Lattice cryptography: isomorphisms & dual attacks

Pulles, L.N.

Chapter 6

Solving Module-LIP using Automorphisms

This chapter is based on joint work with D.M.H van Gent published in IACR's Communications in Cryptology [vGP25].

The search rank-2 module Lattice Isomorphism Problem, or smLIP, over a cyclotomic ring of degree a power of two, can be reduced to a LIP instance of at most half the rank if an adversary knows a nontrivial automorphism of the underlying integer lattice. Knowledge of such a nontrivial automorphism speeds up the key recovery attack on Hawk at least quadratically, which would halve the number of security bits.

The initial version of omSVP was trivially broken by the symplectic automorphism, and it was easily made hard again by including the symplectic automorphism as a trivial win. However, the question remains if more easily computable automorphisms exist. This work provides confidence in the soundness of this updated definition, assuming smLIP is hard, since there are plausibly no more trivial automorphisms that allow winning the omSVP game easily.

Although this work does not affect the security of Hawk, it opens up a new attack avenue involving the automorphism group that may be theoretically interesting on its own.

6.1 Introduction

The *Lattice Isomorphism Problem* (LIP) has recently been introduced as a building block for cryptography [DvW22, BGPS23], and was used in Chapter 5 to develop the fast and compact signature scheme named Hawk. LIP for the integer lattice ($\mathbb{Z}$LIP) has been studied well [Szy03, BM21, Duc24, BN24]. Moreover, LIP for *ideal lattices*, i.e., rank-1 module lattices, is solvable in polynomial time using the so-called Gentry–Szydlo algorithm [GS02, LS17, LS19]. However, the SUF-CMA security of Hawk, which uses a rank-2 module lattice for efficiency reasons, is based on a new problem called *one more shortest vector problem* (omSVP). Although omSVP has received some attention so far [DPPvW22a, LJPW24], it needs to be studied more extensively in order to gain confidence in Hawk.

Private key recovery for Hawk has received significantly more attention than omSVP. The former requires solving the so-called *search module LIP* (smLIP) between free modules of rank 2 over a (totally complex) cyclotomic number field. A recent line of work [MPPW24, LJPW24, APvW25] has developed a polynomial-time attack on smLIP over a number field with at least one real embedding, which remarkably is not the case with Hawk. Instead, smLIP for Hawk reduces uniformly to a principal ideal problem in a quaternion algebra [CME+25], and it remains unknown whether Gentry–Szydlo, the algorithm for the case of number fields, generalises to such algebras. Thus far, it seems smLIP is not much easier than $\mathbb{Z}$LIP despite the extra module structure.

The signature distribution of Hawk cannot be simulated as easily as that of NTRU-based signatures see Section 5.6.2, so it is not known how to instantiate the framework of [GPV08], which prevents leakage of the trapdoor basis [NR09, DN12b]. Instead, security of Hawk is based on omSVP, in which an adversary, given an oracle that produces short vectors, needs to output a short vector that is not a "trivial automorphism", such as negation, of any oracle output. The parameters of Hawk are chosen based on practical cryptanalysis. For these parameters, the security reduction to omSVP produces a trivial instance of omSVP. Nevertheless, the reduction shows soundness of the design, and larger parameters reduce to a presumably hard instance of omSVP but are unfavourable for the compactness of Hawk.

The formulation of omSVP is a delicate matter, because it needs

to include all trivial automorphisms that an adversary can compute easily. In the original formulation [DPPvW22a], the trivial automorphisms were believed to be multiplications by roots of unity, provided by the module structure in Hawk. Later, however, another trivial automorphism, stemming from the self-duality of the module lattice, was discovered [LJPW24]. This so-called *symplectic automorphism* solved the original omSVP trivially. Hence, when Hawk advanced to the second round of NIST's standardisation process of additional digital signature schemes [BBD+24], its specification document was amended to include the symplectic automorphism in the definition of omSVP. In other words, two vectors are deemed equivalent if they are in the same orbit under the group generated by the roots of unity and the symplectic automorphism. Although this amendment presumably makes omSVP harder, it does not provide a satisfactory answer to the following question.

