Vanishing Short Integer Solution, Revisited
Reductions, Trapdoors, Homomorphic Signatures for Low-Degree Polynomials

Kalle Jyrkinen and Russell W. F. Lai

Aalto University, Espoo, Finland

Abstract. The vanishing short integer solution (vSIS) assumption [Cini-Lai-Malavolta, Crypto'23], at its simplest form, asserts the hardness of finding a polynomial with short coefficients which vanishes at a given random point. While vSIS has proven to be useful in applications such as succinct arguments, not much is known about its theoretical hardness. Furthermore, without the ability to generate a hard instance together with a trapdoor, the applicability of vSIS is significantly limited.

We revisit the vSIS assumption focusing on the univariate single-point constant-degree setting, which can be seen as a generalisation of the (search) NTRU problem. In such a setting, we show that the vSIS problem is as hard as finding the shortest vector in certain ideal lattices. We also show how to generate a random vSIS instance together with a trapdoor, under the (decision) NTRU assumption. Interestingly, a vSIS trapdoor allows to sample polynomials of short coefficients which evaluate to any given value at the public point. By exploiting the multiplicativity of the polynomial ring, we use vSIS trapdoors to build a new homomorphic signature scheme for low-degree polynomials.

Keywords: vanishing SIS · NTRU · reduction · lattice trapdoors · homomorphic signatures

1 Introduction

The short integer solution (SIS) problem over a ring $\mathcal{R}$ asks to find a non-zero short vector $\mathbf{u} \in \mathcal{R}^m$ such that $\mathbf{A}\mathbf{u} = \mathbf{0} \bmod q$ for a given uniformly random wide matrix $\mathbf{A} \in \mathcal{R}_q^{n \times m}$. The average-case hardness of SIS immediately implies simple constructions of collision-resistant hash functions and computationally binding commitments, both linearly homomorphic. Ajtai [Ajt96] showed that average-case SIS for $\mathcal{R} = \mathbb{Z}$ is as hard as several natural worst-case problems over Euclidean lattices. Subsequently, Gentry, Peikert and Vaikuntanathan (GPV) [GPV08] and Micciancio and Peikert [MP12] showed how to sample a statistically random instance $\mathbf{A}$ of SIS over $\mathbb{Z}$ together with a trapdoor, given which it is efficient to sample a random short preimage $\mathbf{u}$ satisfying $\mathbf{A}\mathbf{u} = \mathbf{v} \bmod q$ for any given target vector $\mathbf{v}$. These lattice trapdoors enable lattice-based constructions of a wide variety of cryptographic primitives, ranging from hash-and-sign signatures [GPV08] to attribute-based encryption schemes [BGG$^+$14]. Motivated by efficiency, the above results are later generalised to other choices of $\mathcal{R}$ (see e.g. [LPR10,LPR13,LS15]), notably the rings of integers of cyclotomic fields, and are nowadays used as standard tools in the constructions increasingly complex primitives.

Vanishing SIS. In this work, we focus on a generalisation of SIS called vanishing SIS (vSIS), proposed by Cini, Lai and Malavolta [CLM23]. The vSIS problem is in fact a family of problems parametrised by numerous parameters. A basic variant asks to find a non-zero polynomial p with at most a prescribed degree and with short coefficients such that $p(\mathbf{a}_i) = 0 \bmod q$ at all n given uniformly random points $\mathbf{a}_1, \ldots, \mathbf{a}_n$. The SIS problem can be seen as a special case by restricting solutions p to lie in the class of m-variate linear polynomials.

As demonstrated in [CLM23,KLNO24], the vSIS assumption can be used to construct commitments which are both linearly and (somewhat) multiplicatively homomorphic, which can in turn be used to construct succinct arguments with efficient provers and verifiers. Furthermore, due to the polynomial structure, the vSIS commitment scheme can be instantiated to have very short commitment keys. For instance, relying on the single-point (i.e. $n = 1$) univariate vSIS assumption, the vSIS commitment allows to commit to a short vector $\mathbf{u} \in \mathcal{R}^m$ against a commitment key consisting of a single element $a \in \mathcal{R}_q$.

However, for the vSIS assumption to serve as a drop-in replacement of the SIS assumption, at least two challenges must be tackled:

1. To better understand the theoretical hardness of the vSIS problem, preferrably by connecting to worst-case lattice problems.
2. To construct algorithms which sample hard vSIS instances together with their trapdoors, given which short polynomials can be sampled that evaluate at the public point to any desired value.

1.1 Our Contributions

In this work, we make initial progress towards tackling both of the aforementioned challenges, focusing on the univariate single-point constant-degree setting. We refer to the vSIS problem in this setting as "simple vSIS". Concretely, over the ring of integers $\mathcal{R}$ of some number field $\mathcal{K}$, a simple vSIS problem asks the following:

> Given a single random point $v \leftarrow\$ \mathcal{R}_q$, find a non-zero polynomial p over $\mathcal{R}$ with degree at most d with coefficients of norm at most β satisfying $p(v) = 0 \bmod q$.

Note that, for $d = 1$, we essentially recover the search NTRU problem, i.e. finding short $f, g \in \mathcal{R}$ such that $fv + g = 0 \bmod q$, with a different instance distribution. Surrounding this problem, our contributions are summarised as follows.

Reduction from IdSVP to vSIS. We show that solving the simple vSIS problem is as hard as finding approximate shortest vectors in ideal lattices of the form

$$I = \langle z^d \rangle \cap \mathcal{R},$$

where $z \in \mathcal{K}$, under certain parameter conditions. Our reduction is a generalisation of a similar reduction by Pellet-Mary and Stehlé [PS21], who showed that the search NTRU problem is as hard as finding short vectors in $\langle z \rangle \cap \mathcal{R}$. We remark that our reduction is worst-case to worst-case and average-case to average-case, where the latter is for certain not necessarily uniform distributions.

vSIS Trapdoors. Assuming the hardness of decision NTRU, i.e. that $v = f/g \bmod q$ is pseudorandom for f, g sampled from a discrete Gaussian distribution over $\mathcal{R}$, we show how to sample a pseudorandom simple vSIS instance v together with a trapdoor. Using the trapdoor with an existing preimage sampling algorithm (e.g. [GPV08]), on input any target point $t \in \mathcal{R}_q$, we can sample a short degree-$d_{\mathcal{K}}$ polynomial p satisfying $p(v) = t \bmod q$. Similar to the above reduction, our vSIS construction is a generalisation of that of NTRU trapdoors [HHP$^+$03,DLP14].

Homomorphic Signatures for Low-Degree Polynomials. Compared to SIS trapdoors [GPV08,MP12], vSIS trapdoors exhibit additional multiplicative properties which could be useful for additional applications. More precisely, due to the ring structure of polynomials over $\mathcal{R}_q$, if p_1, p_2 are polynomials with short coefficients, say sampled by a vSIS trapdoor, such that $p_1(v) = t_1 \bmod q$ and $p_2(v) = t_2 \bmod q$, then $p_1 \cdot p_2$ is still a short-ish polynomial such that $(p_1 \cdot p_2)(v) = t_1 \cdot t_2 \bmod q$.

Indeed, exploiting such a multiplicative property, we present a new construction of homomorphic signatures which allows to evaluate low-degree polynomials with short coefficients over signed data. Although being less expressive that the leveled fully homomorphic signatures of [GVW15], our scheme only uses generic ring arithmetic, without needing to perform non-ring-arithmetic operations such as binary decomposition. Furthermore, our construction admits a natural generalisation to the multi-key setting.

1.2 Limitations and Open Problems

We point out several limitations of our results in the hope that they will inspire improvements by future works.

Constant-degree. Perhaps the most significant limitation is that, for both our reduction from ideal SVP to vSIS and our construction of vSIS trapdoors, the degree d of the vSIS instance is limited to a constant. In particular, for the reduction, it means that the security of vSIS commitments [CLM23], which is meant for committing to poly(λ)-dimensional messages, is not covered. The root cause of this limitation is that, often, we need to argue that the d-th power of a short element is still short (relative to the modulus q). Thus, to avoid picking a too large modulus, we limit the degree d to constants. To get around this limitation, it seems that new fundamental tools for arguing about polynomials with short coefficients is needed, other than just treating them as linear functions with short coefficients over the monomials.

Single-point. All results in this work concern about vSIS in the single-point setting, where the number of points in vSIS corresponds to the module rank of (module-)SIS [LS15]. In other words, elements in the corresponding "vSIS module" are required to satisfy only one linear equation over $\mathcal{R}_q$. In applications [CLM23,KLNO24], it is useful to be able to pick the number of points n flexibly, which allows fine-tuning the vSIS instance to have a desired hardness level (under the heuristic that vSIS is as hard as SIS). Unfortunately, our current techniques of constructing vSIS trapdoors do not seem to allow finding short module elements satisfying not one but several linear equations at once. We note that a similar problem is open even in the simpler case of (module-)NTRU trapdoors [HHP$^+$03,DLP14,CKKS19,CPS$^+$20] as well.

Univariate. In this work, we focus primarily on univariate vSIS instances. As discussed above, the currently predominant use (so far) of vSIS is to commit to a large quantity of data. Using univariate vSIS, it would mean that the degree d of the vSIS instance needs to be as large as the dimensionality of the data to be committed. Since the number of monomials of a multivariate polynomial grows exponentially in the degree, switching to multivariate vSIS would allow considering a much lower degree d, which is intuitively more secure. Unfortunately, the presence of cross terms makes it difficult to adapt our trapdoor construction to multivariate vSIS.

Class of ideal lattices and moduli. In our ideal SVP to vSIS reduction, we are forced to focus on very specific ideal lattices of the form

$$I = \langle z^d \rangle \cap \mathcal{R}$$

and the vSIS instance has modulus q^d. The restriction on the class of ideal lattices is concerning, since ideals of this form becomes exponentially more scarce as d increases. Furthermore, the restriction on the modulus in particular means that the reduction does not cover vSIS instances with prime or product of primes moduli, which are usually the preferred choices. Lifting these restrictions seems to require a new reduction strategy different from that of [PS21] which we generalise.

Worst-case vs. average-case. Our reduction is worst-case to worst-case and average-case to average-case, whereas a worst-case to average-case reduction akin to that of [Ajt96] from SIVP to SIS is more preferable. We note that this is a minor restriction, given the conditional worst-case to average-case reduction sketched in [CLM23] under the decision NTRU or LWE assumption.

2 Preliminaries

Let $\lambda \in \mathbb{N}$ denote the security parameter. For a finite set S or a distribution χ, we write $x \leftarrow\$ S$ and $x \leftarrow\$ \chi$ for sampling from the uniform distribution over S and the distribution χ respectively. We use $:=$ to denote deterministic assignment and $y \leftarrow \mathcal{A}(x)$ to denote the execution of a (potentially probabilistic) algorithm $\mathcal{A}$. Matrices and vectors are represented by bold upper- and lower-case letters respectively, e.g. $\mathbf{M}$ and $\mathbf{v}$.

2.1 Linear Algebra

We recall the notion of Gram-Schmidt orthogonalisation (GSO).

Definition 1 (Orthogonal projection). *Let K be a field and V be an inner product space over K. For a subspace $W \subseteq V$, we denote the orthogonal projection of $\mathbf{v} \in V$ onto W as $\mathrm{proj}_W(\mathbf{v})$. If $\{\mathbf{w}_i\}_{i \in I}$ for some index set I is an orthogonal basis of W, then*

$$\mathrm{proj}_W(\mathbf{v}) = \sum_{i \in I} \frac{\langle \mathbf{v}, \mathbf{w}_i \rangle}{\langle \mathbf{w}_i, \mathbf{w}_i \rangle} \mathbf{w}_i.$$

An orthogonal projection onto a certain subspace is always unique and does not depend on the choice of the basis. Using $\mathrm{proj}_W(\mathbf{v})$ as a subroutine, we can define the Gram-Schmidt orthogonalisation (GSO) procedures.

Definition 2 (Gram-Schmidt orthogonalisation). *Let K, V be as in Definition 1, $\mathbf{b}_1, \ldots, \mathbf{b}_m$ be linearly independent vectors in V and $\mathbf{B} = \begin{bmatrix}\mathbf{b}_1 \ldots \mathbf{b}_m\end{bmatrix}$. We call $\tilde{\mathbf{B}} = \begin{bmatrix}\tilde{\mathbf{b}}_1 \ldots \tilde{\mathbf{b}}_m\end{bmatrix}$ the Gram-Schmidt orthogonalisation of $\mathbf{B}$, and it is defined as*

$$\begin{cases} \tilde{\mathbf{b}}_1 = \mathbf{b}_1 \\ \tilde{\mathbf{b}}_i = \mathbf{b}_i - \mathrm{proj}_{\mathrm{span}_K(\{\mathbf{b}_1, \ldots, \mathbf{b}_{i-1}\})}(\mathbf{b}_i), \ i \geq 2. \end{cases}$$

The Gram-Schmidt norm (GS-norm) of the matrix $\mathbf{B}$ is denoted $\|\mathbf{B}\|_{\mathrm{GS}}$ and defined as the maximum of the ℓ_2 norm over the columns of $\tilde{\mathbf{B}}$, i.e.

$$\|\mathbf{B}\|_{\mathrm{GS}} = \max_i \|\tilde{\mathbf{b}}_i\|.$$

2.2 Lattice and Discrete Gaussian

For a full-rank matrix $\mathbf{B} \in \mathbb{R}^{n \times k}$ with $k \leq n$, the lattice spanned by $\mathbf{B}$ is $\Lambda = \mathbf{B} \cdot \mathbb{Z}^k$ and $\mathbf{B}$ is said to be a basis of Λ. Its dual lattice is $\Lambda^\vee = \{\mathbf{x} \in \mathrm{span}_{\mathbb{R}}(\Lambda) : \langle \Lambda, \mathbf{x} \rangle \subseteq \mathbb{Z}\}$.

