On the Hardness of some Vandermonde Knapsack problems

Dipayan Das[1] and Arindam Mukherjee[2]

[1] Department of Mathematics and Statistics, Florida Atlantic University, United States
dasd@fau.edu
[2] Department of Mathematics, Achhruram Memorial College, India
arindamaths@gmail.com

Abstract. The Vandermonde Knapsack problem comprises a family of algebraic variants of the Knapsack problem. This includes the Partial Vandermonde (PV) Knapsack problem (DCC'15, ACNS'14, ACISP'18, DCC'20, Indocrypt'25), the Vanishing SIS (vSIS)-based commitment problem (Crypto'23, PKC'25), and related assumptions. These problems have played an important role in enabling efficient lattice-based cryptographic constructions.
Recently, two independent works by Boudgoust, Gachon, and Pellet-Mary (Crypto'22), and by Das and Joux (Eurocrypt'24), proposed attacks demonstrating that certain instances of the PV Knapsack problem are weak. In this paper, we present new attacks on the PV Knapsack problem for power-of-two cyclotomic rings. By combining our techniques with the attack of Das and Joux, we show that a substantially larger fraction of keys are weak in this setting than was previously known. We then extend our attack to the integer variant of the vSIS commitment problem, demonstrating that certain instances are also weak for specific parameter regimes.

Keywords: Lattice-based assumptions, Lattice reductions, Vanishing SIS commitment problem, PV Knapsack problem

1 Introduction

The Knapsack problem[3] is predominantly the most used post-quantum assumption till date. However, for efficiency reasons, it is often desirable to use an algebraic variant of the problem. Recently, the Vandermonde variant of the Knapsack problem has found a lot of applications in post-quantum cryptography.

In this paper, we fix q as a prime integer of (sufficiently) large size. Let $\mathfrak{R}_q$ be a ring of characteristics q [4] and $\mathcal{S} \subseteq \mathfrak{R}_q$. Let $\boldsymbol{f}(\boldsymbol{x}) \in \mathfrak{R}_q[\boldsymbol{x}]$ be a hidden polynomial of degree less than n with coefficients being ternary (or binary) elements in $\mathfrak{R}_q$.

[3] In this paper, we use "knapsack problem" to refer the low-density inhomogeneous SIS or LWE problem.
[4] For simplicity, one can assume $\mathfrak{R}_q := \mathbb{Z}_q$.

Let $t = \epsilon \cdot n$, where $0 < \epsilon < 1$. Informally, the Vandermonde Knapsack problem is the following:

Given t uniformly random evaluations of $\boldsymbol{f}$ from $\mathcal{S}$, find the hidden polynomial $\boldsymbol{f}$.

Let $v \in \mathcal{S}$. If we denote f_i as the i-th coefficient of $\boldsymbol{f}$ represented in the polynomial basis in $\boldsymbol{x}$, then the evaluation $\boldsymbol{f}$ at v is denoted by $\boldsymbol{f}(v)$ and we write

$$\boldsymbol{f}(v) := \sum_{i=0}^{n-1} f_i v^i \tag{1}$$

Since evaluations can be expressed in the powers of v and the coefficients of $\boldsymbol{f}$ is small, the recovery of $\boldsymbol{f}$ from the evaluations is thus called the Vandermonde Knapsack problem. Now, based on the choice of parameter selection, we have different variants of the problem, which lie in the family of Vandermonde Knapsack problem. In this paper, we are mostly interested in the PV Knapsack and integer variant of the vSIS commitment problem.

We first give an informal description of the two problems. From the context, the rationale for placing both within the family of Vandermonde knapsack problems will become clear.

In the PV Knapsack problem, we first set $\mathfrak{R}_q := \mathbb{Z}_q$. Then, we fix the integer n as a power of two, $q = 1 \bmod 2n$, and $\mathcal{S} := \{v \in \mathfrak{R}_q : v^n = -1\} \subset \mathfrak{R}_q$. While in the (univariate) integer vSIS commitment problem, we also fix $\mathfrak{R}_q := \mathbb{Z}_q$, but $\mathcal{S} := \mathfrak{R}_q^* \subset \mathfrak{R}_q$.

The PV Knapsack problem, which was initially called the partial Fourier recovery problem, was introduced in [13]. The problem has served as a foundational building block for several cryptographic applications, e.g., digital signatures [13,17], encryption schemes [14,4], trapdoor-free accumulator [5]. The vSIS problem was introduced in two separate works [11,6]. The problem was used for an efficient post-quantum functional commitment scheme for linear relations and efficient lattice-based succinct arguments[5]. Later, the problem was used to construct new homomorphic signatures [15] and lattice-based anonymous credential signatures [8].

Despite the Vandermonde structure, the underlying assumptions are comparable to that of the standard knapsack problem. Indeed, in [11], the authors wrote the following: *We note that similar assumptions have been proposed by prior works (see [BSS22] for a summary). Thus far, there are no known algorithms which solve these variants faster than their non-Vandermonde counterparts.*

Recent work has explored methods for exploiting the algebraic structure of the problem [3,7]. In particular, these attacks identify keys that are easier to solve than those covered by the general hardness assumption. When such a system is deployed by many users, each with their own public key, some users

[5] In [11], this problem is referred to as the Ring-Vandermonde SIS commitment. In this work, we refer to both formulations as the vSIS commitment problem.

may inadvertently select weak keys. However, the proposed attacks apply only to the PV knapsack problem. As far as we know, these results do not affect other problems in the Vandermonde family of knapsack problems.

In this paper, we further investigate the hardness of several variants of the Vandermonde knapsack problem. Given its wide range of cryptographic applications, it is important to examine the underlying hardness assumptions more carefully. In particular, we focus on identifying alternative algorithms for solving the Vandermonde knapsack problem efficiently.

2 Preliminaries

2.1 Notations

For any positive integer N, we denote $\phi(N)$ as the Euler's totient function, i.e., the number of integers between 1 and N that are co-prime to N. We write $\mathbb{Z}_N$ to denote the ring of integers modulo N. We also write $\mathbb{Z}_N^*$ to denote the multiplicative group of units modulo N. Then $|\mathbb{Z}_N^*| = \phi(N)$. In particular, when q is a prime, $\mathbb{Z}_q$ is a finite field, so $|\mathbb{Z}_q^*| = \phi(q) = q - 1$. For any element $\beta \in \mathbb{Z}_q^*$, its (multiplicative) order is defined as the smallest positive integer r such that $\beta^r = 1 \in \mathbb{Z}_q^*$. Since q is prime, r divides $q - 1$. Assuming q is an odd prime, we represent its element by the unique representative in the integer interval $[-\frac{q-1}{2}, \frac{q-1}{2}]$.

All the vectors used in this paper are column vectors and are denoted with bold letters. For any vector $\boldsymbol{v}$, we use $\|\boldsymbol{v}\|$ (resp. $\|\boldsymbol{v}\|_\infty$) to signify the ℓ_2 norm (resp. ℓ_∞) of the vector $\boldsymbol{v}$. Further, we identify polynomials and vectors as being the same data type using the coefficient embedding.

Let $\mathcal{R}_q := \mathbb{Z}_q[\boldsymbol{x}]/(\mathbf{g})$ denote the quotient polynomial ring of $\mathbb{Z}_q[\boldsymbol{x}]$ by the polynomial $\mathbf{g}(\boldsymbol{x})$ of degree n. In this paper, we fix $\mathbf{g}(\boldsymbol{x}) := \boldsymbol{x}^n + 1$ where n is an integer of the form power of two, unless stated otherwise. This is known as the power of two cyclotomic ring. When $q = 1 \bmod 2n$, the polynomial $\mathbf{g}$ factors linearly over $\mathbb{Z}_q$. In such a case, we denote Ω as the set of all the primitive roots of $\mathbf{g}$ in $\mathbb{Z}_q$. In particular, for any primitive root ω, we have

$$\Omega = \{\omega^i : i \in \mathbb{Z}_{2n}^*\} \text{ when } n \text{ is a power of two}$$

To represent an element in $\mathcal{R}_q$, we use the polynomial basis $\{1, \boldsymbol{x}, \ldots, \boldsymbol{x}^{n-1}\}$. Let $\theta_n \in \mathbb{Z}_{2n}^*$. In the quotient polynomial ring $\mathcal{R}_q$, for any polynomial $\boldsymbol{f} \in \mathcal{R}_q$, we can also interpret $\boldsymbol{f}(\boldsymbol{x}^{\theta_n})$ as a polynomial in $\mathcal{R}_q$. More precisely, if we write $\boldsymbol{f}(\boldsymbol{x}) = f_0 + f_1\boldsymbol{x} + \cdots + f_{n-1}\boldsymbol{x}^{n-1}$, then

$$\boldsymbol{f}(\boldsymbol{x}^{\theta_n}) = f_0 + f_1\boldsymbol{x}^{\theta_n} + \cdots + f_{n-1}\boldsymbol{x}^{\theta_n n-1}$$

In particular, since we are working modulo $\mathbf{g}(\boldsymbol{x})$, taking $\theta_n = -1, n+1, n-1 \in \mathbb{Z}_{2n}^*$, we have the following, respectively:

$$\boldsymbol{f}(\boldsymbol{x}^{-1}) = \boldsymbol{f}(1/\boldsymbol{x}) = f_0 - f_1\boldsymbol{x}^{n-1} - f_2\boldsymbol{x}^{n-2} - \cdots - f_{n-1}\boldsymbol{x}$$
$$\boldsymbol{f}(\boldsymbol{x}^{n+1}) = \boldsymbol{f}(-\boldsymbol{x}) = f_0 - f_1\boldsymbol{x} + f_2\boldsymbol{x}^2 - \cdots - f_{n-1}\boldsymbol{x}^{n-1}$$
$$\boldsymbol{f}(\boldsymbol{x}^{n-1}) = \boldsymbol{f}(1/(-\boldsymbol{x})) = f_0 + f_1\boldsymbol{x}^{n-1} - f_2\boldsymbol{x}^{n-2} + \cdots + f_{n-1}\boldsymbol{x}$$

Now let $\mathbb{Z}_q^{<n}[\boldsymbol{x}]$ denotes the set of polynomials of degree $< n$ in the ring $\mathbb{Z}_q[\boldsymbol{x}]$. Any element in $\mathbb{Z}_q^{<n}[\boldsymbol{x}]$ can be represented using the polynomial basis $\{1, \boldsymbol{x}, \ldots, \boldsymbol{x}^{n-1}\}$. Let $\theta_q \in \mathbb{Z}_q^*$. For any polynomial $\boldsymbol{f} \in \mathbb{Z}_q^{<n}[\boldsymbol{x}]$, we can interpret $\boldsymbol{f}(\theta_q\boldsymbol{x})$ as an element in $\mathbb{Z}_q^{<n}[\boldsymbol{x}]$. More precisely, if $\boldsymbol{f}(\boldsymbol{x}) = f_0 + f_1\boldsymbol{x} + \cdots + f_{n-1}\boldsymbol{x}^{n-1}$, then we have

$$\boldsymbol{f}(\theta_q\boldsymbol{x}) = f_0 + f_1\theta_q\boldsymbol{x} + f_2\theta_q^2\boldsymbol{x}^2 + \cdots + f_{n-1}\theta_q^{n-1}\boldsymbol{x}^{n-1}$$

In particular, if we take $\theta_q = -1 \in \mathbb{Z}_q^*$, we have the following:

$$\boldsymbol{f}(-\boldsymbol{x}) = f_0 - f_1\boldsymbol{x} + \cdots - f_{n-1}\boldsymbol{x}^{n-1}$$

2.2 The PV Knapsack problem

Let Ω_t be a random subset of Ω of size $t = \lfloor n/2 \rfloor$ containing t distinct elements. Also, let $\boldsymbol{f}(\boldsymbol{x})$ be an (hidden) element in $\mathcal{R}_q$ whose coefficients are sampled uniformly at random from the set $\{-1, 0, 1\}$.

Definitions 1 (PV Knapsack problem) *Given $\mathcal{R}_q$ and $\{(\omega, \boldsymbol{f}(\omega)) : \omega \in \Omega_t\}$, find the hidden polynomial $\boldsymbol{f}$.*

The problem can also be defined with $\mathbf{g}(\boldsymbol{x}) := \boldsymbol{x}^n - 1$ when n is a prime. In this case, we take $q = 1 \bmod n$ to ensure that the $\mathbf{g}(\boldsymbol{x})$ factors linearly over $\mathbb{Z}_q$. Here, for any primitive root ω, we have

$$\Omega = \{\omega^i : i \in \mathbb{Z}_n^*\} \text{ when } n \text{ is a prime}$$

2.3 The Integer vSIS commitment problem

Assume n is a positive integer of sufficiently large size. Let $\mathcal{S}_t$ be a random subset of $\mathbb{Z}_q^* \setminus \{-1, 1\}$ of size $t = \lfloor n/2 \rfloor$. Also, let $\boldsymbol{f}(\boldsymbol{x})$ be an (hidden) element in $\mathbb{Z}_q^{<n}[\boldsymbol{x}]$ whose coefficients are sampled uniformly at random from the set $\{0, 1\}$.

Definitions 2 (Integer vSIS commitment problem) *Given $\mathbb{Z}_q$ and $\{(\upsilon, \boldsymbol{f}(\upsilon)) : \upsilon \in \mathcal{S}_t\}$, find the hidden polynomial $\boldsymbol{f}$.*

The definition naturally extends from the integer ring $\mathbb{Z}_q$ to the quotient polynomial ring $\mathcal{R}_q$. In [6], the authors introduce the integer variant of the problem in the introductory discussion of the vSIS problem (Section 2 of [6]). However, we note that both the constructions in [6,11] are ultimately carried out in the polynomial ring $\mathcal{R}_q$. We believe that the integer variant is interesting in its own right and deserves a closer examination regarding its hardness assumption.

Note 1. We restrict evaluations to the set $\mathbb{Z}_q^* \setminus \{-1, 1\}$ in order to avoid a simple distinguishing attack. Indeed, evaluation at 1 or -1 yields an immediate distinguishing attack on the problem and must therefore be excluded. We give the details in Section 6.

Note 2. The problem admits a natural extension from the binary to the ternary case, analogous to the PV Knapsack. However, in line with the Vanishing-SIS commitment schemes, we define the binary version.

The knapsack problems above also admit corresponding decision versions, which ask to distinguish between evaluations of $\boldsymbol{f}(\boldsymbol{x})$ whose coefficients are sampled uniformly from $\mathbb{Z}_q$ and those with ternary (or binary) coefficients, evaluated on elements of Ω_t (or $\mathcal{S}_t$).

2.4 Lattice reduction

Let $\boldsymbol{B} \in \mathbb{Z}^{n \times n}$ be a full rank matrix. The lattice generated by $\boldsymbol{B}$ is

$$\mathcal{L}(\boldsymbol{B}) := \{\boldsymbol{B}\boldsymbol{z} : \boldsymbol{z} \in \mathbb{Z}^n\}$$

We denote the dimension and volume of the lattice by $\texttt{Dim}$ and $\texttt{Vol}$, respectively, where $\texttt{Dim} = n$ and $\texttt{Vol} = |\det(\boldsymbol{B})|$. A lattice is called a q-ary lattice if it contains $q\mathbb{Z}^n$ as a sublattice.