Question 6.1. Which automorphisms can be computed easily by an adversary against the omSVP game?

6.1.1 Contributions

Our main result is the following theorem, which regards $\mathbb{Z}^{2n}$ as a module lattice R^2 over the ring $R = \mathbb{Z}[X]/(X^n + 1)$. The group $\mathcal{G}_{\mathbf{Q}}$ is defined in Definition 5.14.

Theorem 6.2. *Given a Hermitian form* $\mathbf{Q}$ *of the module lattice* R^2*, and an automorphism* $\mathbf{U} \in \mathrm{GL}_{2n}(\mathbb{Z}) \setminus \mathcal{G}_{\mathbf{Q}}$ *of* $\mathrm{Rot}(\mathbf{Q})$*, then, using standard lattice reduction heuristics, one can solve smLIP for* $\mathbf{Q}$ *by running BKZ-*β *with* $\beta = n/2 + 1$ *on some sublattice* Λ *constructed in polynomial time from* $\mathbf{U}$ *and that has rank at most* n*.*

Informally, one may think of Theorem 6.2 as follows: if one has a rotation of the module lattice R^2 and a nontrivial automorphism of the underlying integer lattice, one can recover that rotation by finding a shortest vector in a sublattice of half the rank and containing an unusually short vector.

Theorem 6.2 shows that with a nontrivial automorphism one can construct a specific lattice Λ of rank at most n. In Section 6.3, we show that BKZ-β heuristically recovers a shortest vector in such Λ for $\beta = n/2+1$, basically because this lattice is as unusual as (and similar to) $\mathbb{Z}^n$. The heuristics for $\mathbb{Z}^n$ have been experimentally verified [DPPvW22a],

and it is proved [Duc24] that BKZ-β recovers a unit vector in $\mathbb{Z}^n$ when $\beta = n/2 + o(n)$. Thus, Theorem 6.2 provides a *quadratic speed up* in solving smLIP, since BKZ-β runs in time exponential in β. Roughly speaking, if it is easy to find a nontrivial automorphism, then so is smLIP. We view this as evidence that computing a nontrivial automorphism is hard, although it may as well be that smLIP is just easier than what is currently assumed.

If we have an oracle for a *random* automorphism, then [JWL+23] solves SVP significantly faster. Even without the oracle, if the automorphism is random we can with high probability obtain a much better result. It is unreasonable to assume, however, that if an attacker obtains a nontrivial automorphism it will be a truly random one. More likely, it will be highly structured, like the symplectic automorphism. Hence, Theorem 6.2 justifies the choice for $\mathcal{G}$, *the group of trivial automorphisms*, in the definition of omSVP [BBD+26]. We know that $\mathcal{G}$ must include the roots of unity *and* the symplectic automorphism [LJPW24], because one can trivially win the omSVP game otherwise. Conversely, the theorem shows that smLIP and hence omSVP becomes much easier if one knows an additional automorphism $\sigma \notin \mathcal{G}$. If there is any choice of $\mathcal{G}$ that would make omSVP hard, then the current $\mathcal{G}$ is the smallest such choice.

The situation is summarised in Figure 6.1.

6

6.2 Main result

In this section we will prove the main result. Recall the definitions of $\mathcal{G}$ and $\mathcal{G}_{\mathbf{Q}}$ from Definition 5.14, in particular $\mathcal{G} = \langle X, \omega \rangle$.

Lemma 6.3. *The group $G = \mathcal{G}$ satisfies all the following properties.*

1. *G acts regularly, i.e., transitively and freely, on the shortest vectors of the lattice R^2;*
2. *G has exactly one element of order 2, namely $-1 = X^n = \omega^2$;*
3. *G is isomorphic to $(\mu \rtimes \langle\omega\rangle)/\langle -1\rangle \cong (C_{2n} \rtimes C_4)/C_2$, where ω acts on μ by $X \mapsto X^\star = X^{-1}$,*
4. *the multiplication map $\mu \times \{1, \omega\} \to G$ is a bijection.*