For $s > 0$, denote $\rho_s(\mathbf{x}) \coloneqq \exp(-\pi \|\mathbf{x}\|^2 / s^2)$ for any $\mathbf{X} \in \mathbb{R}^n$ the Gaussian function with parameter s. For a lattice $\Lambda \subseteq \mathbb{R}^n$, the discrete Gaussian distribution over a coset A of Λ with parameter s is defined as $\mathcal{D}_{A,s}(\mathbf{x}) \coloneqq \rho_s(\mathbf{x}) / \rho_s(A)$ for any $\mathbf{x} \in A$, where $\rho_s(A) \coloneqq \sum_{\mathbf{x}' \in A} \rho_s(\mathbf{x}')$.

Definition 3 (Smoothing Parameter [MR07]). *Let $\Lambda \subset \mathbb{R}^n$ be a lattice. For real $\epsilon > 0$, the smoothing parameter $\eta_\epsilon(\Lambda)$ is the smallest real $s > 0$ such that $\rho_{1/s}(\Lambda^\vee \setminus \{\mathbf{0}\}) \leq \epsilon$.*

Lemma 1 ([GPV08, eprint Lemma 3.1]). *Let $\Lambda \subset \mathbb{R}^n$ be a full-rank lattice and $\mathbf{B} \in \mathbb{R}^{n \times n}$ be a basis of Λ. For any real $\epsilon > 0$, we have*

$$\eta_\epsilon(\Lambda) \leq \|\mathbf{B}\|_{\mathrm{GS}} \cdot \sqrt{\log(2n \cdot (1 + 1/\epsilon))/\pi}.$$

In particular, $\eta_\epsilon(\Lambda) \leq \|\mathbf{B}\|_{\mathrm{GS}} \cdot \sqrt{n}$ for some ϵ negligible in n.

Lemma 2 ([MR07, Lemma 4.4]). *Let $\Lambda \subset \mathbb{R}^n$ be a lattice. For any $\mathbf{c} \in \mathbb{R}^n$, real $\epsilon \in (0, 1)$ and $s \geq \eta_\epsilon(\Lambda)$, $\Pr[\|\mathbf{x} - \mathbf{c}\| > s\sqrt{n} \mid \mathbf{x} \leftarrow\$ \mathcal{D}_{\Lambda,s,\mathbf{c}}] \leq \frac{1+\epsilon}{1-\epsilon} \cdot 2^{-n}$.*

Lemma 3 ([GPV08, eprint Corollary 2.8]). *Let $\Psi \subseteq \Lambda \subset \mathbb{R}^n$ be full-rank lattices. For any $\epsilon \in (0, 1/2)$, any real $s \geq \eta_\epsilon(\Psi)$, any any $\mathbf{c} \in \mathbb{R}^n$, the distribution $\mathcal{D}_{\Lambda,s,\mathbf{c}} \bmod \Psi$ and the uniform distribution over Λ/Ψ are within statistical distance at most 2ϵ.*

Lemma 4 ([GPV08, eprint Theorem 4.1]). *There exists a PPT algorithm that, given a basis $\mathbf{B} \in \mathbb{R}^{n \times n}$ of a lattice Λ, a parameter $s \geq \|\mathbf{B}\|_{\mathrm{GS}} \cdot \sqrt{n}$, and a centre $\mathbf{c} \in \mathbb{R}^n$, outputs a sample from a distribution that is statistically close in n to $\mathcal{D}_{\Lambda,s,\mathbf{c}}$.*

2.3 Algebraic Number Theory

We use $\mathcal{K}$ to denote a number field of degree $d_{\mathcal{K}}$ and $\mathcal{R}$ its ring of integers. We assume that a $\mathbb{Z}$-basis $\{b_1, \ldots, b_{d_{\mathcal{K}}}\} \subset \mathcal{R}$ is fixed. We write $\sigma = (\sigma_i)_{i=1}^{d_{\mathcal{K}}} : \mathcal{K} \to \mathbb{C}^{d_{\mathcal{K}}}$ for the canonical embedding of $\mathcal{K}$. The notation of the embeddings is extended naturally to vectors by concatenation.

We define the ℓ_p-norm of field vectors via the canonical embedding, i.e. for $\mathbf{x} \in \mathcal{K}^n$, $\|\mathbf{x}\|_p := \|\sigma(\mathbf{x})\|_p$. We omit the subscript p when it is taken as $p = 2$. We write $\beta_{\mathcal{K}} := \max_i \|b_i\|_\infty$ for the ℓ_∞-norm of the longest basis element. We extend the notation to polynomials over $\mathcal{K}$ by defining the norm of a polynomial to be the maximum norm of the coefficients; i.e. for polynomial $f = \sum_{i=0}^{d} f_i X^i \in \mathcal{K}[X]$, we denote $\|f\|_p := \max_{i=0}^{d} \|f_i\|_p$.

Proposition 1. *For any $x = \sum_{i=1}^{d_{\mathcal{K}}} b_i x_i \in \mathcal{K}$ where $x_i \in \mathbb{Q}$, denote $\{x\} := \sum_{i=1}^{d_{\mathcal{K}}} b_i \{x_i\}$ the fractional part of x w.r.t. the $\mathbb{Z}$-basis $\{b_1, \ldots, b_{d_{\mathcal{K}}}\}$. It holds that*

$$\|\{x\}\|_\infty \le d_{\mathcal{K}} \cdot \beta_{\mathcal{K}}/2.$$

Proof. Using elementary properties of σ we get $\|\sum_i b_i x_i\|_\infty \le \sum_i \|b_i x_i\|_\infty \le \sum_i \|b_i\|_\infty \|x_i\|_\infty \le d_{\mathcal{K}} \cdot \beta_{\mathcal{K}}/2$. □

Definition 4. *For $x \in \mathcal{K}$, we define the matrix $M(x) \in \mathbb{R}^{d_{\mathcal{K}} \times d_{\mathcal{K}}}$ as*

$$M(x) = \left[\sigma(xb_1) \cdots \sigma(xb_d)\right].$$

We extend the notation coefficient-wise to vectors in $\mathcal{K}^n$ and matrices in $\mathcal{K}^{m \times n}$.

The discriminant of $\mathcal{K}$ is denoted by $\Delta_{\mathcal{K}}$. For $z \in \mathcal{K}$, write $\langle z \rangle$ for the (fractional) ideal generated by z. The algebraic norm of a (fractional) ideal I of $\mathcal{K}$ is denoted as $\mathcal{N}(I)$.

Note that $\sigma(\mathcal{K})$ is isomorphic to $\mathbb{Q}^{d_{\mathcal{K}}} \subset \mathbb{R}^{d_{\mathcal{K}}}$. Therefore, any $\mathcal{R}$-module $\mathcal{M} \subseteq \mathcal{R}^m$ can be identified $\mathcal{M}$ as a $d_{\mathcal{K}} m$-dimensional lattice via $\sigma(\mathcal{M})$. In particular, for a matrix $\mathbf{A} \in \mathcal{R}_q^{n \times m}$ and a vector $\mathbf{v} \in \mathcal{R}_q^n$, we consider the following $\mathcal{R}$-module (lattice) $\Lambda_q^\perp(\mathbf{A})$ and its coset $\Lambda_q^{\mathbf{v}}(\mathbf{A})$:

$$\Lambda_q^\perp(\mathbf{A}) := \{\mathbf{u} \in \mathcal{R}^m : \mathbf{A}\mathbf{u} = \mathbf{0} \bmod q\},$$
$$\Lambda_q^{\mathbf{v}}(\mathbf{A}) := \{\mathbf{u} \in \mathcal{R}^m : \mathbf{A}\mathbf{u} = \mathbf{v} \bmod q\}.$$

The following is an immediate corollary of Lemma 4 for $\Lambda = \Lambda_q^\perp(\mathbf{A})$.

Corollary 1. *There exists a PPT algorithm* SampPre *that, given a basis $\mathbf{B} \in \mathcal{R}^{m \times m}$ of $\Lambda_q^\perp(\mathbf{A})$, a target $\mathbf{v} \in \mathcal{R}_q^n$ and a parameter $s \ge \|\mathbf{B}\|_{\mathrm{GS}} \cdot \sqrt{d_{\mathcal{K}} m}$, outputs a sample from a distribution that is statistically close in $d_{\mathcal{K}} m$ to $\mathcal{D}_{\Lambda_q^{\mathbf{v}}(\mathbf{A}), s}$.*

2.4 Computational Problems

Definition 5 (Hermite Shortest Vector Problem (HSVP)). *Let $n \in \mathbb{N}$ and $\mu \ge 1$. The* $\mathsf{HSVP}_{n,\mu}$ *problem asks to find, given a basis $\mathbf{B} \in \mathbb{R}^{n \times n}$ of $\mathcal{L} = \mathbf{B} \cdot \mathbb{Z}^n$, a vector $\mathbf{x} \in \mathcal{L} \setminus \{0\}$ satisfying*

$$\|\mathbf{x}\| \le \mu \cdot \sqrt{n} \cdot \det(\mathcal{L})^{\frac{1}{n}}.$$

The $\mathsf{id\text{-}HSVP}_{\mathcal{R},\mu}$ *problem is a restriction of the* $\mathsf{HSVP}_{n,\mu}$ *problem, where $n = d_{\mathcal{K}}$, in that the input lattice is restricted to be an ideal lattice in $\mathcal{R}$.*

Definition 6 (NTRU). *Let $\mathcal{R}, q, \chi$ be parametrised by λ, where $\mathcal{R}$ is a number ring, q is a positive integer, and χ is a distribution over $\mathcal{R}$. The* $\mathsf{dNTRU}_{\mathcal{R},q,\chi}$ *problem is to distinguish the uniform distribution over $\mathcal{R}_q$ from the following distribution:*

$$\left\{ h \in \mathcal{R}_q : \begin{array}{c} f, g \leftarrow\$\ \chi : \langle f, g \rangle = \mathcal{R} \wedge g \in \mathcal{R}_q^\times \\ h = f/g \bmod q \end{array} \right\}$$

The $\mathsf{dNTRU}_{\mathcal{R},q,\chi}$ *assumption states that no PPT algorithm solves the above problem with non-negligible probability in λ.*

We recall the univariate single-point variant of the vanishing SIS assumption introduced in [CLM23].

Definition 7 (Vanishing-SIS (vSIS, [CLM23])). *Let $\mathcal{R}, d, q, \beta$ be parametrised by λ, where d, q are positive integers and $\beta > 0$. The $\mathsf{vSIS}_{\mathcal{R},d,q,\beta}$ assumption states that, for any PPT adversary $\mathcal{A}$, it holds that*

$$\Pr\left[\begin{array}{c} p(v) = 0 \bmod q \\ \wedge\ 0 < \|p\| \leq \beta \end{array} \middle| \begin{array}{l} v \leftarrow\!\!\$\ \mathcal{R}_q^\times \\ p \leftarrow \mathcal{A}(v) \end{array}\right] \leq \mathrm{negl}(\lambda)$$

where in the above $p = \sum_{i=0}^{d} p_i \cdot X^i \in \mathcal{R}[X]$ is a univariate polynomial of degree at most d and $\|p\| := \max_{i=0}^{d} \|p_i\|$.

3 Reduction from Id-SVP to vSIS

The goal of this section is to generalise the results of [PS21], where the authors reduce search-NTRU from ideal-HSVP. As a result, we obtain a reduction from a special distribution of ideal-HSVP to vSIS, under a specific parameter regime.

Before going into the main theorem of the section, let us first state and prove a lemma that represents the core of the reduction. Roughly, the lemma states the following: Given an ideal of the form $I = \langle z^d \rangle \cap \mathcal{R}$, we can define a vSIS instance $v = \lfloor q/z \rceil \bmod q^d$ such that (i) a short solution p exists, and (ii) for any sufficiently short solution p' the leading coefficient must belong to the ideal I.

Lemma 5 (Transforming ideal-HSVP instance to vSIS instance). *Let $q, d \in \mathbb{N}$ such that $q \geq 2$. Also, let $I \subseteq \mathcal{R}$ be a non-zero ideal of form $I = \langle z^d \rangle \cap \mathcal{R}$ where $z \in \mathcal{K}$, and define $v = \lfloor q/z \rceil \bmod q^d$. Then, for every such v we have*

(i) there exists a non-zero polynomial $p \in \mathcal{R}[X]$ such that $\deg(p) \leq d$, $p(v) = 0 \bmod q^d$ and

$$\|p\| \leq d_{\mathcal{K}}^{d+\frac{1}{2}} \cdot \Delta_{\mathcal{K}}^{\frac{1}{2d_{\mathcal{K}}}} \cdot \beta_{\mathcal{K}}^{d} \cdot \mathcal{N}(I)^{\frac{1}{d_{\mathcal{K}}}},$$

and

(ii) for any non-zero polynomial $p' \in \mathcal{R}[X]$ satisfying $\deg(p') \leq d$, $p'(v) = 0 \bmod q^d$ and

$$\frac{q}{\|p'\|_\infty} > \Delta_{\mathcal{K}}^{\frac{1}{2d_{\mathcal{K}}}} \cdot \mathcal{N}(I)^{\frac{1}{d_{\mathcal{K}}}} \cdot \max\left\{1,\ \|z^{-1}\|_\infty^{d-1},\ \left(\frac{\beta_{\mathcal{K}} \cdot d_{\mathcal{K}}}{2}\right)^{d}\right\} \cdot (2^{d+1} - 2),$$

it holds that the leading coefficient of p' is in $I \setminus \{0\}$.