The most important computational problem involving the lattices is the shortest vector problem (SVP). The problem asks to find a shortest non-zero vector in the lattice $\mathcal{L}$, given any basis of the lattice. For any lattice of dimension $\texttt{Dim}$ and volume $\texttt{Vol}$, Minkowski's theorem provides an upper bound on the length of shortest non-zero vector. Let $\boldsymbol{v}$ be a shortest vector in the lattice $\mathcal{L}$, then Minkowski's theorem states that

$$\|\boldsymbol{v}\|_\infty \leq \texttt{Vol}^{1/\texttt{Dim}} \text{ and } \|\boldsymbol{v}\| \leq \sqrt{\texttt{Dim}}\texttt{Vol}^{1/\texttt{Dim}}$$

Definitions 3 (q-ary Kernel Lattice) *For any full-rank matrix $\boldsymbol{A} \in \mathbb{Z}_q^{t \times n}$ with $n > t$, the q-ary kernel lattice of $\boldsymbol{A}$ is defined as*

$$\mathcal{L}^\perp_{\boldsymbol{A},q} = \{\boldsymbol{v} \in \mathbb{Z}^n | \boldsymbol{A}\boldsymbol{v} = \boldsymbol{0} \bmod q\}.$$

Let us write $\boldsymbol{A}$ as $(\boldsymbol{A}_1|\boldsymbol{A}_2)$ where $\boldsymbol{A}_1$ is the submatrix formed using the first t columns of $\boldsymbol{A}$ and $\boldsymbol{A}_2$ is formed using the rest $(n-t)$ columns of $\boldsymbol{A}$. Assuming

$\boldsymbol{A}_1$ is invertible[6],

$$\begin{bmatrix} q\boldsymbol{I}_t & -\boldsymbol{A}_1^{-1}\boldsymbol{A}_2 \\ \boldsymbol{0} & \boldsymbol{I}_{n-t} \end{bmatrix}$$

forms a basis of $\mathcal{L}^{\perp}_{\boldsymbol{A},q}$. Note that this lattice has $\texttt{Dim} = n$ and $\texttt{Vol} = q^t$. Finding a small solution in the lattice $\mathcal{L}^{\perp}_{\boldsymbol{A},q}$ is called the SIS problem.

Let us write λ_i to denote the smallest radius of a closed ball containing at least i linearly independent vectors in the lattice $\mathcal{L}$. In the case when $\lambda_2 > \gamma\lambda_1$ for some $\gamma \geq 1$, then there exits a γ-unique SVP (uSVP_γ) solution in $\mathcal{L}$.

Definitions 4 (γ-unique SVP or uSVP_γ) *Given a lattice $\mathcal{L}$ with the promise that $\lambda_2 > \gamma\lambda_1$ for $\gamma \geq 1$, find vector $\boldsymbol{v} \in \mathcal{L}$ (unique upto sign) such that $\|\boldsymbol{v}\| = \lambda_1$.*

Lattice reduction algorithms are the goto method for finding short vectors in a lattice. A lattice reduction algorithm transforms a given lattice basis to a new basis consisting of vectors that are short and close to orthogonal. The LLL algorithm [16] is a polynomial-time lattice reduction algorithm which finds a small vector in the lattice with an exponential large approximation factor. To find better approximate solutions, a stronger lattice reduction algorithm, such as Block Korkine-Zolotarev (BKZ) algorithm [18], is used. This comes at a cost to the overall runtime of the lattice reduction.

For our experiments, we used an efficient implementation of BKZ, found in the fplll library [19].

The analysis of [12] suggests that the uSVP_γ with $\texttt{Dim} = n$ can be solved using a lattice reduction algorithm that achieves a root Hermite factor close to $\delta := \gamma^{1/n}$. In particular, when γ is large enough, the value of δ becomes feasible with practical lattice reduction algorithms [2,1]. As a result, with high enough values of δ, the uSVP_γ becomes efficiently solvable.

3 Previous Attacks

3.1 The PV Knapsack problem

The PV Knapsack problem was introduced in [13] as an efficient post-quantum, lattice-based assumption. Owing to its rich algebraic structure, the problem has attracted significant interest in cryptography [13,17,4,5]. In the original works, the problem was instantiated with $\boldsymbol{g}(\boldsymbol{x}) = \boldsymbol{x}^n - 1$ for a prime n [13,14]. More recent studies, however, predominantly consider the choice $\boldsymbol{g}(\boldsymbol{x}) = \boldsymbol{x}^n + 1$ for n a power of two [17,4,5]. This shift is likely due to the widespread use of the latter polynomial in lattice-based cryptography. In the following, we recall all previously considered attacks. We emphasize that these attacks apply equally to both choices of $\boldsymbol{g}(\boldsymbol{x})$.

[6] If $\boldsymbol{A}_1$ is not invertible, then we can permute the columns of $\boldsymbol{A}$ to start with an invertible matrix $\boldsymbol{A}_1$.

Direct Primal Attack [13] This attack converts the PV knapsack problem into a structured variant of the knapsack problem. This is achieved by expressing the evaluations using the polynomial basis from the Equation 1.

The attack takes as input the (partial) Vandermonde matrix $\boldsymbol{V} \in \mathbb{Z}_q^{t \times n}$ generated by the $\omega \in \Omega_t$ in any arbitrary order and the corresponding evaluations $\boldsymbol{b} \in \mathbb{Z}_q^t$, and look for a solution of $\boldsymbol{f}$ such that

$$\boldsymbol{V}\boldsymbol{f} = \boldsymbol{b} \bmod q$$

with $\|\boldsymbol{f}\|_\infty \leq 1$.

The authors proposed lattice reduction algorithms to recover the solution. In particular, in this attack we look for the uSVP solution $(\boldsymbol{f}, -1)^T$ in the Kernel lattice

$$\mathcal{L}^{\perp}_{\boldsymbol{V}',q} = \{\boldsymbol{v} \in \mathbb{Z}^{n+1} | \boldsymbol{V}'\boldsymbol{v} = \boldsymbol{0} \bmod q\}$$

where $\boldsymbol{V}' = (\boldsymbol{V}|\boldsymbol{b})$. It is to be noted that the vector $(\boldsymbol{f}, -1)^T$ is unusually short [7] in the lattice $\mathcal{L}^{\perp}_{\boldsymbol{V}',q}$. That is, with high probability, $(\boldsymbol{f}, -1)^T$ is the uSVP solution in $\mathcal{L}^{\perp}_{\boldsymbol{V}',q}$. Consequently, the problem reduces to finding uSVP solution in a (Kernel) lattice of $\texttt{Dim} = n + 1$ and $\texttt{Vol} = q^t$. This attack serves as the reference point for parameter selection, with the goal of ensuring that solving the uSVP instance remains hard in both classical and quantum computation.

For the ease of exposition, we use the notation $\texttt{Solve-uSVP}(\boldsymbol{V}, \boldsymbol{b})$ to represent the procedure that recovers the unique shortest vector $\boldsymbol{f}$ such that

$$\boldsymbol{V}\boldsymbol{f} = \boldsymbol{b} \bmod q$$

Dual Attack In this attack, we look for a small lattice vector in the dual ideal to solve the PV Knapsack problem. Let us first rephrase the PV knapsack problem as a BDD problem.

Let $\boldsymbol{z} \in \mathbb{Z}^n$ be any solution to $\boldsymbol{V}\boldsymbol{z} = \boldsymbol{b} \bmod q$. Then there exists $\boldsymbol{u}$ in the kernel lattice

$$\mathcal{L}^{\perp}_{\boldsymbol{V},q} = \{\boldsymbol{v} \in \mathbb{Z}^n | \boldsymbol{V}\boldsymbol{v} = \boldsymbol{0} \bmod q\}$$

which is very close to the target $\boldsymbol{z}$. Indeed, setting $\boldsymbol{u} = \boldsymbol{z} - \boldsymbol{f} \in \mathcal{L}^{\perp}_{\boldsymbol{V},q}$ provides a very close solution to the target $\boldsymbol{z}$.

Algebraically, any element in $\mathcal{L}^{\perp}_{\boldsymbol{V},q}$ corresponds to an element in the ideal $I_{\Omega_t} = \prod_{\omega \in \Omega_t} (\boldsymbol{x} - \omega) \subset \mathcal{R}_q$. As a result, if we look for a small solution $\boldsymbol{u}'$ in the dual ideal $I_{\Omega \setminus \Omega_t} = \prod_{\omega \in \Omega \setminus \Omega_t} (\boldsymbol{x} - \omega)$, we have

$$\boldsymbol{u}'\boldsymbol{z} = \boldsymbol{u}'\boldsymbol{u} + \boldsymbol{u}'\boldsymbol{f} = \boldsymbol{u}'\boldsymbol{f}$$

It is to be noted that finding a small solution $\boldsymbol{u}'$ in the dual ideal $I_{\Omega \setminus \Omega_t}$ can be achieved by looking for a small vector in the Kernel lattice generated by a

[7] The norm of $(\boldsymbol{f}, -1)^T$ is significantly smaller than that of the shortest vector in a random lattice of the same dimension and volume, which can be approximated using Minkowski's theorem.

Vandermonde matrix with rows as powers of $\omega \in \Omega \setminus \Omega_t$. Also, note that, there are exactly $n - t$ elements in the set $\Omega \setminus \Omega_t$. Consequently, this task reduces to finding a short lattice vector in $\texttt{Dim} = n$ and $\texttt{Vol} = q^{n-t}$. Now, if $\|\boldsymbol{u}'\boldsymbol{f}\|_\infty < q/2$, there is no modular q operations involved in the product $\boldsymbol{u}'\boldsymbol{f}$. As a result, $\boldsymbol{f}$ can be directly recovered.

Dimension reduced Dual attack [3] In this attack, we show how to reduce the lattice cost of finding a small vector in the dual attack. We give a simplified version of the attack as given in [7].

Let $\Omega_{2t_1^{(A)}}$ be the largest subset of Ω_t that remains invariant under the transformation of inverse, where $0 \leq t^{(A)} \leq \lfloor t/2 \rfloor$. Note that all the (primitive) roots in Ω can be paired with its inverse. As a result, all the elements in the set $\Omega \setminus \Omega_{2t_1^{(A)}}$ can also be paired with its inverse.

Now, if $t_1^{(A)}$ is not too small, then PV Knapsack problem reduces to the BDD problem in the kernel lattice

$$\mathcal{L}^{\perp}_{\boldsymbol{V}^*,q} = \{\boldsymbol{v} \in \mathbb{Z}^n | \boldsymbol{V}^*\boldsymbol{v} = \boldsymbol{0} \bmod q\}$$

where $\boldsymbol{V}^*$ is the Vandermonde matrix generated by the $\omega \in \Omega_{2t_1^{(A)}}$.

Now, instead of searching for a small element $\boldsymbol{u}'$ in the ideal $I_{\Omega \setminus \Omega_{2t_1^{(A)}}}$, one can instead look for a $\boldsymbol{u}'$ satisfying

$$\boldsymbol{u}'(\omega) = \boldsymbol{u}'(\omega^{-1}) \quad \text{for all } \omega \in \Omega \setminus \Omega_{2t_1^{(A)}}.$$

This symmetry condition allows us to restrict the search to a small polynomial $\boldsymbol{u}'(x) \in I_{\Omega \setminus \Omega_{2t_1^{(A)}}}$ generated by the basis $\{1, \boldsymbol{x}, \boldsymbol{x}^2, \ldots, \boldsymbol{x}^{\lfloor n/2 \rfloor}\}$. Indeed, for such a $\boldsymbol{u}'$, both $\boldsymbol{x}$ and $\boldsymbol{x}^{-1}$ are roots. As a result, such $\boldsymbol{u}'(x)$ can be recovered from a lattice of reduced dimension.

However, in both the Dual attack and the Dimension reduced Dual attack, the condition $\|\boldsymbol{u}'\boldsymbol{f}\|_\infty < q/2$ must hold for the attack to succeed. This stringent norm requirement restricts both attacks to only a small set of *worst-case* keys.

Note that, for the values of n commonly considered in the literature, one may also exploit higher-order symmetries to further reduce the lattice dimension. However, for a random choice of Ω_t, the number of available evaluations becomes too small on the average. Consequently, the reduction to the BDD no longer applies.

Dimension reduced Primal attack [7] In this attack, we demonstrate how to reduce the lattice dimension using symmetries within the primal strategy. The key advantage is that the primal attack corresponds to the uSVP, which is believed to be easier to solve than the SVP [12,2,1].

We start from $\Omega_{2t_1^{(A)}}$, the largest subset of Ω_t that is invariant under the inversion map. For each $\omega \in \Omega_{2t_1^{(A)}}$, we have $t_1^{(A)}$ distinct evaluations of $\boldsymbol{f}$,

namely at both ω and ω^{-1}. Consequently, this yields $t_1^{(A)}$ distinct evaluations of $\boldsymbol{f}(\boldsymbol{x}) \pm \boldsymbol{f}(1/\boldsymbol{x})$. Now, recovering $\boldsymbol{f}(\boldsymbol{x})$ from $\boldsymbol{f}(\boldsymbol{x}) \pm \boldsymbol{f}(1/\boldsymbol{x})$ is easy. Therefore, the PV problem reduces to recovering

$$\boldsymbol{\psi}_{\pm}(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{x}) \pm \boldsymbol{f}(1/\boldsymbol{x}).$$

Now, $\boldsymbol{\psi}_{\pm}$ can be generated by a basis of order $\lceil n/2 \rceil$ via the invariant transformation of $\boldsymbol{x} \to \boldsymbol{x}^{-1}$. Indeed, it is easy to verify that $\boldsymbol{\psi}_{\pm}$ is generated by the basis

$$\boldsymbol{x}^i \pm \boldsymbol{x}^{-i} \textbf{ for } 0 \leq i \leq \lfloor n/2 \rfloor$$

Also, since the coefficients of $\boldsymbol{f}(\boldsymbol{x})$ are in $\{-1, 0, 1\}$, the coefficients of $\boldsymbol{\psi}_{\pm}$ are in $\{-2, -1, 0, 1, 2\}$. As a result, similar to direct primal attack, we can look for the uSVP solution $(\boldsymbol{\psi}_{\pm}, -1)^T$ from its evaluations. For this, we look for a very short solution $(\boldsymbol{\psi}_{\pm}, -1)^T$ in the Kernel lattice

$$\mathcal{L}^{\perp}_{\boldsymbol{W}'_{\pm},q} = \{\boldsymbol{v} \in \mathbb{Z}^{\lceil n/2 \rceil + 1} | \boldsymbol{W}'_{\pm}\boldsymbol{v} = \boldsymbol{0} \bmod q\}$$

where $\boldsymbol{W}'_{\pm} = (\boldsymbol{W}_{\pm}|\boldsymbol{b}_{\pm})$. Here, $\boldsymbol{W}_{\pm}$ denotes the matrix with $t_1^{(A)}$ rows and $\lceil n/2 \rceil$ columns, whose entries are given by the evaluations of each of the $\lceil n/2 \rceil$ monomials at an arbitrary choice of $t_1^{(A)}$ representatives of the pairs (ω, ω^{-1}) appearing in $\Omega_{2t_1^{(A)}}$. And $\boldsymbol{b}_{\pm}$ denotes the known evaluations of $\boldsymbol{\psi}_{\pm}$ at these representatives.

Now, when $t_1^{(A)}$ is sufficiently large, which occurs with non-negligible probability, the cost of solving the uSVP is significantly reduced. The volume of the lattice is proportional to the number of distinct evaluations $t_1^{(A)}$, and therefore the problem becomes easier as $t_1^{(A)}$ increases. This gives rise to the existence of weak keys.

As in the dimension-reduced dual attack, one may also attempt to exploit higher-order symmetries. However, this typically further decreases the number of available evaluations. Consequently, the uniqueness gap is substantially reduced, which does not facilitate the efficient recovery of the uSVP solution.

3.2 The vSIS commitment problem

To the best of our knowledge, there is no formal attack strategy given for the vSIS commitment problem. Nevertheless, the analysis in Section 3 of [11] suggests that the only viable approach is a direct primal attack, analogous to that used for the PV knapsack problem.

Here, the attack takes as input the Vandermonde matrix $\boldsymbol{V} \in \mathbb{Z}_q^{t \times n}$ generated by the $\upsilon \in \mathcal{S}_t$ in any arbitrary order and the corresponding evaluations $\boldsymbol{b} \in \mathbb{Z}_q^t$, and look for a solution of $\boldsymbol{f}$ such that

$$\boldsymbol{V}\boldsymbol{f} = \boldsymbol{b} \bmod q$$

with $\|\boldsymbol{f}\|_\infty \leq 1$.

Similar to the PV Knapsack, we can use lattice reduction algorithms to recover the solution $\boldsymbol{f}$. In particular, we can look for the uSVP solution $(\boldsymbol{f}, -1)^T$ in the Kernel lattice

$$\mathcal{L}^{\perp}_{\boldsymbol{V}',q} = \{\boldsymbol{v} \in \mathbb{Z}^{n+1} | \boldsymbol{V}'\boldsymbol{v} = \boldsymbol{0} \bmod q\}$$

where $\boldsymbol{V}' = (\boldsymbol{V}|\boldsymbol{b})$. Consequently, the problem reduces to finding uSVP solution in a (Kernel) lattice of $\texttt{Dim} = n + 1$ and $\texttt{Vol} = q^t$.

