice. HAWK’s designers derive blocksizes $\beta_{\text{key}} \in \{211, 452, 940\}$ for the three parameter sets [HAWK25, Table 7] to achieve different security bounds based on best-known attacks. Ducas [Duc23] provides a powerful tool: a block-reduction algorithm that, given a lattice already known to be isometric to a scaled $\mathbb{Z}^d$, returns its hidden orthogonal basis with calls to an exact-SVP oracle in dimension only $\lfloor d/2 \rfloor + 1$. By itself this yields no HAWK attack: applied to the rank- $2n$ key lattice, it gives oracle dimension $n + 1$, with no improvement on the blocksizes above. Nevertheless, NIST’s second-round status report encourages further analysis of HAWK’s “security assumptions, particularly the smLIP problem within the specific structure of cyclotomic number fields” and notes that prior module-LIP techniques “currently appear inapplicable to the complex cyclotomic number fields utilized by HAWK” [Ala26, §3.4].

Prior module-LIP cryptanalysis [GS02; MPPW24; EP24; Che25; LJPW24] works with complex conjugation c , whose fixed field is totally real. We use a second order-2 Galois involution on the base field, $\tau : \zeta \mapsto -\zeta$, whose fixed field is the *complex* cyclotomic subfield $\mathbb{Q}(\zeta^2)$. We show that the associated τ -cocycle $V_\tau = B^{-1}\tau(B)$ is a shortest vector of a publicly computable lattice of rank only n : 2×2 matrices over R_n cut out by $\mathbb{Q}$ -linear constraints with coefficients depending only on Q .

We prove this lattice is isometric to the near-hypercubic lattice $\sqrt{n/4} \mathbb{Z}^{n/2+1} \oplus \sqrt{n/2} \mathbb{Z}^{n/2-1}$ (Proposition 4.5), so Ducas’s block reduction [Duc23; BN24]¹ finds all its shortest vectors with SVP calls in dimension $n/2 + 1$ (Theorem 5.1, Appendix E). Recovering the key from V_τ is then the descent of van Gent and Pulles [GP25]: their kernel sublattice $\{x \in R_n^2 : \tau(x) = V_\tau^{-1}x\}$ is isometric to a scaled $\mathbb{Z}^n$ (Proposition D.2), and one further run of Ducas’s algorithm on it returns an equivalent key. Combined, this is our main result (Theorem 6.1): HAWK key recovery reduces deterministically to $\text{poly}(n)$ calls to exact SVP in dimension $n/2 + 1$. We implement the attack and recover HAWK-256 secret keys end-to-end in §6.2. The construction does not transfer to Falcon (Appendix F). Conductors with cyclic $(\mathbb{Z}/m)^\times$ (for $m > 4$, those of the form p^k or $2p^k$ with p an odd prime) evade it; we expect many other composite conductors do not (§7).

In the gate-count model of the specification [HAWK25, §5.3], the reduction lowers the key-recovery cost of HAWK-512 from 2^{150} to at most 2^{108} gates and of HAWK-1024 from 2^{288} to at most 2^{182} (Table 1, §6.1). §2 outlines the attack; §3 fixes notation; §§4–6 contain the construction, the reduction, and the experiments; §§7–9 cover the scope of the attack, prior work, and open questions.

¹ The term *near-hypercubic* is Bambury and Nguyen’s [BN24]; the precise class meant throughout this paper is fixed in §3.3.

2 Technical Overview

A HAWK public key is a matrix $Q = B^*B \in \text{GL}_2(R_n)$, where $B \in \text{SL}_2(R_n)$ is the secret key, $K_n = \mathbb{Q}(\zeta)$ for a primitive 2^ℓ -th root of unity ζ , $R_n = \mathbb{Z}[\zeta]$, and $n = 2^{\ell-1}$ is the degree of K_n (§3.1). Key recovery asks for any B' with $B'^*B' = Q$; such a B' signs in place of the secret key [HAWK25, §5.1]. The Galois group of $K_n/\mathbb{Q}$ contains three involutions: complex conjugation c , the map $\tau : \zeta \mapsto -\zeta$, and $\sigma := c\tau$ (§3.2). The attack computes the matrix $V_\tau := B^{-1}\tau(B)$ from Q alone (without needing B), then computes a B' from V_τ . It has three steps.

2.1 The attack in three steps

1. Step 1 (§4) computes from Q a basis of a lattice $\Lambda_B^{(\tau)} \subset M_2(R_n)$ of rank n that contains V_τ as a shortest vector in polynomial time.
2. Step 2 (Theorem 5.1) computes all $2(n/2 + 1)$ shortest vectors of $\Lambda_B^{(\tau)}$. §6.1 bounds the number of SVP-oracle calls in dimension $n/2 + 1$ by $O(n^2 \log n)$. In practice, a single progressive sieve sufficed for HAWK-256, §6.2.
3. Step 3 (Proposition D.2, Appendix D) computes B' from the $2(n/2 + 1)$ candidates: exactly two of them, $\pm V_\tau$, satisfy $Y \equiv I \pmod{2}$, and these are followed by at most two further reduction runs, each with SVP-oracle calls in dimension at most $n/2 + 1$.

Theorem 6.1 combines the three steps into a deterministic reduction: for every power of two $n \geq 4$, HAWK- n key recovery reduces to $\text{poly}(n)$ arithmetic operations and $\text{poly}(n)$ calls to an exact-SVP oracle in dimension at most $n/2 + 1$. Table 1 (§6.1) lists the costs per parameter set; Appendix F explains why Falcon is unaffected.

2.2 Derivation overview

Step 1: the cocycle lattice. The attack needs, from the public key alone, a lattice that contains a vector leading back to the secret key. That vector is the cocycle $V_\tau = B^{-1}\tau(B)$. Although it is built from B , the vector satisfies the cocycle equation and an intertwining relation between the two public Gram matrices Q and $\tau(Q)$. Both become $\mathbb{Q}$ -linear conditions on its entries with coefficients computable from the public key (Lemma 4.1). The integer solutions of the two conditions form the lattice $\Lambda_B^{(\tau)}$, of rank n (Lemma 4.3). A $\mathbb{Z}$ -basis of it is computed from Q in polynomial time (§4.1). V_τ is a shortest vector of $\Lambda_B^{(\tau)}$. Lengths are measured by $Q^{(\tau)}(Y) := \text{Tr}_F(\det Y)$, which is the trace from the degree- $n/4$ field F fixed by c and τ , §3.2. $\det V_\tau = 1$ since $\det B = 1$, and so its squared length is the minimum $\text{Tr}_F(1) = n/4$, regardless of the key-generation distribution (Lemma 4.4). The step uses two properties of HAWK: $\det B = 1$, and the public key is the Gram matrix of B .

Step 2: the isometry class and its reduction. Finding the shortest vectors of $\Lambda_B^{(\tau)}$ requires its isometry class. The change of variables $Y \mapsto BY\tau(B)^{-1}$ carries the cocycle lattice of any key isometrically to that of the trivial key $B = I$ (Lemma 4.2). The map is secret, but only its existence matters, so the class can be read off at $B = I$. There, working in the power basis, $(\Lambda_I^{(\tau)}, Q^{(\tau)}) \cong \sqrt{n/4} \mathbb{Z}^{n/2+1} \oplus \sqrt{n/2} \mathbb{Z}^{n/2-1}$ (Proposition 4.5): up to scaling, the cocycle lattice is in the near-hypercubic class $[\mathbb{Z}^p \oplus \sqrt{2} \mathbb{Z}^q]$ with $p = n/2 + 1$, $q = n/2 - 1$.

Ducas’s block-reduction framework [Duc23; BN24] solves LIP for this class with SVP-oracle calls in dimension $p = n/2 + 1$ (Theorem 5.1, proved in Appendix E), returning all $2(n/2 + 1)$ shortest vectors of $\Lambda_B^{(\tau)}$.

Step 3: from a shortest vector to the key. This step is the automorphism framework of [GP25] (§5, Appendix D). Each shortest vector Y defines an order-2 automorphism of the public quadratic module, with fixed sublattice $\mathcal{M}_Y := \{x \in R_n^2 : \tau(x) = Y^{-1}x\}$ under the form $T(x, y) := \text{Tr}_{K_n}(x^*Qy)$. Only $Y = \pm V_\tau$ lead back to the key, and they are the two candidates with $Y \equiv I \pmod{2}$ (the parity test, Appendix D). For $\pm V_\tau$, $(\mathcal{M}_Y, T)$ is isometric to $\sqrt{n} \mathbb{Z}^n$ with the columns of B^{-1} (up to roots of unity) as its shortest vectors (Proposition D.2). One run of the algorithm of [Duc23] returns the shortest vectors with SVP-oracle calls in dimension at most $n/2 + 1$, and any pair u_1, u_2 with $u_1^*Qu_2 = 0$ gives an equivalent key $[u_1 u_2]^{-1}$. Theorem 6.1 assembles the three steps; §6.1 counts the oracle calls and gates.

3 Preliminaries

3.1 Fields, dimensions, and the key-recovery problem

Throughout, $\ell \geq 3$ is fixed, ζ_m denotes a primitive m -th root of unity and $\zeta := \zeta_{2^\ell}$, and

$$\begin{aligned} K_n &= \mathbb{Q}(\zeta), & R_n &= \mathbb{Z}[\zeta], \\ n &:= [K_n : \mathbb{Q}] = 2^{\ell-1}, & K_n^+ &:= \mathbb{Q}(\zeta + \zeta^{-1}). \end{aligned}$$

Here K_n and R_n are the cyclotomic field and its ring of integers, in the notation of the specification [HAWK25]; n is the ring degree (also the specification’s n), the module lattice R_n^2 has $\mathbb{Z}$ -rank $2n$, K_n^+ is the maximal totally real subfield of K_n , and $[K_n^+ : \mathbb{Q}] = n/2$. For a general number field E , $\mathcal{O}_E$ denotes its ring of integers, and we abbreviate $\text{Tr}_E := \text{Tr}_{E/\mathbb{Q}}$ and $N_E := N_{E/\mathbb{Q}}$ for the trace and norm down to $\mathbb{Q}$. Superscripts on a ring or field carry several meanings, always distinguishable from context:

- the unit group ($R^\times$) and the k -fold Cartesian power (R^k for an integer $k \geq 1$);
- the fixed field or ring of an automorphism, or of the subgroup $\langle \cdots \rangle$ generated by several (E^γ , $E^{(c, \tau)}$, $\mathcal{O}_E^\gamma$, R_n^γ);

- the (-1) -eigenspace of an involution $(E^{\gamma=-1}, R_n^{\gamma=-1})$;
- the maximal totally real subfield (E^+) .

For $X \in M_{a \times b}(K_n)$ (the $a \times b$ matrices over K_n ; we abbreviate $M_a := M_{a \times a}$) write $X^* := \overline{X}^T$ where $\bar{\cdot} = c$ is complex conjugation $\zeta \mapsto \zeta^{-1}$; for $x \in K_n$ we set $|x|^2 := x c(x)$, an element of K_n^+ . The HAWK parameter sets HAWK-256, HAWK-512, and HAWK-1024 are $\ell \in \{9, 10, 11\}$, i.e. $n \in \{256, 512, 1024\}$; the names refer to n , as in the specification.

A HAWK secret key is $B \in \text{SL}_2(R_n)$; the public key is $Q = B^* B \in \text{GL}_2(R_n)$, which is Hermitian with $\det Q = 1$. Two keys B, B' are *equivalent* if $B' = UB$ for some $U \in \text{U}_2(R_n) := \{U \in \text{GL}_2(R_n) : U^* U = I\}$ (equivalently $B'^* B' = Q$). The **key-recovery problem** HAWK-KR(n) (smLIP for the isometry class of R_n^2) is: given Q , output any $B' \in \text{GL}_2(R_n)$ with $B'^* B' = Q$. Any such B' suffices to forge signatures [HAWK25, §5.1].

3.2 The three involutions

$\text{Gal}(K_n/\mathbb{Q}) \cong (\mathbb{Z}/2^\ell)^\times \cong \mathbb{Z}/2 \times \mathbb{Z}/2^{\ell-2}$ is non-cyclic for $\ell \geq 3$. Its 2-torsion subgroup has order 4, with three nontrivial commuting involutions:

$$c : \zeta \mapsto \zeta^{-1}, \quad \tau : \zeta \mapsto -\zeta, \quad \sigma := c\tau = \tau c : \zeta \mapsto -\zeta^{-1}.$$

Their fixed fields are $K_n^c = K_n^+$ (totally real), $K_n^\tau = \mathbb{Q}(\zeta^2) = \mathbb{Q}(\zeta_{2^{\ell-1}})$ (complex cyclotomic — the field underlying HAWK- $(n/2)$), and K_n^σ (a CM field of degree $n/2$, i.e. a totally imaginary quadratic extension of a totally real field). Set $F := K_n^{(c,\tau)} = K_n^+ \cap K_n^\tau$, totally real of degree $n/4$.

3.3 Lattice conventions

For a Euclidean lattice Λ , $\lambda_1(\Lambda)$ denotes the length of a shortest nonzero vector, $\Lambda^\vee := \{w \in \text{span}(\Lambda) : \langle w, \Lambda \rangle \subseteq \mathbb{Z}\}$ the dual lattice, and $[\Lambda]$ the isometry (LIP) class; between quadratic lattices, $\cong$ always means isometry, never mere $\mathbb{Z}$ -module isomorphism. For $s > 0$, $s \cdot \Lambda$ denotes Λ with all lengths scaled by s (equivalently, its Gram matrix scaled by s^2). We write $A \oplus B$ for the orthogonal direct sum of quadratic lattices, and call a lattice *hypercubic* if it is isometric to a scaling $s \cdot \mathbb{Z}^m$ of the integer lattice for some $s > 0$. Bambury and Nguyen [BN24] call a lattice *near-hypercubic* when its first minimum is almost as short as the integer lattice's, the operational hypothesis being $\lambda_1(\Lambda) \lambda_1(\Lambda^\vee) < 1 - \varepsilon$ for some $\varepsilon = 1/\text{poly}(\text{rank } \Lambda)$ [BN24, Thm 1]. The term names a broad family, and in this paper it always refers to the specific class arising from HAWK, the lattices isometric to a scaling of $\mathbb{Z}^p \oplus \sqrt{2} \mathbb{Z}^q$ (for which $\lambda_1(\Lambda) \lambda_1(\Lambda^\vee) = 1/\sqrt{2}$, for any $p, q \geq 1$ and any scaling).