Proof. Towards (i), let α be the shortest non-zero element in I, measured in the infinity norm. Let

$$p(X) = \alpha \cdot \left(X + \left\{\frac{q}{z}\right\}\right)^{d}.$$

We show that p is a short degree-d polynomial in $\mathcal{R}[X]$ vanishing at v modulo q^d.

Observe that $\deg(p) = d$ by construction. To check the vanishing property, notice that $v = \frac{q}{z} - \left\{\frac{q}{z}\right\}$, and thus

$$p(v) = \alpha \cdot \left(\frac{q}{z}\right)^{d} = 0 \bmod q^d$$

where the second equality follows from our assumptions: since $\alpha \in \langle z^d \rangle$ there exist $r \in \mathcal{R}$ such that $\alpha = rz^d$ and thus $\alpha/z^d \in \mathcal{R}$.

Next, we show that the coefficients of p are in $\mathcal{R}$. Observe that they are given by the binomial expansion

$$p_i = \alpha \cdot \binom{d}{i} \cdot \left\{\frac{q}{z}\right\}^{d-i}, \quad i \in \{0, \ldots, d\}.$$

We have

$$\alpha \cdot \left\{\frac{q}{z}\right\}^{d-i} = \alpha \cdot \left(\frac{q}{z} - v\right)^{d-i} = \sum_{j=0}^{d-i} \alpha \cdot \binom{D-i}{j} \cdot \left(\frac{q}{z}\right)^{j} \cdot (-v)^{d-i-j}.$$

Since $\alpha \in \langle z^d \rangle \cap \mathcal{R}$, all of the summands are in $\mathcal{R}$ and hence $\alpha \left\{q/z\right\}^{d-i} \in \mathcal{R}$ for all $i \in \{0, \ldots, d\}$. This implies $p \in \mathcal{R}[X]$.

Lastly, let us compute an upper bound on the norm of the solution. Minkowski's bound implies $\|\alpha\|_\infty \leq \Delta_{\mathcal{K}}^{1/(2d_{\mathcal{K}})} \cdot \mathcal{N}(I)^{1/d_{\mathcal{K}}}$ and by equivalence of norms, $\|\alpha\| \leq \sqrt{d_{\mathcal{K}}} \cdot \Delta_{\mathcal{K}}^{1/(2d_{\mathcal{K}})} \cdot \mathcal{N}(I)^{1/d_{\mathcal{K}}}$. Also, by Proposition 1 we have $\| \left\{q/z\right\} \|_\infty \leq \beta_{\mathcal{K}} \cdot d_{\mathcal{K}}/2$. Thus,

$$\begin{aligned}
\|p\| &\leq \|\alpha\| \cdot \max_{i \in \{0,\ldots,d\}} \left\{ \binom{d}{i} \right\} \cdot \left\| \left\{\frac{q}{z}\right\} \right\|_\infty^{d} \\
&\leq \sqrt{d_{\mathcal{K}}} \cdot \Delta_{\mathcal{K}}^{\frac{1}{2d_{\mathcal{K}}}} \cdot \mathcal{N}(I)^{\frac{1}{d_{\mathcal{K}}}} \cdot \max_{i \in \{0,\ldots,d\}} \left\{ \binom{d}{i} \right\} \cdot \left(\frac{\beta_{\mathcal{K}} \cdot d_{\mathcal{K}}}{2}\right)^{d} \\
&\leq d_{\mathcal{K}}^{d+\frac{1}{2}} \cdot \Delta_{\mathcal{K}}^{\frac{1}{2d_{\mathcal{K}}}} \cdot \beta_{\mathcal{K}}^{d} \cdot \mathcal{N}(I)^{\frac{1}{d_{\mathcal{K}}}}
\end{aligned}$$

where we used the bound $\max_{i \in \{0,\ldots,d\}} \left\{ \binom{d}{i} \right\} \leq \sum_{i=0}^{d} \binom{d}{i} = (1+1)^d = 2^d$.

Next, we prove (ii). Without loss of generality we may assume that $\deg(p') = d$; if $\deg(p') = d^* < d$, simply consider the polynomial $X^{d-d^*} p'$ instead. We have

$$p'(v) = \sum_{i=0}^{d} p'_i \cdot \left(\frac{q}{z} - \left\{\frac{q}{z}\right\}\right)^{i} = \sum_{i=0}^{d} p'_i \cdot \sum_{j=0}^{i} \binom{i}{j} \cdot \left(\frac{q}{z}\right)^{j} \cdot \left(-\left\{\frac{q}{z}\right\}\right)^{i-j} = q^d r$$

for some $r \in \mathcal{R}$. Multiplying both sides by α/q^d and reordering yields

$$\frac{\alpha p'_d}{z^d} = \alpha r - \frac{\alpha}{q} \cdot \left(\begin{aligned} &p'_d \cdot \sum_{j=0}^{d-1} \binom{d}{j} \cdot \frac{q^{j-d+1}}{z^j} \cdot \left(-\left\{\frac{q}{z}\right\}\right)^{d-j} \\ &\quad + \sum_{i=0}^{d-1} p'_i \cdot \sum_{j=0}^{i} \binom{i}{j} \cdot \frac{q^{j-d+1}}{z^j} \cdot \left(-\left\{\frac{q}{z}\right\}\right)^{i-j} \end{aligned} \right).$$

Denote the second term on the right-hand side as θ; then, observe that $\|\theta\|_\infty$ is at most

$$\begin{aligned}
&\frac{\|\alpha\|_\infty}{q} \cdot \|p'\|_\infty \cdot \max\left\{1,\ \|z^{-1}\|_\infty^{d-1},\ \left\|\left\{\frac{q}{z}\right\}\right\|_\infty^{d}\right\} \cdot \left(\sum_{j=0}^{d-1}\binom{d}{j} + \sum_{i=0}^{d-1}\sum_{j=0}^{i}\binom{i}{j}\right) \\
= &\frac{\Delta_{\mathcal{K}}^{\frac{1}{2d_{\mathcal{K}}}} \cdot \mathcal{N}(I)^{\frac{1}{d_{\mathcal{K}}}}}{q} \cdot \|p'\|_\infty \cdot \max\left\{1,\ \|z^{-1}\|_\infty^{d-1},\ \left(\frac{\beta_{\mathcal{K}} \cdot d_{\mathcal{K}}}{2}\right)^{d}\right\} \cdot (2^{d+1} - 2)
\end{aligned}$$

where we used Minkowski's bound on $\|\alpha\|_\infty$, as well as properties of geometric sums to get $\sum_{j=0}^{d-1} \binom{d}{j} + \sum_{i=0}^{d-1} \sum_{j=0}^{i} \binom{i}{j} = 2^d - 1 + \sum_{i=0}^{d-1} 2^i = 2^d - 1 + \frac{1-2^d}{1-2} = 2^{d+1} - 2$. Therefore, by the assumption on $\|p'\|_\infty$, we have that $\|\theta\|_\infty < 1$. Because $\alpha \in \langle z^d \rangle$, $\alpha p'_d/z^d \in \mathcal{R}$. We also have $\alpha r \in \mathcal{R}$, and $\mathcal{R}$ being an additive group hence implies $\theta \in \mathcal{R}$. Since the infinity norm of any non-zero element in $\mathcal{R}$ is greater than or equal to 1, we conclude that $\theta = 0$.

As a result, $\alpha p'_d/z^d = \alpha r$. Dividing both sides by α we get that $\frac{p'_d}{z^d} = r \in \mathcal{R}$ and therefore $p'_d \in \langle z^d \rangle$, concluding the proof. □

Now we are ready to prove the main theorem.

Theorem 1. *Let $d, N_0, N_1, q, \beta, z_0 \in \mathbb{N}$ such that $N_0 \le N_1$,*

$$\beta \ge d_{\mathcal{K}}^{d+\frac{1}{2}} \cdot \mathit{\Delta}_{\mathcal{K}}^{\frac{1}{2d_{\mathcal{K}}}} \cdot \beta_{\mathcal{K}}^{d} \cdot N_1^{\frac{1}{d_{\mathcal{K}}}}$$

and

$$q > \beta \cdot \mathit{\Delta}_{\mathcal{K}}^{\frac{1}{2d_{\mathcal{K}}}} \cdot N_1^{\frac{1}{d_{\mathcal{K}}}} \cdot \max\left\{1,\ z_0^{d-1},\ \left(\frac{\beta_{\mathcal{K}} \cdot d_{\mathcal{K}}}{2}\right)^{d}\right\} \cdot (2^{d+1} - 2).$$

Also, define

$$\mu = \frac{\beta}{\sqrt{d_{\mathcal{K}}} \cdot \mathit{\Delta}_{\mathcal{K}}^{1/(2d_{\mathcal{K}})} \cdot N_0^{1/d_{\mathcal{K}}}}.$$

There is a PPT (with respect to $\mathrm{size}(z)$*,* $\log q$ *and* d*) reduction from worst-case* $\mathsf{id\text{-}HSVP}_{\mathcal{R},\mu}$ *to worst-case* $\mathsf{vSIS}_{\mathcal{R},d,q^d,\beta}$ *for ideals* $I \subseteq \mathcal{R}$ *that*

- *satisfy* $\mathcal{N}(I) \in [N_0, N_1]$*,*
- *are of form* $I = \langle z^d \rangle \cap \mathcal{R}$ *where* $z \in \mathcal{K}$ *satisfying* $\|z^{-1}\|_\infty \le z_0$ *and*
- *are assumed to be represented by the element* z*.*

Moreover, let $\mathcal{D}^{\mathsf{id\text{-}HSVP}}$ *be a distribution over ideals satisfying the above conditions. Then, there exists a distribution* $\mathcal{D}^{\mathsf{vSIS}}$ *over* $\mathsf{vSIS}_{\mathcal{R},d,q^d,\beta}$ *instances and a PPT (w.r.t.* $\mathrm{size}(z)$*,* $\log q$ *and* d*) reduction from average-case* $\mathsf{id\text{-}HSVP}_{\mathcal{R},\mu}$ *(for ideals sampled from* $\mathcal{D}^{\mathsf{id\text{-}HSVP}}$*) to average-case* $\mathsf{vSIS}_{\mathcal{R},d,q^d,\beta}$ *(for instances sampled from* $\mathcal{D}^{\mathsf{vSIS}}$ *over* $\mathsf{vSIS}_{\mathcal{R},d,q^d,\beta}$*).*

Proof. Let $\mathcal{A}$ be a PPT worst-case $\mathsf{vSIS}_{\mathcal{R},d,q^d,\beta}$ oracle. Define the reduction $\mathsf{id\text{-}HSVP\text{-}to\text{-}vSIS}^{\mathcal{A}}_{\mathcal{R},d,q^d,\beta}$ that takes as input an ideal $I = \langle z^d \rangle \cap \mathcal{R}$ satisfying $\mathcal{N}(I) \in [N_0, N_1]$.

```
$\mathsf{id\text{-}HSVP\text{-}to\text{-}vSIS}^{\mathcal{A}}_{\mathcal{R},d,q^d,\beta}(I)$
$v := \left\lfloor \frac{q}{z} \right\rceil \bmod q^d$
$p \leftarrow \mathcal{A}(v)$
Let $p^*$ be the leading coefficient of the polynomial $p$
**return** $p^*$
```

Due to the lower bound on β, v is a valid $\mathsf{vSIS}_{\mathcal{R},d,q^d,\beta}$ instance by the first claim of Lemma 5. Therefore p is a non-zero polynomial in $\mathcal{R}[X]$ satisfying $\deg(p) \le d$, $p(v) = 0 \bmod q^d$ and $\|p\| \le \beta$. Since $\|p\|_\infty \le \|p\|$ and thanks to q being bounded from below, the second claim of Lemma 5 implies that $p^* \in I \setminus \{0\}$. Observe that

$$\|p^*\| \le \beta \le \mu \cdot \sqrt{d_{\mathcal{K}}} \cdot \mathit{\Delta}_{\mathcal{K}}^{\frac{1}{2d_{\mathcal{K}}}} \cdot N_0^{\frac{1}{d_{\mathcal{K}}}} \le \mu \cdot \sqrt{d_{\mathcal{K}}} \cdot \mathit{\Delta}_{\mathcal{K}}^{\frac{1}{2d_{\mathcal{K}}}} \cdot \mathcal{N}(I)^{\frac{1}{d_{\mathcal{K}}}}$$

and hence p^* is a solution to the $\mathsf{id\text{-}HSVP}_{\mathcal{R},\mu}$ instance I.

To conclude the proof of the first part, it remains to bound the running time of the reduction. Notice that we query the oracle once and the rest of the operations consist of division, rounding, taking residue and finding the leading coefficient. All of these can be done in time $\mathrm{poly}(\mathrm{size}(z), \log q, d)$.

To prove the second part of the theorem, consider the same reduction as before but let $\mathcal{A}$ now be a PPT average-case oracle with non-negligible success probability. By a similar argument that was used when proving the first part, the reduction does not decrease the success probability and the running time is still polynomial. □

The next corollary further emphasises the approximation factor of the reduction and follows from combining the definition of μ and the lower bound on β in the previous theorem.