4 Our Contributions

In this paper, we introduce new evaluation attacks on the Vandermonde Knapsack problem. Our approach evaluates carefully chosen elements to construct lower-dimensional Knapsack instances. The resulting dimensional reduction allows certain Vandermonde-based Knapsack problems to be solved more efficiently than previously assumed.

4.1 New attack on PV Knapsack problem over power-of-two cyclotomics

Let $n = 2^k$ for some positive integer k. In the ring $\mathcal{R}_q := \mathbb{Z}_q[\boldsymbol{x}]/(\boldsymbol{x}^n + 1)$, the automorphisms correspond to the map $\boldsymbol{x} \to \boldsymbol{x}^j$ for $j \in \mathbb{Z}^*_{2n}$. Hence, the symmetries of the ring are governed by the structure of the multiplicative group $\mathbb{Z}^*_{2n}$.

We start from the observation that in $\mathbb{Z}^*_{2n}$, there are exactly three subgroups of order 2, namely the subgroups generated by $-1, n \pm 1 \in \mathbb{Z}^*_{2n}$.[8] We show that all these generators allow us to have different symmetries of order 2 in the power-of-two cyclotomic ring. In particular, the symmetries are achieved via the following invariant transformation in the ring $\mathcal{R}_q$:

$$\boldsymbol{x} \to \boldsymbol{x}^{-1}, \qquad \boldsymbol{x} \to \boldsymbol{x}^{n+1} = -\boldsymbol{x}, \qquad \boldsymbol{x} \to \boldsymbol{x}^{n-1} = -\boldsymbol{x}^{-1}$$

The invariant transformation $\boldsymbol{x} \to \boldsymbol{x}^{-1}$ was the only symmetry exploited in [7]. In particular, it relied exclusively on the reciprocal symmetry of order 2 induced by $\boldsymbol{x} \to \boldsymbol{x}^{-1}$. In contrast, we additionally consider signed symmetry, corresponding to the transformation $\boldsymbol{x} \to -\boldsymbol{x}$, as well as signed reciprocal symmetry, corresponding to the transformation $\boldsymbol{x} \to -\boldsymbol{x}^{-1}$. We show that exploiting all three symmetries adaptively allows a significantly larger weak-key set for the PV Knapsack problem than those identified in prior attacks.

[8] This is easy to verify. Assume $k \geq 1$, then

$$\begin{aligned}(-1)^2 &= 1 \bmod 2n\\ (n-1)^2 &= (n^2 + 1 - 2n) \bmod 2n = (n^2+1) \bmod 2n = 1 \bmod 2n\\ (n+1)^2 &= (n^2 + 1 + 2n) \bmod 2n = (n^2+1) \bmod 2n = 1 \bmod 2n\end{aligned}$$

Recall that in the dimension-reduced primal attack (Section 3.1), the cost depends on the number of available pairs $t_1^{(A)}$ in the set Ω_t. In particular, when sufficiently many pairs are present, the resulting uSVP cost can be reduced substantially. However, such a large number of pairs occurs only for a small fraction of keys for a given symmetry, thereby limiting the overall applicability of this strategy. By contrast, exploiting the new symmetries in an adaptive manner enables us to attack a significantly larger fraction of keys.

Let $\Omega_{2t_1^{(A)}}$, $\Omega_{2t_1^{(B)}}$, and $\Omega_{2t_1^{(C)}}$ denote the largest subsets of Ω_t consisting of all inverse-invariant pairs (we call it Type-A pairs), negation-invariant pairs (we call it Type-B pairs), and inverse-negation-invariant pairs (we call it Type-C pairs), respectively. Observe that these invariant subsets can be computed efficiently from Ω_t. We initiate the attack by determining the maximum among the available numbers of pairs, namely $t_1^{(A)}$, $t_1^{(B)}$, and $t_1^{(C)}$. Based on this value, we adaptively select the corresponding symmetry to mount the dimension reduced primal attack on the PV Knapsack problem.

We then analyze the joint probability distribution that governs the occurrence of the three different forms of pairs arising from random Ω_t. We show that, with very high probability, a large number of evaluations fall into at least one of these three types of pairs. As a result, we obtain a significantly larger set of weak keys for the PV Knapsack problem than previously identified in [3,7].

However, the new symmetries exploited in this paper apply only when $\boldsymbol{g}(\boldsymbol{x}) = \boldsymbol{x}^n + 1$ with n power of two, and do not extend to the case when $\boldsymbol{g}(\boldsymbol{x}) = \boldsymbol{x}^n - 1$ with n prime. To see this, consider the ring $\mathbb{Z}_q[\boldsymbol{x}]/(\boldsymbol{x}^n - 1)$ for some prime n. Here, the automorphisms correspond to the map $\boldsymbol{x} \to \boldsymbol{x}^j$ for $j \in \mathbb{Z}_n^*$. Hence, the symmetries of the ring are governed by the structure of the multiplicative group $\mathbb{Z}_n^*$. In the group $\mathbb{Z}_n^*$, there is exactly one subgroup of order 2, which is generated by $-1 \in \mathbb{Z}_n^*$. As a result, the ring has a unique symmetry of order 2, which is induced by the invariant transformation given by $\boldsymbol{x} \to \boldsymbol{x}^{-1}$. This coincides with the symmetry of order 2 in [7].

We note that the dimension-reduced dual attack (described in Subsection 3.1) can be generalized for all the symmetries when $\mathbf{g}(\boldsymbol{x}) = \boldsymbol{x}^n + 1$. In particular, [3] considers the largest subset of Ω stabilized by an automorphism subgroup of order 2, which captures all the symmetries discussed above. Nevertheless, as discussed in Section 5 of [3], the resulting key-recovery attack remains applicable to worst-case keys. We show that, in this ring, the additional symmetries yield a substantially larger class of weak keys, thereby highlighting a clear distinction in hardness depending on the choice of $\mathbf{g}$ used in the literature.

4.2 New attacks on the Integer vSIS commitment problem

We first show that, for certain choices of the prime modulus q, the distinguishing problem can be solved efficiently. This undermines the hiding property of the underlying commitment scheme. The key observation is that the factorization of $q-1$ can be exploited to obtain useful algebraic relations of the hidden coefficients of the polynomial $\boldsymbol{f}$. In particular, when $q - 1$ contains small factors, we can

evaluate $\boldsymbol{f}$ at elements of correspondingly small order in $\mathbb{Z}_q^*$. These evaluations yield algebraic relations that allow us to construct a low-dimensional lattice containing a short vector related to the coefficients of $\boldsymbol{f}$. Thus, by efficiently recovering a short vector in this lattice, we obtain a distinguisher for the integer vSIS commitment problem.

Next, we propose a key recovery attack that applies to arbitrary prime moduli q. Our approach is based on a symmetry-based attack, analogous to the one used against PV Knapsack, which we adapt to the integer vSIS commitment problem. In particular, we exploit the symmetries induced by the choice of the prime modulus q to recover $\boldsymbol{f}$. We show that, for any prime modulus q of small order, key recovery can be performed efficiently for certain instances of the integer vSIS problem.

In the ring $\mathbb{Z}_q[\boldsymbol{x}]$, the automorphisms correspond to the map $\boldsymbol{x} \to a\boldsymbol{x} + b$ for $a \in \mathbb{Z}_q^*$ and $b \in \mathbb{Z}_q$. Hence, the symmetries of the ring are governed by the structure of the multiplicative group $\mathbb{Z}_q^*$. Now, in the group $\mathbb{Z}_q^*$, there is exactly one subgroup of order 2, namely the subgroup generated by $-1 \in \mathbb{Z}_q^*$. We show that this generator allows us to have a symmetry of order 2 in the ring $\mathbb{Z}_q[\boldsymbol{x}]$. In particular, the symmetry is achieved via the following invariant transformation in the ring $\mathbb{Z}_q[\boldsymbol{x}]$:

$$\boldsymbol{x} \to -\boldsymbol{x}$$

We start with the fact that all the elements in $\mathbb{Z}_q^*$ can be paired with their negation. Now consider $\mathcal{S}_{2t_1}$ to be the largest subset of $\mathcal{S}_t$ containing all the negation-invariant pairs, i.e., $\mathcal{S}_{2t_1}$ contains all the pairs of the form $(v, -v)$. Also, in the ring $\mathbb{Z}_q[\boldsymbol{x}]$, the evaluation of $\boldsymbol{f}(\boldsymbol{x})$ at $-v$ coincides with the evaluation of $\boldsymbol{f}(-\boldsymbol{x})$ at v. As a result, we can compute the evaluation of $\boldsymbol{f}$ both at v and $-v$ when $v \in \mathcal{S}_{2t_1}$. Consequently, this yields t_1 distinct evaluations of $\boldsymbol{f}(\boldsymbol{x}) \pm \boldsymbol{f}(-\boldsymbol{x})$.

Similar to the primal attacks on the PV Knapsack problem, we hope to recover

$$\boldsymbol{f}(\boldsymbol{x}) \pm \boldsymbol{f}(-\boldsymbol{x})$$

efficiently using a lattice of reduced dimension if the number of pairs t_1 is sufficiently large. This gives back $\boldsymbol{f}(\boldsymbol{x})$.

It only remains to see when the value of t_1 becomes large for this attack to be effective. For this, we show that when the modulus $q = \alpha n$ for some small constant α, as in the small-modulus variant of Knapsack, the value of t_1 is large with non-negligible probability. The small modulus variant of Knapsack has found applications in several lattice-based constructions; for example, it underlies the signature scheme proposed in [10].

As a result, for this particular case of the problem, a fraction of the keys are weak. As far as we know, there are no easy instances of the small modulus variant of the Knapsack problem. [9]

[9] It is worth noting that [9] has made progress on efficiently solving the small-modulus variant of the SIS problem. Nevertheless, the proposed methods recover only non-trivial solutions and do not provide the uniqueness guarantee sought in the Knapsack problem.

5 New attack on the PV Knapsack problem

In this section, we start by revisiting the main result of the order-2 reciprocal symmetry attack introduced by Das and Joux in [7].

Lemma 1 (Lemma 1 of [7]). *Let $\boldsymbol{f}(\boldsymbol{x})$ be any polynomial in $\mathcal{R}_q$. Now define $\boldsymbol{\psi}_{\pm}(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{x}) \pm \boldsymbol{f}(1/\boldsymbol{x})$. Then $\boldsymbol{\psi}_{\pm}$ can be generated by a basis of order $n/2$. Additionally, if the coefficients of $\boldsymbol{f}(\boldsymbol{x})$ are sampled uniformly from the set $\{-1, 0, 1\}$, then the expected squared-norm of $\boldsymbol{\psi}_{\pm}$ is upper-bounded by $2n/3$ in the new basis representation.*

The result follows from the fact that $\boldsymbol{\psi}_{+}$ is generated by the basis

$$\{2, (\boldsymbol{x} - \boldsymbol{x}^{n-1}), \ldots, (\boldsymbol{x}^{n/2-1} - \boldsymbol{x}^{n/2+1})\}$$

and $\boldsymbol{\psi}_{-}$ is generated by the basis

$$\{(\boldsymbol{x} + \boldsymbol{x}^{n-1}), \ldots, (\boldsymbol{x}^{n/2-1} + \boldsymbol{x}^{n/2+1}), 2\boldsymbol{x}^{n/2}\}$$

which both has order $n/2$ as required.

Now, if the coefficients of $\boldsymbol{f}$ are sampled uniformly at random from $\{-1, 0, 1\}$, then the coefficients of $\boldsymbol{\psi}_{\pm}$ are in $\{-2, -1, 0, 1, 2\}$ and follow the probability distribution from Table 1, except for the special coefficients. The special coefficients associated with 2 and $2\boldsymbol{x}^{n/2}$ only follow the initial uniform distribution in the set $\{-1, 0, 1\}$. Therefore, using this distribution of the coefficients and using linearity of expectation, the expected squared-norm of $\boldsymbol{\psi}_{\pm}$ is upper-bounded by $2n/3$ in the new basis representation.

Coefficient	-2	-1	0	1	2
Probability	1/9	2/9	3/9	2/9	1/9

Table 1: Probability distribution of the coefficients of $\boldsymbol{\psi}_{\pm}$.

We now describe two new symmetries of order 2 that arise naturally in the power-of-two cyclotomic ring, namely the sign symmetry and reciprocal-sign symmetry. Leveraging these structural properties, we develop a new adaptive attack against the PV Knapsack problem.

Lemma 2. *Let $\boldsymbol{f}(\boldsymbol{x})$ be any polynomial in $\mathcal{R}_q$. Now define $\boldsymbol{\phi}_{\pm}(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{x}) \pm \boldsymbol{f}(-\boldsymbol{x})$. Then $\boldsymbol{\phi}_{\pm}$ can be generated by a basis of order $n/2$. Additionally, if the coefficients of $\boldsymbol{f}(\boldsymbol{x})$ are sampled uniformly from the set $\{-1, 0, 1\}$, then the expected squared-norm of $\boldsymbol{\phi}_{\pm}$ is $n/3$ in the new basis representation.*

Proof. Consider the mapping

$$\boldsymbol{x}^i \to \boldsymbol{x}^i \pm (-\boldsymbol{x})^i \text{ for } 0 \leq i < n$$

It is easy to verify that $\boldsymbol{x}^i + (-\boldsymbol{x})^i$ vanishes when $i \bmod 2 \neq 0$ and $2\boldsymbol{x}^i$ when $i \bmod 2 = 0$. Similarly, $\boldsymbol{x}^i - (-\boldsymbol{x})^i$ vanishes when $i \bmod 2 = 0$ and $2\boldsymbol{x}^i$ when $i \bmod 2 \neq 0$. As a result, $\boldsymbol{\phi}_\pm$ can be generated by a basis of order $n/2$. In particular, $\boldsymbol{\phi}_+$ can be generated by the basis

$$\{2, 2\boldsymbol{x}^2, \ldots, 2\boldsymbol{x}^{n-2}\}$$

which is of order $n/2$. Similarly, $\boldsymbol{\phi}_-$ can be generated by the basis

$$\{2\boldsymbol{x}, 2\boldsymbol{x}^3, \ldots, 2\boldsymbol{x}^{n-1}\}$$

which is also of order $n/2$.

Now if the coefficients of $\boldsymbol{f}$ are sampled from the set $\{-1, 0, 1\}$, the coefficients of $\boldsymbol{\phi}_\pm$ are also in the set $\{-1, 0, 1\}$ and follow the initial uniform distribution. As a result, the expected squared-norm of $\boldsymbol{\phi}_\pm$ is $n/3$, as claimed. [10]

Lemma 3. *Let $\boldsymbol{f}(\boldsymbol{x})$ be any polynomial in $\mathcal{R}_q$. Now define $\boldsymbol{\chi}_\pm(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{x}) \pm \boldsymbol{f}(1/(-\boldsymbol{x}))$. Then $\boldsymbol{\chi}_\pm$ can be generated by a basis of order $n/2$. Additionally, if the coefficients of $\boldsymbol{f}(\boldsymbol{x})$ are sampled uniformly from the set $\{-1, 0, 1\}$, then the expected squared-norm of $\boldsymbol{\chi}_\pm$ is upper-bounded by $2n/3$ in the new basis representation.*

The result follows combining from the above two lemmas. In particular, the polynomial $\boldsymbol{\chi}_+$ is generated by the basis

$$\{2, (\boldsymbol{x} + \boldsymbol{x}^{n-1}), (\boldsymbol{x}^2 - \boldsymbol{x}^{n-2}), \ldots, (\boldsymbol{x}^{n/2-1} + \boldsymbol{x}^{n/2+1})\}$$

and $\boldsymbol{\chi}_-$ is generated by the basis

$$\{(\boldsymbol{x} - \boldsymbol{x}^{n-1}), (\boldsymbol{x}^2 + \boldsymbol{x}^{n-2}), \ldots, (\boldsymbol{x}^{n/2-1} - \boldsymbol{x}^{n/2+1}), 2\boldsymbol{x}^{n/2}\}$$

which has both order $n/2$, as required. The expected squared-norm condition of $\boldsymbol{\chi}_\pm$ follows from Lemma 1.