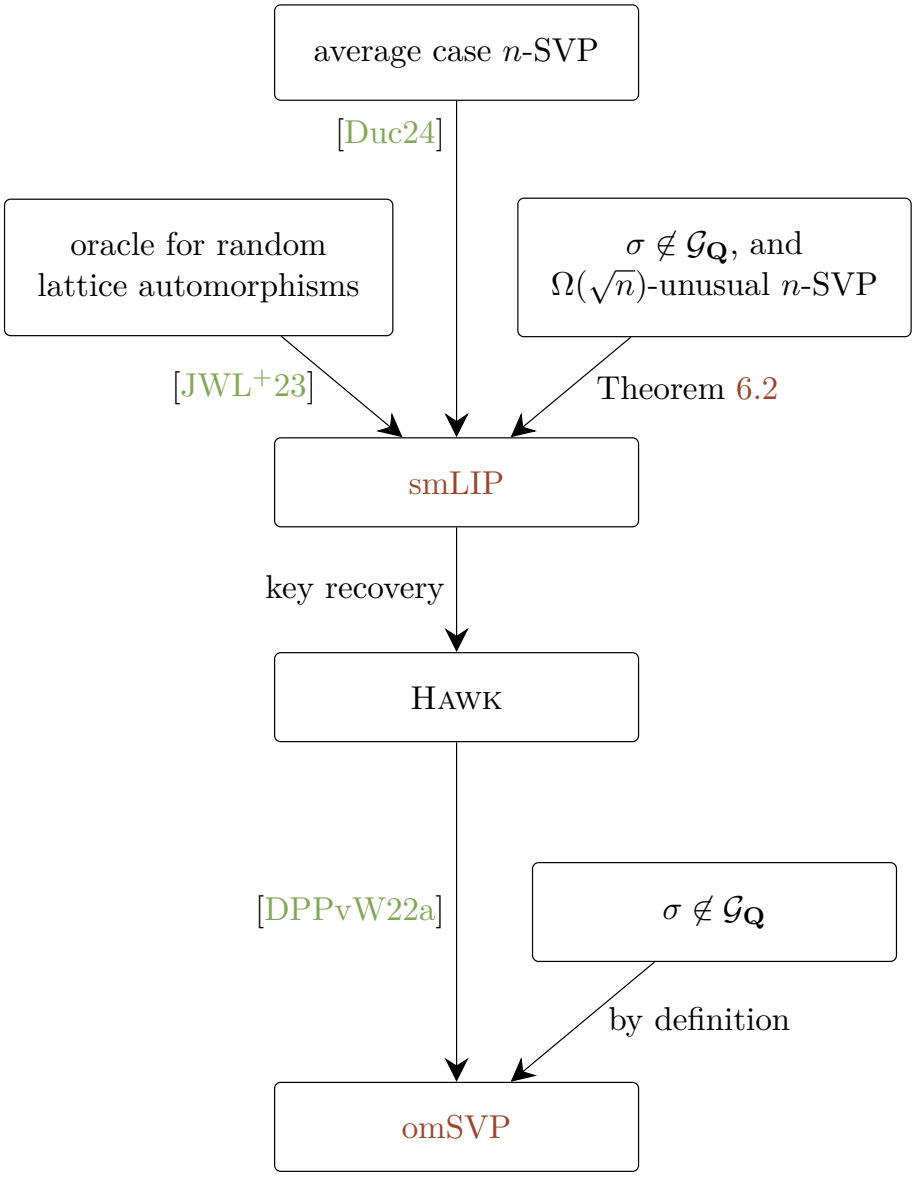

Figure 6.1: Reductions between problems related to the SUF-CMA security of Hawk. An arrow $A \to B$ indicates B reduces to A, i.e., B can be solved in polynomial time by making multiple queries to an algorithm that solves A.

