Vandermonde meets Regev: public key encryption schemes based on partial Vandermonde problems

Katharina Boudgoust, Amin Sakzad, Ron Steinfeld

▶ **To cite this version:**

Katharina Boudgoust, Amin Sakzad, Ron Steinfeld. Vandermonde meets Regev: public key encryption schemes based on partial Vandermonde problems. Designs, Codes and Cryptography, 2022, 90 (8), pp.1899-1936. ⟨10.1007/s10623-022-01083-7⟩. ⟨hal-04535099⟩

HAL Id: hal-04535099

https://hal.science/hal-04535099v1

Submitted on 21 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

HAL Authorization

Vandermonde meets Regev: Public Key Encryption Schemes Based on Partial Vandermonde Problems

Katharina Boudgoust*[1], Amin Sakzad[2], and Ron Steinfeld[2]

katharina.boudgoust@cs.au.dk, amin.sakzad@monash.edu, ron.steinfeld@monash.edu,

[1] Dept. Computer Science, Aarhus University, Aarhus, Denmark
[2] Faculty of Information Technology, Monash University, Australia

Abstract. PASS Encrypt is a lattice-based public key encryption scheme introduced by Hoffstein and Silverman in 2015. The efficiency and algebraic properties of PASS Encrypt and of the underlying partial Vandermonde knapsack problem (PV-Knap) make them an attractive starting point for building efficient post-quantum cryptographic primitives. Recall that PV-Knap asks to recover a polynomial of small norm from a partial list of its Vandermonde transform. Unfortunately, the security foundations of PV-Knap-based encryption are not well understood, and in particular, no security proof for PASS Encrypt is known. In this work, we make progress in this direction. First, we present a modified version of PASS Encrypt with a security proof based on decision PV-Knap and a leaky variant of it, named the PASS problem. We next study an alternative approach to build encryption based on PV-Knap. To this end, we introduce the partial Vandermonde LWE problem (PV-LWE), which we show is computationally equivalent to PV-Knap. Following Regev's design for LWE-based encryption, we use PV-LWE to construct an efficient encryption scheme. Its security is based on PV-LWE and a hybrid variant of PV-Knap and Polynomial LWE. Finally, we give a refined analysis of the concrete security of both schemes against best known lattice attacks.

Keywords: Lattice-Based Cryptography · Public Key Encryption · Partial Vandermonde Problem

1 Introduction

Lattice-based cryptography is a relatively young research field of cryptography that uses conjectured hard problems on Euclidean lattices as the theoretical foundation for cryptographic constructions. One of its main advantages over its number-theoretic counterparts, that rely on the conjectured hardness of integer factorization or the discrete logarithm problem, is its assumed resistance

* Most of the research has been conducted while the first author was still affiliated to Univ Rennes, CNRS, IRISA, Rennes, France.

against quantum attacks. It was initiated at the end of the 1990s by two different branches. On the one hand, there have been proposals benefiting from strong theoretical connections to presumed hard worst-case lattice problems [Ajt96,AD97], leading to the development of cryptography based on the SIS and LWE problems, see for instance Peikert's survey [Pei16]. On the other hand, however, very efficient schemes basing their security on average-case structured lattice problems have been introduced, the most popular among them is the NTRU encryption scheme by Hoffstein et al. [HPS98].

Following the latter approach, Hoffstein et al. [HPS+14] proposed the signature scheme PASS Sign, whose security is based on the difficulty of recovering a polynomial of small norm having access only to a partial list of its Fourier transform. Later, Lu et al. [LZA18] complemented the proposal by moving from the partial Fourier transform (evaluation at all roots of unity) to the partial Vandermonde transform (evaluation only at the *primitive* roots of unity) and by giving a rigorous proof of security.

The problem that underlies PASS Sign, as given in [LZA18], can be presented as follows. Let q be a prime, $R = \mathbb{Z}[x]/\langle\phi(x)\rangle$ denote the ring of integers of the ν-th cyclotomic number field of degree $\deg(\phi(x)) = n$, such that $\phi(x)$ splits into linear factors mod q. More precisely, $\phi(x) = \prod_{j\in[n]}(x-\omega_j) \bmod q$, where $\{\omega_j\}_{j\in[n]}$ are the ν-th primitive roots of unity in $\mathbb{Z}_q$. Let $\mathbf{V} = (\omega_j^{k-1})_{j,k\in[n]} \in \mathbb{Z}_q^{n\times n}$ denote the Vandermonde matrix for $\{\omega_j\}_{j\in[n]}$ and $\mathbf{V}_\Omega \in \mathbb{Z}_q^{t\times n}$ be the partial Vandermonde matrix consisting of $t \leq n$ subrows of $\mathbf{V}$ specified by a random subset of t roots $\Omega \subseteq \{\omega_j\}_{j\in[n]}$. Let $\mathbf{f}$ be a ring element of small norm, sampled from some distribution χ_f. Given $\mathbf{V}_\Omega \cdot \mathbf{f} \bmod q$, the problem asks to find $\mathbf{f}$. We call this the partial Vandermonde knapsack problem (PV-Knap).[3] A related problem is partial Vandermonde SIS (PV-SIS), where given $\mathbf{V}_\Omega$ one asks to find a ring element $\mathbf{f}$ of small norm such that $\mathbf{V}_\Omega \cdot \mathbf{f} = \mathbf{0} \bmod q$. This can be formulated as a shortest vector problem over a structured lattice. In parallel, Hoffstein and Silverman [HS15] introduced PASS Encrypt, a public key encryption (PKE) scheme whose computational building blocks are closely related to the ones of PASS Sign. It is very efficient and fulfills additive and (somewhat) multiplicative homomorphic properties. The algebraic structure of PASS Encrypt and of the underlying partial Vandermonde problems makes them a natural starting point for the design of efficient cryptographic primitives. For example, such properties were recently exploited in the context of PASS Sign to construct compact aggregate signature schemes [DHSS20], and it is plausible that combining PASS Encrypt with PASS Sign may form the basis for various compact and efficient privacy-preserving primitives such as group signatures. Unfortunately, the main problem with PASS Encrypt to date is that its security is not well understood, no proof of security was given in [HS15] with respect to the hardness of explicit compu-