Corollary 2. *Use the same notations as in Theorem 1. Also, suppose that there exists a set of ideals of $\mathcal{R}$ satisfying the conditions of the theorem and denote that set by S.*

There exists a reduction from worst-case (respectively, average-case) $\mathsf{id\text{-}HSVP}_{\mathcal{R},\mu}$ *for ideals in S to worst-case (resp. average-case)* $\mathsf{vSIS}_{\mathcal{R},d,q^d,\beta}$*, with*

$$\mu = O\left(d_{\mathcal{K}}^{d} \cdot \beta_{\mathcal{K}}^{d} \cdot (N_1/N_0)^{\frac{1}{d_{\mathcal{K}}}}\right).$$

Assuming that $(N_1/N_0)^{1/d_{\mathcal{K}}} = \mathrm{poly}(d_{\mathcal{K}}, \beta_{\mathcal{K}})$, the approximation factor is polynomial in $d_{\mathcal{K}}$ and $\beta_{\mathcal{K}}$ (for a constant d).

Remark 1. Assuming $(N_1/N_0)^{1/d_{\mathcal{K}}}$ to be polynomial in $d_{\mathcal{K}}$ is not restrictive by itself. Note that we can choose a lower bound N_0 such that there exists sufficiently many ideals $I \subseteq \mathcal{R}$ with $\mathcal{N}(I) \leq N_0$. Then, we claim that if we set $N_1 = 2^{d_{\mathcal{K}}} N_0$ (such that $(N_1/N_0)^{1/d_{\mathcal{K}}} = 2$), there are at least equally many ideals having their norm within the interval $[N_0, N_1]$. This is because we can scale ideals with small norm up as discussed in Section 4.1 of [PS21]. However, when combined with the other two conditions for the ideals, the situation is not as clear. As such, we are currently unable to make an informed comment on whether the restrictions are reasonable or not.

4 Trapdoors for vSIS

In this section, we construct trapdoors for pseudorandom univariate single-point constant-degree vSIS instances under the decision NTRU assumption. Viewing NTRU as univariate single-point degree-1 vSIS, our construction strategy is to extend the NTRU trapdoor generation algorithm to the setting of univariate single-point constant-degree vSIS. In a nutshell, we construct a trapdoor generation algorithm $\mathsf{TrapGen}$ which samples a pseudorandom point $v \in \mathcal{R}_q$ together with a trapdoor td. Then, using the trapdoor preimage sampling algorithm $\mathsf{SampPre}$ from [GPV08] (recalled in Corollary 1), we can sample a polynomial f of a prescribed degree with short coefficients such that $f(v) = y \bmod q$ for any given image $y \in \mathcal{R}_q$.

4.1 Trapdoor for Univariate Single-Point Constant-Degree vSIS

We extend NTRU trapdoors [HHP+03,DLP14] to the setting of univariate single-point constant-degree vSIS, which we simply refer to as vSIS in this subsection.

vSIS Lattices. We begin by generalising the notion of NTRU lattices to that of vSIS lattices; these are free $\mathcal{R}$-modules that correspond to lattices via the canonical embedding.

Definition 8 (vSIS module). *Let $q, d \in \mathbb{N}$ and $v \in \mathcal{R}$. We denote the corresponding vSIS module as $\mathcal{M}_{v,d,q}$ and define it as*

$$\mathcal{M}_{v,d,q} = \left\{ (p_d, \ldots, p_0) \in \mathcal{R}^{d+1} \;\middle|\; \sum_{i=0}^{d} p_i v^i = 0 \bmod q \right\}.$$

Lemma 6 associates any vSIS module with a structured basis matrix.

Lemma 6. *$\mathcal{M}_{v,d,q}$ is a free $\mathcal{R}$-module of rank $d+1$ generated by the matrix*

$$\mathbf{B} = \begin{bmatrix} 1 & & & \\ -v & \ddots & & \\ & \ddots & 1 & \\ & & -v & q \end{bmatrix} \in \mathcal{R}^{(d+1)\times(d+1)}.$$

Proof. Observe that all columns of $\mathbf{B}$ are in $\mathcal{M}_{v,d,q}$. Thus, it remains to show that any $(p_d, \ldots, p_0) \in \mathcal{M}_{v,d,q}$ can be expressed as $\mathbf{Bx}$ for some $\mathbf{x} \in \mathcal{R}^{d+1}$. Indeed, the components of such $\mathbf{x}$ can be given as

$$x_i = \begin{cases} p_d, & i = 1 \\ x_{i-1}v + p_{d+1-i}, & i \in \{2, \ldots, d\} \\ r, & i = d+1. \end{cases} \qquad \square$$

The trapdoor basis. Our trapdoor construction is driven by a simple observation: if $v = f/g \bmod q$ as in NTRU, v is computationally indistinguishable from uniform under the dNTRU-assumption. Moreover, $(g, -f, 0, \ldots, 0), \ldots, (0, \ldots, 0, g, -f) \in \mathcal{R}^{d+1}$ is a set of d $\mathcal{K}$-linearly independent elements in the module $\mathcal{M}_{v,d,q}$. A straight-forward adaptation of the techniques of [PP19] gives an efficient, recursive algorithm to find one more module element to complete the basis, assuming that $\gcd(\mathcal{N}(f), \mathcal{N}(g)) = 1$. The concrete trapdoor generation algorithm $\mathsf{TrapGen}$ is defined in Figure 1.

```
$\mathsf{TrapGen}(1^\lambda, 1^1, 1^{d+1}, q)$
repeat
  repeat
    $f, g \leftarrow\!\!\$\ \chi$
  until $g \in \mathcal{R}_q^\times$, $\gcd(\mathcal{N}(f), \mathcal{N}(g)) = 1$
  $v := fg^{-1} \bmod q$
  Compute $(a_0, \ldots, a_d)$ s.t. $\sum_{i=0}^{d} f^{d-i} g^i a_i = q$ (see Sec. 4 of [PP19])
  $\mathbf{T} := \begin{bmatrix} g & & & a_0 \\ -f & \ddots & & a_1 \\ & \ddots & \ddots & \vdots \\ & & g & a_{d-1} \\ & & -f & a_d \end{bmatrix}$
until $\|M(\mathbf{T})\|_{\mathrm{GS}} \leq B$
return $(v, \mathbf{T})$
```

Fig. 1: vSIS trapdoor generation algorithm parametrised by χ and B, where χ is a distribution over $\mathcal{R}$ and B is an upper bound on $\|M(\mathbf{T})\|_{\mathrm{GS}}$; we will discuss selection of χ and B in Section 4.2. The choice of $(a_0, \ldots, a_d)$ ensures that $\det(\mathbf{T}) = \det(\mathbf{B}) = q$ (where $\mathbf{B}$ is as in Lemma 6), which is necessary for $\mathbf{T}$ to generate the correct module.

Theorem 2. *Use the same notation as in Figure 1 and suppose that the algorithm* $\mathsf{TrapGen}$ *returns* $(v, \mathbf{T})$*. Then,* $\mathbf{T}$ *generates the vSIS module* $\mathcal{M}_{v,d,q}$*.*

Proof. Let $\mathbf{B}$ be as in Lemma 6; we need to show that the modules $\mathcal{M}_{\mathbf{B}}$ and $\mathcal{M}_{\mathbf{T}}$ (the $\mathcal{R}$-modules generated by the columns of $\mathbf{B}$ and $\mathbf{T}$, respectively) are equal. Using the condition imposed on the coefficients $(a_0, \ldots, a_d)$ we obtain

$$\det(\mathbf{T}) = \sum_{i=0}^{d} f^{d-i} g^i a_i = q = \det(\mathbf{B}).$$

Since both determinants are non-zero and the dimensions of the matrices are equal, we conclude that the modules are of equal rank. Since the ranks are equal and $\det(\mathbf{T}) = \det(\mathbf{B})$, it suffices to show that $\mathcal{M}_{\mathbf{T}} \subseteq \mathcal{M}_{\mathbf{B}}$; this can be seen from

$$\mathbf{T} = \mathbf{B} \begin{bmatrix} g & & & & a_0 \\ vg - f & \ddots & & & va_0 + a_1 \\ \vdots & \ddots & \ddots & & \vdots \\ v^{d-2}(vg-f) & \cdots\cdots & vg - f & g & \sum_{i=0}^{d-1} v^{d-1-i} a_i \\ v^{d-1}(vg-f)/q & \cdots & v(vg-f)/q & (vg-f)/q & \left(\sum_{i=0}^{d} v^{d-i} a_i\right)/q \end{bmatrix}. \qquad \square$$

In the TrapGen algorithm constructed in Figure 1, the norm of the last column of $\mathbf{T}$ can in theory be almost arbitrarily large. However, it can be reduced utilising a $\mathcal{R}$-module analogue of Babai's nearest plane algorithm [Bab86]. A similar approach is used with NTRU trapdoors in [HHP$^+$03] and [PP19]. In the formal description of TrapGen in Figure 1, we handle this by only letting TrapGen output $\mathbf{T}$ until it finds one with small enough Gram-Schmidt norm. It remains to be seen whether the algorithm TrapGen is efficient or not. Similarly as with NTRU trapdoors (see e.g. [DLP14]), we can only present heuristic evidence. Our numerical experiments, detailed in Section 4.2, suggest that the expected number of iterations is constant when $B = \Omega(d \cdot \Delta_{\mathcal{K}}^{1/(2d_{\mathcal{K}})} \cdot q^{1/(d+1)})$.

4.2 Empirical Results

Bounding the Gram-Schmidt norm. We study the Gram-Schmidt norm of $M(\mathbf{T})$. Obtaining an upper bound seems infeasible, as such a result has not been established even in the less complicated case of NTRU trapdoors. This is why we can ultimately only provide heuristic results for the expected behavior.

Lemma 7. *Let $\mathcal{K}$ be a cyclotomic field of degree $d_{\mathcal{K}}$ and let $\mathbf{x} \in \mathcal{K}^n$ and $\{\mathbf{v}_1, \ldots, \mathbf{v}_m\} \subseteq \mathcal{K}^n$. Furthermore, denote $S_{\mathcal{K}} = \operatorname{span}_{\mathcal{K}}\{\mathbf{v}_1, \ldots, \mathbf{v}_m\}$ and $S_{\mathbb{C}} = \operatorname{span}_{\mathbb{C}} \left\{ \sigma(\mathbf{v}_i \zeta^{j-1}) \;\middle|\; i \in [m],\, j \in [d] \right\}$. Then, we have*

$$\sigma(\operatorname{proj}_{S_{\mathcal{K}}}(\mathbf{x})) = \operatorname{proj}_{S_{\mathbb{C}}}(\sigma(\mathbf{x})),$$

where projection onto $S_{\mathcal{K}}$ is defined by equipping $\mathcal{K}$ with the inner product inherited from $\mathbb{C}$, i.e. $\langle \mathbf{x}, \mathbf{y} \rangle = \bar{\mathbf{x}}^{\mathsf{T}} \cdot \mathbf{y}$ where $\bar{\cdot}$ denotes the complex conjugate.[1]

Proof. Write $\mathbf{x} = \mathbf{y} + \mathbf{z}$ where $\mathbf{y} = \operatorname{proj}_{S_{\mathcal{K}}}(\mathbf{x})$. Then, $\sigma(\mathbf{x}) = \sigma(\mathbf{y}) + \sigma(\mathbf{z})$. Since $\mathbf{y} \in S_{\mathcal{K}}$, there exists $c_{ij} \in \mathbb{Q}$ such that $\mathbf{y} = \sum_{i \in [m]} \sum_{j \in [d]} c_{ij} \zeta^{j-1} \mathbf{v}_i$ and hence $\sigma(\mathbf{y}) = \sum_{i \in [m]} \sum_{j \in [d]} c_{ij} \sigma(\zeta^{j-1} \mathbf{v}_i)$. Thus, $\sigma(\mathbf{y}) \in S_{\mathbb{C}}$.

Secondly, by construction we have $\mathbf{z} \perp S_{\mathcal{K}}$ and thus $\langle \mathbf{z}, \mathbf{v}_i \zeta^{j-1} \rangle = 0$ for all $i \in [m]$, $j \in [d]$. This implies $\langle \sigma(\mathbf{z}), \sigma(\mathbf{v}_i \zeta^{j-1}) \rangle = 0$ for $i \in [m]$, $j \in [d]$. Therefore $\sigma(\mathbf{z}) \perp S_{\mathbb{C}}$, concluding the proof. □

The following result is inspired by Lemma 2.2 of [CPS$^+$20]. Observe that our setting is different due to using the canonical embedding instead of the coefficient embedding.

Lemma 8. *Let $\mathcal{K}, d_{\mathcal{K}}$ be as in Lemma 7, and denote $\mathbf{B} = \begin{bmatrix} \mathbf{b}_1 \cdots \mathbf{b}_n \end{bmatrix} \in \mathcal{K}^{n \times n}$ and $M(\mathbf{B}) = \begin{bmatrix} \mathbf{r}_1 \cdots \mathbf{r}_{nd_{\mathcal{K}}} \end{bmatrix}$. Then,*

- *$\tilde{\mathbf{r}}_{(i-1)d_{\mathcal{K}}+1} = \sigma\left(\tilde{\mathbf{b}}_i\right)$ for all $i \in [n]$, and*
- *$\|M(\mathbf{B})\|_{\mathrm{GS}} = \|\mathbf{B}\|_{\mathrm{GS}}$.*

Proof. By definition of GSO we have

$$\tilde{\mathbf{r}}_{(i-1)d+1} = \mathbf{r}_{(i-1)d+1} - \operatorname{proj}_{\operatorname{span}_{\mathcal{K}}\left\{\mathbf{r}_1, \ldots, \mathbf{r}_{(i-1)d}\right\}}(\mathbf{r}_{(i-1)d+1})$$

[1] Note that any cyclotomic field $\mathcal{K}$ is closed under complex conjugation.

for all $i \in [n]$. Using the fact that $\mathbf{r}_{(i-1)d+j} = \sigma(\mathbf{b}_i\zeta^{j-1})$ for all $i \in [n]$, $j \in [d]$, yields the first claim together with Lemma 7.