We equip R_n with the *trace form* $\langle x, y \rangle_{K_n} := \text{Tr}_{K_n}(x c(y))$; the power-of-two trace formula (stated in the proof of Proposition 4.5) makes the power basis

(ζ^j) orthogonal for it, with all squared norms equal to n . The block-reduction notation (primitive sublattices, volumes, projected blocks) used by Algorithm 1, §9, and Appendices D–E is set up at the start of Appendix E.

4 The τ -Cocycle Lattice

The matrix $V_\tau := B^{-1}\tau(B)$ measures the failure of the secret basis to be τ -equivariant. Since $\det B = 1$, its determinant $(\det B)^{-1}\tau(\det B)$ equals 1, and applying τ once more gives the *cocycle equation* $V_\tau \tau(V_\tau) = I$: V_τ is the Galois 1-cocycle of B for the quadratic extension K_n/K_n^τ . On its own this equation pins down nothing: every $A \in \mathrm{GL}_2(K_n)$ yields a solution $A^{-1}\tau(A)$. Two further constraints, both publicly checkable, cut V_τ down to a shortest vector of a rank- n lattice.

4.1 Definition

For $Y \in M_2(K_n)$ write $\mathrm{adj} Y = \begin{pmatrix} Y_{22} & -Y_{12} \\ -Y_{21} & Y_{11} \end{pmatrix}$, so $Y \cdot \mathrm{adj} Y = (\det Y)I$. Define the τ -cocycle lattice of the public key Q , and the quadratic form on it, as

$$\Lambda_B^{(\tau)} := \left\{ Y \in M_2(R_n) : \underbrace{\tau(Y) = \mathrm{adj} Y}_{(C1)} \text{ and } \underbrace{QY = \sigma(Y)^T \tau(Q)}_{(C2)} \right\},$$

$$Q^{(\tau)}(Y) := \mathrm{Tr}_F(\det Y),$$

where Q is the public Gram matrix and $F = K_n^{(c,\tau)}$ the totally real field of degree $n/4$ from §3.2.

Constraint (C1) linearizes the cocycle equation $\tau(Y) = Y^{-1}$ by replacing Y^{-1} with $\mathrm{adj} Y$; constraint (C2) imposes the relation $QY = \sigma(Y)^T \tau(Q)$ between the public Gram matrices Q and $\tau(Q)$ with Y as the unknown.

Both constraints are $\mathbb{Q}$ -linear in the $4n$ rational coordinates of Y (with τ and σ acting entrywise): (C1) has constant coefficients and (C2) has coefficients determined by Q and $\tau(Q)$, both publicly computable. Hence $\Lambda_B^{(\tau)}$ is the integer kernel of a public linear system, and a $\mathbb{Z}$ -basis is computable in polynomial time via Hermite normal form.

The form $Q^{(\tau)}$ is $\mathbb{Q}$ -valued on $\Lambda_B^{(\tau)}$: for $Y \in \Lambda_B^{(\tau)}$, (C1) gives $\tau(\det Y) = \det(\mathrm{adj} Y) = \det Y$, and taking $\det$ of (C2) with $\det Q = \det \tau(Q) = 1$ gives $\sigma(\det Y) = \det Y$, so $\det Y \in K_n^\tau \cap K_n^\sigma = K_n^{(c,\tau)} = F$, hence $\det Y \in \mathcal{O}_F$. (Neither constraint alone suffices: (C1) places $\det Y$ in K_n^τ , (C2) in K_n^σ .) Shortness is measured with respect to $Q^{(\tau)}$ throughout. Explicitly, $\det Y$ is quadratic in the $4n$ rational coordinates of Y , so $Q^{(\tau)}$ is a quadratic form in the usual sense; the associated bilinear form is $(Y, Y') \mapsto \frac{1}{2} \mathrm{Tr}_F(\det(Y+Y') - \det Y - \det Y')$, and the Gram matrices of $\Lambda_B^{(\tau)}$ used elsewhere are taken with respect to it.

4.2 The secret cocycle and the secret-basis isometry

Lemma 4.1. *Let $V_\tau := B^{-1}\tau(B) \in \text{SL}_2(R_n)$. Then $V_\tau \in \Lambda_B^{(\tau)}$ and $\det V_\tau = 1$.*

Proof. $\det V_\tau = (\det B)^{-1} \cdot \tau(\det B) = 1$. **(C1):** $\tau(V_\tau) = \tau(B)^{-1}B = V_\tau^{-1} = \text{adj } V_\tau$ (since $\det V_\tau = 1$). **(C2):** $QV_\tau = B^*B \cdot B^{-1}\tau(B) = c(B)^T\tau(B)$, and

$$\begin{aligned} \sigma(V_\tau)^T \tau(Q) &= [c\tau(B^{-1}) \cdot c(B)]^T \cdot \tau(B^*B) \\ &= c(B)^T (c\tau)(B)^{-T} \cdot (c\tau)(B)^T \tau(B) \\ &= c(B)^T \tau(B). \quad \square \end{aligned}$$

The defining constraints involve Q , so the lattice itself depends on the key; the following (secret) isometry shows that its rank, minimum, and isometry class do not.

Lemma 4.2. *The map $\varphi_\tau : Y \mapsto BY\tau(B)^{-1}$ is a $\mathbb{Z}$ -linear bijection $M_2(R_n) \rightarrow M_2(R_n)$ carrying $\Lambda_B^{(\tau)}$ onto $\Lambda_I^{(\tau)}$ (the lattice of §4.1 for the trivial key $B = I$, i.e. $Q = I$) and preserving $\det$ (hence $Q^{(\tau)}$).*

Proof. $B, \tau(B) \in \text{SL}_2(R_n)$, so φ_τ is an R_n -bimodule automorphism. For $W = \varphi_\tau(Y)$: $\tau(W) = \tau(B)\tau(Y)B^{-1}$, and $\text{adj } W = \text{adj}(\tau(B)^{-1}) \cdot \text{adj } Y \cdot \text{adj } B = \tau(B) \cdot \text{adj } Y \cdot B^{-1}$ (using $\text{adj } X = X^{-1}$ for $X \in \text{SL}_2$ and $\text{adj}(XYZ) = \text{adj } Z \cdot \text{adj } Y \cdot \text{adj } X$). These agree iff (C1) holds for Y . Rewriting (C2) as $QY\tau(Q)^{-1} = \sigma(Y)^T$ and substituting $Q = c(B)^T B$, $\tau(Q)^{-1} = \tau(B)^{-1}(c\tau)(B)^{-T}$ gives: (C2) $\Leftrightarrow W = \sigma(W)^T$, which is (C2) at $Q = I$. Finally $\det W = \det B \cdot \det Y \cdot \tau(\det B)^{-1} = \det Y$. $\square$

The attacker does not know φ_τ (it depends on B). Its *existence* shows that $(\Lambda_B^{(\tau)}, Q^{(\tau)})$ and $(\Lambda_I^{(\tau)}, Q^{(\tau)})$ are isometric as quadratic $\mathbb{Z}$ -lattices, so rank, minimum, shortest-vector count, and isometry class can be read off from the $Q = I$ case.

4.3 Rank and explicit parametrization

At $B = I$ the two constraints decouple entrywise, yielding an explicit parametrization and the rank. Write $R_n^\sigma := R_n \cap K_n^\sigma$ for the fixed ring of σ and $R_n^{\tau=-1} := \{x \in R_n : \tau(x) = -x\}$ for the (-1) -eigenspace of τ . We call the intersection of R_n with an eigenspace of an involution an *eigenlattice*.

Lemma 4.3. $\text{rank}_{\mathbb{Z}} \Lambda_B^{(\tau)} = n$, and $\Lambda_I^{(\tau)} = \left\{ \begin{pmatrix} a & b \\ \sigma(b) & \tau(a) \end{pmatrix} : a \in R_n^\sigma, b \in R_n^{\tau=-1} \right\}$.

Proof. By Lemma 4.2 it suffices to parametrize $\Lambda_I^{(\tau)} = \{W \in M_2(R_n) : \tau(W) = \text{adj } W, W = \sigma(W)^T\}$. Writing $W = \begin{pmatrix} a & b \\ b' & a' \end{pmatrix}$, constraint (C1) gives $\tau(a) = a'$, $\tau(b) = -b$, $\tau(b') = -b'$ entrywise, and constraint (C2) gives $a = \sigma(a)$, $a' = \sigma(a')$, $b' = \sigma(b)$. Thus $a \in R_n^\sigma$ is free ($n/2$ $\mathbb{Z}$ -dimensions since $[K_n : K_n^\sigma] = 2$) and $a' = \tau(a)$ is determined, consistently since $\sigma(a') = \sigma\tau(a) = \tau\sigma(a) = a'$. Similarly $b \in R_n^{\tau=-1}$ is free ($n/2$ $\mathbb{Z}$ -dimensions, basis $\{\zeta^{2j+1} : 0 \leq j < n/2\}$) and

$b' = \sigma(b)$ is determined, consistently since $\tau(b') = \tau\sigma(b) = \sigma\tau(b) = -\sigma(b) = -b'$. The total rank is n . $\square$

This computation is where the hypothesis $\tau \neq c$ is used: for $\tau = c$ the composite $\sigma = c\tau$ is the identity, constraint (C2) degenerates, and the analogous lattice has rank $3n/2$ (§8.3, Appendix C).

We write $W(a, b) := \begin{pmatrix} a & b \\ \sigma(b) & \tau(a) \end{pmatrix}$ for the general element of $\Lambda_I^{(\tau)}$, so that $\Lambda_I^{(\tau)} = \{W(a, b) : a \in R_n^\sigma, b \in R_n^{\tau=-1}\}$. Its determinant $a\tau(a) - b\sigma(b)$ simplifies: since $\sigma(a) = a$, $\tau(a) = \tau\sigma(a) = c(a)$ and $a\tau(a) = |a|^2$; since $\tau(b) = -b$, $\sigma(b) = c\tau(b) = -c(b)$ and $-b\sigma(b) = |b|^2$. Hence $\det W(a, b) = |a|^2 + |b|^2$, with no cross term between a and b .

4.4 Positive-definiteness and shortest vectors

Lemma 4.4. (i) $Q^{(\tau)}$ is positive-definite on $\Lambda_B^{(\tau)}$; (ii) $\min_{Y \neq 0} Q^{(\tau)}(Y) = n/4$, attained exactly at the $2(n/2 + 1)$ vectors Y with $Y\tau(Y) = I$ (equivalently $\det Y = 1$).

Proof. By Lemma 4.2 work in $\Lambda_I^{(\tau)}$. By §4.3, $\det W(a, b) = |a|^2 + |b|^2$, a sum of two elements of the form $x c(x)$, totally ≥ 0 in F , and $= 0$ iff $W = 0$. Therefore $Q^{(\tau)}(W) = \text{Tr}_F(|a|^2 + |b|^2) > 0$ for $W \neq 0$.

By the AM–GM inequality over the $[F : \mathbb{Q}] = n/4$ real embeddings, together with $N_F(\det W) \in \mathbb{Z}_{\geq 1}$ (nonzero algebraic integer), we have $Q^{(\tau)}(W) \geq (n/4) N_F(\det W)^{4/n} \geq n/4$. Equality throughout holds iff $\det W = 1$: equality in the AM–GM forces $\det W$ to be constant across the real embeddings of F , and a totally positive constant of integer norm 1 is 1. Conversely every Y with $\det Y = 1$ has $Q^{(\tau)}(Y) = \text{Tr}_F(1) = n/4$. Hence the minimum is $n/4$, attained exactly at the Y with $\det Y = 1$, equivalently (since $Y \cdot \text{adj } Y = (\det Y)I$ and $\text{adj } Y = \tau(Y)$ by (C1)) at the Y with $Y\tau(Y) = I$.

For the count, a shortest vector is a $W(a, b)$ with $|a|^2 + |b|^2 = 1$, both summands totally non-negative algebraic integers in $\mathcal{O}_F$: if both were nonzero, each would lie in $(0, 1)$ at every real embedding, hence have norm in $(0, 1)$, which is impossible. So one summand vanishes, and the nonzero one of a, b is an algebraic integer all of whose archimedean absolute values equal 1, hence a root of unity by Kronecker’s theorem. Either $b = 0$ and a is a root of unity of K_n^σ , i.e. $a = \pm 1$ (since $\sigma(\zeta^k) = \zeta^k$ forces $\zeta^{2k} = (-1)^k$, i.e. $k \in \{0, n\}$), giving the 2 vectors $\pm I$; or $a = 0$ and b is a root of unity of K_n with $\tau(b) = -b$, i.e. $b \in \{\zeta^k : k \text{ odd}\}$, giving n vectors. Total $2(n/2 + 1)$. $\square$

4.5 The isometry class

By Lemma 4.2 the isometry class of $(\Lambda_B^{(\tau)}, Q^{(\tau)})$ is that of the trivial-key lattice:

Proposition 4.5. $(\Lambda_I^{(\tau)}, Q^{(\tau)}) \cong \sqrt{n/4} \mathbb{Z}^{n/2+1} \oplus \sqrt{n/2} \mathbb{Z}^{n/2-1}$.