[3] Note that $\mathbf{V}_\Omega \cdot \mathbf{f} \bmod q$ is the vector of evaluations $(\mathbf{f}(\omega_j))_{\omega_j\in\Omega} \bmod q$ of the polynomial $\mathbf{f}$ at the roots in Ω; the full vector of evaluations $(\mathbf{f}(\omega_j))_{j\in[n]}$ is also known as the Number Theoretic Transform (NTT) of $\mathbf{f}$ and $\mathbf{f}(\omega_j)$ is also referred to as the j-th NTT *slot* of $\mathbf{f}$.

tational problems, and the scheme is deterministic and thus does not satisfy the standard notion of IND-CPA security.

Contributions In this paper, we make progress on understanding the security of PASS Encrypt and the construction of efficient PKE schemes based on partial Vandermonde problems. We now present a summary of our main contributions. A summary of all our results is also depicted in Figure 1.

Provable Secure PASS Encrypt We present a modification of PASS Encrypt in Section 4 together with a security proof based on the decision PV-Knap problem and a leaky variant of it, that we call the PASS problem. The latter problem captures the fact that a ciphertext of PASS Encrypt consists of several partial Vandermonde transforms of *related* elements. In other words, a successful attacker against PV-Knap can be used to win the IND-CPA security game, but a successful attacker against the IND-CPA security of PASS Encrypt may not be powerful enough to solve PV-Knap. This issue was not addressed before in the original version of PASS Encrypt [HS15]. Furthermore, the original scheme is deterministic and thus cannot be IND-CPA secure. Additionally, it used the Fourier transform similar to older versions of PASS Sign. In our slightly modified version of PASS Encrypt, we first move to the Vandermonde transform, as done for PASS Sign by Lu et al. [LZA18]. This change is motivated by the fact that the discrete Fourier transform always maps the all-1 vector to zero and thus partial Fourier SIS[4] is trivially easy. In contrast, our setting does not allow the same trivial solution to partial Vandermonde SIS. Second, we make the scheme probabilistic by adding random terms to the message. We then give a proof of correctness (Lemma 12) for well-chosen parameters and a proof of security (Lemma 13), assuming the hardness of dec-PV-Knap and PASS. A refined analysis of the concrete security of PASS Encrypt is provided in Section 6.1. In particular, we show a novel attack that we call *plaintext recovering using hints* attack, which takes the structure of PASS Encrypt into account. It is inspired by the recent work of Dachman-Soled et al. [DDGR20] on exploiting hints that are given on a LWE secret or noise. Our complexity estimates for this attack show that it does not reduce the attack complexity below that of previously known lattice attacks on PASS, which increases our confidence in its claimed security against best known lattice attacks.

Partial Vandermonde LWE On the other hand, since the leaky PASS problem is somewhat ad-hoc and may turn out to be easier than PV-Knap, it is natural to ask whether it is possible to construct efficient PKE schemes relying only on the hardness of PV-Knap. Towards this goal, we propose a more natural approach to build PKE based on partial Vandermonde problems. As a first

[4] In Fourier SIS, the Fourier matrix, instead of the Vandermonde matrix, is used. The Fourier matrix consists of the powers of *all* roots of unity, whereas the Vandermonde matrix only contains the powers of all *primitive* roots of unity.

step, we enlarge the landscape of problems related to the partial Vandermonde transform and introduce partial Vandermonde LWE (PV-LWE), the dual problem to PV-Knap. In the following we provide an informal definition of PV-LWE, for more details see Section 3.2. Given a partial Vandermonde matrix $\mathbf{V}_\Omega$, an instance of PV-LWE is given by $(\mathbf{V}_\Omega, \mathbf{b} = \mathbf{V}_\Omega^T \cdot \mathbf{s} + \mathbf{e} \bmod q)$, where $\mathbf{s}$ is an element of $\mathbb{Z}_q^t$ and $\mathbf{e}$ is sampled from a distribution over $\mathbb{Z}^n$ that provides elements of small norm. The search variant asks to find $\mathbf{s}$ and the decision variant asks to distinguish a sample of PV-LWE from $(\mathbf{V}_\Omega, \mathbf{b})$, where $\mathbf{b}$ is sampled uniformly at random over $\mathbb{Z}_q^n$. We prove the equality of PV-Knap and PV-LWE in Section 3.3. Note that when replacing the partial Vandermonde matrix $\mathbf{V}_\Omega$ by a random matrix $\mathbf{A} \in \mathbb{Z}_q^{t \times n}$, we obtain an instance of the standard LWE problem [Reg05] and when replacing it by a (square) matrix of multiplication with respect to a quotient ring $\mathbb{Z}_q[x]/\langle \phi(x) \rangle$ for some polynomial $\phi(x)$, we get an instance of Polynomial LWE (P-LWE) [SSTX09,LPR10].