For the second claim we need to prove that $\|\tilde{\mathbf{r}}_{(i-1)d+j}\| \leq \|\sigma(\tilde{\mathbf{b}}_i)\|$ for all $i \in [n]$, $j \in \{2, \dots, d\}$. Towards this, denote

$$\begin{aligned} S_1 &= \operatorname{span}_{\mathbb{C}}\{\tilde{\mathbf{r}}_1, \dots, \tilde{\mathbf{r}}_{(i-1)d}\}, \\ S_2 &= \operatorname{span}_{\mathbb{C}}\{\tilde{\mathbf{r}}_{(i-1)d+1}, \dots, \tilde{\mathbf{r}}_{(i-1)d+j-1}\}, \\ \mathbf{p}_1 &= \operatorname{proj}_{S_1}(\mathbf{r}_{(i-1)d+j}), \text{ and} \\ \mathbf{p}_2 &= \operatorname{proj}_{S_2}(\mathbf{r}_{(i-1)d+j}) \end{aligned}$$

such that $\tilde{\mathbf{r}}_{(i-1)d+j} = \mathbf{r}_{(i-1)d+j} - \mathbf{p}_1 - \mathbf{p}_2$. By the properties of GSO we have $S_1 = \operatorname{span}_{\mathbb{C}}\{\mathbf{r}_1, \dots, \mathbf{r}_{(i-1)d}\}$ and hence

$$\begin{aligned} \mathbf{p}_1 = \operatorname{proj}_{\operatorname{span}_{\mathbb{C}}\{\mathbf{r}_1,\dots,\mathbf{r}_{(i-1)d}\}}(\mathbf{r}_{(i-1)d+j}) &= \sigma\left(\operatorname{proj}_{\operatorname{span}_{\mathcal{K}}\{\mathbf{b}_1,\dots,\mathbf{b}_{i-1}\}}(\mathbf{b}_i\zeta^{j-1})\right) \\ &= \sigma\left(\operatorname{proj}_{\operatorname{span}_{\mathcal{K}}\{\mathbf{b}_1,\dots,\mathbf{b}_{i-1}\}}(\mathbf{b}_i)\zeta^{j-1}\right) \\ &= \sigma\left(\mathbf{b}_i - \tilde{\mathbf{b}}_i\right) \circ \left(\sigma\left(\zeta^{j-1}\right), \dots, \sigma\left(\zeta^{j-1}\right)\right) \end{aligned}$$

where we applied Lemma 7 for the second equality, the third equality follows from linearity of the projection and $\circ$ denotes the Hadamard product. Since all elements of $\sigma(\zeta^{j-1})$ are of magnitude 1, we conclude that $\|\mathbf{p}_1\| = \|\sigma(\mathbf{b}_i - \tilde{\mathbf{b}}_i)\|$. By a similar argument we also have $\|\mathbf{r}_{(i-1)d+j}\| = \|\sigma(\mathbf{b}_i)\|$. Finally, by orthogonality we obtain

$$\begin{aligned} \left\|\tilde{\mathbf{r}}_{(i-1)d+j}\right\|^2 &= \left\|\mathbf{r}_{(i-1)d+j}\right\|^2 - \|\mathbf{p}_1\|^2 - \|\mathbf{p}_2\|^2 \\ &= \|\sigma(\mathbf{b}_i)\|^2 - \left\|\sigma(\mathbf{b}_i - \tilde{\mathbf{b}}_i)\right\|^2 - \|\mathbf{p}_2\|^2 \\ &= \left\|\sigma(\tilde{\mathbf{b}}_i)\right\|^2 - \|\mathbf{p}_2\|^2 \leq \left\|\sigma(\tilde{\mathbf{b}}_i)\right\|^2. \qquad \square \end{aligned}$$

In light of Lemma 8, to estimate the GS-norm of $M(\mathbf{T})$ (with GSO performed over $\mathbb{C}^{(d+1)d_{\mathcal{K}}}$), it suffices to consider the GS-norm of $\mathbf{T}$ (with GSO performed over $\mathcal{K}^{d+1}$). This makes the task easier. Also, thanks to the lemma, we can talk about the GS-norm of $\mathbf{T}$ interchangeably with that of $M(\mathbf{T})$ without ambiguity.

With that out of the way, let us study $\|\mathbf{T}\|_{\mathrm{GS}}$. During this process, we generalise a series of results from [DLP14]. We begin by noticing that as long as f, g are short ring elements (as in NTRU), $\tilde{\mathbf{t}}_1, \dots, \tilde{\mathbf{t}}_d$ are indeed short. However, we also need to estimate the norm of $\tilde{\mathbf{t}}_{d+1}$. The following result formalises this observation.

Corollary 3. *Use the same definitions as in Theorem 2; in addition, let $\mathcal{K}, d_{\mathcal{K}}$ be as in Lemma 7. We have*

$$\|\mathbf{T}\|_{\mathrm{GS}} = \max\{\|\mathbf{t}_1\|, \|\tilde{\mathbf{t}}_{d_{\mathcal{K}}+1}\|\}.$$

Proof. By the structure of $\mathbf{T}$, the vectors $\mathbf{t}_1, \dots, \mathbf{t}_d$ are permutations of each other and hence their norms are equal. Moreover, by Lemma 8,

$$\sigma(\tilde{\mathbf{t}}_i) = \sigma(\mathbf{t}_i) - \operatorname{proj}_S(\mathbf{t}_i)$$

where $S = \operatorname{span}_{\mathbb{C}}\left\{ \mathbf{t}_j\zeta^{k-1} \,\middle|\, j \in [i-1], k \in [d_{\mathcal{K}}] \right\}$. Combining these two observations we get $\|\tilde{\mathbf{t}}_i\| \leq \|\mathbf{t}_i\| = \|\mathbf{t}_1\|$ for all $i \in [d]$. □

Next, we derive an explicit expression for $\tilde{\mathbf{t}}_{d+1}$.

Lemma 9. *Use the same definitions as in Corollary 3. Furthermore, let $\tilde{\mathbf{T}} = [\tilde{\mathbf{t}}_1 \cdots \tilde{\mathbf{t}}_{d+1}]$ be the GSO of $\mathbf{T}$. We have*

$$\tilde{\mathbf{t}}_{d+1} = \frac{q}{\sum_{i=0}^{d} f^{d-i}g^i\overline{f^{d-i}g^i}} \left(\overline{f^d}, \overline{f^{d-1}g}, \dots, \overline{fg^{d-1}}, \overline{g^d}\right).$$

Proof. Denote $\mathbf{c} = \frac{q}{\sum_{i=0}^{d} f^{d-i}g^i\overline{f^{d-i}g^i}}\left(\overline{f^d},\, \overline{f^{d-1}g},\, \ldots,\, \overline{fg^{d-1}},\, \overline{g^d}\right)$. To prove that $\tilde{\mathbf{t}}_{d+1} = \mathbf{c}$ it suffices to show that $\mathbf{c}$ is of form $\mathbf{t}_{d+1} - \mathbf{p}$ where $\mathbf{p}$ is the orthogonal projection $\mathrm{proj}_{\mathrm{span}_{\mathcal{K}}\{\mathbf{t}_1,\ldots,\mathbf{t}_d\}}(\mathbf{t}_{d+1})$. This can be split into two separate statements:

(i) $\mathbf{p} = \mathbf{t}_{d+1} - \mathbf{c} \in \mathrm{span}_{\mathcal{K}}\{\mathbf{t}_1, \ldots, \mathbf{t}_d\}$, and
(ii) $\mathbf{c} \perp \mathrm{span}_{\mathcal{K}}\{\mathbf{t}_1, \ldots, \mathbf{t}_d\}$.

By the structure of $\mathbf{T}$, $\{\mathbf{t}_1, \ldots, \mathbf{t}_d\}$ is a linearly independent set and hence (i) is equivalent to showing that the determinant of the matrix $\mathbf{T}' = \left[\mathbf{t}_1 \cdots \mathbf{t}_d\ \mathbf{t}_{d+1} - \mathbf{c}\right]$ is zero. Applying the Laplace expansion yields

$$\begin{aligned}
\det\left(\mathbf{T}'\right) &= \sum_{j=0}^{d}\left(a_j - \frac{q}{\sum_{i=0}^{d} f^{d-i}g^i\overline{f^{d-i}g^i}}\overline{f^{d-j}g^j}\right) f^{d-j}g^j \\
&= \sum_{j=0}^{d} a_j f^{d-j}g^j - \frac{q}{\sum_{i=0}^{d} f^{d-i}g^i\overline{f^{d-i}g^i}}\sum_{j=0}^{d}\overline{f^{d-j}g^j}f^{d-j}g^j \\
&= \sum_{j=0}^{d} a_j f^{d-j}g^j - q = 0.
\end{aligned}$$

For (ii), notice that for all $i \in [D]$ we have

$$\langle \mathbf{t}_i, \left(\overline{f^d},\, \overline{f^{d-1}g},\, \ldots,\, \overline{fg^{d-1}},\, \ldots,\, \overline{g^d}\right)\rangle = gf^{d-i+1}g^{i-1} - ff^{d-i}g^i = 0.$$ □

The following result combines Corollary 4 and Lemma 9; it provides a way to compute $\|\mathbf{T}\|_{\mathrm{GS}}$ given f, g without going through the GSO process. In practice it allows one to quickly identify bad choices of f, g.

Corollary 4. *Use the same definitions as in Corollary 3. We have*

$$\|\mathbf{T}\|_{\mathrm{GS}} = \max\left\{\|\mathbf{t}_1\|,\, \left\|\frac{q}{\sum_{i=0}^{d} f^{d-i}g^i\overline{f^{d-i}g^i}}\left(\overline{f^d},\, \overline{f^{d-1}g},\, \ldots,\, \overline{fg^{d-1}},\, \overline{g^d}\right)\right\|\right\}.$$

Our construction also admits a lower bound for the norm of $\tilde{\mathbf{t}}_{d+1}$ that will aid in the analysis.

Lemma 10. *Use the same definitions as in Corollary 3. Then,*

$$\left\|\tilde{\mathbf{t}}_{d+1}\right\| \geq \Delta_{\mathcal{K}}^{(d+1)/(2d_{\mathcal{K}})} \cdot \frac{q}{\|\mathbf{t}_1\|^d}.$$

Proof. Denote the GSO of $M(\mathbf{T})$ as $\left[\tilde{\mathbf{r}}_1 \cdots \tilde{\mathbf{r}}_{(d+1)d_{\mathcal{K}}}\right]$. By orthogonality of the vectors $\tilde{\mathbf{r}}_i$ we have

$$\left|\det\left(M(\mathbf{T})\right)\right| = \prod_{i=1}^{(d+1)d_{\mathcal{K}}} \|\tilde{\mathbf{r}}_i\| \leq \|\mathbf{t}_1\|^{d\cdot d_{\mathcal{K}}} \cdot \left\|\tilde{\mathbf{t}}_{d+1}\right\|^{d_{\mathcal{K}}} \tag{1}$$

where the inequality follows from Lemma 8 and Corollary 3. By Theorem 2, $M(\mathbf{T})$ generates the same lattice as

$$\begin{bmatrix}
M(1) & & & \\
M(-v) & \ddots & & \\
 & \ddots & M(1) & \\
 & & M(-v) & M(q)
\end{bmatrix}.$$

Thus, $\det(M(\mathbf{T})) = \det(M(1))^d \cdot \det(M(q)) = \det(M(1))^{d+1} \cdot q^{d_{\mathcal{K}}} = \Delta_{\mathcal{K}}^{(d+1)/2} \cdot q^{d_{\mathcal{K}}}$. Substituting this to Equation (1) yields the desired bound. □

The first lower bound of the following corollary is an immediate result of Corollary 3 and Lemma 10; the second one is obtained by minimising the first expression with respect to $\|\mathbf{t}_1\|$.

Corollary 5. *Use the same definitions as in Corollary 3. We have*

$$\|\mathbf{T}\|_{\mathrm{GS}} \geq \max\left\{ \|\mathbf{t}_1\|,\ \Delta_{\mathcal{K}}^{(d+1)/(2d_{\mathcal{K}})} \cdot \frac{q}{\|\mathbf{t}_1\|^d} \right\}.$$

Thus, the theoretical lower bound over different choices of $\|\mathbf{t}_1\|$ *is*

$$\|\mathbf{T}\|_{\mathrm{GS}} \geq \Delta_{\mathcal{K}}^{1/(2d_{\mathcal{K}})} \cdot q^{1/(d+1)}.$$

In the algorithm TrapGen we have control over the norm of $\mathbf{t}_1$. Therefore, the latter bound is important, asserting a concrete bound on the quality of the trapdoor. For ease of notation in the following sections, let us define the shorthand $c_{\mathcal{K},q,d} = \Delta_{\mathcal{K}}^{1/(2d_{\mathcal{K}})} \cdot q^{1/(d+1)}$.