Proof. By §4.3, $\Lambda_I^{(\tau)} = \{W(a, b) : a \in R_n^\sigma, b \in R_n^{\tau=-1}\}$ with $\det W(a, b) = |a|^2 + |b|^2$, so the two summands are $Q^{(\tau)}$ -orthogonal and $Q^{(\tau)}(W(a, b)) = \text{Tr}_F(|a|^2) + \text{Tr}_F(|b|^2) = \frac{1}{4}\langle a, a \rangle_{K_n} + \frac{1}{4}\langle b, b \rangle_{K_n}$, where $\langle \cdot, \cdot \rangle_{K_n}$ is the trace form of §3.3 and we used $|a|^2, |b|^2 \in F$ and $[K_n : F] = 4$. It remains to compute the restriction of $\langle \cdot, \cdot \rangle_{K_n}$ to the eigenlattices $R_n^{\tau=-1} = \ker(\tau+1)|_{R_n}$ and $R_n^\sigma = \ker(\sigma-1)|_{R_n}$.

By the power-of-two trace formula ($\text{Tr}_{K_n}(\zeta^m)$ equals n if $m \equiv 0$, $-n$ if $m \equiv n$, else $0 \pmod{2^\ell}$), the power basis $(\zeta^j)_{0 \leq j < n}$ is $\langle \cdot, \cdot \rangle_{K_n}$ -orthogonal with all squared norms equal to n : $(R_n, \langle \cdot, \cdot \rangle_{K_n}) \cong \sqrt{n} \mathbb{Z}^n$. Every $\gamma \in \text{Gal}(K_n/\mathbb{Q})$ is an isometry of this lattice ($\text{Tr}_{K_n}(\gamma(x)\gamma(y)) = \text{Tr}_{K_n}(\gamma(x\bar{y})) = \text{Tr}_{K_n}(x\bar{y})$) and maps each ζ^j to $\pm \zeta^{j'}$, i.e. acts on the orthogonal basis by a signed permutation. For an involution the cycles have length 1 or 2, and the ε -eigenlattice ($\varepsilon = \pm 1$) decomposes orthogonally over the cycles: a basis vector fixed with sign ε contributes itself (squared norm n); a basis vector fixed with sign $-\varepsilon$ contributes nothing; a 2-cycle $\zeta^j \mapsto s\zeta^{j'} \mapsto \zeta^j$ contributes the single vector $\zeta^j + \varepsilon s\zeta^{j'}$ (squared norm $2n$).

For τ : $\tau(\zeta^j) = (-1)^j \zeta^j$ — every basis vector is fixed up to sign, there are no 2-cycles, and the (-1) -eigenlattice is the span of the odd powers: $R_n^{\tau=-1} = \bigoplus_{j \text{ odd}} \mathbb{Z}\zeta^j \cong \sqrt{n} \mathbb{Z}^{n/2}$.

For σ : $\sigma(\zeta^j) = (-1)^j \zeta^{-j} = (-1)^{j+1} \zeta^{n-j}$ for $1 \leq j \leq n-1$ (using $\zeta^{-j} = -\zeta^{n-j}$), and $\sigma(1) = 1$. The underlying permutation $j \mapsto n-j$ has two fixed points, $j = 0$ (sign $+1$; this is the case $\sigma(1) = 1$ treated separately above) and $j = n/2$ (sign $(-1)^{n/2+1} = -1$ since $n/2$ is even), and $n/2 - 1$ two-cycles $\{j, n-j\}$; on the rank-2 sublattice $\mathbb{Z}\zeta^j \oplus \mathbb{Z}\zeta^{n-j}$ the involution acts by $\begin{pmatrix} 0 & s \\ s & 0 \end{pmatrix}$ with $s = (-1)^{j+1}$, whose integer fixed vectors are exactly $\mathbb{Z}(\zeta^j + s\zeta^{n-j})$. Hence $R_n^\sigma = \mathbb{Z} \cdot 1 \oplus \bigoplus_{j=1}^{n/2-1} \mathbb{Z}(\zeta^j + (-1)^{j+1} \zeta^{n-j}) \cong \sqrt{n} \mathbb{Z} \oplus \sqrt{2n} \mathbb{Z}^{n/2-1}$.

Dividing all squared norms by 4 ($n/4$ and $n/2$) and combining the two $Q^{(\tau)}$ -orthogonal summands gives the claim. The $n/2 + 1$ short basis vectors (squared norm $n/4$) are $W(1, 0) = I$ and $W(0, \zeta^j)$ for j odd. $\square$

Corollary 4.6. *After rescaling by $\sqrt{4/n}$ (i.e. $2/\sqrt{n}$), the public lattice $\Lambda_B^{(\tau)}$ is in the LIP class $[A']$ [DvW21] of the near-hypercubic lattice*

$$A' := \mathbb{Z}^p \oplus \sqrt{2} \mathbb{Z}^q, \quad p := n/2 + 1, \quad q := n/2 - 1, \quad p + q = n.$$

The Gram representative $D = \text{diag}(I_p, 2I_q)$ is integral; the dual class is $[\mathbb{Z}^p \oplus \frac{1}{\sqrt{2}} \mathbb{Z}^q]$. (The integers p, q fixed here are used throughout §5 and Appendix E; they are not to be confused with the odd prime p of §7.)

Remark 4.7. The proof of Proposition 4.5 uses $\text{Tr}_{K_n}(\zeta^m) = 0$ for all $0 < |m| < n$. For non-power-of-two conductors this fails (e.g., $\text{Tr}_{\mathbb{Q}(\zeta_p)/\mathbb{Q}}(\zeta_p^m) = -1$ for an odd prime p and all $m \not\equiv 0$), and the analogous lattice need not be isometric to any $a\mathbb{Z}^{p'} \oplus b\mathbb{Z}^{q'}$, $0 < a < b$; see §7.

5 Recovering the Key from $\Lambda_B^{(\tau)}$

[Corollary 4.6](#) places the rescaled cocycle lattice in the LIP class $[A'] = [\mathbb{Z}^p \oplus \sqrt{2} \mathbb{Z}^q]$, $p = n/2 + 1$, $q = n/2 - 1$. We run Ducas’s two-block reduction [[Duc23](#), Alg. 1] with cut $k := p - 1$:

Algorithm 1 Ducas’s two-block reduction [[Duc23](#), Alg. 1] on $[A']$, cut $k = p - 1$. The reduction operations are defined in §E.4.

Require: integral Gram matrix of a basis $B = (b_0, \dots, b_{p+q-1})$ of $A \in [A']$ (this B is unrelated to the secret key of §3.1)

Ensure: $U \in \text{GL}_{p+q}(\mathbb{Z})$ with $U^T(\text{input})U = \text{diag}(I_p, 2I_q)$ up to signed permutation

- 1: **while** $\text{vol}(B_{[0:k-1]})^2 > 1$ **do**
 - 2: dual-SVP reduce $B_{[0:k]}$; primal-SVP reduce $B_{[k:p+q-1]}$
 - 3: HKZ-reduce $B_{[0:k-1]}$; HKZ-reduce $B_{[k:p+q-1]}$; size-reduce B
-

Ducas’s progress lemma [[Duc23](#), Lemma 4] fails on $[A']$ (§E.1); Appendix E adapts the product-progress framework of [[BN24](#), Lemma 1, Thm 1] to $[A']$, supplying the no-abort bound ([Lemma E.2](#); see [Remark E.5](#)), and proves:

Theorem 5.1. *On input an integral Gram matrix H of a basis of a lattice in $[A']$, Algorithm 1 terminates after at most $p^2 \ln p + p \ln(2(p-1) \ln \max_i \|b_i\|) + 3p$ iterations (any basis of $[A']$ has $\max_i \|b_i\| \geq \sqrt{2}$) and outputs $U \in \text{GL}_{p+q}(\mathbb{Z})$ with $U^T H U = \text{diag}(I_p, 2I_q)$ up to signed permutation, using SVP-oracle calls in dimension $\leq p = n/2 + 1$ only.*

Proof: Appendix E.

Algorithm 1 returns all $2(n/2 + 1)$ shortest vectors of $\Lambda_B^{(\tau)}$. For each, $x \mapsto Y \tau(x)$ is an order-2 automorphism of the public form (R_n^2, T) outside the trivial group of [[GP25](#)] ([Lemma D.1](#)), so the descent of van Gent and Pulles [[GP25](#), Props. 1–2] applies: a short vector of the kernel sublattice $\mathcal{M}_Y$ recovers an equivalent key.

Appendix D sharpens the analysis to an unconditional bound: the parity test $Y \equiv I \pmod{2}$ selects the two candidates $Y = \pm V_\tau$, for which the kernel is exactly $\sqrt{n} \mathbb{Z}^n$ ([Proposition D.2](#)), so one further run of Ducas’s algorithm [[Duc23](#)] on it returns the key with SVP-oracle calls in dimension at most $n/2 + 1$.

6 Main Theorem and Security Consequences

Theorem 6.1. *For every power of two $n \geq 4$ there is a deterministic algorithm that, given a HAWK- n public key $Q = B^* B$ with $B \in \text{SL}_2(R_n)$, outputs $B' \in \text{GL}_2(R_n)$ with $B'^* B' = Q$, making $\text{poly}(n)$ arithmetic operations and $\text{poly}(n)$ calls to an exact-SVP oracle in dimension at most $n/2 + 1$.*

Proof.

1. Compute a $\mathbb{Z}$ -basis and Gram matrix $G^{(\tau)}$ of $(\Lambda_B^{(\tau)}, Q^{(\tau)})$ from Q . By [Lemma 4.2](#) and [Proposition 4.5](#), $\frac{4}{n}G^{(\tau)}$ is an integral Gram matrix of a lattice in $[\Lambda']$.
2. LLL-reduce $\frac{4}{n}G^{(\tau)}$ and apply [Algorithm 1 \(Theorem 5.1\)](#): the p output basis vectors of minimal norm, pulled back to $M_2(R_n)$ and taken with both signs, are the $2(n/2 + 1)$ shortest vectors of $\Lambda_B^{(\tau)}$. This uses $\text{poly}(n) \times \text{SVP}(n/2 + 1)$.
3. For each of the two Y with $Y \equiv I \pmod{2}$ (exactly $\pm V_\tau$, by the parity test of [Appendix D](#)): compute a $\mathbb{Z}$ -basis and T -Gram of $\mathcal{M}_Y$, $T(x, y) := \text{Tr}_{K_n}(x^* Q y)$, by Hermite normal form on the constraint $\tau(x) = Y^{-1}x$ (with $Y^{-1} = \text{adj } Y \in M_2(R_n)$ since $\det Y = 1$); LLL-reduce $\frac{1}{n}T$ -Gram and solve $\mathbb{Z}^n$ -LIP on it via [\[Duc23\]](#), using $\text{poly}(n) \times \text{SVP}(n/2 + 1)$ ([Appendix D](#)). For each pair (u_1, u_2) from the resulting shortest-vector set with $u_1^* Q u_2 = 0$ (at most $\binom{n}{2}$ pairs up to sign), test whether $B' := [u_1 \ u_2]^{-1}$ satisfies $B' \in \text{GL}_2(R_n)$ and $B'^* B' = Q$; output the first such B' .

By [Proposition D.2](#), step (3) succeeds for $Y = \pm V_\tau$. $\square$

6.1 Cost estimates

[Table 1](#) compares the per-call cost of the SVP oracle at the designers' blocksize β_{key} and at the oracle dimension $n/2 + 1$ of [Theorem 6.1](#), under both the Core-SVP heuristic ($2^{0.292\beta}$ [\[BDGL16\]](#)) and the gate-count model of Albrecht, Gheorghiu, Postlethwaite, and Schanck [\[AGPS20\]](#). Asymptotically the improvement is a factor of two in the oracle dimension: $\beta_{\text{key}}/n \rightarrow 1$ under the model of [\[HAWK25, §5.1\]](#) while ours is $n/2 + 1$. With the provable sieve of [\[ADRS15\]](#) as the oracle this gives time $2^{(n/2+1)+o(n)}$ unconditionally.

Table 1. Cost of one SVP-oracle call. β is the oracle dimension: β_{key} [\[HAWK25, Table 7\]](#) for the spec, $n/2 + 1$ for the attack. *Core-SVP is our conversion $2^{0.292\beta}$ (exponent rounded to the nearest integer); the specification reports blocksizes and AGPS20 gate counts only. AGPS20 gates are the classical list-decoding gate counts of [\[AGPS20\]](#) at sieve dimension $\beta_{\text{eff}} = \beta - \text{d4f}(\beta)$, with dimensions-for-free [\[Duc18\]](#) computed and rounded as in [\[HAWK25, §5.3\]](#). [†]HAWK-256 is a challenge parameter set [\[HAWK25, §3.2\]](#), for which the specification states no gate count (the entry 2^{74} at β_{key} is our computation); [§6.2](#) reports actual recoveries.

	n	β		Core-SVP*		AGPS20 gates	
		Spec	Attack	Spec	Attack	Spec	Attack
HAWK-256 [†]	256	211	129	2^{62}	2^{38}	2^{74}	2^{52}
HAWK-512	512	452	257	2^{132}	2^{75}	2^{141}	2^{86}
HAWK-1024	1024	940	513	2^{274}	2^{150}	2^{278}	2^{158}

The HAWK specification reports total key-recovery costs of 2^{150} and 2^{288} gates for HAWK-512 and HAWK-1024 [\[HAWK25, §5.3\]](#), computed as $n \cdot C_{\text{gate}}(\beta_{\text{eff}})$:

C_{gate} is the AGPS20 gate count of one sieve call [AGPS20], $\beta_{\text{eff}} = \beta_{\text{key}} - \text{d4f}(\beta_{\text{key}})$, and n is the spec’s lower bound on the number of oracle calls in one BKZ tour. Under the same model at our dimension $n/2 + 1$ (so $\beta_{\text{eff}} = 224$ and 464 after dimensions-for-free), the per-call costs are $2^{86.0}$ and $2^{157.8}$; multiplying by the call count below gives totals of at most 2^{108} and 2^{182} .