Partial Vandermonde Regev Encrypt Next, we show how to use PV-LWE to construct a PKE scheme following Regev's design for LWE-based encryption schemes [Reg05]. We call it PV Regev Encrypt and present it in Section 5. In analog to standard Regev-like PKE, the public key is an instance of PV-LWE. To encrypt a message $\mathbf{m}$, a random vector $\mathbf{r}$ of small norm is chosen. The first part $\mathbf{u}$ of a ciphertext $\mathbf{c} = (\mathbf{u}, \mathbf{v})$ is given as an instance of PV-Knap in order to mask the vector $\mathbf{r}$. The second part $\mathbf{v}$ uses the public key and $\mathbf{r}$ to hide the message $\mathbf{m}$ via a sample of P-LWE. Using the algebraic structure of the matrix $\mathbf{V}_\Omega$, we can encrypt an n-bit message. We give a proof of correctness (Lemma 14) and a proof of security (Lemma 15). An analysis of the concrete security of PV Regev Encrypt against best known attacks is provided in Section 6.2. In particular, we show three different attacks that show similarities with those against PASS Encrypt.

Security of PV Regev Encrypt In our security proof of PV Regev Encrypt (Lemma 15), we show that its security is simultaneously based on the hardness of two problems. First, the hardness of the decision variant of PV-LWE (which is equivalent to PV-Knap by our duality result above), and second the hardness of a hybrid variant, which we call the Hybrid-PV-P problem. It consists of an instance of PV-Knap together with an instance of P-LWE, where the underlying secrets are related to each other. We show a sequence of (quantum) hardness reductions that demonstrate, up to a search-to-decision reduction that we leave as an open problem, that the (quantum) hardness of Hybrid-PV-P is (modulo the above search-to-decision relation) implied by the hardness of the PV-Knap problem and the hardness of the standard decision NTRU problem. This gives evidence, as summarized in Figure 1, that the security of PV Regev Encrypt may not require any additional ad-hoc assumption besides the hardness of PV-Knap and the standard NTRU problem.

Partial-NTT P-SIS and P-LWE Problems As an intermediate byproduct of potentially independent interest, we also introduce a natural problem that

we call Partial-NTT P-SIS and its LWE dual counterpart problem Partial-NTT P-LWE, whose hardness we show are implied by the hardness of the PV-Knap and NTRU problems. The Partial-NTT P-SIS problem is a natural relaxation of the standard P-SIS problem, where, given a random $\mathbf{b} \in R_q$, the attacker must find short non-zero $\mathbf{z}_1, \mathbf{z}_2$ such that $\mathbf{z}_1 + \mathbf{b} \cdot \mathbf{z}_2 = \mathbf{0} \bmod \mathcal{I}_{\Omega,q}$ (where Ω is a subset of primitive roots of unity as in the definition of PV-SIS and $\mathcal{I}_{\Omega,q} = \prod_{\omega_j \in \Omega} \langle q, x - \omega_j \rangle$). That is, rather than requiring $\mathbf{z}_1 + \mathbf{b} \cdot \mathbf{z}_2$ to be fully zero in R_q as in the standard P-SIS problem, the requirement for Partial-NTT P-SIS is relaxed to the polynomial $\mathbf{z}_1 + \mathbf{b} \cdot \mathbf{z}_2 \in R_q$ evaluating to zero at the subset of roots of unity in Ω or, equivalently, having zero entries in the NTT slots defined by Ω.

Techniques We now provide some more details on our techniques used in this work. In our design of PV Regev Encrypt, a main technical challenge was to overcome the following leakage problem in PASS Encrypt: for a public key $\mathbf{V}_\Omega \mathbf{f}$, the PASS Encrypt ciphertext contains information not only about $\mathbf{V}_\Omega \mathbf{r}$ for an encryption randomness $\mathbf{r}$, but also about the *complement* Vandermonde transform $\mathbf{V}_{\Omega^c} \mathbf{r}$ of $\mathbf{r}$, where $\Omega^c := \{\omega^j\}_{j \in [n]} \setminus \Omega$. This stems from the fact that knowing the secret key in PASS Encrypt is not sufficient to decrypt the first ciphertext part, so additional information on the chosen randomness has to be provided via the complement partial Vandermonde transform. This additional information, however, can be interpreted as some leakage on the underlying secret randomness. Instead, in our PV Regev Encrypt scheme we use the partial Vandermonde LWE instance and the standard Regev-like PKE design for LWE to overcome this issue. However, a straightforward adaptation of Regev PKE encounters another challenge: for a public key $\mathsf{pk} = \mathbf{V}_\Omega^T \cdot \mathbf{s} + \mathbf{e}$, we can encapsulate a shared key $K \approx \mathbf{r}^T \cdot \mathbf{V}_\Omega^T \cdot \mathbf{s}$ by sending the PV-Knap instance $\mathbf{c} = \mathbf{r}^T \cdot \mathbf{V}_\Omega^T$ in the ciphertext, but this only gives a short one-dimensional shared key K (and hence a short encrypted message). To overcome this efficiency issue, our PV Regev Encrypt scheme uses instead the n-dimensional shared key $K \approx \mathrm{Rot}(\mathbf{r})^T \cdot \mathbf{V}_\Omega^T \cdot \mathbf{s}$, where $\mathrm{Rot}(\mathbf{r})$ is the matrix of multiplication of $\mathbf{r}$. Here, to allow decryption of this n-dimensional shared key, we naively would need to send a long matrix $\mathbf{C} = \mathrm{Rot}(\mathbf{r})^T \cdot \mathbf{V}_\Omega^T$ in the ciphertext. Fortunately, we observe that this is not needed and the short ciphertext $\mathbf{c} = \mathbf{r}^T \cdot \mathbf{V}_\Omega^T$ is sufficient to decrypt. Here, we use the fact that the partial Vandermonde matrix $\mathbf{V}_\Omega$ and the matrix $\mathrm{Rot}(\mathbf{r})$ interact nicely: as shown in Lemma 3, knowing $\mathbf{V}_\Omega \cdot \mathbf{r}$ suffices to efficiently compute $\mathbf{V}_\Omega \cdot \mathrm{Rot}(\mathbf{r})$.