Numerical results. The results of the previous section provide a solid basis for understanding the GS norm of $\mathbf{T}$. In what follows, we attempt to complete the picture by providing results from numerical experiments. It will turn out that, with careful choice of parameters, we can sample the trapdoor from a distribution where $\|M(\mathbf{T})\|_{\mathrm{GS}}$ is (heuristically) expected to be little over the theoretical lower bound, $c_{\mathcal{K},q,d}$.

Our experiments were motivated by the numerical results of [DLP14] for NTRU trapdoors, as well as those of [CKKS19,CPS+20] for module-NTRU trapdoors. In all of these, the norm of the last column of the trapdoor was found to be not much greater than the theoretical lower bound. We expected to observe a similar pattern for the vSIS trapdoors, i.e., that the norm of $\tilde{\mathbf{t}}_{d+1}$ would be somewhat close to $\Delta_{\mathcal{K}}^{(d+1)/(2d_{\mathcal{K}})} \cdot q/\|\mathbf{t}_1\|^d$ as given by Lemma 10.

We conducted our experiments using SageMath. In the experiments, we sampled several independent $f \in \mathcal{R}_q$, $g \in \mathcal{R}_q^\times$ by picking their coefficient vectors from a discrete Gaussian distribution. The Gaussian parameter was varied to obtain range of different values for $\|\mathbf{t}_1\|$. Then, for each pair f, g, we computed $\|\tilde{\mathbf{t}}_{d+1}\|$ using Lemma 9 and plotted $\|\tilde{\mathbf{t}}_{d+1}\|$ against $\|\mathbf{t}_1\|$ to analyse how the former depends on the latter.

This process was repeated for all different combinations of $\mathcal{K} = \mathbb{Q}(\zeta_{\mathfrak{f}})$ for

$$\mathfrak{f} \in \{64, 128\}, \qquad q \in \{1000193, 1000000513\}, \qquad \text{and} \qquad d \in [6],$$

sampling 200 independent pairs f, g for each of them. The choices of q were to ensure that q is totally split in $\mathbb{Q}(\zeta_{\mathfrak{f}})$, which is a popular choice for practical applications.

The results of the experiments for a subset of parameters are shown in Figure 2. It is in line with our expectations, since most samples result in $\|\tilde{\mathbf{t}}_{d+1}\|$ close to the lower bound. Some $\|\tilde{\mathbf{t}}_{d+1}\|$ are still significantly larger (especially for larger d), which might seem alarming. However, in practice we may simply reject such f, g if the proportion of such outliers is sufficiently small; this is similar to what is done in Algorithm 2 of [DLP14].

To understand the average behavior of $\|\tilde{\mathbf{t}}_{d+1}\|$, we fitted a linear least squares model to the data points $(\ln(\|\mathbf{t}_1\|), \ln(\|\mathbf{t}_{d+1}\|))$. Heuristically, we assume that the optimal choice of $\|\mathbf{t}_1\|$ is roughly where the model intersects the line $y = \|\mathbf{t}_1\|$. In this region, approximately one half of the resulting $\|\tilde{\mathbf{t}}_{d+1}\|$ are less than or equal to $\|\mathbf{t}_1\|$, in which case $\|\mathbf{T}\|_{\mathrm{GS}} \leq \|\mathbf{t}_1\|$.

In light of this discussion, the y-coordinate of the intersection between the least squares model and the line $y = \|\mathbf{t}_1\|$ can be viewed as an estimate of the practically achievable $\|\mathbf{T}\|_{\mathrm{GS}}$. Inspired by [CPS+20], we express this value as gs_slack $\cdot\, c_{\mathcal{K},q,d}$ where gs_slack is a constant. According to Figure 3, gs_slack appears to grow roughly linearly with respect to d. In contrast, our experiments do not suggest that the degree of $\mathcal{K}$ or the choice of modulus q would significantly affect gs_slack.

Remark 2. Increasing d leads to better trapdoors only as long as (i) the decrease of $q^{1/(d+1)}$ subdues the (presumably) linear growth of gs_slack and (ii) we are able to sample $\mathbf{t}_1$ with norm $c_{\mathcal{K},q,d}$. In other words, $\|\mathbf{T}\|_{\mathrm{GS}}$ cannot be improved indefinitely by increasing the vSIS degree and it may even get worse after (i) is no longer satisfied.

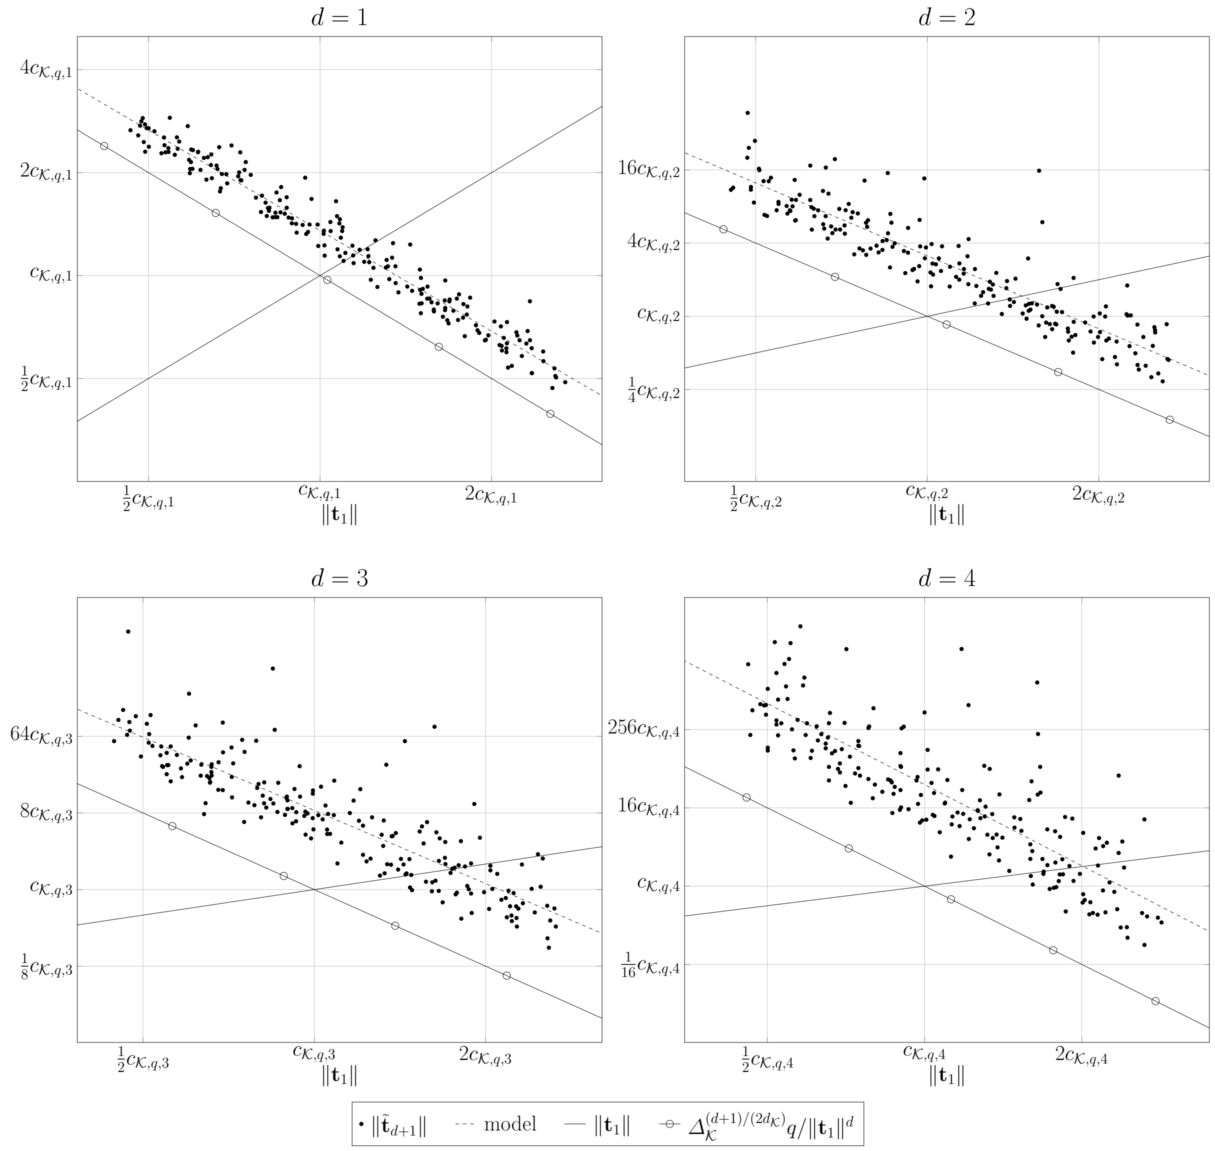

Fig. 2: The norm of the last column of $\tilde{\mathbf{T}}$ as a function of $\|\mathbf{t}_1\|$ for $\mathcal{K} = \mathbb{Q}(\zeta_{128})$, $q = 1000193$ and different values of d. The plots use logarithmic scales for both the x and y axes, with the least squares model estimating the expected behavior.

5 Homomorphic Signatures

To demonstrate the utility of vSIS trapdoors, we present a simple construction of homomorphic signatures (HS) for constant-degree polynomials over $\mathcal{R}$ with short coefficients, where a signature is itself a polynomial $\mathsf{s} \in \mathcal{R}[X]$ with short coefficients. For simplicity, we first focus our discussion on a scheme which allows signing a single dataset with selective security, and defer discussions about full (i.e. adaptive) and multi-dataset security to Section 5.3. We then briefly discuss our construction in relation to existing schemes in Section 5.4.

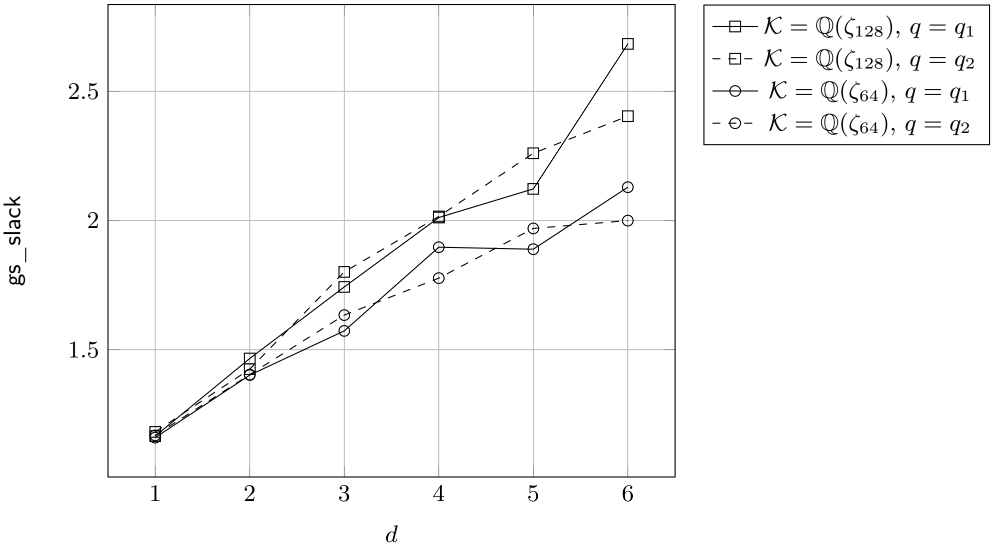

Fig. 3: Growth of gs_slack as a function of d. Here we used the moduli $q_1 = 1000193$ and $q_2 = 1000000513$.