The reduction makes $O(n^2 \log n)$ oracle calls in at most three reduction runs: one Algorithm 1 run in step (2) and at most two $\mathbb{Z}^n$ -LIP runs in step (3), each preceded by LLL. After LLL every basis vector has squared norm at most 2^n , since $\|b_i\|^2 \leq 2^{D-1} \lambda_i(\Lambda)^2$ for rank D [LLL82, p. 518] and the successive minima λ_i are key-independent (at most $\sqrt{2}$ and 1 for steps (2) and (3) respectively, by Propositions 4.5 and D.2). At this input size, the iteration bound of Theorem 5.1 is $(n/2+1)^2 \ln(n/2+1) + O(n \log n)$ for step (2), and that of Remark D.3 is $(n+1)(n/2+1) \ln(n/2+1) + O(n \log n)$ per step-(3) run; each iteration makes two oracle calls, and the terminal HKZ reductions add $O(n)$. The totals come to at most $2^{21.9}$ and $2^{24.0}$ calls for HAWK-512 and HAWK-1024, with the input basis entering only through a doubly-logarithmic term contributing under 1%.

Only the SVP-oracle calls are priced, on both sides. The polynomial-time arithmetic outside the oracle (HNF, LLL, Gram updates, the parity tests of step (3)) is omitted, as is the non-oracle cost of a BKZ tour in the specification’s count. Neither side prices memory access, the progressivity factor of roughly 2^5 [SAB22, §5.2.1], or the overheads of [Duc22]; the latter two are explicitly set aside in [HAWK25, §5.3].

A heuristic estimate in the specification’s own methodology goes further: under the GSA-intersect model of Appendix C, progressive BKZ reaches the minimum of the rescaled cocycle lattice — a τ -cocycle, by Lemma 4.4 — at blocksize 205 (HAWK-512) and 432 (HAWK-1024) rather than the proven oracle dimension $n/2 + 1$, giving heuristic key-recovery costs of about $2^{80.8}$ and $2^{146.5}$ gates against the specification’s 2^{150} and 2^{288} . The provable totals above are thus conservative by roughly a further 2^{27} and 2^{35} ; Appendix C gives the conversion and its caveats.

6.2 Implementation and key recovery for HAWK-256

We implemented the attack as a released end-to-end program [code] and recovered the secret keys of two HAWK-256 public keys generated by the reference key generator [HAWK25]. The implementation follows steps (1)–(3) of Theorem 6.1, replacing the provable LIP solvers with heuristic lattice reduction and taking two shortcuts in step (3); the theorem’s cost bound does not depend on either.

Step (1) constructs the 256×256 integer Gram matrix of $\Lambda_B^{(\tau)}$ via an integer Hermite-normal-form kernel computation followed by LLL reduction of the resulting basis.

Step (2) BKZ-reduces this Gram matrix with fpylll to blocksize 44 and runs a progressive BGJ-type sieve [BGJ15], using the AMX implementation of Zhao, Ding, and Yang [ZDY25], to a maximum sieving dimension of 118. The sieve

returns two shortest vectors Y, Y' of $\Lambda_B^{(\tau)}$ not related by sign. A sieve database contains some shortest vectors long before it provably contains all of them; [Remark 6.2](#) is what makes two enough.

Remark 6.2 (intertwiner sublattice). Let Y, Y' be shortest vectors of $\Lambda_B^{(\tau)}$ with $Y \neq \pm Y'$ and neither equal to $\pm V_\tau$, and set $M := Y \cdot \text{adj}(Y')$. The rank- $n/2$ sublattice

$$\Lambda_a := \{X \in \Lambda_B^{(\tau)} : MX = X\tau(M)\}$$

is cut out by one further public linear system (its coefficients involve only Q, Y, Y'), and its two shortest vectors are $\pm V_\tau$ (proof in Appendix D). The hypothesis that neither sieve output equals $\pm V_\tau$ is testable by the parity test of Appendix D; when it fails the construction is unnecessary.

The implementation locates $\pm V_\tau$ directly from Λ_a : it computes an explicit $\mathbb{Z}$ -basis of this rank-128 lattice and runs LLL followed by progressive BKZ; blocksize at most 24 sufficed in every run. It then reaches an equivalent key by a τ -tower descent from $\ell = 9$ down to $\ell = 3$ (Appendix D, after [Lemma D.4](#)) in place of the single $\mathbb{Z}^n$ -LIP solve of step (3); each level re-runs steps (1)–(2) at half the degree, and the resulting rank-128 cocycle lattice at the $\ell = 8$ level was solved at BKZ blocksize at most 16, every lower level by LLL alone. We make no provable claim for the tower descent; the provable path remains the single $\mathbb{Z}^n$ -LIP solve.

Every attempt succeeded. For one of the two keys, two independent runs each produced a solution: distinct but unitarily-equivalent secret keys, both passing the reference implementation’s sign/verify check. A recovery completes within a few hours on a single 96-core Sapphire Rapids server, dominated by the step-(2) sieve. The release is a single-command pipeline with per-stage checkpoints that verifies its output against the HAWK reference implementation as its final step; Appendix A lists the lemma checks and the (heuristic) end-to-end runs at small ℓ . For HAWK-512 the proven oracle dimension is 257; we have not attempted it.

7 Scope of the Attack

The attack requires two ingredients: a τ -cocycle lattice (§4), and a half-dimension reduction for its isometry class (§5, Appendix E). The first is settled by:

Theorem 7.1. *An order-2 Galois element $\tau \neq c$ of $\mathbb{Q}(\zeta_m)$ — and hence the construction of §4 — exists iff $(\mathbb{Z}/m)^\times$ is non-cyclic, iff $m \notin \{1, 2, 4, p^k, 2p^k\}$ for p an odd prime.*

Proof. §4 requires an order-2 Galois element $\tau \neq c$; by [Lemma B.1](#) one exists iff $(\mathbb{Z}/m)^\times$ is non-cyclic. $\square$

For the second, [Theorem 6.1](#) is proven only for powers of two, but for every non-cyclic m the cocycle lattice has V_τ as a shortest vector with minimum $\varphi(m)/4$ (the proof of [Lemma 4.4](#) uses only $\tau^2 = 1, \tau \neq c$), and we expect heuristic lattice

reduction to find it at blocksize somewhat below $\varphi(m)/2$ (roughly 0.36–0.42 $\varphi(m)$ at the power-of-two parameter sets under the model of Appendix C) whenever m has a small prime factor: preliminary computations at small parameters for $m = 4p$ and $m = 3 \cdot 2^\ell$ suggest the cocycle lattice satisfies the hypothesis of [BN24, Thm 1] with the same constant $\lambda_1 \lambda_1^\vee \rightarrow 1/\sqrt{2}$ as the power-of-two case.

Conductors $m \in \{p^k, 2p^k\}$ (and the degenerate $m \leq 4$) evade the construction. Over those fields the rank- $3\varphi(m)/2$ c -cocycle of §8.3 still applies, with blocksize roughly $0.72 \beta_{\text{key}}$ under the same model (Appendix C).

8 Relation to Prior Work

8.1 Subfield attacks

The NTRU subfield attack of Albrecht, Bai, and Ducas [ABD16] norms the public key $h \mapsto N_{K_n/E}(h)$ to a subfield E ; the subring variant of Kirchner and Fouque [KF17] instead restricts the multiplication-by- h matrix to a subring of the same index. Kirchner and Fouque prove a complexity bound [KF17, Thm 7] that depends on the secret width only through $n \log \sigma_{\text{sk}}$ (where σ_{sk} is the width of the NTRU secret), a quantity they observe is invariant under subfield descent [KF17, §4.2], so descent gains nothing at non-overstretched parameters. By contrast, the cocycle V_τ of §4 has $\det V_\tau = 1$ and $Q^{(\tau)}(V_\tau) = n/4$ regardless of the key-generation distribution (Lemma 4.4), so the conservation argument has no analogue; the construction of §4 involves neither a modulus nor a subfield descent.

8.2 Gentry–Szydlo and quaternion methods

Gentry and Szydlo [GS02, §8] propose extending their algorithm to recover f from a basis of (f) and the relative norm of f over an index-2 subfield only “where the degree-2 extension is complex conjugation”. Every subsequent module-LIP reduction we are aware of — to a relative- or reduced-norm principal-ideal problem [MPPW24; Che25; APvW25], via the symplectic automorphism [LJPW24], or via the Hermitian square-basis oracle of [EP24] — likewise uses no Galois involution other than complex conjugation. The construction of this paper instead uses the involution τ , whose fixed field $K_n^\tau = \mathbb{Q}(\zeta_{2^\ell-1})$ is complex cyclotomic; the resulting public lattice has rank n rather than $2n$, and the reduction target is SVP in dimension $n/2 + 1$ rather than a principal-ideal problem.

8.3 Comparison with the c -cocycle

Chevignard et al. [Che25, §1] construct a rank- $2n$ lattice from the public key whose shortest vector is the secret (their d is our n); for HAWK it is hypercubic, and they show a single SVP call in dimension $2n$ breaks the scheme. Replacing

τ by c throughout §4 gives a related c -cocycle lattice of rank $3n/2$: constraint (C2) degenerates to “ QY symmetric” (since $\sigma = c^2 = \text{id}$), so it is not the lattice of [Che25], which imposes neither (C1) nor (C2). For $\tau \neq c$, $\sigma = c\tau \neq \text{id}$ is itself an order-2 automorphism and (C2) is an independent $n/2$ -dimensional restriction; this rank saving is the source of the factor-2 blocksize improvement. The third involution σ gives the same rank n , since the proofs use only $\tau^2 = 1$, $c\tau \neq 1$, $[K_n : K_n^\tau] = 2$, and the power-of-two trace orthogonality, all of which hold with $\tau \leftrightarrow \sigma$. Higher-order Galois elements admit no adj-linearized analogue of (C1), since the linearization of $\gamma(Y) = Y^{-1}$ requires $\gamma^2 = \text{id}$; we found no useful replacement.

8.4 The reductions of Ducas and of Bambury–Nguyen

Ducas [Duc23] solves $\mathbb{Z}^d$ -LIP with polynomially many SVP-oracle calls in dimension $\lfloor d/2 \rfloor + 1$, by a two-block reduction whose progress lemma [Duc23, Lemma 4] bounds the per-round drop of an integer potential. Bambury and Nguyen [BN24] extend the early-abort form of the reduction to any sublattice of $\mathbb{Z}^D$ with $\lambda_1(A)\lambda_1(A^\vee) < 1 - \varepsilon$, returning one short primal or dual vector at the same oracle dimension. As noted in §5, we run Ducas’s algorithm on the class $[A']$, where his progress lemma fails and Appendix E supplies the replacement; Remark E.5 compares the two routes in detail.

8.5 Lattice automorphisms and van Gent–Pulles

Van Gent and Pulles [GP25] show that, given any $\mathbb{Z}$ -automorphism of the public form outside the trivial group generated by the roots of unity and the symplectic automorphism of [LJPW24], smLIP reduces to BKZ on a kernel sublattice of half the rank: their Propositions 1–2 construct the kernel and recover the key from a short vector of it, deterministically and in polynomial time, and their Theorem 1 estimates the BKZ cost heuristically. §5 and Appendix D run their descent on the automorphisms $x \mapsto Y\tau(x)$ that the shortest vectors of the cocycle lattice supply (Lemma D.1); Appendix D identifies the kernel exactly as $\sqrt{n}\mathbb{Z}^n$, which lets Ducas’s algorithm replace the heuristic BKZ step.

9 Discussion and Open Questions

One attempt to halve the oracle dimension again, to $n/4 + 1$, by running Algorithm 1 at cut $\lfloor p/2 \rfloor$, is blocked: the loop reaches the non-terminal fixed point $S = \mathbb{Z}^{\lfloor p/2 \rfloor} \subset \mathbb{Z}^p$ (where $\text{vol}(S) = 1$ and the product-progress quantity of Lemma E.2 equals 1), and for $n > 8$ the residual back-block LIP still requires $\text{SVP}(n/2 - 1)$ by duality and Corollary E.4. As a generic $[A']$ -LIP problem this would require a gap- $\sqrt{2}$ Rankin solver below half-dimension, for which no technique is known.

The HAWK instance carries structure a generic member of $[A']$ lacks: $A_B^{(\tau)}$ is publicly a rank-4 quadratic $\mathcal{O}_F$ -lattice (for $f \in \mathcal{O}_F$, both (C1) and (C2) commute

with multiplication by f , and $\varphi_\tau(fY) = f\varphi_\tau(Y)$), with $\mathcal{O}_F$ -valued form $\det Y$. Since F is totally real its only roots of unity are ± 1 , which appears to rule out the known structured-sieve speedups; we leave open whether this rank-4 module structure admits a solver with SVP-oracle dimension below $n/2$.

10 Acknowledgements and Disclosures

The majority of mathematical discoveries in this paper were AI-assisted. Human author contribution mainly consisted of directing, organizing and verifying AI work. The write-up was partially AI-assisted but was reviewed and edited by the authors.

We thank Thomas Pornin and the rest of the HAWK team for reviewing the work, and especially for feedback on proper attribution for some of our attack’s components — something AI-assisted work makes more difficult.

References

- [ABD16] M. R. Albrecht, S. Bai, and L. Ducas. “A subfield lattice attack on overstretched NTRU assumptions: Cryptanalysis of some FHE and graded encoding schemes”. In: *CRYPTO 2016, Part I*. Ed. by M. Robshaw and J. Katz. Vol. 9814. LNCS. Springer, 2016, pp. 153–178. URL: https://doi.org/10.1007/978-3-662-53018-4_6.
- [ADRS15] D. Aggarwal, D. Dadush, O. Regev, and N. Stephens-Davidowitz. “Solving the shortest vector problem in 2^n time using discrete Gaussian sampling”. In: *STOC 2015*. ACM, 2015, pp. 733–742.
- [AGPS20] M. R. Albrecht, V. Gheorghiu, E. W. Postlethwaite, and J. M. Schanck. “Estimating quantum speedups for lattice sieves”. In: *ASIACRYPT 2020*. Vol. 12492. LNCS. Springer, 2020, pp. 583–613.
- [Ala26] G. Alagic, M. Bros, P. Ciadoux, Q. Dang, T. H. Dang, J. Kelsey, J. Lichtinger, Y.-K. Liu, C. Miller, D. Moody, R. Peralta, R. Perlner, A. Robinson, H. Silberg, D. Smith-Tone, and N. Waller. *Status Report on the Second Round of the Additional Digital Signature Schemes for the NIST Post-Quantum Cryptography Standardization Process*. NIST Internal Report 8610. NIST, 2026. URL: <https://doi.org/10.6028/NIST.IR.8610>.
- [APvW25] B. Allombert, A. Pellet-Mary, and W. van Woerden. “Cryptanalysis of rank-2 module-LIP: A single real embedding is all it takes”. In: *EUROCRYPT 2025*. Vol. 15602. LNCS. Springer, 2025, pp. 184–212. URL: https://doi.org/10.1007/978-3-031-91124-8_7.
- [BDGL16] A. Becker, L. Ducas, N. Gama, and T. Laarhoven. “New directions in nearest neighbor searching with applications to lattice sieving”. In: *SODA 2016*. SIAM, 2016, pp. 10–24.