A second technical challenge was to provide a sequence of reductions to assess the hardness of Hybrid-PV-P based on the more standard PV-Knap and NTRU problems. Whereas the proofs of Lemma 6, 7 and 11 (syndrome decoding), of Lemma 8 (dual attack) and Lemma 10 ([SSTX09]) are straightforward adaptions of existing proofs, a novel technique was necessary for the proof of Lemma 9, the reduction from PV-SIS to Partial-NTT P-SIS, assuming the hardness of decision NTRU. The NTRU instance acts as a lossy argument, in order to transform a successful adversary A against the Partial-NTT P-SIS problem to a successful

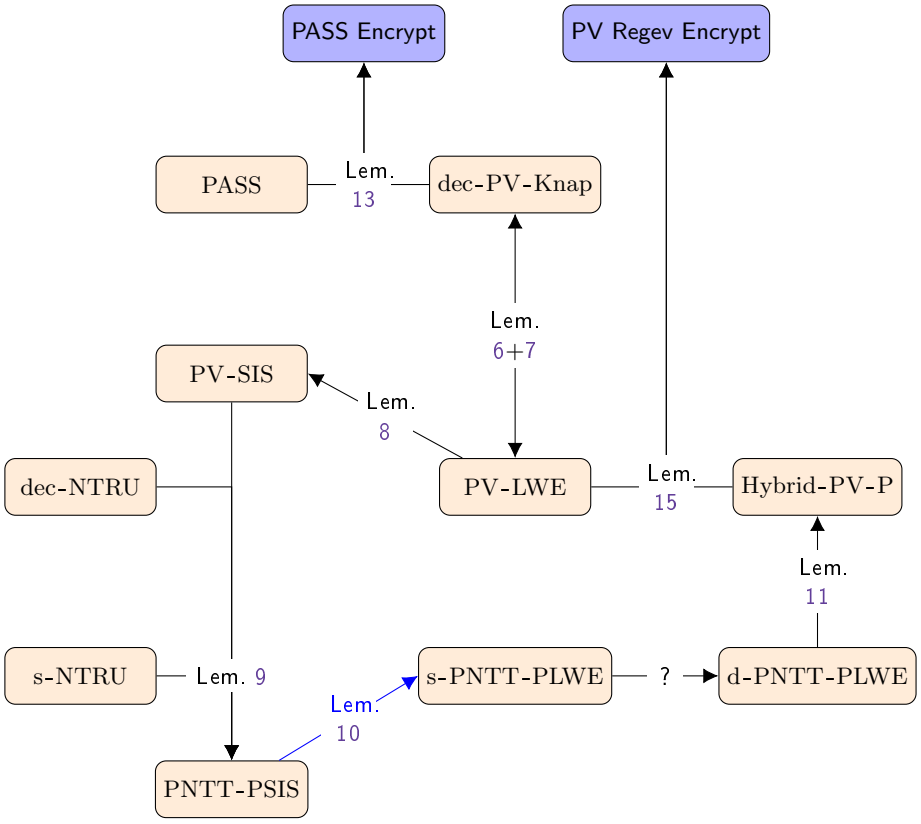

Fig. 1: A summary of all reductions in this paper. The two blue boxes are our proposed encryption schemes. The orange boxes are all different problems being used in this paper. A question mark is used to denote a desired open reduction.

In particular, we show the equivalence of dec-PV-Knap$_\chi$ and PV-LWE$_\chi$, for an error distribution χ in Lemmata 6 and 7. The IND-CPA security of PASS Encrypt is proved based on the hardness of dec-PV-Knap$_{\chi_1}$, dec-PV-Knap$_{\chi_2}$ and PASS$_{\chi',\chi_1,\chi_2}$, where χ' is independent of χ (Lemma 13). The IND-CPA security of PV Regev Encrypt is proved assuming the hardness of PV-LWE$_\chi$ and Hybrid-PV-P$_{\chi,\chi}$ (Lemma 15).

In a sequence of reductions, we show that (A) PV-LWE$_{\chi_\alpha}$ reduces to PV-SIS$_\beta$ in Lemma 8 with β being approximately $\alpha^{-1}/2$ and $\chi_\alpha = D_{\mathbb{Z},\alpha q}$, that (B) PV-SIS$_\beta$, s-NTRU$_{\chi_{\alpha'},\beta'}$, and dec-NTRU$_{\chi_{\alpha'}}$ reduce to PNTT-PSIS$_{\beta'}$ (Lemma 9) with β' being approximately equal to $\beta/(\sqrt{2}\alpha' qn)$ and $\chi_{\alpha'} = D_{\mathbb{Z},\alpha' q}$, that (C) PNTT-PSIS$_{\beta'}$ and s-PNTT-PLWE$_{\chi''_\alpha}$ are related in Lemma 10 with $\alpha''\beta' = \sqrt{n/2}$ and $\chi_{\alpha''} = D_{\alpha'' q}$ and that (D) 'd-PNTT-PLWE$_{\chi''_\alpha}$ reduces to Hybrid-PV-P$_{\chi_{\alpha''},\chi_{\alpha''}}$ (Lemma 11). Colored in blue, Lemma 10 gives a quantum reduction between these problems. All the other reductions are classical.