Using the above results, we adaptively construct a new knapsack problem in reduced dimension to solve the PV Knapsack problem. To this end, we first select the symmetry type that provides the maximum number of available pairs in Ω_t. Let $t_1^{(\max)} = \max(t_1^{(A)}, t_1^{(B)}, t_1^{(C)})$ and let $\sigma \in \{A, B, C\}$ be the corresponding symmetry type such that $t_1^{(\sigma)} = t_1^{(\max)}$. Now we define a mapping $\rho(\boldsymbol{x})$ such that

$$\rho(\boldsymbol{x}) = \begin{cases} \boldsymbol{x}^{-1} & \text{if } \sigma = A \\ -\boldsymbol{x} & \text{if } \sigma = B \\ -\boldsymbol{x}^{-1} & \text{if } \sigma = C \end{cases}$$

[10] The motivation for using the constant 2 in the basis is to reduce the squared norm of $\boldsymbol{\phi}_\pm$. In our attack, we aim to recover $\boldsymbol{\phi}_\pm$ via lattice reduction. The problem becomes easier when the underlying uSVP secret has a smaller squared norm, increasing the likelihood of successful recovery.

and construct the polynomials

$$\boldsymbol{\tau}_{\pm}(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{x}) \pm \boldsymbol{f}(\rho(\boldsymbol{x})) = \begin{cases} \boldsymbol{\psi}_{\pm}(\boldsymbol{x}) & \text{if } \sigma = A \quad \text{(see Lemma 1)} \\ \boldsymbol{\phi}_{\pm}(\boldsymbol{x}) & \text{if } \sigma = B \quad \text{(see Lemma 2)} \\ \boldsymbol{\chi}_{\pm}(\boldsymbol{x}) & \text{if } \sigma = C \quad \text{(see Lemma 3)} \end{cases}.$$

The new reduced Knapsack instances are designed using the $t_1^{(\max)}$ many evaluations of the polynomials $\boldsymbol{\tau}_{\pm}(\boldsymbol{x})$. We start by creating matrix $\mathbf{B}_{\pm}$ with $t_1^{(\max)}$ many rows and $\frac{n}{2}$ many columns, whose entries are the evaluations of each of the $\frac{n}{2}$ basis elements of each of the representatives of the $t_1^{(\max)}$ many distinct pairs of type-σ in Ω_t. We also construct vector $\boldsymbol{b}_{\pm}$ with entries as the known evaluations of $\boldsymbol{\tau}_{\pm}$ at each of the representatives in the set Ω_t. Finally, we look for the solution of $\texttt{Solve-uSVP}(\boldsymbol{B}_{\pm}, \boldsymbol{b}_{\pm}.)$

Now, similar to the direct primal attack, we convert the problem to uSVP on the kernel lattice

$$\mathcal{L}^{\perp}_{\boldsymbol{B}'_{\pm},q} = \{\boldsymbol{v} \in \mathbb{Z}^{\frac{n}{2}+1} | \boldsymbol{B}'_{\pm}\boldsymbol{v} = \mathbf{0} \bmod q\}$$

where $\boldsymbol{B}'_{\pm} = (\boldsymbol{B}_{\pm}|\boldsymbol{b}_{\pm})$. Similar to the direct primal attack, we also expect $(\boldsymbol{\tau}_{\pm}| - 1)^T$ to be an unusually short vector in the lattice. As a result, we expect to yield the uSVP solution. This short vector can be recovered using lattice reduction, such as BKZ.

5.1 Analysis of the attack

Here we do a similar analysis to that in [7]. Note that the concrete analysis of the attack depends on two major facts. First, for a given $t_1^{(\max)}$ what is the cost of solving the reduced uSVP instance, and does this computation outperform the computation of the uSVP instance in the full lattice? Second, for a random Ω_t what is the probability of occurrence of a given value of $t_1^{(\max)}$? Note that the value of $t_1^{(\max)}$ can be computed efficiently for a given Ω_t. Hence, this probability corresponds to the fraction of keys that can be attacked using the reduced uSVP instance.

Comparison of cost for solving uSVP In this section, we compare the cost of solving the uSVP instance reduced by our method with that of solving the full uSVP instance. Recall that the cost of solving a uSVP mostly depends on the root Hermite factor $\delta = \gamma^{1/n}$ [12], where γ is the uniqueness gap. For the direct primal attack [13], the reciprocal symmetry-based attack [7], and our adaptive symmetry-based attack, we consider only lattices in which the existence of an unusually short vector is guaranteed. We also assume the length of the other vectors in the lattice are upper-bounded by the Minkowski's bound. As a result, we estimate the value of $\lambda_2 \approx \sqrt{d}V^{1/d}$, where d is the dimension and V is the volume. On the other hand, to estimate λ_1, we use the square-root of the expected

squared-norm. With these two estimations, the root Hermite factor δ can be estimated by $(\frac{\lambda_2}{\lambda_1})^{1/d}$.

Now, the direct primal attack gives a uSVP in dimension $n+1$ and volume q^t. Here, the expected squared-norm of the shortest vector is $2n/3+1$. We estimate the expected norm by $\sqrt{2n/3+1}$. Hence, the root Hermite factor for the direct primal attack can be estimated by

$$\delta_{\text{full}} = \left(\frac{\sqrt{n+1} q^{t/(n+1)}}{\sqrt{2n/3+1}} \right)^{1/(n+1)}.$$

For our adaptive symmetry-based attack, both the lattices that we attain after reduction phase have dimension $\frac{n}{2}+1$ and volume $q^{t_1^{(\max)}}$. Note that the expected squared-norm of the shortest vector is $2n/3+1$ if $\sigma \in \{A, C\}$ and $n/3+1$ otherwise. Let $\eta+1$ denote the expected squared-norm of the shortest vector that we look for depending on the chosen symmetry. We estimate the expected norm of the shortest vector by $\sqrt{\eta+1}$. Therefore, the corresponding root Hermite factor can be estimated by

$$\delta_{\text{reduced}} = \left(\frac{\sqrt{n/2+1} q^{t_1^{(\max)}/(n/2+1)}}{\sqrt{\eta+1}} \right)^{1/(n/2+1)}$$

Now, to identify the regime where our adaptive symmetry-based attack outperforms the direct primal attack on the full lattice, we compare the corresponding root Hermite factors δ_{reduced} and δ_{full}, respectively. By simplifying the expressions and ignoring the small constant factors, we expect our attack to outperform the direct primal attack on the full lattice whenever

$$(q^{t/n})^{1/n} < (q^{2t_1^{(\max)}/n})^{2/n}, \text{i.e. whenever } t < 4t_1^{(\max)}.$$

When $t = 4t_1^{(\max)}$, the root Hermite factors of both the attacks are equal. However, due to its smaller dimensions, the new attack will outperform the direct primal attack.

We provide the detailed outline of the algorithm in Algorithm 1.

Note 3. It is to be noted that the expected squared norm of $\boldsymbol{\phi}_\pm$ is smaller than that of $\boldsymbol{\psi}_\pm$ and $\boldsymbol{\chi}_\pm$ (See Lemma 2). As a result, one should always consider $\sigma = B$ whenever $t_1^{(B)} \geq \max(t_1^{(A)}, t_1^{(C)})$. This increases the uniqueness gap of the underlying uSVP and thus it is easier to recover $\boldsymbol{\phi}_\pm$ using lattice reductions.

5.2 Joint probability distribution of pairs

To analyze the effectiveness of our new algorithm, we now analyze the probability distribution of $t_1^{(\max)}$.

We first recall the probability distribution of $t_1^{(\sigma)}$ from [7]. For $0 \leq k_\sigma \leq t/2$, the probability that there are exactly k_σ many pairs of type-σ in the set Ω_t is given by

Algorithm 1: Adaptive order 2 Symmetry Attack on PV Knapsack

Input : Ring parameters n, q; sample size $t = \frac{n}{2}$;
Evaluations $E = \{(\omega, \boldsymbol{f}(\omega)) \mid \omega \in \Omega_t\}$.
Output : The secret polynomial $\boldsymbol{f}(\boldsymbol{x})$ or failure $\perp$.

Initialize pair sets: $P_A \leftarrow \emptyset$, $P_B \leftarrow \emptyset$, $P_C \leftarrow \emptyset$
foreach $\omega \in \Omega_t$ **do**
  // Symmetry Detection Phase
  if $\omega^{-1} \in \Omega_t$ ***and*** $(\omega^{-1}, \omega) \notin P_A$ **then**
    $P_A \leftarrow P_A \cup \{(\omega, \omega^{-1})\}$ // Check $x \leftrightarrow 1/x$
  end
  if $-\omega \in \Omega_t$ ***and*** $(-\omega, \omega) \notin P_B$ **then**
    $P_B \leftarrow P_B \cup \{(\omega, -\omega)\}$ // Check $x \leftrightarrow -x$
  end
  if $-1/\omega \in \Omega_t$ ***and*** $(-1/\omega, \omega) \notin P_C$ **then**
    $P_C \leftarrow P_C \cup \{(\omega, -1/\omega)\}$ // Check $x \leftrightarrow -1/x$
  end
end
Calculate pair counts: $t_1^{(A)} \leftarrow |P_A|$, $t_1^{(B)} \leftarrow |P_B|$, $t_1^{(C)} \leftarrow |P_C|$
$t_1^{(\max)} \leftarrow \max(t_1^{(A)}, t_1^{(B)}, t_1^{(C)})$
if $t_1^{(\max)} < t/4$ **then**
  return $\perp$ // Fail: Not enough symmetry pairs found
end
Select symmetry type $\sigma \in \{A, B, C\}$ such that $|P_\sigma| = t_1^{(\max)}$ // Strategy Selection
Define mapping $\rho(\boldsymbol{x})$ based on σ: $\begin{cases} \boldsymbol{x}^{-1} & \text{if } \sigma = A \\ -\boldsymbol{x} & \text{if } \sigma = B \\ -\boldsymbol{x}^{-1} & \text{if } \sigma = C \end{cases}$
Initialize target vectors $\mathbf{b}_\pm \in \mathbb{Z}_q^{t_1^{(\max)}}$ // Dimension Reduction Phase
foreach $(\omega, \omega') \in P_\sigma$ **do**
  $v_\pm \leftarrow (\boldsymbol{f}(\omega) \pm \boldsymbol{f}(\omega')) \bmod q$
  Append $v_\pm$ to $\mathbf{b}_\pm$
end
Construct matrices $\mathbf{B}_\pm$ using reduced bases for $\boldsymbol{f}(\boldsymbol{x}) \pm \boldsymbol{f}(\rho(\boldsymbol{x}))$
$\hat{\boldsymbol{\tau}}_\pm \leftarrow$ Solve-uSVP$(\mathbf{B}_\pm, \mathbf{b}_\pm)$ // Apply Lattice Reduction
Recover candidate $\boldsymbol{f}(\boldsymbol{x})$: $2\boldsymbol{f}(\boldsymbol{x}) \equiv \hat{\boldsymbol{\tau}}_+(\boldsymbol{x}) + \hat{\boldsymbol{\tau}}_-(\boldsymbol{x})$ // Lifting
return $\boldsymbol{f}(\boldsymbol{x})$

$$\Pr(t_1^{(\sigma)} = k_\sigma) = \frac{\binom{n/2}{k_\sigma}\binom{n/2-k_\sigma}{t-2k_\sigma}2^{t-2k_\sigma}}{\binom{n}{t}}.$$

When $t = n/2$, the expected value of the random variable $t_1^{(\sigma)}$ can be approximated by $t/4$. This value was empirically computed in [7]; for completeness we provide a detailed analysis in Appendix A.

Next, we compute the joint probability of having k_A many pairs of type-A, k_B many pairs of type-B and k_C many pairs of type-C in Ω_t. Note that, for every $\omega \in \Omega$, the elements ω^{-1}, $-\omega^{-1}$, and $-\omega$ also belong to Ω. Now, based on the algebraic structure of the power-of-two-cyclotomic ring, the derivation of the joint probability formula is based on the decomposition of Ω into quadruples of the form $\{\omega, -\omega, \omega^{-1}, -\omega^{-1}\}$. Note that the set Ω can be partitioned into $n/4$ disjoint quadruples.

Now, within each quadruplet, the pairs are formed as follows:

- **Type A (Sign):** Pairs are $(\omega, -\omega)$ and $(\omega^{-1}, -\omega^{-1})$. (2 pairs total)
- **Type B (Reciprocal):** Pairs are (ω, ω^{-1}) and $(-\omega, -\omega^{-1})$. (2 pairs total)
- **Type C (Reciprocal-sign):** Pairs are $(\omega, -\omega^{-1})$ and $(-\omega, \omega^{-1})$. (2 pairs total)

When Ω_t is chosen randomly, each quadruplet can fall into exactly one of three scenarios: all elements of the quadruplet are in Ω_t, some elements are in Ω_t, or none of the elements are in Ω_t. Assume that, a random subset Ω_t of size t intersects a specific quadruplet Q in s elements ($0 \leq s \leq 4$). The number of pairs from the specific quadruplet that contributed to the total number of pairs k_A, k_B, k_C depends on the value of s in the following manner:

Vertices Chosen (s)	Contribution to (k_A, k_B, k_C)	Description	Multiplicity
0	$(0, 0, 0)$	Empty	$\binom{4}{0} = 1$
1	$(0, 0, 0)$	Single element	$\binom{4}{1} = 4$
2	$(1, 0, 0)$	Type A pair	2
2	$(0, 1, 0)$	Type B pair	2
2	$(0, 0, 1)$	Type C pair	2
3	$(1, 1, 1)$	Any 3 elements contain 1 of each	$\binom{4}{3} = 4$
4	$(2, 2, 2)$	Full set contains 2 of each	$\binom{4}{4} = 1$

Now, assume that the number of quadruples in each specific selection state of Ω_t is given by:

- m: Number of quadruples for which all four elements are selected.
- j: Number of quadruples for which three of its elements are selected.
- u_A, u_B, u_C: Number of quadruples for which two of its elements are selected.
- v: Number of quadruples for which only one of its element is selected.
- z: Number of quadruples for which none of its elements are selected.

We have the following relations for the number of pairs:

$$\begin{aligned} k_A = 2m + j + u_A &\implies u_A = k_A - 2m - j \\ k_B = 2m + j + u_B &\implies u_B = k_B - 2m - j \\ k_C = 2m + j + u_C &\implies u_C = k_C - 2m - j. \end{aligned}$$

The total number of vertices are

$$t = 4m + 3j + 2(u_A + u_B + u_C) + v.$$

Substituting the values of u_A, u_B, u_C, we get

$$v = t - 2(k_A + k_B + k_C) + 3j + 8m.$$

As a result, the total number of quadruples is given by

$$M = m + j + u_A + u_B + u_C + v + z.$$

Hence, the joint probability of having exactly k_A many pairs of type-A, k_B many pairs of type-B and k_C many pairs of type-C in Ω_t is given by

$$\Pr(t_1^{(A)} = k_A, t_1^{(B)} = k_B, t_1^{(C)} = k_C) = \frac{N_3(k_A, k_B, k_C)}{\binom{n}{t}},$$

where $N_3(k_A, k_B, k_C)$ is the count of favorable cases and is given by

$$N_3(k_A, k_B, k_C) = \sum_{m,j} \frac{(n/4)!}{m!j!u_A!u_B!u_C!v!z!} (1)^m (4)^j (2)^{u_A} (2)^{u_B} (2)^{u_C} (4)^v (1)^z.$$

The sum runs over all non-negative integers m, j (free variables) such that all factorial arguments remain non-negative.