- [BGJ15] A. Becker, N. Gama, and A. Joux. *Speeding-up lattice sieving without increasing the memory, using sub-quadratic nearest neighbor search*. Cryptology ePrint Archive, Report 2015/522. 2015. URL: <https://eprint.iacr.org/2015/522>.
- [BN24] H. Bambury and P. Q. Nguyen. “Improved provable reduction of NTRU and hypercubic lattices”. In: *Post-Quantum Cryptography – PQCrypto 2024*. Vol. 14772. LNCS. Springer, 2024, pp. 343–370.
- [Che25] C. Cheviguard, G. Mureau, T. Espitau, A. Pellet-Mary, H. Pliatsok, and A. Wallet. *A reduction from Hawk to the principal ideal problem in a quaternion algebra*. EUROCRYPT 2025. 2025.
- [code] hawk_recover. *hawk_recover: attack implementation and lemma-verification scripts for §6.2*. github.com/anthropics/cryptography-research-demo.
- [DPPvW22] L. Ducas, E. W. Postlethwaite, L. N. Pulles, and W. van Woerden. “HAWK: Module LIP makes lattice signatures fast, compact and simple”. In: *Advances in Cryptology – ASIACRYPT 2022, Part IV*. Ed. by S. Agrawal and D. Lin. Vol. 13794. LNCS. Springer, 2022, pp. 65–94. URL: https://doi.org/10.1007/978-3-031-22972-5_3.
- [Duc18] L. Ducas. “Shortest vector from lattice sieving: a few dimensions for free”. In: *EUROCRYPT 2018*. Vol. 10820. LNCS. Springer, 2018, pp. 125–145.
- [Duc22] L. Ducas. “Estimating the hidden overheads in the BDGL lattice sieving algorithm”. In: *PQCrypto 2022*. Springer, 2022, pp. 480–497.
- [Duc23] L. Ducas. “Provable lattice reduction of $\mathbb{Z}^n$ with blocksize $n/2$ ”. In: *Des. Codes Cryptogr.* 92.4 (2024), pp. 909–916.
- [DvW21] L. Ducas and W. van Woerden. *On the lattice isomorphism problem, quadratic forms, remarkable lattices, and cryptography*. EUROCRYPT 2022. 2022.
- [EP24] T. Espitau and H. Pliatsok. *On Hermitian decomposition lattices and the module-LIP problem in rank 2*. Cryptology ePrint Archive, Report 2024/1148. 2024. URL: <https://eprint.iacr.org/2024/1148>.
- [GP25] D. M. H. van Gent and L. N. Pulles. “HAWK: Having automorphisms weakens key”. In: *IACR Communications in Cryptology 2.2* (2025). art. 20. URL: <https://doi.org/10.62056/a3qjp2w9p>.
- [GS02] C. Gentry and M. Szydlo. “Cryptanalysis of the revised NTRU signature scheme”. In: *EUROCRYPT 2002*. Ed. by L. R. Knudsen. Vol. 2332. LNCS. Springer, 2002, pp. 299–320.
- [HAWK25] J. W. Bos, O. Bronchain, L. Ducas, S. Fehr, Y.-H. Huang, T. Pornin, E. W. Postlethwaite, T. Prest, L. N. Pulles, and W. van Woerden. *HAWK*. Tech. rep. version 1.1, submission to the NIST PQC additional-signatures call. National Institute of Standards and Technology, 2025. URL: <https://hawk-sign.info>.

- [KF17] P. Kirchner and P.-A. Fouque. “Revisiting lattice attacks on over-stretched NTRU parameters”. In: *EUROCRYPT 2017, Part I*. Vol. 10210. LNCS. Springer, 2017, pp. 3–26.
- [LJPW24] H. Luo, K. Jiang, Y. Pan, and A. Wang. “Cryptanalysis of rank-2 module-LIP with symplectic automorphisms”. In: *ASIACRYPT 2024*. Springer, 2024, pp. 359–385.
- [LLL82] A. K. Lenstra, H. W. Lenstra Jr., and L. Lovász. “Factoring polynomials with rational coefficients”. In: *Math. Ann.* 261.4 (1982), pp. 515–534.
- [MPPW24] G. Mureau, A. Pellet-Mary, H. Pliatsok, and A. Wallet. *Cryptanalysis of rank-2 module-LIP in totally real number fields*. EUROCRYPT 2024. 2024.
- [NIST22] NIST. *Call for additional digital signature schemes for the post-quantum cryptography standardization process*. 2022. URL: <https://csrc.nist.gov/projects/pqc-dig-sig>.
- [PFH20] T. Prest, P.-A. Fouque, J. Hoffstein, P. Kirchner, V. Lyubashevsky, T. Pornin, T. Ricosset, G. Seiler, W. Whyte, and Z. Zhang. *FALCON*. Tech. rep. Specification v1.2 (2020-10-01). NIST Post-Quantum Cryptography Standardization Process, 2020. URL: <https://falcon-sign.info>.
- [SAB22] P. Schwabe, R. Avanzi, J. Bos, L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, J. M. Schanck, G. Seiler, and D. Stehlé. *CRYSTALS-Kyber*. Tech. rep. v3.02. NIST PQC submission, 2021.
- [ZDY25] Z. Zhao, J. Ding, and B.-Y. Yang. “Sieving with streaming memory access”. In: *IACR Trans. Cryptogr. Hardw. Embed. Syst.* 2025.2 (2025). Software: <https://github.com/zhaoziyu0008/BGJ-Sieve-AMX>, pp. 362–384. URL: <https://doi.org/10.46586/tches.v2025.i2.362-384>.

A Computational Verification

The released code [code] is the heuristic implementation of §6.2, not the provable reduction of §5: its SVP stage is progressive BKZ plus a sieve. An accompanying script re-derives the §4 statements exactly on a random key at $\ell \leq 6$: Lemmas 4.1–4.4 and Proposition 4.5, including the isometry class $\sqrt{n/4} \mathbb{Z}^{n/2+1} \oplus \sqrt{n/2} \mathbb{Z}^{n/2-1}$ and the shortest-vector count $2(n/2 + 1)$ by exact enumeration. End-to-end, the same code recovers keys on synthetic instances at $\ell = 5, 6, 7, 8$ in 0.7 s, 1.4 s, 19 s, and 8.6 minutes respectively (single core; BKZ alone reaches the minimum for $n \leq 128$), and on HAWK-256 keys at $\ell = 9$. Instances at $\ell \leq 8$ are synthetic (a random $\text{SL}_2(R_n)$ basis from a reimplemented generator); at $\ell = 9$ keys are generated, and every recovered key is checked by a sign/verify round-trip, using the official HAWK submission package [HAWK25].

B Involutions of $(\mathbb{Z}/m)^\times$

The following lemma is the group-theoretic input to Theorem 7.1; we include its proof for completeness.

Lemma B.1. *For $m \geq 1$ the following are equivalent: (i) $(\mathbb{Z}/m)^\times$ contains an element of order 2 other than -1 ; (ii) $(\mathbb{Z}/m)^\times$ is non-cyclic; (iii) $m \notin \{1, 2, 4, p^k, 2p^k\}$ for p an odd prime.*

Proof. By the Chinese Remainder Theorem, $(\mathbb{Z}/m)^\times \cong \prod_i (\mathbb{Z}/p_i^{k_i})^\times$ over the prime-power factors $p_i^{k_i}$ of m . Each factor with p_i odd is cyclic (the primitive root theorem) of order $p_i^{k_i-1}(p_i - 1)$, which is even; $(\mathbb{Z}/2)^\times$ is trivial, $(\mathbb{Z}/4)^\times$ is cyclic of order 2, and $(\mathbb{Z}/2^k)^\times \cong \mathbb{Z}/2 \times \mathbb{Z}/2^{k-2}$ is non-cyclic for $k \geq 3$.

(iii) $\Rightarrow$ (i): if $m \notin \{1, 2, 4, p^k, 2p^k\}$ then either $8 \mid m$, in which case the factor $(\mathbb{Z}/2^k)^\times$ with $k \geq 3$ already contains $(\mathbb{Z}/2)^2$ in its 2-torsion, or the decomposition has at least two factors of even order (this happens when two distinct odd primes divide m , or when one odd prime divides m and $4 \mid m$). When there are at least two even-order factors, the element that reduces to -1 in one of them and to $+1$ in every other factor has order 2 and differs from -1 in a second even-order factor. Either way the 2-torsion subgroup has order at least 4 and contains an involution other than -1 .

(i) $\Rightarrow$ (ii): for $m > 2$ the class of -1 has order 2, so (i) provides two distinct elements of order 2; a cyclic group has at most one. (For $m \leq 2$ the group is trivial and (i) is false.)

(ii) $\Rightarrow$ (iii): contrapositively, for $m \in \{1, 2, 4, p^k, 2p^k\}$ the decomposition has at most one nontrivial factor, which is cyclic. $\square$

C The $0.72 \beta_{\text{key}}$ estimate of §7

Throughout this appendix m is the conductor and $d_+ := [K_n^+ : \mathbb{Q}] = \varphi(m)/2$ ($d_+ = n/2$ for the power-of-two parameter sets). The blocksize model is the

heuristic GSA-intersect estimate (GSA: Geometric Series Assumption), as used for [HAWK25, Table 8]: an unusually short vector of length ν in a rank- d lattice of volume $\mathcal{V}$ is found by progressive BKZ at the smallest β with

$$\sqrt{\beta/d} \nu \leq \delta_\beta^{2\beta-d-1} \mathcal{V}^{1/d},$$

$$\delta_\beta = \left(\frac{\beta}{2\pi e}\right)^{1/(2(\beta-1))},$$

a condition that depends only on d and the gap $\ln \nu - \frac{1}{d} \ln \mathcal{V}$; the gap-0 instance is [DPPvW22, eq. (4)].

For every conductor, the c -cocycle parametrization at $B = I$ is $W = \begin{pmatrix} a & b \\ b & c(a) \end{pmatrix}$ with $a \in R_n$, $b \in R_n^{c=-1}$, $\det W = |a|^2 + |b|^2$, and $Q^{(c)}(W) = \text{Tr}_{K_n^+}(\det W) = \frac{1}{2}\langle a, a \rangle_{K_n} + \frac{1}{2}\langle b, b \rangle_{K_n}$ (using $|a|^2, |b|^2 \in K_n^+$ and $[K_n : K_n^+] = 2$). As in Lemma 4.4 with K_n^+ in place of F , the minimum is $\text{Tr}_{K_n^+}(1) = d_+$, attained at the secret cocycle $B^{-1}c(B)$, so $\nu = \sqrt{d_+}$.

Power-of-two conductors. The key lattice is a scaled $\mathbb{Z}^{2n}$, gap 0; the model gives $\beta_{\text{key}} = 221/457/933$, within ± 10 of the simulated values $211/452/940$ of [HAWK25, Table 7] and within ± 1 of [HAWK25, Table 8]. For the c -cocycle lattice, the trace-orthogonal power basis (§3.3) gives $(R_n, \langle \cdot, \cdot \rangle_{K_n}) = \sqrt{n} \mathbb{Z}^n$, and c acts on that basis as a signed permutation with one fixed point (1), one antifixed point ($\zeta^{n/2}$), and $d_+ - 1$ sign-2-cycles, giving $(R_n^{c=-1}, \langle \cdot, \cdot \rangle_{K_n}) = \sqrt{n} \mathbb{Z} \oplus \sqrt{2n} \mathbb{Z}^{d_+-1}$ by the per-cycle accounting of the proof of Proposition 4.5. Hence $\Lambda_I^{(c)} \cong \sqrt{d_+} \mathbb{Z}^{2d_++1} \oplus \sqrt{2d_+} \mathbb{Z}^{d_+-1}$, gap $\approx -\frac{1}{6} \ln 2$.

Power-of-two τ -cocycle. The same model applies to the rescaled cocycle lattice of Corollary 4.6, $\Lambda' = \mathbb{Z}^{n/2+1} \oplus \sqrt{2} \mathbb{Z}^{n/2-1}$: rank $d = n$, $\nu = 1$, $\mathcal{V} = 2^{(n/2-1)/2}$, so the gap is $-\frac{n/2-1}{2n} \ln 2 \approx -\frac{1}{4} \ln 2$. The model then gives block sizes $92/205/432$ at HAWK-256/-512/-1024, that is 0.42–0.46 of the same-model β_{key} above and somewhat below the oracle dimension $n/2 + 1$ of Theorem 6.1; the next paragraph converts these into the gate counts quoted in §6.1. The ± 10 calibration against the simulated block sizes is for the key lattice only; a simulator run on the class $[\Lambda']$, as was done for the key lattice in [HAWK25, Table 7], would be the proper counterpart.

Odd-prime-power conductors. The power basis is not orthogonal, so neither lattice has gap 0. The key-lattice minimum is $\varphi(m)$: for $x \in R_n \setminus \{0\}$, $\text{Tr}_{K_n}|x|^2 = \sum_\sigma |\sigma(x)|^2 \geq \varphi(m) |N(x)|^{2/\varphi(m)} \geq \varphi(m)$ by AM–GM and $|N(x)| \geq 1$, with equality iff x is a root of unity. The volumes follow from $|\text{disc } K_n|$ and a closed form for $\det(R_n^{c=-1}, \langle \cdot, \cdot \rangle_{K_n})$, both checked against exact linear algebra at $m \in \{5, 7, 9, 11, 13, 25, 27\}$.