adversary B against PV-SIS. More precisely, given $\mathbf{V}_\Omega$, B generates an instance of NTRU $\mathbf{b} = \mathbf{g}/\mathbf{f}$, where $\mathbf{g}, \mathbf{f}$ are ring elements of small norm. It then runs A on the input $(\mathbf{V}_\Omega, \mathbf{b})$. The successful adversary A returns a solution $(\mathbf{z}_1, \mathbf{z}_2)$ of small norm such that $\mathbf{z}_1 + \mathbf{b} \cdot \mathbf{z}_2 = \mathbf{0} \bmod \mathcal{I}_{\Omega,q}$. The adversary B can now use the trapdoor $\mathbf{f}$ to compute $\mathbf{z} = \mathbf{f} \cdot (\mathbf{z}_1 + \mathbf{b} \cdot \mathbf{z}_2) = \mathbf{f} \cdot \mathbf{z}_1 + \mathbf{g} \cdot \mathbf{z}_2$. It fulfills $\mathbf{z} = \mathbf{0} \bmod \mathcal{I}_{\Omega,q}$, or equivalently $\mathbf{V}_\Omega \cdot \mathbf{z} = \mathbf{0} \bmod q$ and thus is a valid solution to PV-SIS.

Open Questions From a general perspective, it would be interesting to deepen our understanding of the different partial Vandermonde problems and their relations. In particular, a search-to-decision reduction for PNTT-PLWE (the question mark in Figure 1) would complete the proof of security for PV Regev Encrypt based only on dec-PV-Knap and NTRU problems. An interesting related question is whether NTRU hardness is really necessary for PV Regev Encrypt and how the hardness of NTRU is related to the hardness of dec-PV-Knap. In order to gain more confidence in the PASS problem, it would also be helpful to relate it to a more standard lattice problem. Another interesting open problem is to find worst-case to average-case reductions to PV-Knap and to PV-LWE, respectively. An orthogonal direction is to further investigate the homomorphic properties of PASS Encrypt. In its current state, it is not bootstrappable as the decryption circuit is too deep. A possible lead could be to adapt the flattening technique of Gentry et al. [GSW13] that was used in the setting of LWE-based fully homomorphic encryption. We see at least one obstacle to overcome here: Flattening techniques (relying on (reverse) bit-decomposition) are usually directly applied to the ciphertext, whereas in our case we would like to apply them not on the ciphertext (which consists of elements that lie in the range of $\mathbf{V}_\Omega$ or $\mathbf{V}_{\Omega^c}$) but on the corresponding preimages.

Organization In Section 2 we recall the notions that are necessary for the rest of the paper. In Section 3 we recall known and introduce new problems related to the partial Vandermonde transformation matrix. We propose a provable secure variant of PASS Encrypt in Section 4 and introduce a new encryption scheme that we call PV Regev Encrypt in Section 5. We conclude with a detailed section on the concrete security of both schemes in Section 6.

2 Preliminaries

For any $n \in \mathbb{N}$, we represent the set $\{0, \ldots, n-1\}$ by $[n]$. Let q be a positive integer, then $\mathbb{Z}_q$ is the ring of integers modulo q. Vectors are denoted in bold lowercase and matrices in bold capital letters. We refer to the column vectors and entries of a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ as $\mathbf{A} = (\mathbf{a}_j)_{j \in [n]} = (a_{kj})_{k \in [m], j \in [n]}$. By $\mathbf{A}^T$ we denote the transpose of the matrix $\mathbf{A}$ and by $\mathbf{0}$ the zero vector/matrix. The identity matrix of order n is denoted by $\mathbf{I}_n$. If $\mathbf{a} \in \mathbb{R}^n$, then $\|\mathbf{a}\|_2$ (resp. $\|\mathbf{a}\|_1$ and $\|\mathbf{a}\|_\infty$) denotes the ℓ_2 (resp. ℓ_1 and infinity) norm of $\mathbf{a}$. For two vectors $\mathbf{a}_1, \mathbf{a}_2$ we denote by $\|(\mathbf{a}_1, \mathbf{a}_2)\|_2$ the norm of their (vertically) concatenated vector.

For a set S and a distribution χ over S, we denote by $a \leftarrow \chi$ the process of sampling $a \in S$ according to χ. By $U(S)$ we denote the uniform distribution over S.

For $n \in \mathbb{N}$, we denote by $D_s(\mathbf{x}) = \exp(-\pi\|\mathbf{x}\|_2^2/s^2)/s^n$, the Gaussian distribution with parameter s (and standard deviation $s/\sqrt{2\pi}$) on $\mathbb{R}^n$, and we denote by $D_{\mathbb{Z}^n,s}(\mathbf{x}) = D_s(\mathbf{x})/D_s(\mathbb{Z}^n)$ the corresponding discrete Gaussian distribution on $\mathbb{Z}^n$. For any $a \in \mathbb{R}$, we denote by $\lfloor a \rfloor$ the greatest integer less than or equal to a, which extends component-wise to vectors over $\mathbb{R}$.

2.1 Number Theory

For $\nu \in \mathbb{Z}$, let $K = \mathbb{Q}(\zeta)$ be the ν-th cyclotomic number field of degree $n = \varphi(\nu)$, where $\zeta = \exp(-2\pi i/\nu)$ is a complex primitive ν-th root of unity and φ denotes Euler's totient function. Further, let $R = \mathbb{Z}[\zeta]$ be its ring of integers. The ring R is isomorphic to $\mathbb{Z}[x]/\langle\phi(x)\rangle$, where $\phi(x)$ is the minimal polynomial of ζ. We further denote $R_q = R/qR$. Throughout this work, we identify every ring element $a \in R$ with the coefficient vector $\mathbf{a} = (a_j)_{j\in[n]}$ of the corresponding polynomial $a(x)$ in $\mathbb{Z}[x]/\langle\phi(x)\rangle$, which we call the *coefficient embedding*. Multiplication by $\mathbf{a}$ in R can be represented by a matrix multiplication, where we denote the corresponding matrix by $\mathrm{Rot}(\mathbf{a}) \in \mathbb{Z}^{n\times n}$. For $j \in [n]$, the j-th column of $\mathrm{Rot}(\mathbf{a})$ is given by $x^j \cdot a(x) \bmod \phi(x)$. In particular, for ν a power of two with $n = \nu/2$, it holds $R \cong \mathbb{Z}[x]/\langle x^n + 1\rangle$ and for any $\mathbf{a} = \sum_{j\in[n]} a_j x^j \in R$ it yields