5.1 Definitions

Definition 9 (Single-dataset HS). *A single-dataset homomorphic signature (HS) scheme for message space $\mathcal{X}$ and admissible functions $\mathcal{G}$ over $\mathcal{X}$ consists of a tuple of algorithms* $(\mathsf{PrmsGen}, \mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Eval}, \mathsf{Process}, \mathsf{Verify})$ *with the following syntax.*

- $\mathsf{pp} \leftarrow \mathsf{PrmsGen}(1^\lambda, 1^N)$*: Gets the security parameter λ and a data-size bound N. Generates the public parameters* pp.
- $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda, \mathsf{pp})$*: Gets the security parameter along with the public parameters. Generates the public key and the secret key.*
- $(\mathsf{s}_1, \ldots, \mathsf{s}_N) \leftarrow \mathsf{Sign}_{\mathsf{sk}}(x_1, \ldots, x_N)$*: Signs a tuple of data $(x_1, \ldots, x_N) \in \mathcal{X}^N$.*
- $\mathsf{s}^* \leftarrow \mathsf{Eval}_{\mathsf{pp}}(g, (x_1, \mathsf{s}_1), \ldots, (x_l, \mathsf{s}_l))$*: Homomorphically evaluates function $g \in \mathcal{G}$, outputting a new signature* s^**.*
- $\alpha_g \leftarrow \mathsf{Process}_{\mathsf{pp}}(g)$*: Computes a "public key" of g that is later used in the verification step.*
- $b \leftarrow \mathsf{Verify}_{\mathsf{pk}}(\alpha_g, y, \mathsf{s}^*)$*: Uses the signature s^* to check that y is equal to $g(x_1, \ldots, x_l)$; outputs 1 if that is the case, 0 otherwise.*

Definition 10 (Correctness). *Let Σ be a single-dataset HS scheme as defined in Definition 9. We say that Σ is correct if it satisfies the following two notions for* $\mathsf{pp} \leftarrow \mathsf{PrmsGen}(1^\lambda, 1^N)$ *and* $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda, 1^N)$*:*

- *Signing correctness. Let $\mathsf{id}_i : \mathcal{X}^N \to \mathcal{X}$ be defined as the canonical extension of $\mathsf{id}_i(x_1, \ldots, x_N) = x_i$. We require that for any $i \in [N]$, $(x_1, \ldots, x_N) \in \mathcal{X}^N$ and $(\mathsf{s}_1, \ldots, \mathsf{s}_N) \leftarrow \mathsf{Sign}_{\mathsf{sk}}(x_1, \ldots, x_N)$, it holds that*

$$\mathsf{Verify}_{\mathsf{pk}}(\mathsf{Process}_{\mathsf{pp}}(\mathsf{id}_i), x_i, \mathsf{s}_i) = 1$$

  except with negligible probability in λ, with probability taken over the randomness of $\mathsf{PrmsGen}, \mathsf{KeyGen}$ *and* $\mathsf{Sign}_{\mathsf{sk}}$.
- *Evaluation correctness. For any $g \in \mathcal{G}$, $(x_1, \ldots, x_N) \in \mathcal{X}^N$, $(\mathsf{s}_1, \ldots, \mathsf{s}_N) \leftarrow \mathsf{Sign}_{\mathsf{sk}}(x_1, \ldots, x_N)$ and $\mathsf{s}^* \leftarrow \mathsf{Eval}_{\mathsf{pp}}(g, (x_1, \mathsf{s}_1), \ldots, (x_N, \mathsf{s}_N))$, it holds that*

$$\mathsf{Verify}_{\mathsf{pk}}(\mathsf{Process}_{\mathsf{pp}}(g), g(x_1, \ldots, x_N), \mathsf{s}^*) = 1$$

  except with negligible probability in λ, with probability taken over the randomness of $\mathsf{PrmsGen}, \mathsf{KeyGen}$*,* $\mathsf{Sign}_{\mathsf{sk}}$ *and* $\mathsf{Eval}_{\mathsf{pp}}$.

Remark 3. In the above, we only define single-hop evaluation correctness for simplicity. It is clear how the definition can be extended to capture multi-hop evaluation correctness.

Definition 11 (Security). *Let N be a function of λ. We define the security of single-dataset HS via the following security games:*

```
$\mathsf{Selective\text{-}security}_{\Sigma,\mathcal{A}}(1^\lambda)$
$(x_1, \ldots, x_N) \leftarrow \mathcal{A}$
$\mathsf{pp} \leftarrow \mathsf{PrmsGen}(1^\lambda, 1^N)$
$(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda, \mathsf{pp})$
$(\mathsf{s}_1, \ldots, \mathsf{s}_N) \leftarrow \mathsf{Sign}_{\mathsf{sk}}(x_1, \ldots, x_N)$
$(g, y, \mathsf{s}) \leftarrow \mathcal{A}(\mathsf{pp}, \mathsf{pk}, (\mathsf{s}_1, \ldots, \mathsf{s}_N))$
$b_0 = \mathsf{Verify}_{\mathsf{pk}}(\mathsf{Process}_{\mathsf{pp}}(g), y, \mathsf{s})$
$b_1 = (g(x_1, \ldots, x_N) \neq y)$
$b_2 = (g \in \mathcal{G})$
return $b_0 \wedge b_1 \wedge b_2$
```

```
$\mathsf{Full\text{-}security}_{\Sigma,\mathcal{A}}(1^\lambda)$
$\mathsf{pp} \leftarrow \mathsf{PrmsGen}(1^\lambda, 1^N)$
$(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda, \mathsf{pp})$
$(x_1, \ldots, x_N) \leftarrow \mathcal{A}(\mathsf{pp}, \mathsf{pk})$
$(\mathsf{s}_1, \ldots, \mathsf{s}_N) \leftarrow \mathsf{Sign}_{\mathsf{sk}}(x_1, \ldots, x_N)$
$(g, y, \mathsf{s}) \leftarrow \mathcal{A}((\mathsf{s}_1, \ldots, \mathsf{s}_N))$
$b_0 = \mathsf{Verify}_{\mathsf{pk}}(\mathsf{Process}_{\mathsf{pp}}(g), y, \mathsf{s})$
$b_1 = (g(x_1, \ldots, x_N) \neq y)$
$b_2 = (g \in \mathcal{G})$
return $b_0 \wedge b_1 \wedge b_2$
```

An HS scheme is said to be selectively secure if, for any polynomially bounded adversary $\mathcal{A}$, the probability $Pr\left[\mathsf{Selective\text{-}security}_{\Sigma,\mathcal{A}}(1^\lambda) = 1\right]$ is negligible in λ. The scheme is called fully secure if the same holds for $Pr\left[\mathsf{Full\text{-}security}_{\Sigma,\mathcal{A}}(1^\lambda) = 1\right]$.

5.2 Construction

In this section let $q, N, d_{\text{init}}, d_{\text{eval}}, d_{\text{max}} \in \mathbb{N}$, $\beta_{\text{init}}, \beta_{\text{eval}}, \beta_{\text{max}} > 0$ be some norm bounds, $s > 0$ be a Gaussian parameter, Using this notation, Figure 4 presents a vSIS-based single-dataset HS scheme over message space $\mathcal{X}$ that allows evaluation of degree-d_{eval} N-variate polynomials of norm at most β_{eval}.

```
$\mathsf{PrmsGen}(1^\lambda, 1^N)$
$\mathsf{pp} := (r_1, \ldots, r_N) \leftarrow\!\$\ \mathcal{R}_q^N$
return pp

$\mathsf{KeyGen}(1^\lambda, \mathsf{pp})$
$(v, \mathsf{td}) \leftarrow \mathsf{TrapGen}(1^\lambda, 1^1, 1^{d_{\text{init}}}, q)$
$(\mathsf{pk}, \mathsf{sk}) := (v, \mathsf{td})$
return $(\mathsf{pk}, \mathsf{sk})$

$\mathsf{Sign}_{\mathsf{sk}}((x_i)_{i \in [N]})$
for $i \in [N]$ do
    $f_i'(X) \leftarrow \mathsf{SampPre}\left(\mathsf{td}, \frac{r_i - x_i}{v}, s\right)$
    $f_i(X) := X \cdot f_i' + x_i$
return $(f_1, \ldots, f_N)$

$\mathsf{Eval}_{\mathsf{pp}}(g, f_1, \ldots, f_N)$
$h := g(f_1, \ldots, f_N)$
return $h$

$\mathsf{Process}_{\mathsf{pp}}(g)$
assert $\deg(g) \leq d_{\text{eval}}$ and $\|g\| \leq \beta_{\text{eval}}$
$\alpha_g := g(r_1, \ldots, r_N) \bmod q$
return $\alpha_g$

$\mathsf{Verify}_{\mathsf{pk}}(\alpha_g, y, h)$
if $h(v) = \alpha_g \bmod q$ and $h(0) = y$ and
 $\deg(h) \leq d_{\max}$ and $\|h\| \leq \beta_{\max}$
    return 1
return 0
```

Fig. 4: A vSIS-based HS where TrapGen is as constructed in Figure 1 and SampPre is the algorithm described in Corollary 1 where an output is interpreted as a polynomial f of degree at most $d_{\text{init}} - 1$.

Remark 4. In the construction of Figure 4, we must require that v is a unit in $\mathcal{R}_q$. This means that we have to use a slightly modified version of the algorithm $\mathsf{TrapGen}$. However, over a suitable choice of $\mathcal{R}$ and q, e.g. when $\mathcal{R}_q$ splits into super-polynomial-size fields, the overwhelming majority of elements in $\mathcal{R}_q$ are units. Thus, neither the output distribution nor the efficiency changes drastically.

Theorem 3. *Let $q, N, d_{\text{init}}, d_{\text{eval}}, d_{\max} \in \mathbb{N}$ with $d_{\text{init}} \geq 2$ and $d_{\max} \geq d_{\text{init}} \cdot d_{\text{eval}}$. Let $\beta_{\text{td}}, \beta_{\text{init}}, \beta_{\text{eval}}, \beta_{\max}, s > 0$ be such that $\beta_{\text{td}} \geq \mathsf{gs\_slack} \cdot \Delta_{\mathcal{K}}^{1/(2d_{\mathcal{K}})} \cdot q^{1/d_{\text{init}}}$ for some constant $\mathsf{gs\_slack} > 0$, $\beta_{\text{init}} \geq s\sqrt{d_{\mathcal{K}} \cdot d_{\text{init}}}$,*

$$\beta_{\max} \geq \binom{d_{\text{eval}} + N}{d_{\text{eval}}} \cdot (d_{\text{init}} + 1)^{d_{\text{eval}}} \cdot \beta_{\text{init}}^{d_{\text{eval}}} \cdot \beta_{\text{eval}},$$

and $s \geq \beta_{\text{td}} \cdot \sqrt{d_{\mathcal{K}} \cdot d_{\text{init}}}$. Let $\mathcal{R}$ be the ring of integers of a power-of-2 cyclotomic field of degree $d_{\mathcal{K}}$ with $\mathcal{R}_q$ splitting into super-polynomial-size fields. Let $\mathsf{TrapGen}$ be such that it returns $\mathsf{td} = \mathbf{B}$ with $\|\mathbf{B}\|_{\text{GS}} \leq \beta_{\text{td}}$ with overwhelming probability in λ. Let the message space $\mathcal{X}$ be the set of all $\mathcal{R}$ elements of norm at most β_{init}. Let the set of admissible functions $\mathcal{G}$ be the set of all degree-d_{eval} N-variate polynomials of norm at most β_{eval}. Then, the HS scheme of Figure 4 is correct. Furthermore, it is selectively secure under the $\mathsf{dNTRU}_{\mathcal{R},\chi}$ and $\mathsf{vSIS}_{\mathcal{R},d_{\max},q,2\beta_{\max}}$ assumptions, where $\chi = \mathcal{D}_{\mathcal{R},s}$.

Proof. Signing correctness follows directly from the definition of the algorithm Sign, the guarantee of $\|\mathbf{B}\|_{\text{GS}}$ by the assumption on $\mathsf{TrapGen}$, the guarantee of $\mathsf{SampPre}$ by Corollary 1, and the Gaussian tail bound Lemma 2. In particular, fresh signatures f_i are of norm $\|f_i\| \leq \beta_{\text{init}}$ with overwhelming probability in λ.

For evaluation correctness, let $(f_i)_{i\in[N]} \leftarrow \mathsf{Sign}_{\mathsf{pk},\mathsf{sk}}((x_i)_{i\in[N]})$, $g : \mathcal{X}^N \to \mathcal{X}$ be a multivariate polynomial satisfying $\deg(g) \leq d_{\text{eval}}$, $\|g\| \leq \beta_{\text{eval}}$ and let $h \leftarrow \mathsf{Eval}(g, (f_i)_{i\in[N]})$. By signing correctness we have $h(v) = g(f_1(v), \ldots, f_N(v)) = g(r_1, \ldots, r_N) \bmod q$ and $h(0) = g(x_1, \ldots, x_N) \bmod q$, and by the degree bounds $\deg(g) \leq d_{\text{eval}}$ (by assumption) and $\deg(f_i) \leq d_{\text{init}}$ (by construction) we have $\deg(h) \leq d_{\max}$. We are left to bound the norm of h; writing $g(X_1, \ldots, X_N) = \sum_{\alpha:|\alpha|\leq d_{\text{eval}}} g_\alpha \cdot \prod_{i\in[N]} X_i^{\alpha_i}$ where $\alpha = (\alpha_1, \ldots, \alpha_N)$ ranges over tuples of non-negative integers with ℓ_1-norm at most d_{eval} and $f_i(X) = \sum_{j=0}^{d_{\text{init}}} f_{i,j} X^j$, observe that

$$\begin{aligned}
\|h\| &= \left\| \sum_{\alpha:|\alpha|\leq d_{\text{eval}}} g_\alpha \prod_{i\in[N]} \left(\sum_{j=0}^{d_{\text{init}}} f_{i,j} X^j \right)^{\alpha_i} \right\| \\
&\leq \sum_{\alpha:|\alpha|\leq d_{\text{eval}}} \left\| g_\alpha \prod_{i\in[N]} \left(\sum_{j=0}^{d_{\text{init}}} f_{i,j} X^j \right)^{\alpha_i} \right\| \quad \text{// triangle inequaliy} \\
&\leq \sum_{\alpha:|\alpha|\leq d_{\text{eval}}} \beta_{\text{eval}} \cdot \left\| \prod_{i\in[N]} \left(\sum_{j=0}^{d_{\text{init}}} f_{i,j} X^j \right)^{\alpha_i} \right\| \quad \text{// } \|g\| \leq \beta_{\text{eval}} \\
&\leq \sum_{\alpha:|\alpha|\leq d_{\text{eval}}} \beta_{\text{eval}} \cdot \left\| \prod_{i\in[N]} \beta_{\text{init}}^{\alpha_i} \cdot \left(\sum_{j=0}^{d_{\text{init}}} X^j \right)^{\alpha_i} \right\| \quad \text{// } \|f_i\| \leq \beta_{\text{init}} \\
&= \sum_{\alpha:|\alpha|\leq d_{\text{eval}}} \beta_{\text{eval}} \cdot \beta_{\text{init}}^{|\alpha|} \cdot \underbrace{\left\| \left(\sum_{j=0}^{d_{\text{init}}} X^j \right)^{|\alpha|} \right\|}_{\leq (d_{\text{init}}+1)^{|\alpha|} \text{ by induction on } |\alpha|} \\
&\leq \sum_{\alpha:|\alpha|\leq d_{\text{eval}}} \beta_{\text{eval}} \cdot \beta_{\text{init}}^{d_{\text{eval}}} \cdot (d_{\text{init}} + 1)^{d_{\text{eval}}} \quad \text{// } |\alpha| \leq d_{\text{eval}} \\
&= \binom{d_{\text{eval}} + N}{d_{\text{eval}}} \cdot \beta_{\text{eval}} \cdot \beta_{\text{init}}^{d_{\text{eval}}} \cdot (d_{\text{init}} + 1)^{d_{\text{eval}}} \quad \text{// counting monomials} \\
&\leq \beta_{\max}.
\end{aligned}$$

Towards security, suppose that $\mathcal{A}$ is a PPT adversary that breaks the selective security of the scheme. Given a vSIS instance v, we can construct a reduction which queries $\mathcal{A}$ for the messages $(x_i)_{i\in[N]} \in \mathcal{X}^N$. Then, for all $i \in [N]$ it generates a degree-d_{init} polynomial f_i by setting the constant coefficient equal to x_i and sampling the rest of the coefficients at random from $\mathcal{D}_{\mathcal{R},s}$. Moreover, for all $i \in [N]$ it computes $r_i = f_i(v) \bmod q$. The reduction gives the tuple $((f_i)_{i\in[N]}, v, (r_i)_{i\in[N]})$ to $\mathcal{A}$ and receives (g, y, h) in return.