Proof. For all $x, y \in R$ we have

$$\omega X \begin{pmatrix} x \\ y \end{pmatrix} = \omega \begin{pmatrix} Xx \\ Xy \end{pmatrix} = \begin{pmatrix} +X^\star \cdot y^\star \\ -X^\star \cdot x^\star \end{pmatrix} = X^\star \begin{pmatrix} +y^\star \\ -x^\star \end{pmatrix} = X^\star \omega \begin{pmatrix} x \\ y \end{pmatrix},$$

so $\omega X \omega^{-1} = X^\star$. It follows that $\mu \rtimes \langle \omega \rangle \to G$ is a (surjective) group homomorphism. Its kernel consists of all pairs (X^i, ω^j) such that $X^i = \omega^{-j}$. Since $\omega^2 = -1 \in \mu$ and $\omega \notin \mu$, we conclude that the kernel is generated by $(-1, -1)$, establishing the group structure. It follows that the multiplication map $\mu \times \{1, \omega\} \to G$ is a bijection.

In μ, the only element of order 2 is -1. As $(X^i\omega)^2 = X^i \omega X^i \omega = X^i X^{-i} \omega\omega = \omega^2 = -1$, we conclude that -1 is the only element of G of order 2.

The set S of shortest vectors of R^2 equals $(\mu \times \mathbf{0}) \sqcup (\mathbf{0} \times \mu)$. It is easy to see that the action of G on S is transitive: under $\mu \subseteq G$ there are the two orbits $\mu \times \mathbf{0}$ and $\mathbf{0} \times \mu$, while ω interchanges the two. Suppose $\mathbf{x} \in S$ is fixed by $X^i\omega^j$ with $j \in \{0, 1\}$. Then $j = 0$, otherwise $\mathbf{x}$ and $\omega^j \mathbf{x}$ are in different orbits under μ. It follows that $X^i \mathbf{x} = \mathbf{x}$, which is only possible if $X^i = 1$. Hence, the action is free. □

Lemma 6.4. *There exists a polynomial time algorithm that, given a Hermitian form* $\mathbf{Q}$ *of* R^2 *and* $\mathbf{x} \in R^2$ *of length* 1 *or* $\sqrt{2}$*, computes all shortest vectors of* $\mathbf{Q}$*.*

Proof. By [LJPW24] we may compute $G = \mathcal{G}_\mathbf{Q}$.

Case $\|\mathbf{x}\| = 1$: We may compute by Lemma 6.3 all shortest vectors as the orbit of $\mathbf{x}$ under G.

Case $\|\mathbf{x}\| = \sqrt{2}$: Compute for each $\sigma \in G \setminus \{1\}$ the element

$$\mathbf{y}_\sigma = \Big(\sum_{i=0}^{k_\sigma - 1} \sigma^i \Big)(\mathbf{x}) \ ,$$

where $k_\sigma = \mathrm{ord}(\sigma)/2$. We will show that $\mathbf{y}_\sigma$ is twice a shortest vector for some σ, in which case we may reduce to the previous case with $\mathbf{x} \leftarrow \mathbf{y}_\sigma/2$. Note that $\mathbf{x} = \mathbf{x}_1 - \mathbf{x}_2$ for some shortest vectors $\mathbf{x}_1 \neq \mathbf{x}_2$. By Lemma 6.3 there is some $\sigma \in G$ such that $\sigma(\mathbf{x}_1) = \mathbf{x}_2$, and $\sigma^{k_\sigma} = -1$. Then

$$\begin{aligned} \mathbf{y}_\sigma &= \sum_{i=0}^{k_\sigma - 1} \sigma^i(\mathbf{x}_1) - \sum_{i=0}^{k_\sigma - 1} \sigma^i(\mathbf{x}_2) = \sum_{i=0}^{k_\sigma - 1} \sigma^i(\mathbf{x}_1) - \sum_{i=0}^{k_\sigma - 1} \sigma^{i+1}(\mathbf{x}_1) \\ &= \mathbf{x}_1 - \sigma^{k_\sigma}(\mathbf{x}_1) = 2\mathbf{x}_1, \end{aligned}$$

as was to be shown. □

Proposition 6.5. *There is a polynomial-time reduction from smLIP on a Hermitian form* $\mathbf{Q}$ *of* R^2 *to finding a nonzero vector of length at most* $\sqrt{2}$ *of* $\mathbf{Q}$.