Now, we have all the tools ready to investigate the probability distribution of $t_1^{(\max)}$. Observe that, the event $t_1^{(\max)} \leq k$ occurs if and only if all three individual pair counts are less than or equal to k simultaneously, i.e.,

$$t_1^{(A)} \leq k \quad \text{and} \quad t_1^{(B)} \leq k \quad \text{and} \quad t_1^{(C)} \leq k.$$

Given the joint probability $\Pr(t_1^{(A)} = k_a, t_1^{(B)} = k_b, t_1^{(C)} = k_c)$, the cumulative probability $\Pr(t_1^{(\max)} \leq k)$ is found by summing over the cubic region defined by $[0, k]^3$. Hence we have

$$\Pr(t_1^{(\max)} \leq k) = \sum_{k_a=0}^{k} \sum_{k_b=0}^{k} \sum_{k_c=0}^{k} \Pr(t_1^{(A)} = k_a, t_1^{(B)} = k_b, t_1^{(C)} = k_c).$$

The probability can be computed as

$$\Pr(t_1^{(\max)} = k) = \Pr(t_1^{(\max)} \leq k) - \Pr(t_1^{(\max)} \leq k-1) = \sum_{\substack{0 \leq k_a, k_b, k_c \leq k \\ \max(k_a, k_b, k_c) = k}} \frac{N_3(k_a, k_b, k_c)}{\binom{n}{t}}.$$

Figure 1 compares the distributions of $t_1^{(\max)}$ and $t_1^{(\sigma)}$. In the weak-key region, $t_1^{(\max)}$ assigns substantially more probability mass than $t_1^{(\sigma)}$, and its expectation

n	$\Pr(t_1^{(\sigma)} \geq t/4)$ [7]	$\Pr(t_1^{(\max)} \geq t/4)$ [This work]
256	0.5527	**0.9226**
512	0.5373	**0.9093**
1024	0.5264	**0.8996**

Table 2: Comparison between $\Pr(t_1^{(\max)} \geq t/4)$ and $\Pr(t_1^{(\sigma)} \geq t/4)$ for different values of n.

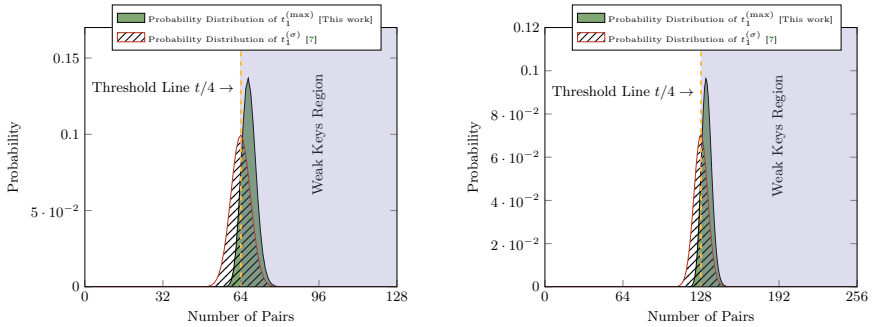

(a) $n = 512$ and $t = 256$ (b) $n = 1024$ and $t = 512$

Fig. 1: Comparison between the Probability Distribution of $t_1^{(\sigma)}$ and $t_1^{(\max)}$ intersecting with the weak keys region.

lies well inside the weak region. As a result, our attack applies to a significantly larger set of weak keys for the PV Knapsack problem than that of [7]. The details of the concrete values are given in Table 2.

To better analyze the distribution of $t_1^{(\max)}$, we generated 100000 random instances of the problem and empirically computed the probability distribution of the $t_1^{(\max)}$. We observed that it precisely matched our theoretical probability computation (see Figure 2).

5.3 Experimental Results

In this paper, all experiments were conducted on a Mac Studio desktop equipped with an Apple M4 Max chip and 64 GB of memory. We conducted experiments on two parameter sets proposed in the literature. In the first case, we take $n = 512, q = 205207553$ which originates in [5][11]. In the second case, we take $n = 512, q = 2^{16}+1$ which originates in [17]. Both parameter sets claim to achieve 128-bit security in their respective papers.[12]

[11] The implementation details of the scheme is available in `https://github.com/khsjds/COMPASS`

[12] At this stage, we do not compare the costs using lattice estimators; instead, we focus on practical comparisons based on lattice reduction experiments.

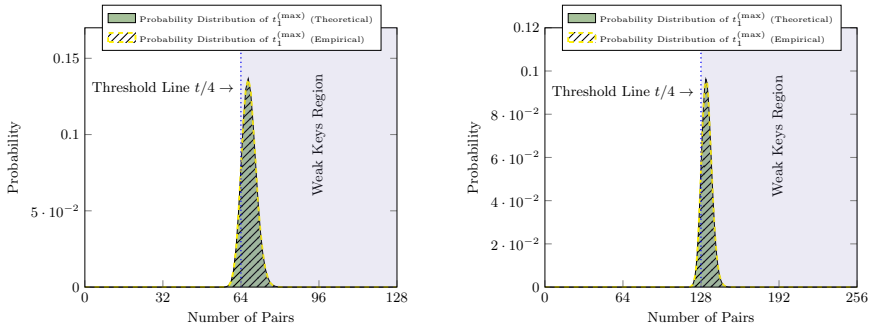

(a) $n = 512$ and $t = 256$ (b) $n = 1024$ and $t = 512$

Fig. 2: Comparison between the Probability Distribution of $t_1^{(\max)}$ computed theoretically vs. empirically.

In both cases, we first apply the direct primal attack and then use our adaptive order-2 symmetry attack for a fair comparison. In our experiments, we apply BKZ lattice reduction using the fplll software [19] with block size 65. Additionally, for our dimension reduced attack, we use the lattice reduction algorithm in parallel. The running time denotes the time required by the lattice reduction algorithm to recover the secret.

In the first case, with parameter set from [5], the direct primal attack provides the uSVP secret in approximately 59 hours. Then we conducted 20 trials with our dimension reduced attack, and in 95 % of the instances, the secrets were recovered in under 3 hours. We further observed that the attack was also successful when $t_1^{(\max)} < t/4$. We believe this improvement is due to the significantly larger modulus q, which increases the uniqueness gap of the underlying uSVP. As a result, the secrets were recovered more efficiently and consistently across all instances.

In the second case, with parameter set from [17], we never recovered the secret in the direct primal attack running lattice reduction for 15 days. We then performed 30 trials and selected the instance with the highest value of $t_1^{(\max)}$, recovering the secret within 105 hours of lattice reduction with our dimension reduced attack. Then we increased the number of trials to 100; selecting the instance with the highest $t_1^{(\max)}$ allowed us to recover the secret within 36 hours. Finally, we increased the number of trials to 200; selecting the instance with the highest $t_1^{(\max)}$ allowed us to recover the secret within 16 hours. [13]

[13] As reported in [7], key recovery for this parameter set was successful with $t_1^{(A)} = 80$ corresponding to 32768 trials, and required a total computation time of approximately 117 hours on an Intel Xeon CPU E5-2683 v4 @ 2.10 GHz (1200 MHz).

5.4 Higher order symmetry

Now we generalize the attack for higher order symmetry. We show the attack using symmetry of order 4; extending to any higher order symmetry is then similar. First note that there are exactly two (distinct) cyclic subgroups of order 4 in $\mathbb{Z}_{2n}^*$, namely the subgroup generated by $(n/2) \pm 1 \in \mathbb{Z}_{2n}^*$. Let us define $\theta_{n[1]} = (n/2) - 1$ and $\theta_{n[2]} = (n/2) + 1$, the two generators of order 4 in $\mathbb{Z}_{2n}^*$[14]. Consequently, we obtain the following results.

Lemma 4. *Let $n = 2^k$ and $k \geq 3$. Also, let $\theta_{n[1]} = (n/2) - 1 \in \mathbb{Z}_{2n}^*$. Then $\theta_{n[1]}$ is an element of order 4, i.e., $\theta_{n[1]}^4 = 1 \in \mathbb{Z}_{2n}^*$. Let $\boldsymbol{f}(\boldsymbol{x})$ be any polynomial in $\mathcal{R}_q$. Now define $\boldsymbol{\phi}_{\theta_{n[1]}}(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{f}(\boldsymbol{x}^{\theta_{n[1]}}) + \boldsymbol{f}(\boldsymbol{x}^{\theta_{n[1]}^2}) + \boldsymbol{f}(\boldsymbol{x}^{\theta_{n[1]}^3})$. Then $\boldsymbol{\phi}_{\theta_{n[1]}}$ can be generated by a basis of order $n/4$. Additionally, if the coefficients of $\boldsymbol{f}(\boldsymbol{x})$ are sampled uniformly from the set $\{-1, 0, 1\}$, then the expected squared-norm of $\boldsymbol{\phi}_{\theta_{n[1]}}$ is upper bounded by $n/3$ in the new basis representation.*

Proof. Consider the mapping

$$\boldsymbol{x}^i \to \boldsymbol{x}^i + (\boldsymbol{x}^{\theta_{n[1]}})^i + (\boldsymbol{x}^{\theta_{n[1]}^2})^i + (\boldsymbol{x}^{\theta_{n[1]}^3})^i \text{ for } 0 \leq i < n/2$$

We first claim that $\boldsymbol{x}^i \to \boldsymbol{x}^i + (\boldsymbol{x}^{\theta_{n[1]}})^i + (\boldsymbol{x}^{\theta_{n[1]}^2})^i + (\boldsymbol{x}^{\theta_{n[1]}^3})^i$ vanishes when $i \bmod 2 \neq 0$ and $2(\boldsymbol{x}^i + \boldsymbol{x}^{\theta_{n[1]} i})$ when $i \bmod 2 = 0$. To see this, first note that

$$\begin{aligned}\theta_{n[1]}^2 =& ((n/2) - 1)^2 \bmod 2n \\ =& n + 1 \bmod 2n \text{ when } k \geq 3\end{aligned}$$

Also,

$$\theta_{n[1]}^3 = (n+1)\theta_{n[1]} \bmod 2n$$

Since we are working modulo $\boldsymbol{g}(\boldsymbol{x}) = \boldsymbol{x}^n + 1$, we have

$$\begin{aligned}\boldsymbol{x}^{\theta_{n[1]}^2} =& \boldsymbol{x}^{n+1} = -\boldsymbol{x} \\ \boldsymbol{x}^{\theta_{n[1]}^3} =& \boldsymbol{x}^{(n+1)\theta_{n[1]}} = -\boldsymbol{x}^{\theta_{n[1]}}\end{aligned}$$

As a result, we have

$$\boldsymbol{x}^i + (\boldsymbol{x}^{\theta_{n[1]}})^i + (\boldsymbol{x}^{\theta_{n[1]}^2})^i + (\boldsymbol{x}^{\theta_{n[1]}^3})^i = \boldsymbol{x}^i + \boldsymbol{x}^{\theta_{n[1]} i} + (-\boldsymbol{x})^i + (-\boldsymbol{x}^{\theta_{n[1]}})^i$$

So the mapping $\boldsymbol{x}^i + (\boldsymbol{x}^{\theta_{n[1]}})^i + (\boldsymbol{x}^{\theta_{n[1]}^2})^i + (\boldsymbol{x}^{\theta_{n[1]}^3})^i$ vanishes when $i \bmod 2 \neq 0$ and $2(\boldsymbol{x}^i + \boldsymbol{x}^{\theta_{n[1]} i})$ when $i \bmod 2 = 0$, as claimed. Thus $\boldsymbol{\phi}_{\theta_{n[1]}}$ is generated by the basis

$$\{4, 2(\boldsymbol{x}^2 + \boldsymbol{x}^{2\theta_{n[1]}}), 2(\boldsymbol{x}^4 + \boldsymbol{x}^{4\theta_{n[1]}}), \ldots, 2(\boldsymbol{x}^{n/2-2} + \boldsymbol{x}^{(n/2-2)\theta_{n[1]}})\}$$

[14] Note that, when $n = 2^k$, there is exactly two (distinct) cyclic subgroups of order 2^ℓ in $\mathbb{Z}_{2n}^*$ for $1 < \ell < k$, namely the subgroups generated by $(n/2^{\ell-1}) \pm 1 \in \mathbb{Z}_{2n}^*$.

which is of order $n/4$.

Finally, if the coefficients of $\boldsymbol{f}$ are sampled uniformly at random from the set $\{-1, 0, 1\}$, the coefficients $f_i + f_{\theta_{n[1]}i}$ of $\boldsymbol{\phi}_{\theta_{n[1]}}$ lie in the set $\{-2, -1, 0, 1, 2\}$ and follow the distribution as in Table 1, except for the special coefficient associated with 4. This special coefficient follows the initial uniform distribution in the set $\{-1, 0, 1\}$. As a result, the expected squared norm is upper-bounded by $n/3$.

Lemma 5. *Let $n = 2^k$ and $k \geq 3$. Also, let $\theta_{n[2]} = (n/2) + 1 \in \mathbb{Z}^*_{2n}$. Then $\theta_{n[2]}$ is also an element of order 4, i.e., $\theta^4_{n[2]} = 1 \in \mathbb{Z}^*_{2n}$. Let $\boldsymbol{f}(\boldsymbol{x})$ be any polynomial in $\mathcal{R}_q$. Now define $\boldsymbol{\phi}_{\theta_{n[2]}}(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{f}(\boldsymbol{x}^{\theta_{n[2]}}) + \boldsymbol{f}(\boldsymbol{x}^{\theta^2_{n[2]}}) + \boldsymbol{f}(\boldsymbol{x}^{\theta^3_{n[2]}})$. Then $\boldsymbol{\phi}_{\theta_{n[2]}}$ can be generated by a basis of order $n/4$. Additionally, if the coefficients of $\boldsymbol{f}(\boldsymbol{x})$ are sampled uniformly from the set $\{-1, 0, 1\}$, then the expected squared-norm of $\boldsymbol{\phi}_{\theta_{n[2]}}$ is $n/6$ in the new basis representation.*

Proof. Since $\theta_{n[2]}$ is also an element of order 4, using the same argument as in Lemma 4, the mapping

$$\boldsymbol{x}^i \to \boldsymbol{x}^i + (\boldsymbol{x}^{\theta_{n[2]}})^i + (\boldsymbol{x}^{\theta^2_{n[2]}})^i + (\boldsymbol{x}^{\theta^3_{n[2]}})^i \text{ for } 0 \leq i < n$$

vanishes when $i \bmod 2 \neq 0$ and $2(\boldsymbol{x}^i + \boldsymbol{x}^{\theta_{n[1]}i})$ when $i \bmod 2 = 0$.

Moreover, assuming i is even, we claim that for this particular choice of $\theta_{n[2]}$, $2(\boldsymbol{x}^i + \boldsymbol{x}^{i\theta_{n[2]}})$ vanishes when $i \bmod 4 = 2$ and $4\boldsymbol{x}^i$ when $i \bmod 4 = 0$. To see this, first let $i = 2j$ for some odd j, then we have

$$\begin{aligned}
2(\boldsymbol{x}^{2j} + \boldsymbol{x}^{2j\theta_{n[2]}}) &= 2(\boldsymbol{x}^{2j} + \boldsymbol{x}^{2j(n/2)+1}) \\
&= 2(\boldsymbol{x}^{2j} + \boldsymbol{x}^{nj}\boldsymbol{x}^{2j}) \\
&= 2(\boldsymbol{x}^{2j} + \boldsymbol{x}^{n}\boldsymbol{x}^{2j}) \\
&= 2(\boldsymbol{x}^{2j} - \boldsymbol{x}^{2j}) \text{ since we are working modulo } \boldsymbol{g}(\boldsymbol{x}) = \boldsymbol{x}^n + 1 \\
&= 0
\end{aligned}$$

But, taking $i = 4j$, we have

$$\begin{aligned}
2(\boldsymbol{x}^{4j} + \boldsymbol{x}^{4j\theta_{n[2]}}) &= 2(\boldsymbol{x}^{4j} + \boldsymbol{x}^{4j((n/2)+1)})) \\
&= 2(\boldsymbol{x}^{4j} + \boldsymbol{x}^{2nj}\boldsymbol{x}^{4j}) \\
&= 2(\boldsymbol{x}^{4j} + \boldsymbol{x}^{2n}\boldsymbol{x}^{4j}) \\
&= 2(\boldsymbol{x}^{4j} + \boldsymbol{x}^{4j}) \\
&= 4\boldsymbol{x}^{4j}
\end{aligned}$$

Thus $\boldsymbol{\phi}_{\theta_{n[2]}}$ is generated by the basis

$$\{4, 4\boldsymbol{x}^4, \ldots 4(\boldsymbol{x}^{n-4})\}$$

which is of order $n/4$. Finally, if the coefficients of $\boldsymbol{f}$ are sampled uniformly at random from the set $\{-1, 0, 1\}$, the coefficients of $\boldsymbol{\phi}_{\theta_{n[2]}}$ are also in the set $\{-1, 0, 1\}$ and follow the initial uniform distribution. As a result, we have the expected squared norm as $n/6$.