Ratio. Across every conductor we computed, $\beta_{c\text{-cocycle}}/\beta_{\text{key}} \in [0.69, 0.73]$: 0.71–0.72 at HAWK-512 and HAWK-1024, 0.72–0.73 for odd-prime-power conductors of comparable degree ($m = 769, 1153, 3^7, 5^5, 2309$), and 0.69–0.73 at HAWK-256 depending on whether the denominator is the simulated or same-model β_{key} . The ratio $3/4$ suggested by the ranks alone is corrected downward

chiefly by the $2^{1/6}$ volume excess: the cocycle target is relatively shorter in its lattice than the unit vectors are in the key lattice.

Gate-count conversion for §6.1. At the power-of-two τ -cocycle block sizes above (205 for HAWK-512, 432 for HAWK-1024), the recipe of [HAWK25, §5.3] converts as follows: dimensions-for-free $\text{d4f}(\beta) = \beta \ln(4/3) / \ln(\beta/(2\pi e))$, rounded, with β_{eff} lowered to a multiple of 8, gives 176 and 392; the classical list-decoding gate counts of [AGPS20] are $2^{71.8}$ and $2^{136.5}$ per call; and multiplying by n for one tour gives the heuristic totals $2^{80.8}$ and $2^{146.5}$ quoted in §6.1. Isolating $\pm V_\tau$ from the shortest vectors (Remark 6.2) is cheaper by the same recipe. The estimate counts only the top-level step-(2) reduction and prices the final descent by the τ -tower of §6.2; the alternative, a direct reduction on the $\mathbb{Z}^n$ -LIP kernel of step (3), would cost more ($2^{85.5}$ and $2^{153.6}$). Each tower level repeats the top-level computation at half the degree and is cheaper by the same recipe (the $n = 256$ level costs about 2^{51} , the $n = 512$ level about 2^{81}), so the top level dominates the tower. The block sizes are GSA-intersect estimates rather than the simulated ones behind [HAWK25, Table 7], a difference of about ± 10 in β , or ± 2.4 bits; memory, the progressivity factor, and the overheads of [Duc22] remain unpriced on both sides, as in §6.1.

D The Descent as a Strengthened Analysis of van Gent–Pulles

Throughout, Y denotes one of the $2(n/2 + 1)$ shortest vectors of $\Lambda_B^{(\tau)}$, so that $Y\tau(Y) = I$ and $\det Y = 1$ (Lemma 4.4(ii)), hence $Y^{-1} = \tau(Y)$, and $Y^*QY = \tau(Q)$: by (C2), $Y^*QY = Y^*\sigma(Y)^T\tau(Q)$, and $Y^*\sigma(Y)^T = (\sigma(Y)c(Y))^T = \sigma(Y\tau(Y))^T = I$ using $c\tau = \sigma$ and (C1).

Lemma D.1. *The map $U_Y : R_n^2 \rightarrow R_n^2$, $x \mapsto Y\tau(x)$, is an order-2 automorphism of the quadratic $\mathbb{Z}$ -module (R_n^2, T) , $T(x, y) = \text{Tr}_{K_n}(x^*Qy)$, lying outside the trivial group G_n of [GP25] (generated by the roots of unity and the symplectic automorphism, the latter included after [LJPW24]); and $\ker(U_Y - 1) = \mathcal{M}_Y$, $\ker(U_Y + 1) = \mathcal{M}_{-Y}$.*

Proof. Order 2: $U_Y^2(x) = Y\tau(Y)x = x$ by the cocycle equation. Isometry: $T(U_Yx, U_Yy) = \text{Tr}_{K_n}(\tau(x)^*Y^*QY\tau(y)) = \text{Tr}_{K_n}(\tau(x)^*\tau(Q)\tau(y)) = \text{Tr}_{K_n}(\tau(x^*Qy)) = T(x, y)$, using $Y^*QY = \tau(Q)$, $\tau(x)^* = \tau(x^*)$ (the involutions c and τ commute), and the Galois-invariance of the trace. Exclusion: U_Y is τ -semilinear ($U_Y(\alpha x) = \tau(\alpha)U_Y(x)$ for $\alpha \in R_n$), while every element of G_n is R_n -linear and $\tau \neq \text{id}$. Kernels: $U_Yx = \pm x \iff \tau(x) = (\pm Y)^{-1}x \iff x \in \mathcal{M}_{\pm Y}$, using $Y^{-1} = \tau(Y)$. $\square$

In the terminology of [GP25], U_Y lies in the orthogonal group of the public form outside their trivial group, $\mathcal{M}_{\pm Y} = \ker(U_Y \mp 1)$ are the kernel sublattices of [GP25, Prop. 2], and [GP25, Prop. 1] — proved there, deterministic polynomial time — recovers an equivalent key from any sufficiently short kernel vector. [GP25, Lem. 3 and Prop. 2] prove that these kernels are of the form $\mathbb{Z}^s \oplus \sqrt{2}\mathbb{Z}^t$

for some $s, t \geq 0$ with $\lambda_1 \leq \sqrt{2}$ — enough for the heuristic pricing of [GP25, Thm. 1], but not enough to determine the split (s, t) , which differs between $Y = \pm V_\tau$ and the other candidates. The proofs below identify the kernels exactly — $(\mathcal{M}_{\pm V_\tau}, T) \cong \sqrt{n} \mathbb{Z}^n$, with the columns of B^{-1} up to roots of unity as shortest vectors — and this exactness is what admits the provable solver of [Duc23] (Remark D.3) and so upgrades the endgame from the heuristic pricing of [GP25, Thm. 1] to the unconditional accounting of Theorem 6.1.

Proposition D.2. *Let $Y \in \Lambda_B^{(\tau)}$ be a shortest vector with $\varphi_\tau(Y) = \pm I$. Then:*

1. $\mathcal{M}_Y = B^{-1} \mathcal{O}_{K_n^\tau}^2$ (resp. $B^{-1}(\zeta \mathcal{O}_{K_n^\tau})^2$ for $-I$), a $\mathbb{Z}$ -lattice of rank n ;
2. $(\mathcal{M}_Y, T) \cong \sqrt{n} \mathbb{Z}^n$;
3. the shortest vectors of $(\mathcal{M}_Y, T)$ are $\{\pm B^{-1} \zeta^{2j} e_k : 0 \leq j < n/2, k \in \{1, 2\}\}$ (resp. $\{\pm B^{-1} \zeta^{2j+1} e_k\}$ in the $-I$ case), where e_1, e_2 is the standard basis of R_n^2 ; any pair (u_1, u_2) from this set with $u_1^* Q u_2 = 0$ satisfies $[u_1 \ u_2]^{-1} \in U_2(R_n) \cdot B$, an equivalent HAWK secret key. (Here $[u_1 \ u_2]$ is the matrix with columns u_1, u_2 .)

Proof. (1) For $Y = V_\tau$: $\tau(x) = V_\tau^{-1}x = \tau(B)^{-1}Bx \Leftrightarrow \tau(Bx) = Bx \Leftrightarrow Bx \in (K_n^\tau)^2$. Intersecting with R_n^2 and using $BR_n^2 = R_n^2$ ($B \in \text{SL}_2(R_n)$) and $R_n \cap K_n^\tau = \mathcal{O}_{K_n^\tau}$: $\mathcal{M}_Y = B^{-1} \mathcal{O}_{K_n^\tau}^2$. For $Y = -V_\tau$, $\tau(Bx) = -Bx$ gives $\mathcal{M}_Y = B^{-1}(\zeta \mathcal{O}_{K_n^\tau})^2$. (2) For $x = B^{-1}y$: $x^* Q x = y^* (B^{-1})^* Q B^{-1} y = y^* y$, so $T(x, x') = \text{Tr}_{K_n}(y^* y')$. On the basis $\{\zeta^{2j} e_k : 0 \leq j < n/2, k \in \{1, 2\}\}$ of $\mathcal{O}_{K_n^\tau}^2$: $\text{Tr}_{K_n}(\zeta^{2j} \zeta^{2j'}) = n \delta_{j,j'}$. Hence Gram $= nI_n$. (3) The shortest vectors of a lattice in class $[\mathbb{Z}^m]$ are its orthogonal $\mathbb{Z}$ -basis up to sign, so here $\{\pm B^{-1} \zeta^{2j} e_k\}$. For any two of these, $(B^{-1} \zeta^{2j} e_k)^* Q (B^{-1} \zeta^{2j'} e_{k'}) = \zeta^{2(j'-j)} \delta_{k,k'}$, which vanishes iff $k \neq k'$. For such a pair, $[u_1 \ u_2] = B^{-1}M$ with M a monomial matrix whose nonzero entries are $2n$ -th roots of unity, hence $M \in U_2(R_n)$ and $[u_1 \ u_2]^{-1} \in U_2(R_n) \cdot B$. $\square$

Selecting $\pm V_\tau$: the parity test. The other n shortest vectors do not lead to the key: for Y with $\varphi_\tau(Y) = W(0, \omega)$, $\omega = \zeta^j$ with j odd (proof of Lemma 4.4(ii)), the substitution $y = Bx$ of Proposition D.2's proof turns the defining condition of $\mathcal{M}_Y$ into $y_2 = -\omega^{-1} \tau(y_1)$ with $y_1 \in R_n$ free, and $T(x, x') = 2 \text{Tr}_{K_n}(\overline{y_1} y'_1)$, so $(\mathcal{M}_Y, T) \cong \sqrt{2n} \mathbb{Z}^n$, whose shortest vectors carry no orthonormal pair. The two useful candidates are recognized by reduction modulo 2: since $\tau(\zeta^i) = (-1)^i \zeta^i \equiv \zeta^i$, the involution τ acts trivially on $R_n/2R_n$, so $\tau(B) \equiv B$ and $\pm V_\tau = \pm B^{-1} \tau(B) \equiv I \pmod{2}$; whereas any other shortest vector $Y = B^{-1}W(0, \omega) \tau(B) \equiv B^{-1}W(0, \omega) B \not\equiv I \pmod{2}$, because B is invertible modulo 2 (as $\det B = 1$), so conjugation by it is a bijection of $M_2(R_n/2R_n)$ fixing I , and $W(0, \omega)$ has zero diagonal. Thus exactly the two candidates $\pm V_\tau$ satisfy $Y \equiv I \pmod{2R_n}$, an $O(n)$ -coefficient check. (The reduction does require the exact shortest vectors: $\mathcal{M}_Y \neq 0$ iff $\det Y = 1$, so an approximate-SVP output Y' with $Q^{(\tau)}(Y') > n/4$ gives $\mathcal{M}_{Y'} = \{0\}$.)

Remark D.3 (improved iteration count for [Duc23]). The two-bound termination analysis of Theorem 5.1 (in the notation of its proof, §E.4: potential W_t , primal factor α) applies to [Duc23, Alg. 1] on a lattice in $[\mathbb{Z}^d]$,

$d = 2k' + 1$: after the primal step the front block $\mathcal{L}(B_{[0:k']})$ is a primitive rank- $(k'+1)$ sublattice of volume $W_t^{1/2}\alpha$ with $\alpha \leq 1$ [Duc23, Lemma 3], its dual has covolume $(W_t^{1/2}\alpha)^{-1}$, and Minkowski's theorem gives the same recursion $W_{t+1} \leq (k'+1)W_t^{1-1/(k'+1)}$; the basic per-iteration factor is $1 - 1/d$ [Duc23, Lemma 4]. The iteration count of [Duc23, Thm 2] therefore improves to $d(1 + (k'+1)\ln(k'+1)) + (k'+1)\ln(2k'\ln\max_i \|b_i\|) + 3(k'+1)$ for any input basis, again doubly logarithmic in the input. Applied through the even-rank adjustment $\Lambda \mapsto \Lambda \oplus \mathbb{Z}$ of [Duc23] (the rank- n lattice $(\mathcal{M}_Y, \frac{1}{n}T)$ becomes the odd rank $n+1$), it gives $(n+1)(n/2+1)\ln(n/2+1) + O(n\log n)$ iterations per descent run under the LLL input bound of §6.1.