Proof. Let $\mathbf{Q}$ be a Hermitian form of R^2 and suppose we have such a vector. By Lemma 6.4 we may compute the shortest vectors. The action of μ partitions them into two orbits, and let $\mathbf{b}_1$ and $\mathbf{b}_2$ be elements of the two orbits. With $\mathbf{B} = [\mathbf{b}_1, \mathbf{b}_2]$ we have $\mathbf{B}^\star\mathbf{Q}\mathbf{B} = \mathbf{I}_2$, so we may return $\mathbf{B}^{-1}$. □

In the lemma below, recall A_k is the root lattice of type A from Example 2.14.

Lemma 6.6. *Let* $\sigma \in \mathrm{O}(\mathbb{Z}^n)$ *with order dividing a prime* p *and let* $\Phi_p(Z) = Z^{p-1} + Z^{p-2} + \cdots + 1$. *Then the sublattices* $\Lambda_1 = \ker(\sigma - 1)$ *and* $\Lambda_2 = \ker(\Phi_p(\sigma))$ *of* $\mathbb{Z}^n$ *satisfy* $\mathrm{rk}(\Lambda_1) + \mathrm{rk}(\Lambda_2) = n$. *Moreover, we have*

$$\Lambda_1 \cong \mathbb{Z}^a \times \sqrt{p} \cdot \mathbb{Z}^c \quad \text{and} \quad \Lambda_2 \cong \mathbb{Z}^b \times A_{p-1}^c, \tag{6.1}$$

for some $a, b, c \in \mathbb{N}$, *with* $b = 0$ *if* $p > 2$. *If* $\sigma \neq \pm 1$, *we have* $\Lambda_1, \Lambda_2 \neq \mathbf{0}$.

Proof. Since σ is a signed permutation, it partitions the coordinates of $\mathbb{Z}^n$. This gives a decomposition $\mathbb{Z}^n = L_1 \oplus \cdots \oplus L_k$ into pairwise orthogonal sublattices, with each L_i isomorphic to either $\mathbb{Z}$ or $\mathbb{Z}^p$. With $\pi_i : \mathbb{Z}^n \to L_i$ the corresponding projections, we have $\sigma \circ \pi_i = \pi_i \circ \sigma$. It follows from linear algebra that

$$\Lambda_j = (L_1 \cap \Lambda_j) \oplus \cdots \oplus (L_k \cap \Lambda_j) \ ,$$

holds for $j = 1$ and $j = 2$. Hence, it is sufficient to prove the lemma for the case $\mathbb{Z}^n = L_1$.

If $n = 1$, then $\sigma = \pm 1$. Hence, (Λ_1, Λ_2) equals $(\mathbb{Z}, \mathbf{0})$, or if $p = 2$ possibly $(\mathbf{0}, \mathbb{Z})$. In particular, they are of the prescribed forms.

Suppose now that $n = p$. Additionally assume that σ is just a permutation without sign flips, which is automatic when p is odd. Let

$$B_1 = \mathbb{Z}\mathbf{z} \quad \text{and} \quad B_2 = \{\mathbf{x} \in \mathbb{Z}^n \mid \mathbf{z}^\mathsf{T}\mathbf{x} = 0\},$$

where $\mathbf{z}$ is the all-1 vector in $\mathbb{Z}^n$. It is easy to verify that $B_1 \subseteq \Lambda_1$ and $B_2 \subseteq \Lambda_2$. Since $\mathbb{Q}B_1 + \mathbb{Q}B_2 = \mathbb{Q}^n$, we have by linear algebra that

$\Lambda_j = (\mathbb{Q}B_j) \cap \mathbb{Z}^n = B_j$ $(j = 1, 2)$. Hence, $\Lambda_1 = \mathbb{Z}\mathbf{z} \cong \sqrt{p} \cdot \mathbb{Z}$ and $\Lambda_2 = A_{p-1}$.