Finally, in the group $\mathbb{Z}_{2n}^*$, there also exist a (non-cyclic) subgroup of order 4 isomorphic to the Klein 4-group. This subgroup can be explicitly written as $V_{\theta_{n[3]}} = \{\pm 1, n \pm 1\} \subset \mathbb{Z}_{2n}^*$. We show that this subgroup also yields a symmetry of order 4 in $\mathcal{R}_q$.

Lemma 6. *Let $n = 2^k$ and $k \geq 3$. Let $\boldsymbol{f}(\boldsymbol{x})$ be any polynomial in $\mathcal{R}_q$. Now define $\boldsymbol{\phi}_{\theta_{n[3]}}(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{f}(\boldsymbol{x}^{-1}) + \boldsymbol{f}(\boldsymbol{x}^{n+1}) + \boldsymbol{f}(\boldsymbol{x}^{n-1})$. Then $\boldsymbol{\phi}_{\theta_{n[3]}}$ can be generated by a basis of order $n/4$. Additionally, if the coefficients of $\boldsymbol{f}(\boldsymbol{x})$ are sampled uniformly from the set $\{-1, 0, 1\}$, then the expected squared-norm of $\boldsymbol{\phi}_{\theta_{n[3]}}$ is upper bounded by $n/3$ in the new basis representation.*

Proof. Consider the mapping

$$\boldsymbol{x}^i \to \boldsymbol{x}^i + (\boldsymbol{x}^{-1})^i + (\boldsymbol{x}^{n-1})^i + (\boldsymbol{x}^{n+1})^i \text{ for } 0 \leq i < n/2$$

We first claim that $\boldsymbol{x}^i \to \boldsymbol{x}^i + (\boldsymbol{x}^{-1})^i + (\boldsymbol{x}^{n-1})^i + (\boldsymbol{x}^{n+1})^i$ vanishes when $i \bmod 2 \neq 0$ and $2(\boldsymbol{x}^i + \boldsymbol{x}^{-i})$ when $i \bmod 2 = 0$. To see this, recall

$$\begin{aligned} \boldsymbol{x}^{n-1} &= -\boldsymbol{x}^{-1}, \\ \boldsymbol{x}^{n+1} &= -\boldsymbol{x}. \end{aligned}$$

As a result, we have

$$\boldsymbol{x}^i + (\boldsymbol{x}^{-1})^i + (\boldsymbol{x}^{n-1})^i + (\boldsymbol{x}^{n+1})^i = \boldsymbol{x}^i + \boldsymbol{x}^{-i} + (-\boldsymbol{x})^{-i} + (-\boldsymbol{x})^i$$

So the mapping $\boldsymbol{x}^i + (\boldsymbol{x}^{-1})^i + (\boldsymbol{x}^{n-1})^i + (\boldsymbol{x}^{n+1})^i$ vanishes when $i \bmod 2 \neq 0$ and $2(\boldsymbol{x}^i + \boldsymbol{x}^{-i})$ when $i \bmod 2 = 0$, as claimed. Thus $\boldsymbol{\phi}_{\theta_{n[3]}}$ is generated by the basis

$$\{4, 2(\boldsymbol{x}^2 + \boldsymbol{x}^{-2}), 2(\boldsymbol{x}^4 + \boldsymbol{x}^{-4}), \ldots, 2(\boldsymbol{x}^{n/2-2} + \boldsymbol{x}^{-(n/2-2)})\}$$

which is of order $n/4$.

Finally, if the coefficients of $\boldsymbol{f}$ are sampled uniformly at random from the set $\{-1, 0, 1\}$, the coefficients of $\boldsymbol{\phi}_{\theta_{n[3]}}$ lie in the set $\{-2, -1, 0, 1, 2\}$ and follow the distribution as in Table 1, except for the special coefficient associated with 4. This special coefficient follows the initial uniform distribution in the set $\{-1, 0, 1\}$. As a result, the expected squared norm is upper-bounded by $n/3$.

We can similarly formulate a low-dimensional Knapsack problem to solve the PV Knapsack problem using the symmetries of order 4. But it only allows to recover a quarter of the information about $\boldsymbol{f}$. The $\boldsymbol{\phi}_{\theta_{n[k]}}$ for $k \in \{1, 2, 3\}$ is equivalent to $\boldsymbol{\phi}_+$ in order 2. However, we cannot directly construct an equivalent of $\boldsymbol{\phi}_-$ in order 4. Below we present the full recovery for $k = 1$; the cases $k = 2, 3$ follow straightforwardly.

In order to recover the complete solution of $\boldsymbol{f}$, we define three additional polynomials

$$\boldsymbol{f}'(\boldsymbol{x}) = \boldsymbol{x}\boldsymbol{f}(\boldsymbol{x}), \qquad \boldsymbol{f}''(\boldsymbol{x}) = \boldsymbol{x}^2\boldsymbol{f}(\boldsymbol{x}), \qquad \boldsymbol{f}'''(\boldsymbol{x}) = \boldsymbol{x}^3\boldsymbol{f}(\boldsymbol{x})$$

Now, define $\boldsymbol{\phi}'_{\theta_{n[1]}}(\boldsymbol{x}) = \boldsymbol{f}'(\boldsymbol{x}) + \boldsymbol{f}'(\boldsymbol{x}^{\theta_{n[1]}}) + \boldsymbol{f}'(\boldsymbol{x}^{\theta^2_{n[1]}}) + \boldsymbol{f}'(\boldsymbol{x}^{\theta^3_{n[1]}})$. It is easy to verify that the polynomial $\boldsymbol{\phi}'_{\theta_n[1]}(\boldsymbol{x})$ can be generated by a basis of order $n/4$. Also, both $\boldsymbol{\phi}_{\theta_n[1]}$ and $\boldsymbol{\phi}'_{\theta_n[1]}$ have the same expected squared norm. Indeed, since we are working modulo $\boldsymbol{g}(\boldsymbol{x}) = \boldsymbol{x}^n + 1$, multiplication by $\boldsymbol{x}$ only gives different shifts of the coefficients (in absolute value) represented in the polynomial basis. Finally, writing

$$\boldsymbol{\phi}'_{\theta_n[1]}(\boldsymbol{x}) = \boldsymbol{x}\boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{x}^{\theta_{n[1]}}\boldsymbol{f}(\boldsymbol{x}^{\theta_{n[1]}}) + \boldsymbol{x}^{\theta^2_{n[1]}}\boldsymbol{f}(\boldsymbol{x}^{\theta^2_{n[1]}}) + \boldsymbol{x}^{\theta^3_{n[1]}}\boldsymbol{f}(\boldsymbol{x}^{\theta^3_{n[1]}})$$

we can see that if the evaluations of $\boldsymbol{f}$ are known in a quadruplet $(\omega, \omega^{\theta_{n[1]}}, \omega^{\theta^2_{n[1]}}, \omega^{\theta^3_{n[1]}})$ that appear in Ω_t, we know the evaluation of $\boldsymbol{\phi}'_{\theta_n[1]}$. As a result, we can use lattice reduction, similar to $\boldsymbol{\phi}_{\theta_n[1]}$, to recover $\boldsymbol{\phi}'_{\theta_n[1]}$. This allows us to get another quarter information of $\boldsymbol{f}$. Combining the solutions of $\boldsymbol{\phi}_{\theta_{n[1]}}, \boldsymbol{\phi}'_{\theta_{n[1]}}, \boldsymbol{\phi}''_{\theta_{n[1]}}, \boldsymbol{\phi}'''_{\theta_{n[1]}}$ allows to recover the complete $\boldsymbol{f}$.

However, for a random Ω_t, the number of quadruples that arise in this setting is significantly reduced on the average. Consequently, the number of weak keys that can be targeted using order-4 symmetries is much smaller. Experimentally, we find that for the concrete parameters used in the literature, less than 1 % of the keys can be attacked using order 4. Nevertheless, since detecting the presence of such symmetries is computationally inexpensive, it is advisable to check for their existence prior to initiating the lattice-reduction phase of the attack.

6 New attacks on the integer vSIS commitment problem

In this section, we present new attacks on the integer vSIS problem. We first show that, for certain prime moduli q, the distinguishing problem can be solved efficiently. We then investigate a key-recovery attack applicable to arbitrary prime moduli of small order.

6.1 Attack using small order elements

Recall in the integer vSIS commitment problem, we are given $t = \lfloor n/2 \rfloor$ random evaluations from the set $\mathbb{Z}_q^* \setminus \{-1, 1\}$ of a hidden polynomial $\boldsymbol{f}(\boldsymbol{x}) \in \mathbb{Z}_q^{<n}[\boldsymbol{x}]$ with coefficients in $\{0, 1\}$. We have excluded the elements $\{-1, 1\}$ from the evaluation set to avoid a simple key distinguishing attack on the Knapsack problem.

To see the attack more clearly, first note that the probability of occurrence of 1 (or -1) in the random subset $\mathcal{S}_t$ of $\mathbb{Z}_q^*$ is $t/(q-1)$. Now, let us assume we know the evaluation of $\boldsymbol{f}$ at 1 (or -1).

Note that $0 \leq \boldsymbol{f}(\pm 1) \leq n$ when the coefficients of $\boldsymbol{f}$ are in the set $\{0, 1\}$. If this evaluation returns a value in the integer interval $[0, n]$, the distinguisher outputs 1 when the original distribution is integer vSIS with probability 1, and outputs 1 when it comes from the uniform distribution with probability at most $(n+1)/q$. As a result, including the evaluation of 1 (or -1) $\in \mathbb{Z}_q^*$ provides a distinguishing attack with advantage at least $(1 - (n+1)/q)(t/(q-1))$, which is non-negligible when $q > n + 1$.

The exclusion of $\pm 1 \in \mathbb{Z}_q^*$ removes the elements of order 1 and 2 in the ring $\mathbb{Z}_q^*$. However, for some choices of q, $\mathbb{Z}_q^*$ can also have other small order elements. Indeed, depending on the choice of q, the group $\mathbb{Z}_q^*$ may contain elements of varying order. In particular, the order structure of $\mathbb{Z}_q^*$ is determined by the divisors of $|\mathbb{Z}_q^*| = q - 1$. If the factorization of $q - 1$ is known and includes small order elements, we show that this can also be exploited to mount a distinguishing attack against the integer vSIS problem.

Before presenting the details of the attack, we first make the coefficient distribution of the hidden polynomial symmetric around 0. As we will see later, this provides a better approximation of the secret and increases the effectiveness and applicability of our small-order attack.

For this, we define

$$\boldsymbol{f}'(\boldsymbol{x}) := 2(\boldsymbol{f}(x) - \sum_{i=0}^{n-1}(1/2)\boldsymbol{x}^i) \in \mathbb{Z}_q^{<n}$$

Note that, when $f_i = 0$ (similarly, when $f_i = 1$), the coefficient of the polynomial $\boldsymbol{f}'$ given by $f'_i = 2(f_i - \frac{1}{2}) = -1$ ($f'_i = 1$). Consequently, if the coefficients of $\boldsymbol{f}$ are sampled uniformly at random from $\{0, 1\}$, then the coefficients of $\boldsymbol{f}'$ are uniformly distributed over $\{-1, 1\}$. Finally, if the evaluation of $\boldsymbol{f}$ is known at $\upsilon \in \mathcal{S}_t$, we know the evaluation of $\boldsymbol{f}'$ at υ.[15] As a result, recovering $\boldsymbol{f}$ from its evaluations can be formulated as recovering $\boldsymbol{f}'$ from its evaluations. In the remainder of the attack in this subsection, we therefore focus on recovering $\boldsymbol{f}'$, after which recovering $\boldsymbol{f}$ is straightforward.

We now present the small-order attack, beginning with the following lemma.

Lemma 7. *Let $2 < r < n$ be an integer satisfying $r = 0 \bmod q - 1$. Let $\beta \in \mathbb{Z}_q^*$ be an element of order r, i.e., $\beta^r = 1 \in \mathbb{Z}_q^*$. Also, let $\boldsymbol{f}'(\boldsymbol{x})$ be any polynomial in $\mathbb{Z}_q^{<n}[\boldsymbol{x}]$. Then the evaluation of $\boldsymbol{f}'(\boldsymbol{x})$ at β can be generated by a basis of order r. Additionally, if the coefficients of $\boldsymbol{f}'(\boldsymbol{x})$ are sampled uniformly from the set $\{-1, 1\}$, then the expected squared-norm of the new basis representation is $r\lceil n/r \rceil$.*

Proof. Since $r = 0 \bmod q - 1$, there exists $\beta \in \mathbb{Z}_q^*$ of order r. In fact, there are exactly $\phi(r)$ many such $\beta \in \mathbb{Z}_q^*$. Then for any β, the evaluation $\boldsymbol{f}'(\beta)$ can be written as

$$\boldsymbol{f}'(\beta) := \sum_{i=0}^{n-1} f'_i \beta^i = \sum_{i=0}^{r-1} f''_i \beta^i \text{ since } \beta^r = 1 \bmod q$$

where $f''_i = \sum_{k=1}^{\lceil n/r \rceil} f'_{i+kr}$. As a result, $\boldsymbol{f}'(\beta)$ can be generated by the basis

$$\{1, \beta, \ldots, \beta^{r-1}\}$$

[15] It is to be noted that the evaluation of the polynomial $\sum_{i=0}^{n-1}(1/2)\boldsymbol{x}^i \in \mathbb{Z}_q^{<n}[\boldsymbol{x}]$ at the point $\upsilon \in \mathcal{S}_t$ can be computed simply by $\sum_{i=0}^{n-1}(1/2)\upsilon^i$. So the evaluation of $\boldsymbol{f}'(\upsilon)$ is given by $2(\boldsymbol{f}(\upsilon) - \sum_{i=0}^{n-1}(1/2)\upsilon^i)$.

which is of order r.

Now, if the coefficients of $\boldsymbol{f}'$ are uniform at random from the set $\{-1, 1\}$, the expected squared value of f''_i is $\lceil n/r \rceil$. This is because the coefficients f'_i are symmetric about 0 and therefore have an expected value of 0. Hence, by linearity of expectation, the expected squared norm of the coefficients of $\boldsymbol{f}''$ is $r\lceil n/r \rceil$, as claimed.

Next we show that the knowledge of some evaluations of small order elements β helps to mount a distinguishing attack against the integer vSIS problem, thanks to the Lemma 7.

Let us assume t_r evaluations of β (each of order r) are known, where $0 \leq t_r \leq \phi(r)$. Now, consider a Knapsack problem using t_r many evaluations of β that appear in $\mathcal{S}_t$. For this, we define a matrix $\boldsymbol{B}_r$ with t_r many rows and r many columns. The rows of $\boldsymbol{B}_r$ consist of the evaluations of each of the r monomials at the t_r elements of order r. We also construct a vector $\boldsymbol{b}_r$ with entries as the known evaluations of $\boldsymbol{f}'(\beta)$. Now we search for the small solutions on the kernel lattice of dimension $r + 1$ and volume q^{t_r} defined by

$$\mathcal{L}^{\perp}_{\boldsymbol{B}'_r,q} = \{\boldsymbol{y} \in \mathbb{Z}^{r+1} | \boldsymbol{B}'_r \boldsymbol{y} = \boldsymbol{0} \bmod q\}$$

where $\boldsymbol{B}'_r = (\boldsymbol{B}_r | \boldsymbol{b}_r)$.

For simplicity, assume r is a prime integer, then for each order r elements β, we have $\sum_{i=0}^{r-1} \beta^i = 0 \bmod q$. As a result, the trivial vector $(\boldsymbol{1}, 0)^T$ belongs to the lattice $\mathcal{L}^{\perp}_{\boldsymbol{B}'_r,q}$, which we expect to be unusually short in this lattice.