$$\mathrm{Rot}(\mathbf{a}) = \begin{pmatrix} a_0 & -a_{n-1} & \cdots & -a_1 \\ a_1 & a_0 & \ddots & \vdots \\ \vdots & \vdots & \ddots & -a_{n-1} \\ a_{n-1} & a_{n-2} & \cdots & a_0 \end{pmatrix} \in \mathbb{Z}^{n\times n}.$$

We need the following upper bound on the norm of the product of two ring elements in power-of-two cyclotomics.

Lemma 1. *Let n be a power of 2. Further, let $R = \mathbb{Z}[x]/\langle x^n + 1\rangle$. For any two elements $\mathbf{a}, \mathbf{b} \in R$ it yields $\|\mathbf{a} \cdot \mathbf{b}\|_\infty \le 2 \cdot \|\mathbf{a}\|_1 \cdot \|\mathbf{b}\|_\infty$.*

Proof. Using the notion of expansion factor, as introduced in [LM06], we know that in R it holds $\|\mathbf{a} \cdot \mathbf{b}\|_\infty \le 2 \cdot \|\mathbf{a} \star \mathbf{b}\|_\infty$, where $\star$ denotes the convolution product of two polynomials in the ring $\mathbb{Z}[x]$, without modulo $x^n + 1$. It yields

$$\|\mathbf{a} \star \mathbf{b}\|_\infty \le \max_{k\in[2n-1]} \sum_{j\in[k+1]} |a_j| \cdot \|\mathbf{b}\|_\infty \le \|\mathbf{a}\|_1 \cdot \|\mathbf{b}\|_\infty. \quad \square$$

For $d \le n$, we denote $T_n(d)$ the set of ternary polynomials in R with exactly d coefficients that equal 1, d that equal -1 and $n - 2d$ coefficients that equal 0.

Let q be a prime such that $q = 1 \bmod \nu$. In this case, the minimal polynomial $\phi(x)$ of ζ completely splits in $\mathbb{Z}_q[x]$, i.e., $\phi(x) = \prod_{j\in[n]}(x - \omega_j) \bmod q$, where every ω_j is a distinct primitive n-th root of unity in $\mathbb{Z}_q$. Further, we have $\langle q\rangle =$

$\prod_{j\in[n]}\langle q, x-\omega_j\rangle$. In this case, the ν-th cyclotomic number field $K=\mathbb{Q}(\zeta)$ of degree n possesses exactly n field homomorphisms $\sigma_j\colon K\to\mathbb{Z}_q$ for $j\in[n]$ that map ζ to each of the distinct roots ω_j of the minimal polynomial $\phi(x)$. The *canonical embedding*[5] σ is the field homomorphism from K to $\mathbb{Z}_q^n$ defined as $\sigma(a)=(\sigma_j(a))_{j\in[n]}$, where addition and multiplication of vectors in $\mathbb{Z}_q^n$ are performed component-wise. The canonical and coefficient embedding are linked through the Vandermonde matrix of the roots of $\phi(x)$. It yields $\sigma(a)=\mathbf{V}\cdot\mathbf{a} \bmod q$, where $\mathbf{V}=(\omega_j^k)_{j,k\in[n]}\in\mathbb{Z}_q^{n\times n}$.

We divide the set $\{\omega_j\}_{j\in[n]}$ into two random and disjoint subsets Ω and Ω^c of roughly the same size, say $|\Omega|=t=\lfloor n/2\rfloor$ and $\Omega^c=n-t$. The partial Vandermonde transformation matrix $\mathbf{V}_\Omega\in\mathbb{Z}_q^{t\times n}$ and its complement $\mathbf{V}_{\Omega^c}\in\mathbb{Z}_q^{(n-t)\times n}$ are given by $\mathbf{V}_\Omega=(\omega_{i_j}^k)_{j\in[t],k\in[n]}$ and $\mathbf{V}_{\Omega^c}=(\omega_{i_{t+\ell}}^k)_{\ell\in[n-t],k\in[n]}$,

where $\omega_{i_j}\in\Omega$ for $j\in[t]$ and $\omega_{i_{t+k}}\in\Omega^c$ for $k\in[n-t]$. Note that $\mathbf{V}_\Omega$ and $\mathbf{V}_{\Omega^c}$ both have maximal row rank t and $n-t$, respectively. Multiplying the coefficient vector of an element $\mathbf{a}\in R$ by $\mathbf{V}_\Omega$ (resp. $\mathbf{V}_{\Omega^c}$) corresponds to the evaluation of $\mathbf{a}$ at the points ω_{i_j} for $j\in[t]$ (resp. at the points $\omega_{i_{t+k}}$ for $k\in[n-t]$). To ease notations, we omit most of the time the product syntax $\cdot$ and simply write $\mathbf{V}_\Omega\mathbf{a}$ for the matrix-vector product.