Suppose now that $n = 2$ and σ is not a permutation. Then $-\sigma$ is a permutation, otherwise $\sigma^2 \neq 1$. Compared to σ, this results in interchanging the corresponding lattices Λ_1 and Λ_2. Since $A_1 \cong \sqrt{2} \cdot \mathbb{Z}$, we are done. □

Proposition 6.7. *There exists a polynomial-time algorithm that, given a Hermitian form* $\mathbf{Q}$ *of* R^2 *and* $\sigma \in \mathrm{O}(\mathrm{Rot}(\mathbf{Q})) \setminus \mathcal{G}_{\mathbf{Q}}$*, computes a nonzero sublattice* Λ *of* $\mathrm{Rot}(\mathbf{Q})$ *of rank at most* n *such that* $\lambda_1(\Lambda) \leq \sqrt{2}$*.*

Proof. We may compute in polynomial time, e.g. using the characteristic polynomial of σ, the multiplicative order of σ. In particular, we may test whether $-1 \in \langle\sigma\rangle$.

Assume first that this is not the case. Then by replacing σ by some power, we may assume σ has prime order p and $\sigma \neq \pm 1$. We may now use Lemma 6.6 with $\sigma \in \mathrm{O}\,\mathbb{Z}^{2n}$, which allows us to obtain the lattices Λ_1 and Λ_2 from this lemma using linear algebra. Let Λ to be the one of minimal rank. Since $\mathrm{rk}(\Lambda_1) + \mathrm{rk}(\Lambda_2) = 2n$, we have $\mathrm{rk}\,\Lambda \leq n$. It remains to show that $\lambda_1(\Lambda_1), \lambda_1(\Lambda_2) \leq \sqrt{2}$. For Λ_2, this immediately follows from Lemma 6.6 and $\lambda_1(A_{p-1}) = \sqrt{2}$. If $p = 2$, then the same follows for Λ_1. Suppose p is odd. Then $\Lambda_1 \cong \mathbb{Z}^a \times \sqrt{p} \cdot \mathbb{Z}^c$ and $\Lambda_2 \cong A_{p-1}^c$ for some a, c. As $a + pc = \mathrm{rk}(\Lambda_1) + \mathrm{rk}(\Lambda_2) = 2n$ is a power of 2 and p is odd, we conclude that $a > 0$. Hence, $\lambda_1(\Lambda_1) = 1$ and we are done.

Now consider the case for general σ. By [LJPW24], we may compute $G = \mathcal{G}_{\mathbf{Q}}$. It suffices to show that $-1 \notin \langle\tau\rangle$ for some $\tau \in G \cdot \sigma$, since we may then apply the above argument on τ in the rôle of σ. Let $\mathbf{x} \in R^2$ be a shortest vector. Then by Lemma 6.3 there exists some $\rho \in G$ such that $\rho(\mathbf{x}) = \sigma(\mathbf{x})$. Hence, $\tau = \rho^{-1} \cdot \sigma$ fixes $\mathbf{x}$, and $\tau \neq 1$ because otherwise $\sigma \notin G$. Since -1 fixes no shortest vectors, we conclude that $-1 \notin \langle\tau\rangle$. □

6.3 Heuristic hardness of SVP

In this section, we analyze the hardness of recovering a shortest vector from a lattice Λ appearing in Proposition 6.7. Specifically, we determine the minimal block size β such that BKZ-β finds a shortest vector in Λ, using standard heuristics in lattice reduction.

We use the methodology based on the "2016 estimates" which were phrased for lattices related to *Learning with Errors* [ADPS16], and verified experimentally [AGVW17, DDGR20, PV21]. Here, we apply this methodology to a lattice of the form as in Equation (6.1). These 2016 estimates have already been used to motivate the experimentally verified estimates that BKZ-β recovers the shortest vector in a rotation of $\mathbb{Z}^n$ when $\beta = n/2 + o(n)$ [DPPvW22a].

Definition 6.8 ([DvW22])**.** A rank-n lattice Λ is an *f-unusual-SVP* instance if we have

$$f \cdot \lambda_1(\Lambda) \leq \mathrm{GH}(n) \cdot \det(\Lambda)^{1/n}.$$

For example, $\mathbb{Z}^n$ and A_n are $\Omega(\sqrt{n})$-unusual-SVP instances, as n tends to infinity.