Now, in the integer vSIS distribution, using the Lemma, we know the vector $(\boldsymbol{f}'', -1)^T \in \mathcal{L}^{\perp}_{\boldsymbol{B}'_r,q}$ has an expected squared norm of $r\lceil n/r \rceil + 1$.

But, in the uniform distribution, we hope there is no non-trivial vector of squared norm $r\lceil n/r \rceil + 1$ in the lattice. Then the non-trivial uSVP solution in the lattice yields a distinguishing attack.[16] The complexity of the attack is determined by the cost of recovering the uSVP solutions. Since we assume r is small, we expect these solutions to be recoverable efficiently in this lattice, for example, using the LLL algorithm [16].

Now let us analyze the effectiveness of the attack. Since the lattice we consider here come from a cryptographic problem, we can follow a standard heuristic approach and assume that they behave as randomly as it can. More precisely, as in primal lattice attacks, we consider a lattice that is guaranteed to contain a distinguished short vector corresponding to the secret. We then heuristically assume that all other linearly independent lattice vectors have lengths consistent with those of a random lattice, which we estimate using Minkowski's bound.

As a result, the attack is successful when $(\boldsymbol{f}'', -1)^T$ is a guaranteed uSVP solution in the lattice, i.e., the norm of $(\boldsymbol{f}'', -1)^T$ is significantly shorter than

[16] Since $(\boldsymbol{1}, 0)^T$ and $(\boldsymbol{f}'', -1)^T$ belong to the lattice, any integer linear combination of them also lies in the lattice. In some of our trials, the recovered uSVP solution was a very small linear combination of these vectors, which is still sufficient for a distinguishing attack because it reveals information about $\boldsymbol{f}''$. Moreover, the exact vector $(\boldsymbol{f}'', -1)^T$ can be recovered by adding further linear constraints, such as guessing $\sum_{i=0}^{n-1} f'_i = \sum_{i=0}^{r-1} f''_i$.

any vector that should normally occur in this lattice. Then assuming the norm of the shortest vector by the square-root of the expected squared-norm and the length of the other vectors in the lattice has the norm from Minkowski's bound (along with replacing the value $\sqrt{n+1}$ by $\sqrt{n}$), this happens when

$$\sqrt{n} < \sqrt{r} q^{t_r/r}$$

i.e., whenever $t_r > r \log_q(\sqrt{\frac{n}{r}})$.

When $\mathcal{S}_t$ is randomly chosen, the value of t_r is randomized, and follows the hypergeometric distribution with parameter $(q-3, \phi(r), t)$. Indeed, the probability of getting t_r many order r elements in a random subset $\mathcal{S}_t$ is given by $\frac{\binom{\phi(r)}{t_r}\binom{q-3-\phi(r)}{t-t_r}}{\binom{q-3}{t}}$. As a result, the expected value of t_r is $t\phi(r)/(q-3)$.

Now, to ensure that the distinguishing problem remains hard, we require the expected value to be significantly smaller than the target value of t_r for every small divisors r of $q-1$. Under this condition, the above distinguishing attack achieves only negligible advantage. As an alternative, one may choose the prime q such that $q-1$ has all (odd) prime divisors of the size order n.

Finally, to illustrate the effectiveness of the small-order attack, we provide a concrete example. We choose the parameters $n = 500$, $t = 250$, and $q = 1009$. Note that

$$q - 1 = 2^4 \cdot 3^2 \cdot 7.$$

We first apply the direct primal attack to this parameter set, using both LLL and BKZ with block size 10. However, neither approach was able to recover any distinguishing information about the problem. In contrast, by using the new attack with elements of order up to

$$168 = 2^3 \cdot 3 \cdot 7,$$

we obtain a distinguishing advantage of 0.31 using LLL alone as the underlying lattice-reduction algorithm. The distinguishing advantage increases to 0.40 when using BKZ with block size 10. The results are based on 1000 trials, and the total running time was only a few seconds in the Mac desktop.

Note 4. If the coefficients of the hidden polynomial are binary, as in the integer vSIS assumption, the corresponding norm in the new basis of order r can be roughly estimated as $n/\sqrt{r}$. As a result, the required threshold t_r increases to $r \log_q(n/r)$, thereby reducing the effectiveness and applicability of our small-order attack.

Note 5. In the PV Knapsack problem, the evaluation set is restricted to primitive roots ω: those of order $2n$ when $\boldsymbol{g}(\boldsymbol{x}) = \boldsymbol{x}^n + 1$ and those of order n when $\boldsymbol{g}(\boldsymbol{x}) = \boldsymbol{x}^n - 1$. This restriction prevents the small-order attack from being applicable to the PV Knapsack problem.

6.2 Key recovery using symmetry of order 2

We now examine the key-recovery attack on the integer vSIS commitment problem for an arbitrary prime modulus q. Recall, in the case of the vSIS problem, we work in the ring $\mathbb{Z}_q[\boldsymbol{x}]$, instead of the quotient polynomial ring $\mathcal{R}_q$. As a result, we look for the symmetries in $\mathbb{Z}_q^*$ to reduce the lattice dimension to solve the problem.

We begin with symmetries of order 2 in $\mathbb{Z}_q^*$. Since every odd prime q satisfies $q = 1 \bmod 2$, such a symmetry always exists.

Lemma 8. *Let $\boldsymbol{f}(\boldsymbol{x})$ be any polynomial in $\mathbb{Z}_q^{<n}[\boldsymbol{x}]$. Now define $\boldsymbol{\phi}_\pm(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{x}) \pm \boldsymbol{f}(-\boldsymbol{x})$. Then $\boldsymbol{\phi}_\pm$ can be generated by a basis of order $n_\pm$, where $n_+ = \lceil n/2 \rceil$ and $n_- = \lfloor n/2 \rfloor$. Additionally, if the coefficients of $\boldsymbol{f}(\boldsymbol{x})$ are sampled uniformly from the set $\{0, 1\}$, then the expected squared-norm of $\boldsymbol{\phi}_\pm$ is $n_\pm/2$ in the new basis representation.*

The proof is similar to the Lemma 2. In particular, the polynomial $\boldsymbol{\phi}_+$ can be generated by the basis

$$\{2, 2\boldsymbol{x}^2, \ldots, 2\boldsymbol{x}^{2\lceil n/2 \rceil - 2}\}$$

which is of order $n_+ = \lceil n/2 \rceil$. Similarly, the polynomial $\boldsymbol{\phi}_-$ can be generated by the basis

$$\{2\boldsymbol{x}, 2\boldsymbol{x}^3, \ldots, 2\boldsymbol{x}^{2\lfloor n/2 \rfloor - 1}\}$$

which is of order $n_- = \lfloor n/2 \rfloor$. Also, if the coefficients of $\boldsymbol{f}$ are sampled from the set $\{0, 1\}$, the coefficients of $\boldsymbol{\phi}_\pm$ are also in the set $\{0, 1\}$ and follow the initial uniform distribution. As a result, the expected squared-norm of $\boldsymbol{\phi}_\pm$is $n_\pm/2$ in the new basis representation.

Now using Lemma 8, we can construct a new knapsack problem to solve the problem. Like in the PV Knapsack, we look for recovering $\boldsymbol{\phi}_\pm(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{x}) \pm \boldsymbol{f}(-\boldsymbol{x})$ to solve the integer vSIS commitment problem.

Let $\mathcal{S}_{2t_1}$ be the largest subset of $\mathcal{S}_t$ containing all the negation-invariant pairs, i.e., $\mathcal{S}_{2t_1}$ contains all pairs of the form $(\upsilon, -\upsilon)$. Then we can construct new Knapsack problems in reduced dimensions using t_1 many evaluations of the polynomials $\boldsymbol{\phi}_\pm$. Now, similar to the PV knapsack case, we define matrices $\boldsymbol{B}_\pm$ with t_1 many rows and $n_\pm$ many columns. The rows of $\boldsymbol{B}_\pm$ are the evaluations of each of the $n_\pm$ monomials in this t_1 representative of pairs in $\mathcal{S}_{2t_1}$. We also construct vector $\boldsymbol{b}_\pm$ with entries as the known evaluations of $\boldsymbol{\phi}_\pm$ at each of these representatives. Now we search for the solution $\texttt{Solve-uSVP}(\boldsymbol{B}_\pm, \boldsymbol{b}_\pm)$. We achieve this by solving the uSVP on the kernel lattices

$$\mathcal{L}^\perp_{\boldsymbol{B}'_\pm, q} = \{\boldsymbol{y} \in \mathbb{Z}^{n_\pm + 1} | \boldsymbol{B}'_\pm \boldsymbol{y} = \boldsymbol{0} \bmod q\}$$

where $\boldsymbol{B}'_\pm = (\boldsymbol{B}_\pm | \boldsymbol{b}_\pm)$. We also expect $(\boldsymbol{\phi}_\pm | -1)^T$ to be an unusually short vector in the lattice, which can then be recovered using lattice reduction algorithms. The pseudo-code of our attack is given in Algorithm 2.

6.3 Analysis of the key recovery attack

Similar to the PV Knapsack, we start by comparing the cost of solving vSIS problem by looking at the root Hermite factor δ.

In the direct primal attack, we solve a uSVP instance in dimension $n+1$ and volume q^t. Here, the expected squared-norm of the shortest vector that we look for is $n/2+1$. We estimate the expected norm by $\sqrt{n/2+1}$. Hence, the root Hermite factor can be estimated by

$$\delta_{\text{full}} = \left(\frac{\sqrt{n+1} q^{t/(n+1)}}{\sqrt{n/2+1}} \right)^{1/(n+1)}.$$

For our symmetry-based attack, we solve uSVP instance in dimension $n_{\pm}+1$ and volume q^{t_1}. Here, the expected squared-norm of the shortest vector that we look for is $n_{\pm}/2+1$. We estimate the expected norm by $\sqrt{n_{\pm}/2+1}$. Therefore, the corresponding root Hermite factor can be estimated by

$$\delta_{\text{reduced}} = \left(\frac{\sqrt{(n_{\pm}+1)} q^{t_1/(n_{\pm}+1)}}{\sqrt{n_{\pm}/2+1}} \right)^{1/(n_{\pm}+1)}.$$

Like before, we compare the corresponding root Hermite factors δ_{full} and δ_{reduced} to identify when our new attack outperforms the direct primal attack. By simplifying the expressions and ignoring the constant factors, this happens whenever

$$(q^{t/n})^{1/n} \leq (q^{2t_1/n})^{2/n}, \text{i.e. whenever } t \leq 4t_1.$$

We now analyze the condition $t_1 \geq t/4$ for a random $\mathcal{S}_t \subset \mathbb{Z}_q^* \setminus \{-1, 1\}$ under different values of q. For this, let the modulus of the vSIS commitment problem be $q = \alpha n$ for some $\alpha > 1$. Under this assumption, we can approximate the expected number of pairs in $\mathcal{S}_t$ by $t/(4\alpha)$. The detailed analysis of the expected number of pairs in $\mathcal{S}_t$ is provided in Appendix B.

When α is large, the expected number of pairs for a random $\mathcal{S}_t$ is significantly smaller than $t/4$. Consequently, only a small fraction of keys are vulnerable to the proposed order-2 attack. In contrast, when α is a small constant-particularly when α is close to 1-the expected number of pairs approaches $t/4$. This regime represents the crossover point between the direct primal attack and our symmetry-based attack. As a result, a non-negligible fraction of keys becomes susceptible to the order-2 attack. Therefore, to avoid weak instances, the assumption should exclude the small-modulus setting.

Finally, to demonstrate the effectiveness of the key-recovery attack, we consider a concrete example with parameters $n = 260$, $t = 130$, and $q = 521$. As in the experiments of Section 5.3, we first apply the direct primal attack and then our order-2 symmetry attack, using a block size of 65 in both cases. Since our attack depends only on the value of t_1, rather than on the particular choice of elements in $\mathcal{S}_{2t_1}$, we fix t_1 and randomly select t_1 distinct pairs in $\mathcal{S}_{2t_1}$. The probability of obtaining a given value of t_1 directly determines the expected number of trials

Algorithm 2: Order 2 Symmetry Attack on Integer vSIS commitment

Input : parameters n, q; sample size t;
Evaluations $E = \{(\upsilon, \boldsymbol{f}(\upsilon)) \mid \upsilon \in \mathcal{S}_t\}$.
Output : The secret polynomial $\boldsymbol{f}(\boldsymbol{x}) \in \mathbb{Z}_q^{<n}[x]$ or failure $\perp$.

Initialize pair set: $P \leftarrow \emptyset$ `// Initialization`
`// Symmetry Detection Phase (Only Sign Symmetry exists in` $\mathbb{Z}_q^*$`)`
foreach $\upsilon \in \mathcal{S}_t$ **do**
  if $-\upsilon \in \mathcal{S}_t$ ***and*** $(-\upsilon, \upsilon) \notin P$ **then**
    $P \leftarrow P \cup \{(\upsilon, -\upsilon)\}$ `// Check pair` $\upsilon \leftrightarrow -\upsilon$
  end
end
Calculate pair count: $t_1 \leftarrow |P|$ `// Check pair count against threshold`
if $t_1 < t/4$ **then**
  return $\perp$ `// Fail: Not enough symmetry pairs found`
end
Initialize $\boldsymbol{b}_\pm \in \mathbb{Z}_q^{t_1}$ `// Dimension Reduction Phase`
foreach $(\upsilon, -\upsilon) \in P$ **do**
  $u_\pm \leftarrow (\boldsymbol{f}(\upsilon) \pm \boldsymbol{f}(-\upsilon)) \bmod q$
  Append $u_\pm$ to $\boldsymbol{b}_\pm$
end
Construct Reduced Bases for $\boldsymbol{\phi}_\pm(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{x}) \pm \boldsymbol{f}(-\boldsymbol{x})$
Construct matrix $\boldsymbol{B}_+$ using basis $\{2, 2\boldsymbol{x}^2, \ldots, 2\boldsymbol{x}^{2\lceil n/2 \rceil - 2}\}$ evaluated at $\upsilon \in P$
Construct matrix $\boldsymbol{B}_-$ using basis $\{2\boldsymbol{x}, 2\boldsymbol{x}^3, \ldots, 2\boldsymbol{x}^{2\lfloor n/2 \rfloor - 1}\}$ evaluated at $\upsilon \in P$
$\hat{\boldsymbol{\phi}}_\pm \leftarrow$ `Solve-uSVP`$(\boldsymbol{B}_\pm, \boldsymbol{b}_\pm)$ `// Apply Lattice Reduction`
Recover candidate $\boldsymbol{f}(x)$: $2\boldsymbol{f}(x) \equiv \hat{\boldsymbol{\phi}}_+(\boldsymbol{x}) + \hat{\boldsymbol{\phi}}_-(\boldsymbol{x})$ `// Lifting`
return $\boldsymbol{f}(\boldsymbol{x})$

required to obtain the corresponding set of pairs. This value can be obtained from the distribution of $t_1^{(\sigma)}$ given in Section 5.2.

For the direct primal attack, the uSVP secret is recovered in approximately 92 hours. In contrast, our dimension-reduced attack with $t_1 = 33$, which requires approximately 2^{23} expected trials, recovers the secret in less than one hour. For $t_1 = 32$ and $t_1 = 31$, corresponding to approximately 2^{21} and 2^{19} expected trials, respectively, the secret is recovered in approximately 10 and 66 hours, respectively.

6.4 Key recovery using higher order symmetry

Now, for some specific values of q, we can also use the higher order symmetry. In the next lemma, we show the symmetry of order 4 when $q = 1 \bmod 4$. It then extends naturally to any arbitrary symmetry of order $k < n$ when $q = 1 \bmod k$.