Furthermore, $\mathbf{V}_\Omega\mathbf{a}$ is a partial canonical embedding vector of $\mathbf{a}$, as introduced in Section 2.1. Knowing $\mathbf{V}_\Omega\mathbf{a}$ and $\mathbf{V}_{\Omega^c}\mathbf{a}$ gives the complete canonical embedding vector $\sigma(a)=\mathbf{V}\cdot\mathbf{a}\in\mathbb{Z}_q^n$ and thus uniquely identifies it modulo q. The matrix $\mathbf{V}_\Omega$ defines a ring homomorphism from R to $\mathbb{Z}_q^t$, where the latter is equipped with component-wise addition and multiplication, denoted by $+$ and $\circ$.

2.2 Lattices

A (full-rank) n-dimensional lattice Λ is the set of all linear integer combinations of some linearly independent vectors $\mathbf{B}=\{\mathbf{b}_j\}_{j\in[n]}\subseteq\mathbb{R}^n$, i.e., $\Lambda=\sum_{j\in[n]}\mathbb{Z}\cdot\mathbf{b}_j$. The (first) minimum $\lambda_1(\Lambda)$ of the lattice Λ is the Euclidean norm of a shortest non-zero vector in Λ. The dual lattice Λ^* is defined by $\Lambda^*=\{\mathbf{x}\in\mathbb{R}^n\colon\langle\mathbf{x},\mathbf{y}\rangle\in\mathbb{Z},\forall\mathbf{y}\in\Lambda\}$. If $\mathbf{B}$ is a basis of Λ, then $\mathbf{B}^*=(\mathbf{B}^T)^{-1}$ is a basis of Λ^*. One of the most studied lattice problems is the shortest vector problem (SVP). Within this paper we are using the approximate variant of its search version.

Definition 1 (SVP_γ). *Let $\gamma=\gamma(n)\geq 1$ be a function in n. An input to the* approximate shortest vector problem SVP_γ *is a basis $\mathbf{B}$ of an n-dimensional lattice Λ. The goal is to find a vector $\mathbf{z}\neq\mathbf{0}$ such that $\|\mathbf{z}\|_2\leq\gamma\cdot\lambda_1(\Lambda)$.*

Using a suitable embedding, any ideal $\mathcal{I}$ of the ring R (of degree n) defines a lattice in $\mathbb{R}^n$ which we call an *ideal lattice*. The approximate shortest vector problem restricted to ideal lattices is denoted by Id-SVP_γ for γ the approximation factor. We remark that for large approximation factors, as $\gamma=2^{\tilde{O}(\sqrt{n})}$, the Id-$\mathrm{SVP}_\gamma$ problem has been shown to be solvable in polynomial

[5] In other works, this term may also refer to the complex embeddings from K to $\mathbb{C}$.

time [CDPR16,CDW17,PHS19]. However, in this paper we are interested in much smaller approximation factors that are only polynomial in n, for which no efficient attacks are known. We further need a promise lattice problem.

Definition 2 ($\mathbf{BDD}_\delta$). *Let $\mathbf{B}$ be a basis of an n-dimensional lattice Λ and $0 < \delta < \lambda_1(\Lambda(\mathbf{B}))/2$ be a positive real. An input to the* bounded distance decoding problem BDD_δ *is a point $\mathbf{y} \in \mathbb{R}^n$ of the form $\mathbf{y} = \mathbf{x} + \mathbf{e}$, where $\mathbf{x} \in \Lambda(\mathbf{B})$ and $\|\mathbf{e}\|_\infty \leq \delta$. The problem asks to find $\mathbf{x}$.*

2.3 The P-LWE, P-SIS and NTRU Problems

We now recall the Polynomial LWE (P-LWE) problem, as introduced in [SSTX09].

Definition 3 (P-LWE). *Let $q \geq 2$, $n > 0$ and $\phi(x)$ be a polynomial of degree n, defining $R_q = \mathbb{Z}_q[x]/\langle\phi(x)\rangle$. Let χ_e be a distribution over $\mathbb{Z}$ and fix $\mathbf{s} \in R_q$. Let $A_{\phi,\mathbf{s},\chi_e}$ denote the* Polynomial LWE distribution *over $R_q \times R_q$, obtained by sampling $\mathbf{a} \leftarrow U(R_q)$, $\mathbf{e} \leftarrow \chi_e^n$ and returning $(\mathbf{a}, \mathbf{b} = \mathbf{a}\cdot\mathbf{s} + \mathbf{e} \bmod q)$. The* Polynomial LWE problem *comes in two variants:*

> P-SLWE$_{\phi,\chi_e}$*: Let $\mathbf{s} \in R_q$. Given arbitrarily many samples of $A_{\phi,\mathbf{s},\chi_e}$, find $\mathbf{s}$.*
> P-LWE$_{\phi,\chi_e,\chi_s}$*: Let χ_s be a distribution over $\mathbb{Z}$ and sample $\mathbf{s} \leftarrow \chi_s^n$. Distinguish between arbitrarily many independent samples from $A_{\phi,\mathbf{s},\chi_e}$ and the same number of independent samples from $U(R_q \times R_q)$.*

Worst-to-average case reductions guarantee that P-LWE is at least as hard as Id-SVP$_\gamma$ for some approximation factor γ specified by the corresponding reduction [SSTX09,LPR10].

We now recall the Polynomial SIS (P-SIS) problem, as introduced in [LM06] and [PR06]. We specify it to its Hermite normal form and to 2 summands only.

Definition 4 (P-SIS). *Let $q \geq 2$, $n > 0$ and $\phi(x)$ be a polynomial of degree n, defining $R_q = \mathbb{Z}_q[x]/\langle\phi(x)\rangle$. Further, let β be a positive real. The problem* P-SIS$_{q,\beta}$ *is as follows: Given $\mathbf{a} \leftarrow U(R_q)$, find $\mathbf{z}_1, \mathbf{z}_2 \in R$ such that $\mathbf{z}_1 + \mathbf{a}\cdot\mathbf{z}_2 = \mathbf{0} \bmod q$ and $0 < \|(\mathbf{z}_1, \mathbf{z}_2)\|_2 \leq \beta$.*

Worst-to-average case reductions guarantee that P-SIS is at least as hard as Id-SVP$_\gamma$ for some approximation factor γ specified by the reduction.