The hardness of recovering a shortest vector in an unusual-SVP lattice is mainly determined by the ratio $\mathrm{GH}(n)\det(\Lambda)^{1/n}/\lambda_1(\Lambda)$, rather than the UniqueSVP factor $\lambda_2(\Lambda)/\lambda_1(\Lambda)$ [AD21]. For such f-unusual-SVP instance with an unusually short vector $\mathbf{v}$, a *threshold phenomenon* occurs in the terminal block of a BKZ-β tour once the projection of $\mathbf{v}$ onto this block is much shorter than the expected first minimum if that terminal block had been a random lattice. Now when this $\mathbf{v}$ is part of the basis at position $\mathrm{rk}(\Lambda) - \beta + 1$, subsequent tours of BKZ-β will then move $\mathbf{v}$ to the front of the basis with steps of $\beta - 1$ per tour [PV21].

The lattice computed in Proposition 6.7 is of the form $\mathbb{Z}^a \times \sqrt{p}\mathbb{Z}^c$ (with $a > 0$) or $\mathbb{Z}^b \times A_{p-1}^c$, and thus is an instance of the $\Omega(\sqrt{k})$-unusual-SVP, where $k = \mathrm{rk}(\Lambda) \leq n$. In the following heuristic, let $\delta_\beta = \mathrm{GH}(\beta)^{1/(\beta-1)}$ be the *root Hermite factor*, and note $\delta_\beta \approx (\beta/(2\pi e))^{\frac{1}{2(\beta-1)}}$ for $\beta > 50$.

Heuristic 6.9. *Suppose $\Lambda \subseteq R^2$ is a nonzero rank-k lattice with $\lambda_1(\Lambda) \leq \sqrt{2}$. If $\beta \in \mathbb{Z}_{\geq 2}$ satisfies*

$$\sqrt{2\beta/k} \leq \delta_\beta^{2\beta-k-1}, \tag{6.2}$$

then BKZ-β will recover a shortest vector of Λ. In particular, this condition holds asymptotically for $\beta = k/2 + 1$.

Heuristic Justification. Because Λ can be identified with a sublattice of $\mathbb{Z}^{2n}$, by [Duc24, Lem. 2], its volume is at least 1. Note that the projection of a shortest vector of Λ onto the terminal block of a BKZ-β tour has

an expected norm of $\lambda_1(\Lambda) \cdot \sqrt{\beta/k} \leq \sqrt{2\beta/k}$. Moreover, if the terminal block during a BKZ-β tour would be a random block, its first minimum would be $\delta_\beta^{2\beta-k+1} \cdot \det(\Lambda)^{1/k} \geq \delta_\beta^{2\beta-k+1}$. Thus, if Equation (6.2) holds, then the [ADPS16] success condition

$$\lambda_1(\Lambda) \cdot \sqrt{\beta/k} \leq \delta_\beta^{2\beta-k+1} \cdot \det(\Lambda)^{1/k}$$

is satisfied, and we expect a threshold phenomenon to occur. For the last statement, note the left hand side of Equation (6.2) is equal to $\sqrt{1+2/k} \leq e^{1/k}$, and the right hand side of Equation (6.2) is at least $(k/(4\pi e))^{1/k}$, which is larger than $e^{1/k}$ once $k > 100$. △

Note that we use the *Geometric Series Assumption* (GSA) for this estimate, i.e., we assume the lengths of the Gram–Schmidt basis vectors follow a geometric series, and there are better simulators for these lengths [CN11, BSW18]. Still, the GSA gives a decent approximation. In addition, there is a more refined simulator that can take multiple shortest vectors of a lattice (e.g. n for $\mathbb{Z}^n$) into account [DDGR20], which expects a slightly lower required block size. However, because we are interested in the worst-case instance of Λ occurring in Proposition 6.7, such a refined simulator does not help in cases when there is a single shortest vector.

6

6.4 Group-theoretic heuristics

Our results are strong in that they impose no conditions on the automorphism, besides being nontrivial, but they 'only' halve the security parameter of Hawk. As [JWL+23] shows, under the stronger assumption that an attacker has access to an oracle giving automorphisms of a lattice, SVP can be (probabilistically) solved in polynomial time. If instead of having an oracle, the attacker generates a single uniformly random automorphism, then we can, as we will argue, heuristically break Hawk with high probability.