Lemma D.4 (the τ -tower step). *Let $Y \in \Lambda_B^{(\tau)}$ be a shortest vector with $\varphi_\tau(Y) = \pm I$ (so $Y\tau(Y) = I$, $\det Y = 1$, and $Y \equiv I \pmod{2}$), and write $K' := K_n^\tau$, $\mathcal{O}' := \mathcal{O}_{K'}$, so that $R_n = \mathcal{O}' \oplus \zeta\mathcal{O}'$. Then (i) $\mathcal{M}_Y$ is the $\mathcal{O}'$ -module generated by the columns of $(I+Y)/2$ and $(I-Y)/(2\zeta)$ (both in $M_2(R_n)$), and (ii) the alternating $\mathcal{O}'$ -bilinear form $\mu(x, y) := \frac{1}{2}\text{Tr}_{K_n/K'}(\det[x \ y])$ on $\mathcal{M}_Y$ is unimodular: $\mu(f_1, f_2) \in \mathcal{O}'^\times$ for every $\mathcal{O}'$ -basis (f_1, f_2) . Consequently any pair $(x, y) \in \mathcal{M}_Y^2$ with $\mu(x, y) = 1$ is a free $\mathcal{O}'$ -basis of $\mathcal{M}_Y$.*

Proof. (i) The matrices $(I+Y)/2$ and $(I-Y)/(2\zeta)$ are integral because $Y \equiv I \pmod{2}$ and ζ is a unit. For $a \in \mathcal{O}^2$: $\tau(\frac{I+Y}{2}a) = \frac{I+Y^{-1}}{2}a = Y^{-1}\frac{I+Y}{2}a$, and likewise for $(I-Y)/(2\zeta)$ using $\tau(\zeta) = -\zeta$, so both columns lie in $\mathcal{M}_Y$. Conversely, for $x \in \mathcal{M}_Y$ write $x = a + \zeta b$ with $a, b \in \mathcal{O}^2$; averaging $x = a + \zeta b$ with $x = Y\tau(x) = Ya - \zeta Yb$ gives $x = \frac{I+Y}{2}a + \frac{I-Y}{2\zeta}(\zeta^2 b)$ with $\zeta^2 b \in \mathcal{O}^2$. (ii) By Proposition D.2(1), $\mathcal{M}_Y = B^{-1}\Omega\mathcal{O}^2$ with $\Omega = 1$ (resp. $\Omega = \zeta$ for $-I$). For $x = B^{-1}\Omega u$, $y = B^{-1}\Omega w$ with $u, w \in \mathcal{O}^2$, since $\det B = 1$ we get $\det[x \ y] = \Omega^2 \det[u \ w] \in K'$, hence $\mu(x, y) = \Omega^2 \det[u \ w]$ ($\det[x \ y]$ is already τ -fixed). For an $\mathcal{O}'$ -basis $(f_1, f_2) = (B^{-1}\Omega u_1, B^{-1}\Omega u_2)$ we have $[u_1 \ u_2] \in \text{GL}_2(\mathcal{O}')$, so $\mu(f_1, f_2) = \Omega^2 \det[u_1 \ u_2] \in \mathcal{O}'^\times$. Finally, $\mu(x, y) = \det(C)\mu(f_1, f_2)$ where $C \in M_2(\mathcal{O}')$ expresses (x, y) in the basis; if $\mu(x, y) = 1$ then $\det C \in \mathcal{O}'^\times$ and (x, y) is a basis. $\square$

For the other shortest vectors (the case $\varphi_\tau(Y) = W(0, \omega)$, $\omega = \zeta^j$, above) the same computation gives μ with values in $2\zeta^{1-j}\mathcal{O}'$, never a unit — a second reason that only $\pm V_\tau$ descend.

The τ -tower descent (§6.2). The implementation replaces the single $\mathbb{Z}^n$ -LIP solve of step (3) by a descent through the tower of τ -fixed subfields $\mathbb{Q}(\zeta) \supset \mathbb{Q}(\zeta^2) \supset \mathbb{Q}(\zeta^4) \supset \dots$, i.e. from the conductor exponent ℓ down to 3. Each level computes an $\mathcal{O}_{K_n^\tau}$ -basis of $\mathcal{M}_{V_\tau} = B^{-1}\mathcal{O}_{K_n^\tau}^2$: by (i), the module is generated by the columns of $(I+V_\tau)/2$ and $(I-V_\tau)/(2\zeta)$, and by (ii), a free basis is any pair (x, y) with $\mu(x, y) = 1$, obtained by a Bézout step in $\mathcal{O}'$ (no class-group computation). The level then Gauss-reduces this basis under the Q -Hermitian form and recurses on the resulting HAWK public key over K_n^τ , re-running steps (1)–(2) there to obtain the next level's V_τ .

Proof of Remark 6.2. Recall $M := Y \text{adj}(Y')$ and $\Lambda_a := \{X \in \Lambda_B^{(\tau)} : MX = X\tau(M)\}$ for shortest vectors $Y \neq \pm Y'$ of $\Lambda_B^{(\tau)}$, neither equal to $\pm V_\tau$. Transporting

by the isometry φ_τ of Lemma 4.2, the condition $MX = X\tau(M)$ becomes $NW = W\tau(N)$ with $N = \varphi_\tau(Y) \cdot \text{adj}(\varphi_\tau(Y'))$; since $\varphi_\tau(Y), \varphi_\tau(Y')$ are shortest vectors other than $\pm I$, they have the form $W(0, b)$ with $b = \zeta^k$, k odd (§4.3; proof of Lemma 4.4(ii)), so N is diagonal. For $W = W(a, b)$ the diagonal entries of $NW - W\tau(N)$ vanish identically and the off-diagonal entries vanish only for $b = 0$ (using $Y \neq \pm Y'$), so $\varphi_\tau(\Lambda_a) = \{W(a, 0) : a \in R_n^\sigma\}$, the summand of $\Lambda_I^{(\tau)}$ containing $\pm I$, whose two shortest vectors are $\pm I$ by Proposition 4.5. $\square$

E Correctness of Block Reduction on $[\Lambda']$

This appendix proves that LIP for the class $[\Lambda'] = [\mathbb{Z}^p \oplus \sqrt{2}\mathbb{Z}^q]$ is solvable with polynomially many calls to exact SVP in dimension p (Theorem 5.1). The algorithm is Lucas’s two-block reduction [Duc23, Alg. 1], run unchanged on $[\Lambda']$ — the same primal/dual loop at the same oracle dimension, with only the cut fixed at $k := p - 1$ for this even-rank class and a final size-reduction made explicit; what is new is the analysis, since Lucas’s integrality and progress lemmas are proved only for $\mathbb{Z}^d$ (Lemmas E.1–E.2 replace them, and this appendix re-derives termination and correctness). The missing ingredient is the no-abort product-progress bound for $[\Lambda']$. In outline: the algorithm maintains a basis of the lattice and alternates two SVP-oracle steps on projected blocks of dimension at most p : a *primal* step that makes the Gram–Schmidt length $\|b_{p-1}^*\|$ as short as possible and a *dual* step that makes it as long as possible. The squared volume of the first $p - 1$ basis vectors is a positive integer (Lemma E.1) that each primal-plus-dual round multiplies by a factor at most $1 - 1/p$ (Lemma E.2), so after $O(n^2 \log n)$ rounds (for polynomially-bounded input) it reaches 1, at which point the front block spans a copy of $\mathbb{Z}^{p-1}$ and the hidden orthogonal basis can be read off (Theorem 5.1).

Notation. A sublattice S of a lattice Λ is *primitive* if $S = \text{span}(S) \cap \Lambda$; $\text{vol}(S)$ is its covolume (root-determinant of its Gram), π_U (for a subspace U) is orthogonal projection onto U , $\pi_S^\perp$ is orthogonal projection onto $\text{span}(S)^\perp$, and for S primitive in Λ one has $S^\vee = \pi_{\text{span}(S)}(\Lambda^\vee)$. For a basis $B = (b_0, \dots, b_{d-1})$ (unrelated to the secret key of §3.1) of a lattice $\mathcal{L}(B)$, b_i^* denotes the i -th Gram–Schmidt vector (the star is unrelated to the adjoint $*$) and $B_{[i:j]} := (\pi_{\{b_0, \dots, b_{i-1}\}}^\perp(b_i), \dots, \pi_{\{b_0, \dots, b_{i-1}\}}^\perp(b_j))$ the projected block.

Write $\{e_j\}_{j=0}^{p-1}$ for the standard basis of the $\mathbb{Z}^p$ summand (norm 1) and $\{\sqrt{2}f_m\}_{m=0}^{q-1}$ for the $\sqrt{2}\mathbb{Z}^q$ summand (norm $\sqrt{2}$), and $\mathbb{R}^p, \mathbb{R}^q \subset \mathbb{R}^{p+q}$ for the spans of the two families; the algorithm is given only the Gram matrix of an arbitrary basis, so this orthogonal basis exists but is hidden, and finding it is the task. Set $k := p - 1$ for the cut between the two blocks: the cut must satisfy $k < p$ for the front block to be able to terminate as an orthonormal subset of the $\mathbb{Z}^p$ summand (and for Lemma E.2’s Case 1 to find a unit vector outside it), and $k = p - 1$, the largest such cut, keeps both blocks’ SVP dimensions at most p .

E.1 Block reduction for $\mathbb{Z}^d$ and its obstruction

Ducas [Duc23, Thms 1, 2] solves $\mathbb{Z}^d$ -LIP with $\text{poly}(d)$ SVP calls in dimension $\lfloor d/2 \rfloor + 1$, via a 2-block slide reduction with invariant $\text{vol}(B_{[0:k-1]})^2 \in \mathbb{Z}_{>0}$. Its progress lemma is [Duc23, Lemma 4]: if $S \subset \mathbb{Z}^d$ is primitive of rank k with $\text{vol}(S) > 1$, then $\lambda_1(\pi_S^\perp(\mathbb{Z}^d)) \leq \sqrt{1 - 1/d}$. This is the per-round drop of the integer potential described above, for the hypercubic class.

This lemma is **false** for A' : take $S = \mathbb{Z}^{k-1} \oplus \sqrt{2}\mathbb{Z} \subset A'$ (rank k , $\text{vol}^2 = 2 > 1$). Then $e_{k-1} \in A'$, $e_{k-1} \perp S$, $\|e_{k-1}\| = 1$, so $\lambda_1(\pi_S^\perp(A')) = 1 > \sqrt{1 - 1/(p+q)}$, so the lemma fails.

Nor can the input itself be transformed into a hypercubic class: the determinant class in $\mathbb{Q}^\times/(\mathbb{Q}^\times)^2$ of a rational quadratic form of even rank is unchanged by rescaling, dualizing, and passing to a commensurable lattice, and for A' that class is $2^q \equiv 2$ (q odd) while for every rationally scaled hypercubic lattice it is 1. The fix is to bound the **product** of the primal and dual steps' contributions. The dual lattice $A'^\vee = \mathbb{Z}^p \oplus \frac{1}{\sqrt{2}}\mathbb{Z}^q$ has its *short* vectors in the q -block — where A' has its *long* vectors — so when the primal step stalls, the dual step compensates.

E.2 The integrality lemma

Lemma E.1 (Analogue of [Duc23, Lemma 2]). *For any sublattice $S \subset A'$ of rank $r \geq 1$: $\text{vol}(S)^2 \in \mathbb{Z}_{>0}$. If $\text{vol}(S)^2 = 1$ then $S \subset \mathbb{Z}^p$ and $S \cong \mathbb{Z}^r$.*

Proof. A $\mathbb{Z}$ -basis of S has the form $M = \begin{pmatrix} A \\ C \end{pmatrix}$ with $A \in \mathbb{Z}^{p \times r}$, $C \in \mathbb{Z}^{q \times r}$, and Gram $G_S = M^T D M = A^T A + 2C^T C \in \mathbb{Z}^{r \times r}$, so $\text{vol}(S)^2 = \det G_S \in \mathbb{Z}_{>0}$. Suppose $\det G_S = 1$. Reducing mod 2: $\det(A^T A) \equiv 1 \pmod{2}$, so $\det(A^T A)$ is odd, hence A has full rank r and $A^T A$ is positive-definite. The Minkowski determinant inequality $\det(P + P') \geq \det P$ for P positive-definite and P' positive-semidefinite (equivalently $\det(I + P^{-1/2} P' P^{-1/2}) \geq 1$, since all eigenvalues of the identity plus a positive-semidefinite matrix are ≥ 1) gives $\det(A^T A) \leq \det G_S = 1$. As $\det(A^T A) \in \mathbb{Z}_{>0}$, it equals 1. Then $\det(I + 2(A^T A)^{-1/2} C^T C (A^T A)^{-1/2}) = 1$ with positive-semidefinite inner factor forces $C^T C = 0$, i.e. $C = 0$. So $S \subset \mathbb{Z}^p$ with $\text{vol}(S) = 1$; [Duc23, Lemma 2] applied inside $\mathbb{Z}^p$ gives $S \cong \mathbb{Z}^r$. $\square$

E.3 The Product Progress Lemma

One primal step and the subsequent dual step of [Duc23, Alg. 1] jointly multiply $\text{vol}(B_{[0:k-1]})$ by the product bounded in the following lemma.

Lemma E.2. *Let $S \subset A'$ be primitive of rank $k = p - 1$ with $\text{vol}(S) > 1$. Let v be a shortest nonzero vector of $\pi_S^\perp(A')$, let $\tilde{v} \in A'$ be any lift of v , and set $S' := S + \mathbb{Z}\tilde{v}$ (primitive of rank p). Then*

$$\lambda_1(\pi_S^\perp(A')) \cdot \lambda_1(S'^\vee) \leq \sqrt{1 - 1/p}.$$

Proof. S' is primitive: v is shortest in $\pi_S^\perp(A')$, hence primitive there. The map $\pi_S^\perp : A' \rightarrow \pi_S^\perp(A')$ is a surjective lattice homomorphism with kernel $\text{span}(S) \cap A' = S$ (by primitivity of S), so $S' = (\pi_S^\perp)^{-1}(\mathbb{Z}v)$ and $A'/S' \cong \pi_S^\perp(A')/\mathbb{Z}v$ is torsion-free. Hence $S'^\vee = \pi_{\text{span}(S')}(A'^\vee)$.

Case 1: $\text{span}(S) \not\subset \mathbb{R}^p$. Since $\dim \text{span}(S) = k < p$, some $e_j \notin \text{span}(S)$; then $0 \neq \pi_S^\perp(e_j) \in \pi_S^\perp(A')$ with $\|\pi_S^\perp(e_j)\| \leq 1$, so $\lambda_1(\pi_S^\perp(A')) \leq 1$. As $\text{span}(S') \supset \text{span}(S)$ has a vector with nonzero $\mathbb{R}^q$ -component, some f_m satisfies $\pi_{\text{span}(S')}(f_m) \neq 0$. Since $\frac{1}{\sqrt{2}}f_m \in A'^\vee$, we get $0 \neq \frac{1}{\sqrt{2}}\pi_{\text{span}(S')}(f_m) \in S'^\vee$ with norm $\leq \frac{1}{\sqrt{2}}$. Product $\leq 1 \cdot \frac{1}{\sqrt{2}} \leq \sqrt{1 - 1/p}$ for all $p \geq 2$.