Lemma 9. *Let $q = 1 \bmod 4$. Let $\theta_q \in \mathbb{Z}_q^*$ be an element of order 4, i.e., $\theta_q^4 = 1 \in \mathbb{Z}_q^*$. Also, let $\boldsymbol{f}(\boldsymbol{x})$ be any polynomial in $\mathbb{Z}_q^{<n}[\boldsymbol{x}]$. Now define $\boldsymbol{\phi}_{\theta_q}(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{f}(\theta_q\boldsymbol{x}) + \boldsymbol{f}(\theta_q^2\boldsymbol{x}) + \boldsymbol{f}(\theta_q^3\boldsymbol{x})$. Then $\boldsymbol{\phi}_{\theta_q}$ can be generated by a basis of order n_{θ_q}, where $n_{\theta_q} = \lceil n/4 \rceil$. Additionally, if the coefficients of $\boldsymbol{f}(\boldsymbol{x})$ are sampled uniformly from the set $\{0,1\}$, then the expected squared-norm of $\boldsymbol{\phi}_{\theta_q}$ is $n_{\theta_q}/2$ in the new basis representation.*

Proof. Since $|\mathbb{Z}_q^*| = q-1$ and $q = 1 \bmod 4$, there exists $\theta_q \in \mathbb{Z}_q^*$, which is of order 4. Now, if θ_q is of order 4, we have the following:

$$\begin{aligned} 1 + \theta_q^i + \theta_q^{2i} + \theta_q^{3i} &= 0 \bmod q \text{ when } i \neq 0 \bmod 4 \\ &= 4 \bmod q \text{ when } i = 0 \bmod 4 \end{aligned}$$

Indeed, since θ_q is of order 4,

$$\theta_q^i = \theta_q^{2i} = \theta_q^{3i} = 1 \text{ when } i = 0 \bmod 4$$

and

$$\begin{aligned} 1 + \theta_q^i + \theta_q^{2i} + \theta_q^{3i} &= 1 + \theta_q + \theta_q^2 + \theta_q^3 \\ &= 0 \text{ when } i \neq 0 \bmod 4 \end{aligned}$$

Now consider the mapping

$$\boldsymbol{x}^i \to \boldsymbol{x}^i + \boldsymbol{x}^i\theta_q^i + \boldsymbol{x}^i\theta_q^{2i} + \boldsymbol{x}^i\theta_q^{3i} \text{ for } 0 \leq i < n$$

From the above relation on θ_q, the mapping vanishes when $i \bmod 4 \neq 0$ and $4\boldsymbol{x}^i$ when $i \bmod 4 = 0$. As a result, the $\boldsymbol{\phi}_{\theta_q}(\boldsymbol{x})$ is generated by the basis

$$\{4, 4\boldsymbol{x}^4, \ldots, 4\boldsymbol{x}^{4\lceil n/4 \rceil - 4}\}$$

which is of order $n_{\theta_q} = \lceil n/4 \rceil$.

Now, if the coefficients of $\boldsymbol{f}$ are sampled from the set $\{0,1\}$, the coefficients of $\boldsymbol{\phi}_{\theta_q}$ are also in the set $\{0,1\}$ and follow the initial uniform distribution. As a result, the expected squared-norm of $\boldsymbol{\phi}_{\theta_q}$ is $n_{\theta_q}/2$.

Similarly, using the above lemma, we can formulate a low-dimensional Knapsack problem to attack the integer vSIS commitment problem. But it only helps to recover a quarter of the coefficients of the hidden polynomial $\boldsymbol{f}$. Indeed, from Lemma 9, recovering $\boldsymbol{\phi}_{\theta_q}$ allows to recover coefficients of $\boldsymbol{f}$, which is of the form f_{4i} for $0 \leq i < \lceil n/4 \rceil$. In order to recover the missing coefficients of $\boldsymbol{f}$, we define three additional polynomials

$$\boldsymbol{f}'(\boldsymbol{x}) = \boldsymbol{x}\boldsymbol{f}(\boldsymbol{x}), \qquad \boldsymbol{f}''(\boldsymbol{x}) = \boldsymbol{x}^2\boldsymbol{f}(\boldsymbol{x}), \qquad \boldsymbol{f}'''(\boldsymbol{x}) = \boldsymbol{x}^3\boldsymbol{f}(\boldsymbol{x})$$

Also, define $\boldsymbol{\phi}'_{\theta_q}(\boldsymbol{x}) = \boldsymbol{f}'(\boldsymbol{x}) + \boldsymbol{f}'(\theta_q\boldsymbol{x}) + \boldsymbol{f}'(\theta_q^2\boldsymbol{x}) + \boldsymbol{f}'(\theta_q^3\boldsymbol{x})$. It is easy to verify that $\boldsymbol{\phi}'_{\theta_q}(\boldsymbol{x})$ can be generated by a basis of order $\lfloor n/4 \rfloor$. Indeed, $\boldsymbol{\phi}'_{\theta_q}(\boldsymbol{x})$ can be

generated by the basis $\{4\boldsymbol{x}^4, 4\boldsymbol{x}^8, \ldots, 4\boldsymbol{x}^{4\lceil(n+1)/4\rceil-4}\}$. Also, since we are working over the polynomial ring $\mathbb{Z}_q[\boldsymbol{x}]$, $\boldsymbol{\phi}'_{\theta_q}$ has the expected squared-norm $\lfloor n/4 \rfloor/2$.

Finally, writing

$$\boldsymbol{\phi}'_{\theta_q}(\boldsymbol{x}) = \boldsymbol{x}\boldsymbol{f}(\boldsymbol{x}) + (\theta_q\boldsymbol{x})\boldsymbol{f}(\theta_q\boldsymbol{x}) + (\theta_q^2\boldsymbol{x})\boldsymbol{f}(\theta_q^2\boldsymbol{x}) + (\theta_q^3\boldsymbol{x})\boldsymbol{f}(\theta_q^3\boldsymbol{x})$$

we can see that if the evaluations of $\boldsymbol{f}$ are known in a quadruplet $(\upsilon, \theta_q\upsilon, \theta_q^2\upsilon, \theta_q^3\upsilon)$ that appear in $\mathcal{S}_t$, we know the evaluation of $\boldsymbol{\phi}'_{\theta_q}$.

As a result, similar to $\boldsymbol{\phi}_{\theta_q}$, we can use lattice reduction to recover $\boldsymbol{\phi}'_{\theta_q}$ to get the coefficients of the from f_{4i-1} for $i > 0$. Similarly, we get the coefficients of the form f_{4i-2} and f_{4i-3} from recovering $\boldsymbol{\phi}''_{\theta_q}$ and $\boldsymbol{\phi}'''_{\theta_q}$, respectively. This allows to recover the complete solution of $\boldsymbol{f}$.

However, the substantial reduction in evaluations for a random $\mathcal{S}_t$ limits the number of weak keys that can be attacked. Since identifying the presence of such symmetries is straightforward, it is prudent to verify their existence before initiating the lattice-reduction phase of the attack.

7 Conclusion

In this paper, we show that certain variants of the Vandermonde Knapsack problem may be easier to solve than the standard Knapsack problem. To this end, we introduce new symmetry-based attacks that exploit the underlying Vandermonde structure. We further demonstrate that, for some cases–such as power-of-two cyclotomic rings–these symmetries are not unique.

Acknowledgments

We thank Martin Albrecht and Russell Lai for helpful discussions on the Vanishing SIS problem. We also thank the anonymous reviewers of CRYPTO 2026 and ASIACRYPT 2026 for their constructive comments and valuable suggestions.

References

1. Albrecht, M.R., Göpfert, F., Virdia, F., Wunderer, T.: Revisiting the expected cost of solving uSVP and applications to LWE. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017, Part I. LNCS, vol. 10624, pp. 297–322. Springer, Cham (Dec 2017). https://doi.org/10.1007/978-3-319-70694-8_11
2. Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: Post-quantum key exchange - A new hope. In: Holz, T., Savage, S. (eds.) USENIX Security 2016. pp. 327–343. USENIX Association (Aug 2016), https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/alkim
3. Boudgoust, K., Gachon, E., Pellet-Mary, A.: Some easy instances of ideal-SVP and implications on the partial vandermonde knapsack problem. In: Dodis, Y., Shrimpton, T. (eds.) CRYPTO 2022, Part II. LNCS, vol. 13508, pp. 480–509. Springer, Cham (Aug 2022). https://doi.org/10.1007/978-3-031-15979-4_17

4. Boudgoust, K., Sakzad, A., Steinfeld, R.: Vandermonde meets regev: public key encryption schemes based on partial vandermonde problems. DCC **90**(8), 1899–1936 (2022). https://doi.org/10.1007/s10623-022-01083-7
5. Chang, T., Hsu, J., Hsu, H., Tso, R., Mambo, M.: COMPASS: A compact pass-lineage accumulator with succinct proofs. In: Dutta, R., Feo, L.D., Gangopadhyay, S. (eds.) Progress in Cryptology - INDOCRYPT 2025 - 26th International Conference on Cryptology in India, Bhubaneshwar, India, December 14-17, 2025, Proceedings. Lecture Notes in Computer Science, vol. 16372, pp. 145–169. Springer (2025). https://doi.org/10.1007/978-3-032-13301-4_7, https://doi.org/10.1007/978-3-032-13301-4_7
6. Cini, V., Lai, R.W.F., Malavolta, G.: Lattice-based succinct arguments from vanishing polynomials - (extended abstract). In: Handschuh, H., Lysyanskaya, A. (eds.) CRYPTO 2023, Part II. LNCS, vol. 14082, pp. 72–105. Springer, Cham (Aug 2023). https://doi.org/10.1007/978-3-031-38545-2_3
7. Das, D., Joux, A.: Key recovery attack on the partial Vandermonde knapsack problem. In: Joye, M., Leander, G. (eds.) EUROCRYPT 2024, Part VII. LNCS, vol. 14657, pp. 205–225. Springer, Cham (May 2024). https://doi.org/10.1007/978-3-031-58754-2_8
8. Dubois, A., Klooß, M., Lai, R.W.F., Woo, I.K.Y.: Lattice-based proof-friendly signatures from vanishing short integer solutions. In: Jager, T., Pan, J. (eds.) PKC 2025, Part I. LNCS, vol. 15674, pp. 452–486. Springer, Cham (May 2025). https://doi.org/10.1007/978-3-031-91820-9_15
9. Ducas, L., Espitau, T., Postlethwaite, E.W.: Finding short integer solutions when the modulus is small. In: Handschuh, H., Lysyanskaya, A. (eds.) CRYPTO 2023, Part III. LNCS, vol. 14083, pp. 150–176. Springer, Cham (Aug 2023). https://doi.org/10.1007/978-3-031-38548-3_6
10. Espitau, T., Fouque, P.A., Gérard, F., Rossi, M., Takahashi, A., Tibouchi, M., Wallet, A., Yu, Y.: Mitaka: A simpler, parallelizable, maskable variant of falcon. In: Dunkelman, O., Dziembowski, S. (eds.) EUROCRYPT 2022, Part III. LNCS, vol. 13277, pp. 222–253. Springer, Cham (May / Jun 2022). https://doi.org/10.1007/978-3-031-07082-2_9
11. Fisch, B., Liu, Z., Vesely, P.: Orbweaver: Succinct linear functional commitments from lattices. In: Handschuh, H., Lysyanskaya, A. (eds.) CRYPTO 2023, Part II. LNCS, vol. 14082, pp. 106–131. Springer, Cham (Aug 2023). https://doi.org/10.1007/978-3-031-38545-2_4
12. Gama, N., Nguyen, P.Q.: Predicting lattice reduction. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 31–51. Springer, Berlin, Heidelberg (Apr 2008). https://doi.org/10.1007/978-3-540-78967-3_3
13. Hoffstein, J., Pipher, J., Schanck, J.M., Silverman, J.H., Whyte, W.: Practical signatures from the partial Fourier recovery problem. In: Boureanu, I., Owesarski, P., Vaudenay, S. (eds.) ACNS 2014. LNCS, vol. 8479, pp. 476–493. Springer, Cham (Jun 2014). https://doi.org/10.1007/978-3-319-07536-5_28
14. Hoffstein, J., Silverman, J.H.: PASS-encrypt: a public key cryptosystem based on partial evaluation of polynomials. DCC **77**(2-3), 541–552 (2015). https://doi.org/10.1007/s10623-015-0089-z
15. Jyrkinen, K., Lai, R.W.F.: Vanishing short integer solution, revisited - reductions, trapdoors, homomorphic signatures for low-degree polynomials. In: Jager, T., Pan, J. (eds.) PKC 2025, Part II. LNCS, vol. 15675, pp. 273–300. Springer, Cham (May 2025). https://doi.org/10.1007/978-3-031-91823-0_9
16. Lenstra, A.K., Lenstra, H.W.j., Lovász, L.: Factoring polynomials with rational coefficients. Math. Ann. **261**, 515–534 (1982). https://doi.org/10.1007/BF01457454

17. Lu, X., Zhang, Z., Au, M.H.: Practical signatures from the partial Fourier recovery problem revisited: A provably-secure and Gaussian-distributed construction. In: Susilo, W., Yang, G. (eds.) ACISP 18. LNCS, vol. 10946, pp. 813–820. Springer, Cham (Jul 2018). https://doi.org/10.1007/978-3-319-93638-3_50
18. Schnorr, C.P., Euchner, M.: Lattice basis reduction: Improved practical algorithms and solving subset sum problems. In: Fundamentals of computation theory. 8th international conference, FCT '91, Gosen, Germany, September 9–13, 1991. Proceedings, pp. 68–85. Berlin etc.: Springer-Verlag (1991). https://doi.org/10.1007/3-540-54458-5_51
19. development team, T.F.: fplll, a lattice reduction library, Version: 5.5.0 (2023), https://github.com/fplll/fplll, available at https://github.com/fplll/fplll

Appendix

A Expectation of $t_1^{(\sigma)}$ in the PV Knapsack problem

The random variable $t_1^{(\sigma)}$ counts the number of pairs of the form $(\omega, \rho(\omega))$ contained in the set Ω_t. For each of the $n/2$ many pair, define the indicator random variable

$$X_j = \begin{cases} 1 & \text{if both elements of } j^{\text{th}} \text{ pair are selected in } \Omega_t, \\ 0 & \text{otherwise.} \end{cases}$$

Then

$$t_1^{(\sigma)} = \sum_{j=1}^{n/2} X_j.$$

By linearity of expectation,

$$\mathbb{E}[t_1^{(\sigma)}] = \sum_{j=1}^{n/2} \mathbb{E}[X_j].$$

Now, the probability of a given pair to be selected in Ω_t is given by

$$\Pr(X_j = 1) = \frac{\binom{n-2}{t-2}}{\binom{n}{t}} = \frac{t(t-1)}{n(n-1)}.$$

Therefore,

$$\mathbb{E}[t_1^{(\sigma)}] = \frac{n}{2} \cdot \frac{t(t-1)}{n(n-1)} = \frac{t(t-1)}{2(n-1)}.$$

As a result, when $t = \frac{n}{2}$,

$$\mathbb{E}[t_1^{(\sigma)}] \approx t/4$$

for large values of n.

B Expectation of t_1 in the integer vSIS commitment problem

Let $U = \mathbb{Z}_q^* \setminus \{-1, 1\}$ be the set from which we sample the t elements in $\mathcal{S}_t$ for the integer vSIS commitment problem. The probability of a given pair $(\upsilon, -\upsilon)$ to be selected in $\mathcal{S}_t$ is given by $\frac{\binom{|U|-2}{t-2}}{\binom{|U|}{t}}$.

Now we can express t_1 as a sum of indicator random variables as

$$t_1 = \sum_{i=1}^{|U|/2} X_i,$$

where $X_i = 1$, if the i-th pair is fully contained in $\mathcal{S}_t$, and $X_i = 0$ otherwise.

Therefore,

$$\mathbb{E}[t_1] = \mathbb{E}[\sum_{i=1}^{|U|/2} X_i] = \sum_{i=1}^{|U|/2} \mathbb{E}[X_i] = \frac{|U|}{2} \cdot \frac{\binom{|U|-2}{t-2}}{\binom{|U|}{t}} = \frac{|U|}{2} \cdot \frac{t(t-1)}{|U|(|U|-1)} = \frac{t(t-1)}{2(q-4)}.$$

Now taking $q = \alpha n = 2\alpha t$, we get

$$\mathbb{E}[t_1] = \frac{t}{4} \cdot \frac{1 - \frac{1}{t}}{\alpha - \frac{2}{t}} \approx \frac{t}{4\alpha}$$

for large values of n.