Suppose σ is uniformly sampled from $\mathrm{O}(\mathrm{Rot}(\mathbf{Q}))$, and let $g \in S_{2n}$ be the corresponding permutation on the shortest vectors of $\mathrm{Rot}(\mathbf{Q}) \cong \mathbb{Z}^{2n}$ modulo sign, i.e., on the coordinates of the lattice. Then, g is a uniformly random permutation. Each orbit $\Omega/\{\pm 1\}$ of g comes with a sign which is negative if the corresponding action of σ on Ω is transitive. Let c_k^s be the number of orbits of g of length $k \in [2n]$ and sign $s \in \{-, +\}$. One

can show, analogously to Lemma 6.6, that the characteristic polynomial of σ equals

$$\prod_{s\in\{-,+\}}\prod_{k=1}^{2n}(X^{2n}-s)^{c_k^s},$$

and that we have

$$\Lambda_1 = \ker(\sigma-1) \cong \prod_{k=1}^{2n}(\sqrt{k}\cdot\mathbb{Z})^{c_k^+}.$$

We will now argue that with high probability $\Lambda_1 \neq 0$ and $\operatorname{rk}\Lambda_1 = O(\log n)$ holds. In particular, smLIP reduces, given σ, to a comparatively trivial instance of SVP.

The probability that g has no fixed points is approximately $1/e$, so we may assume that g *has* a fixed point. Alternatively, one notes that if we multiply σ by the unique element of $\mathcal{G}_{\mathbf{Q}}$ such that the result g' fixes some preselected shortest vector $\mathbf{x}$, then g' is a uniformly random permutation among all permutations fixing $\mathbf{x}$. Multiplying σ by -1 if necessary, we may assume $c_1^+ > 0$. In particular, $\Lambda_1 \neq \{\mathbf{0}\}$ and $\lambda_1(\Lambda_1) = 1$. The expected number of orbits of g is

$$\sum_{k=1}^{n} 1/k \approx \log(n)\ .$$

Hence, the expected rank of Λ_1 is at most $\log(n)$.

On a less rigorous note, the group generated by $\mathcal{G}_{\mathbf{Q}}$ and σ will often be significantly larger than $|\mathcal{G}_{\mathbf{Q}}|\cdot|\langle\sigma\rangle|$. It is also not unlikely that it maps surjectively to S_{2n}. In this case it becomes possible to sample random-enough automorphisms, bringing us to the regime of [JWL+23].

6.5 Conclusion

Theorem 6.2 now easily follows from combining Proposition 6.7 and Heuristic 6.9. Namely, given a Hermitian form $\mathbf{Q}$ of R^2 and a nontrivial automorphism σ of $\textsc{Rot}(\mathbf{Q})$, Proposition 6.7 computes in polynomial time a basis for a lattice Λ of rank at most n, such that recovering a shortest vector of Λ allows solving smLIP for $\mathbf{Q}$. Then, Heuristic 6.9 shows, heuristically, that BKZ-β recovers a shortest vector of Λ when $\beta = \operatorname{rk}(\Lambda)/2+1 \leq n/2+1$.

Based on the results of Section 6.4, it is reasonable to suspect that Λ is of much lower rank in the average case. Indeed, sampling σ uniformly at random likely gives a lattice of rank at most $\log(n)$, for which directly solving SVP only takes polynomial time in n.

In practice, if one happens to find a nontrivial automorphism σ, it may be worthwhile to take a random composition of σ, X and $\omega_{\mathbf{Q}}$'s until finding a lattice of rank at most $\log(n)$, since $\langle \sigma, X, \omega_{\mathbf{Q}} \rangle = \mathrm{O}(\mathrm{Rot}(\mathbf{Q}))$ may happen.

Although we prove a worst-case result that, given a nontrivial automorphism, block size $n/2 + 1$ is heuristically sufficient, in practice we expect one could find a lattice of extremely low rank and subsequently solve smLIP and break Hawk. In this light, we believe that smLIP is *practically* as hard as finding a nontrivial automorphism, but theoretically not harder than finding a nontrivial automorphism and solving SVP on an easier lattice.