We argue that the reduction faithfully simulates the selective security game for $\mathcal{A}$. We first consider an intermediate distribution of $((f_i)_{i\in[N]}, v, (r_i)_{i\in[N]})$ where v is sampled by $\mathsf{TrapGen}$ and the rest of the tuple is simulated as in the reduction. Define $\Lambda \coloneqq \Lambda_q^\perp(1, v, \ldots, v^{d_{\mathrm{init}}-1})$. By our assumption on $\mathsf{TrapGen}$, with overwhelming probability in λ, we have $\eta_\epsilon(\Lambda) \le s$ for some ϵ negligible in λ. Then, by Lemma 3, the simulated $((f_i)_{i\in[N]}, (r_i)_{i\in[N]})$ are statistically close to those in the real experiment.

To move from the intermediate distribution to that induced by the reduction, we note that a genuine public key v as constructed in Figure 1 is an NTRU instance, which is indistinguishable to a uniformly random $\mathcal{R}_q^\times$ element by the $\mathsf{dNTRU}_{\mathcal{R},\chi}$ assumption and by the assumption that $\mathcal{R}_q$ splits into super-polynomial-size fields.

Next, we analyse the output (g, y, h) received by the reduction from the adversary. With a non-negligible probability (g, y, h) is a forgery, that is, it satisfies

$$\begin{cases} g(x_1, \ldots, x_N) \neq y, \\ h(v) = g(r_1, \ldots, r_N) \bmod q, \\ h(0) = y, \\ \deg(h) \le d_{\max}, \\ \|h\| \le \beta_{\max}. \end{cases} \tag{2}$$

We claim that if this is the case, then $p \coloneqq h - g((f_i)_{i\in[N]})$ is a solution to the $\mathsf{vSIS}_{\mathcal{R},d_{\max},q,2\beta_{\max}}$ problem. Indeed, by the first and third conditions of Equation (2) we have

$$p(0) = h(0) - g(f_1(0), \ldots, f_N(0)) = y - g(x_1, \ldots, x_N) \neq 0,$$

the second one implies

$$p(v) = h(v) - g(f_1(v), \ldots, f_N(v)) = h(v) - g(r_1, \ldots, r_N) = 0 \bmod q$$

and by the fourth one the degree is low enough. Finally, by applying a similar reasoning as in the proof of evaluation correctness we get $\|g(f_1, \ldots, f_N)\| \le \beta_{\max}$. Hence, the fifth condition of Equation (2) yields the desired bound on $\|p\|$. □

Note that our scheme can be easily turned into a multi-key scheme [FMNP16]. In more detail, since a signature $\mathsf{s}_i \in \mathcal{R}[X_i]$ signed by user i can be interpreted as a constant-degree polynomial with formal variable X_i, homomorphically evaluating a constant-degree multivariate polynomial on $(\mathsf{s}_i)_i$ would yield another constant-degree polynomial with formal variable $(X_i)_i$, which can then be verified against the vector of public keys obtained by concatenating the public keys of all signers.

5.3 Adaptive and Multi-Dataset Security

Our construction in Section 5.2 achieves selective security in the single-dataset setting. We next discuss how generic transforms in [GVW14,GVW15] can be applied to obtain fully (i.e. adaptively) secure multi-dataset schemes.

First, we observe that the single-dataset to multi-dataset transform in [GVW14, Theorem 5.1], which preserves adaptive security, also preserves selective security. That is, the same transform turns a selectively secure single-dataset HS into a selectively secure multi-dataset HS. Notably, this transformation does not require the public parameters pp to be succinct, i.e. sublinear in N.

To upgrade a selectively secure single-dataset HS directly to an adaptively secure multi-dataset HS, we can alternatively use the generic transform in [GVW14, Theorem A.1]. A caveat is that this transform requires the public parameters pp of the underlying single-dataset HS to be of size sublinear in N, which is not the case for our construction. However, since our pp only consists of N uniformly random $\mathcal{R}_q$ elements, they can be generated from a short seed using a random oracle. Consequently, we can obtain an adaptively secure multi-dataset HS based on the same assumptions as in Theorem 3 in the random oracle model. Note that, however, as discussed in [GVW14, Appendix A] this transform results in HS schemes which are not suitable for verifiable outsourcing since they do not support verification preprocessing.

Finally, we note that one could potentially cast our single-dataset HS as a homomorphic trapdoor function (HTDF) and combine it with (another instance of) the HS via [GVW14, Theorem 4.3] to obtain an adaptively secure single-dataset HS, which can then be further upgraded to an adaptively secure multi-dataset HS using [GVW14, Theorem 5.1].

5.4 Related Work on Homomorphic Signatures

There is a long line of research on the construction of homomorphic signatures for various classes of functions based on different assumptions. We refer to [ABF24] for a recent summary and focus on comparing our construction with the most related ones in the literature below.

Our HS construction features a simple evaluation algorithm for constant-degree bounded-norm multivariate polynomials. We emphasise that our construction is meant to demonstrate the utility of vSIS trapdoors, but is unlikely to be concretely efficient (at least in its current form). The main bottleneck is that the lower bound of $\beta_{\max}$ in terms of β_{init} in Theorem 3 and the linear dependency of β_{init} on $q^{1/d_{\text{init}}}$ forces the modulus q to be concretely large even for small values of N, β_{eval} and d_{eval}.

Homomorphic signatures for beyond linear functions was first realised by [BF11] who constructed an HS for bounded-degree polynomials from lattices. Their idea is to let signatures be short elements in $\mathcal{R}$ and encode the message in the quotient ring $\mathcal{R}/\mathcal{I}$ for some ideal $\mathcal{I}$, so that signatures can be multiplied using multiplication over $\mathcal{R}$. In our construction signatures instead live in the polynomial ring $\mathcal{R}[X]$ and messages are encoded in $\mathcal{R}$. HS for (bounded polynomial depth arithmetic) circuits was first constructed in [GVW15] from lattices, based on which our construction is adapted. Signature multiplication in that scheme is based on the homomorphic computation technique developed in [GSW13,BGG$^+$14].

Apart from direct constructions, [CFT22] presented a generic construction of HS for a function class from functional commitments (FC) for the same function class. For example, when instantiating with FCs for constant-degree polynomials [ACL$^+$22,CFT22,CLM23], one obtains HS for low-degree polynomials. Furthermore, if the FC is chainable [BCFL23], then the resulting HS supports multi-hop evaluation. However, as discussed in [BCFL23], HS obtained in this way are not arbitrarily composable, unlike [BF11,GVW15] and our construction.

Acknowledgments. We thank the anonymous PKC 2025 reviewers for their helpful comments. The research of Russell W. F. Lai is supported by Research Council of Finland grants 358951.

References

ABF24. Gaspard Anthoine, David Balbás, and Dario Fiore. Fully-succinct multi-key homomorphic signatures from standard assumptions. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part III*, volume 14922 of *LNCS*, pages 317–351. Springer, Cham, August 2024. 20

ACL$^+$22. Martin R. Albrecht, Valerio Cini, Russell W. F. Lai, Giulio Malavolta, and Sri Aravinda Krishnan Thyagarajan. Lattice-based SNARKs: Publicly verifiable, preprocessing, and recursively composable - (extended abstract). In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 102–132. Springer, Cham, August 2022. 20

Ajt96. Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *28th ACM STOC*, pages 99–108. ACM Press, May 1996. 1, 3

Bab86. László Babai. On Lovász' lattice reduction and the nearest lattice point problem. *Combinatorica*, 6:1–13, 1986. 11

BCFL23. David Balbás, Dario Catalano, Dario Fiore, and Russell W. F. Lai. Chainable functional commitments for unbounded-depth circuits. In Guy N. Rothblum and Hoeteck Wee, editors, *TCC 2023, Part III*, volume 14371 of *LNCS*, pages 363–393. Springer, Cham, November / December 2023. 20

BF11. Dan Boneh and David Mandell Freeman. Homomorphic signatures for polynomial functions. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 149–168. Springer, Berlin, Heidelberg, May 2011. 20

BGG$^+$14. Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 533–556. Springer, Berlin, Heidelberg, May 2014. 1, 20

CFT22. Dario Catalano, Dario Fiore, and Ida Tucker. Additive-homomorphic functional commitments and applications to homomorphic signatures. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022, Part IV*, volume 13794 of *LNCS*, pages 159–188. Springer, Cham, December 2022. 20

CKKS19. Jung Hee Cheon, Duhyeong Kim, Taechan Kim, and Yongha Son. A new trapdoor over module-NTRU lattice and its application to ID-based encryption. Cryptology ePrint Archive, Report 2019/1468, 2019. 3, 14

CLM23. Valerio Cini, Russell W. F. Lai, and Giulio Malavolta. Lattice-based succinct arguments from vanishing polynomials - (extended abstract). In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part II*, volume 14082 of *LNCS*, pages 72–105. Springer, Cham, August 2023. 1, 3, 6, 20

CPS$^+$20. Chitchanok Chuengsatiansup, Thomas Prest, Damien Stehlé, Alexandre Wallet, and Keita Xagawa. ModFalcon: Compact signatures based on module-NTRU lattices. In Hung-Min Sun, Shiuh-Pyng Shieh, Guofei Gu, and Giuseppe Ateniese, editors, *ASIACCS 20*, pages 853–866. ACM Press, October 2020. 3, 11, 14

DLP14. Léo Ducas, Vadim Lyubashevsky, and Thomas Prest. Efficient identity-based encryption over NTRU lattices. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 22–41. Springer, Berlin, Heidelberg, December 2014. 2, 3, 9, 11, 12, 14

FMNP16. Dario Fiore, Aikaterini Mitrokotsa, Luca Nizzardo, and Elena Pagnin. Multi-key homomorphic authenticators. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 499–530. Springer, Berlin, Heidelberg, December 2016. 19

GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008. 1, 2, 4, 9

GSW13. Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 75–92. Springer, Berlin, Heidelberg, August 2013. 20

GVW14. Sergey Gorbunov, Vinod Vaikuntanathan, and Daniel Wichs. Leveled fully homomorphic signatures from standard lattices. Cryptology ePrint Archive, Report 2014/897, 2014. 19, 20

GVW15. Sergey Gorbunov, Vinod Vaikuntanathan, and Daniel Wichs. Leveled fully homomorphic signatures from standard lattices. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *47th ACM STOC*, pages 469–477. ACM Press, June 2015. 2, 19, 20

HHP$^+$03. Jeffrey Hoffstein, Nick Howgrave-Graham, Jill Pipher, Joseph H. Silverman, and William Whyte. NTRUSIGN: Digital signatures using the NTRU lattice. In Marc Joye, editor, *CT-RSA 2003*, volume 2612 of *LNCS*, pages 122–140. Springer, Berlin, Heidelberg, April 2003. 2, 3, 9, 11

KLNO24. Michael Klooß, Russell W. F. Lai, Ngoc Khanh Nguyen, and Michal Osadnik. RoK, paper, SISsors toolkit for lattice-based succinct arguments - (extended abstract). In Kai-Min Chung and Yu Sasaki, editors, *ASIACRYPT 2024, Part V*, volume 15488 of *LNCS*, pages 203–235. Springer, Singapore, December 2024. 1, 3

LPR10. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 1–23. Springer, Berlin, Heidelberg, May / June 2010. 1

LPR13. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for ring-LWE cryptography. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 35–54. Springer, Berlin, Heidelberg, May 2013. 1

LS15. Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *DCC*, 75(3):565–599, 2015. 1, 3

MP12. Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Berlin, Heidelberg, April 2012. 1, 2

MR07. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007. 4

PP19. Thomas Pornin and Thomas Prest. More efficient algorithms for the NTRU key generation using the field norm. In Dongdai Lin and Kazue Sako, editors, *PKC 2019, Part II*, volume 11443 of *LNCS*, pages 504–533. Springer, Cham, April 2019. 10, 11

PS21. Alice Pellet-Mary and Damien Stehlé. On the hardness of the NTRU problem. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part I*, volume 13090 of *LNCS*, pages 3–35. Springer, Cham, December 2021. 2, 3, 6, 9