Case 2: $\text{span}(S) \subset \mathbb{R}^p$. Then $S \subset A' \cap \mathbb{R}^p = \mathbb{Z}^p$, and S is primitive in $\mathbb{Z}^p$. [Duc23, Lemma 4] (applied inside $\mathbb{Z}^p$ at rank $k < p$, $\text{vol} > 1$) gives $\lambda_1(\pi_S^\perp(\mathbb{Z}^p)) \leq \sqrt{1 - 1/p}$, and $\pi_S^\perp(\mathbb{Z}^p) \subset \pi_S^\perp(A')$ transfers the bound. For the dual factor: the $p+q$ vectors $\{e_j\} \cup \{\frac{1}{\sqrt{2}}f_m\} \subset A'^\vee$ span $\mathbb{R}^{p+q}$ and each has norm ≤ 1 , so at least one projects nonzero onto the p -dimensional $\text{span}(S')$; that projection witnesses $\lambda_1(S'^\vee) \leq 1$. Product $\leq \sqrt{1 - 1/p} \cdot 1$. $\square$

Remark E.3. In Case 1 the hypothesis $\text{vol}(S) > 1$ is *not* used; the $\sqrt{2}$ -structure supplies the progress. In Case 2 it is used where [Duc23, Lemma 4] needs it. The counterexample from §E.1 falls under Case 1, where the dual step supplies the factor $\frac{1}{\sqrt{2}}$.

E.4 Termination and correctness of Algorithm 1

In Algorithm 1, *primal-SVP reduction* of a projected block inserts a shortest vector of that block at its first position, so that the primal step sets $\|b_k^*\| = \lambda_1(\pi_{\{b_0, \dots, b_{k-1}\}}^\perp(A))$. *Dual-SVP reduction* of $B_{[0:k]}$ applies a unimodular transform to the first $k+1$ basis vectors making $\|b_k^*\|$ as large as possible, namely $1/\lambda_1(\mathcal{L}(B_{[0:k]}))^\vee$ [Duc23]. *Size reduction* translates each basis vector by integer multiples of the preceding ones so that every Gram–Schmidt coefficient $\mu_{ij} := \langle b_i, b_j^* \rangle / \|b_j^*\|^2$ satisfies $|\mu_{ij}| \leq \frac{1}{2}$. *HKZ* (Hermite–Korkine–Zolotarev) *reduction* of a block size-reduces it and SVP-reduces it and, recursively, each of its projected sub-blocks. All SVP calls are in dimension $\leq \max(k+1, p+q-k) = \max(p, q+1) = p = n/2 + 1$. The algorithm reads and updates only the Gram matrix of B ; “outputting a basis” means outputting the unimodular transform from the input basis to it.

Proof of Theorem 5.1. Termination. By Lemma E.1, $W_t := \text{vol}(B_{[0:k-1]})^2 \in \mathbb{Z}_{>0}$ throughout. Let $S_t := \mathcal{L}(B_{[0:k-1]})$ after iteration t ’s dual step; S_t is primitive in A' (as the lattice generated by the first k vectors of a $\mathbb{Z}$ -basis). We account for the drop in the potential W_t across the primal step of iteration t and the dual step of iteration $t+1$. The primal step sets $\|b_k^*\| = \lambda_1(\pi_{S_t}^\perp(A')) =: \alpha$ and $\mathcal{L}(B_{[0:k]}) =: S'_t$ with $\text{vol}(S'_t) = W_t^{1/2}\alpha$. The next dual step sets $\|b_k^*\| = 1/\lambda_1(S_t'^\vee) =: 1/\gamma$, so $W_{t+1}^{1/2} = \text{vol}(S'_t) \cdot \gamma = W_t^{1/2} \cdot \alpha\gamma$. If $W_t > 1$, Lemma E.2 gives $\alpha\gamma \leq \sqrt{1 - 1/p}$, so $W_{t+1} \leq (1 - 1/p)W_t$. From $W_0 \leq (\max_i \|b_i\|)^{2k}$: at most $2pk \ln \max_i \|b_i\|$ iterations.

The dependence on the input basis can be reduced to a doubly logarithmic one. The argument has two phases: first, a recursion proved next brings $\ln W_t$ to at most $1 + p \ln p$; then the factor $1 - 1/p$ of [Lemma E.2](#) brings the potential to 1. In every iteration we additionally have $W_{t+1} \leq p W_t^{1-1/p}$. The primal step's contribution satisfies $\alpha \leq 1$: since $\dim \text{span}(S_t) = k < p$, some e_j lies outside $\text{span}(S_t)$ and projects to a nonzero vector of $\pi_{S_t}^\perp(\Lambda')$ of norm at most 1. For the dual step, $S_t^\vee$ is a lattice of rank p and covolume $1/\text{vol}(S_t) = (W_t^{1/2}\alpha)^{-1}$. The closed cube of side $2 \text{covol}(S_t^\vee)^{1/p}$ has volume $2^p \text{covol}(S_t^\vee)$, so by Minkowski's convex-body theorem it contains a nonzero point of $S_t^\vee$, of Euclidean norm at most $\sqrt{p} \text{covol}(S_t^\vee)^{1/p}$. This gives $\gamma^2 = \lambda_1(S_t^\vee)^2 \leq p (W_t^{1/2}\alpha)^{-2/p}$. Hence $W_{t+1} = W_t(\alpha\gamma)^2 \leq p \alpha^{2(1-1/p)} W_t^{1-1/p} \leq p W_t^{1-1/p}$. Writing $v_t := \ln W_t$, this reads $v_{t+1} \leq \ln p + (1-1/p)v_t$. Hence $v_t \leq (1-1/p)^t v_0 + p \ln p$, and $v_t \leq 1 + p \ln p$ within $\lceil p \ln v_0 \rceil$ iterations. The factor $1 - 1/p$ of [Lemma E.2](#) then brings the integer potential to 1 within $p(1 + p \ln p)$ further iterations. With $v_0 = \ln W_0 \leq 2k \ln \max_i \|b_i\|$ this gives the bound stated in the theorem (the $+3p$ there absorbs the ceiling and the unpaired first dual step).

Correctness. At termination $\text{vol}(B_{[0:k-1]})^2 = 1$, so by [Lemma E.1](#) the front block $S := \mathcal{L}(B_{[0:k-1]})$ is a rank- k sublattice of $\mathbb{Z}^p$ with an orthonormal basis; since the only vectors of $\mathbb{Z}^p$ of norm 1 are the $\pm e_j$, S is the span of k of the p unit vectors. The final step HKZ-reduces the blocks $B_{[0:k-1]}$ (rank k) and $B_{[k:p+q-1]}$ (rank $q+1$) and size-reduces the full basis, using SVP calls in dimension at most $\max(k, q+1) = n/2 < p$.

For the front block: in an HKZ-reduced basis of S every Gram-Schmidt vector is a unit vector $\pm e_j$ of $\mathbb{Z}^p$ (inductively: the first basis vector is a norm-1 vector of $S \subseteq \mathbb{Z}^p$, hence some $\pm e_j$; projecting S orthogonally to it leaves the span of the remaining unit vectors, whose norm-1 vectors are again $\pm e_j$); the size-reduction coefficients $\mu_{ij} = \langle b_i, b_j^* \rangle / \|b_j^*\|^2 = \langle b_i, \pm e_j \rangle$ are integers, hence are reduced to 0; the front block becomes an orthonormal set of k unit vectors.

For the back block: $\pi_S^\perp(\Lambda') = \mathbb{Z}^{p-k} \oplus \sqrt{2}\mathbb{Z}^q = \mathbb{Z} \oplus \sqrt{2}\mathbb{Z}^q$. Its shortest vector is the remaining unit vector, unique up to sign (everything outside $\mathbb{Z}e$ has norm $\geq \sqrt{2}$), so the first HKZ vector of the back block is $\pm e$; the projection orthogonal to it is $\sqrt{2}\mathbb{Z}^q$, whose HKZ basis is its orthogonal basis by the same argument as for the front block (scaled by $\sqrt{2}$). Size reduction again clears the cross terms: against a unit vector $b_j^* = \pm e$ the coefficient is an integer, and against a vector $b_j^* = \pm\sqrt{2}f$ the coefficient is $\langle b_i, \sqrt{2}f \rangle / 2 \in (2\mathbb{Z})/2 = \mathbb{Z}$. The output basis is therefore the hidden orthogonal basis $(e_0, \dots, e_{p-1}, \sqrt{2}f_0, \dots, \sqrt{2}f_{q-1})$ of Λ' up to signed permutation. (These terminal HKZ reductions use SVP dimension $n/2$ rather than the $\lfloor n/4 \rfloor + 1$ that a recursive half-dimension LIP solve would give; the main loop already requires dimension $p = n/2 + 1$, so the theorem's bound is unaffected, and plain HKZ is also how [\[Duc23, Alg. 1\]](#) finishes for $\mathbb{Z}^d$.) $\square$

Corollary E.4. $[\mathbb{Z}^p \oplus \sqrt{s}\mathbb{Z}^q]\text{-LIP} \leq_{\det} \text{poly}(p) \times \text{SVP}(p)$ (*deterministic reduction*) for any integer $s \geq 2$ and $p > q \geq 0$. (Replace “mod 2” by “mod s ” in

[Lemma E.1](#); [Lemma E.2](#) Case 1 gives product $\leq 1/\sqrt{s} \leq \sqrt{1-1/p}$ for $p \geq 2$; Case 2 is unchanged.)

Remark E.5 (relation to [\[BN24\]](#)). Bambury–Nguyen [\[BN24, Thm 1\]](#) prove that for any rank- d lattice $\Lambda \subseteq \mathbb{Z}^d$ with $\lambda_1(\Lambda)\lambda_1(\Lambda^\vee) < 1 - \varepsilon$, a nonzero vector of norm $\leq \lambda_1(\Lambda)$ in Λ or one of norm $\leq \lambda_1(\Lambda^\vee)$ in $\Lambda^\vee$ can be found with polynomially many SVP calls in dimension $\lfloor d/2 \rfloor + 1$, terminating as soon as one is found; their Theorem 4 returns a full orthogonal basis of $\Lambda \cong \mathbb{Z}^d$ with a γ -approximate oracle for any $\gamma < \sqrt{2 - 2/d}$. The class $[\Lambda']$ satisfies the hypothesis of their Theorem 1 with $\lambda_1(\Lambda')\lambda_1(\Lambda'^\vee) = 1/\sqrt{2}$, and [Theorem 5.1](#) can be read as a near-hypercubic instance of their framework that returns all shortest vectors. The analysis here builds on theirs: [\[BN24, Lemma 1\]](#) expresses the per-round volume drop as the product of the primal and dual steps’ contributions, and their Theorem 1 bounds that product by $\lambda_1(\Lambda)\lambda_1(\Lambda^\vee)$ while neither half-block contains a shortest vector. [Lemma E.2](#) bounds the same product under $\text{vol}(S) > 1$ alone, using the structure of $[\Lambda']$, so that the loop runs until the front-block volume reaches 1, where all shortest vectors can be read off (the direct analogue of [\[Duc23, Lemma 4\]](#) fails for Λ' , §E.1). An alternative route iterates [\[BN24, Thm 1\]](#), removing each found primal or dual shortest vector and recursing; the residual lattices remain of the form $\mathbb{Z}^{p'} \oplus \sqrt{2}\mathbb{Z}^{q'}$ until $q' = 0$, where the hypercubic solver of [\[Duc23\]](#) applies. We organize the argument as a single termination bound.

F Falcon and NTRU-Type Schemes Are Not Affected

Falcon [\[PFH20\]](#) uses the same field $K_n = \mathbb{Q}(\zeta_{2^e})$, so τ exists. However, the τ -cocycle construction does not transfer, for two independent structural reasons. Throughout this section we use Falcon’s standard notation: q denotes the NTRU modulus (not the integer $q = n/2 - 1$ of [Corollary 4.6](#)), f, g, F, G denote the four ring elements of Falcon’s secret basis (unrelated to the field F of §§3–4), and σ_{sk} is the Gaussian width of the NTRU secret.

(i) **No public Hermitian Gram.** Falcon publishes $h = gf^{-1} \bmod q \in R_n/q$ [\[PFH20, eq. \(3.17\)\]](#), not the Gram of the secret basis B . Two secret keys with the same h have different Grams (they differ by left-multiplication by an element of $\text{GL}_2(R_n)$ that need not be unitary), so constraint (C2) of §4.1 — which has coefficients depending on the Gram matrix — is not a function of the public key. Without (C2), the lattice of matrices satisfying (C1) alone has rank $2n$, not n , and offers no advantage.

(ii) $\det B = q \neq 1$. [Lemma 4.2](#)’s integral isometry requires $B \in \text{SL}_2(R_n)$, and [Lemma 4.4](#)’s minimum $n/4$ requires $\det V_\tau = 1$. Falcon’s secret basis is $B = \begin{pmatrix} g & -f \\ F & -G \end{pmatrix}$ [\[PFH20, eq. \(2.2\)\]](#), and the NTRU equation $fG - gF = q$ [\[PFH20, eq. \(3.15\)\]](#) is $\det B = q \notin R_n^\times$; the resulting cocycle $qB^{-1}\tau(B)$ has determinant q^2 and $Q^{(\tau)}$ -value $q^2 \cdot n/4$, above the Gaussian heuristic (the expected first minimum of a random lattice of the same volume) of any publicly derived q -ary lattice (one containing $qM_2(R_n)$) — a *worse* target than $(f, g) \in A_h := \{(u, v) \in R_n^2 : uh \equiv v \bmod q\}$ directly.

More generally, consider any polynomial in the entries of B and their Galois conjugates whose norm is independent of σ_{sk} . Each monomial mixes short factors (f, g , of size $\sim \sigma_{\text{sk}}$) with long ones (the completion F, G , of size $\sim q/\sigma_{\text{sk}}$); for the product to be σ_{sk} -independent the two must balance, and any non-constant balanced monomial has norm $\geq \Theta(q)$. Reason (i) and this scale-imbalance argument apply to any NTRU-type scheme, not only to Falcon. The distribution-independent minimum of [Lemma 4.4](#) is possible for HAWK because all four entries of B have the same scale. To the best of our effort, none of the related side-constructions — descent via N_{K_n/K_n^τ} or $N_{K_n/F}$, or the LIP reduction of [Corollary E.4](#) applied to q -ary lattices — helps against Falcon: each either reduces to the standard subfield attacks [\[ABD16; KF17\]](#) or has no Falcon instance.