We now recall the NTRU problem [HPS98] in its decision and search variant. The decision variant has also been called 'Decision Small Polynomial Ratio' (DSPR) problem [LTV12], and 'NTRU Decisional Key Cracking' problem [Ste14]. We remark that for *overstretched* parameter choices, where the modulus q is subexponentially large in the ring dimension, the NTRU problems have been shown to be solvable in polynomial time, e.g. [DvW21]; however, in this paper we are interested in much smaller values of q similar to that used in the original NTRU cryptosystem [HPS98], for which no efficient attacks are known against either search or decision NTRU problems.

Definition 5 (NTRU). *Let $q \geq 2$, $n > 0$ and $\phi(x)$ be a polynomial of degree n, defining $R_q = \mathbb{Z}_q[x]/\langle\phi(x)\rangle$. Let χ be a distribution over $\mathbb{Z}$. Let N_χ denote the* NTRU key distribution *over R_q, obtained by sampling $\mathbf{g} \leftarrow \chi^n$ and $\mathbf{f} \leftarrow (\chi^n \cap R_q^\times)$ (i.e. χ^n restricted to invertible elements of R_q), and returning $\mathbf{b} = \mathbf{g}/\mathbf{f} \in R_q$. The* NTRU problem *comes in two variants:*

s-NTRU$_{\chi,\beta}$*: Given a sample $\mathbf{b} \in R_q$ from N_χ, find $(\mathbf{z}_1, \mathbf{z}_2) \in R^2$ such that $\mathbf{z}_1 + \mathbf{b} \cdot \mathbf{z}_2 = 0 \bmod q$ and $0 < \|(\mathbf{z}_1, \mathbf{z}_2)\|_2 \leq \beta$.*
dec-NTRU$_\chi$*: Given $\mathbf{b} \in R_q$, distinguish between the real case where $\mathbf{b} \leftarrow N_\chi$ and the random case where $\mathbf{b} \leftarrow U(R_q)$.*

2.4 Public Key Encryption

A Public-Key Encryption (PKE) scheme permits two parties to confidentially exchange messages without sharing a common secret key beforehand. In the following, we provide formal definitions of PKE schemes, their correctness and IND-CPA security properties.

Definition 6 (Public Key Encryption). *A* public key encryption *(*PKE*) scheme $\Pi = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ for a message space M and a ciphertext space C is composed of three* PPT *algorithms, specified as follows:*

KGen**:** *The key generation algorithm* KGen *takes as input a security parameter λ and returns a key pair $(\mathsf{sk}, \mathsf{pk})$, called the secret key* sk *and the public key* pk*.*
Enc**:** *The encryption algorithm* Enc *takes as input the public key* pk *and a message $m \in \mathrm{M}$ and returns the ciphertext $c \leftarrow \mathsf{Enc}(\mathsf{pk}, m) \in \mathrm{C}$.*
Dec**:** *The decryption algorithm* Dec *takes as input the secret key* sk *and a ciphertext $c \in C$ and returns $\mathsf{Dec}(\mathsf{sk}, c) \in \mathrm{M} \cup \{\perp\}$, where $\perp$ denotes the failure symbol.*

Definition 7 (Correctness). *We call the* PKE *scheme $\Pi = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$* correct *with correctness error $\delta \in [0, 1)$ if for any message $m \in \mathrm{M}$ it yields*

$$\Pr[\mathsf{Dec}(\mathsf{sk}, \mathsf{Enc}(\mathsf{pk}, m)) = m] \geq 1 - \delta,$$

where the probability is taken over the key pair $(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{KGen}(1^\lambda)$ and the randomness used by the encryption and decryption algorithms. If $\delta = 0$, we say that Π is perfectly correct.

Informally speaking, the security notion IND-CPA captures that no efficient adversary can distinguish between the ciphertext of two messages, where the adversary has even the right to choose the messages by themselves. The acronym stands for *indistinguishable against chosen-plaintext attacks.*

Definition 8 (IND-CPA Security). *We say that the scheme $\Pi = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ is* IND-CPA secure*, if for all* PPT *adversaries $\mathcal{A}$, there is a negligible function $\mathsf{negl}(\cdot)$ such that*

$$\Pr[\mathsf{IND\text{-}CPA}_\Pi^\mathcal{A} = 1] < \frac{1}{2} + \mathsf{negl}(\lambda)\,,$$

where $\mathsf{IND\text{-}CPA}_\Pi^\mathcal{A}$ is the security game from Figure 2.

$\underline{\text{IND-CPA}^{\mathcal{A}}_{\Pi}}$

$b \leftarrow U(\{0,1\})$
$(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KGen}(1^\lambda)$
$(m_0, m_1) \leftarrow \mathcal{A}(1^\lambda, \mathsf{pk})$
$c \leftarrow \mathsf{Enc}(\mathsf{pk}, m_b)$
$b' \leftarrow \mathcal{A}(1^\lambda, \mathsf{pk}, c)$
return $b = b'$

Fig. 2: The IND-CPA security game.

3 Partial Vandermonde Problems

In Section 3.1 we recall the partial Vandermonde knapsack problem (PV-Knap) and its homogeneous counterpart, the partial Vandermonde SIS (PV-SIS) problem. We then define in Section 3.2 the partial Vandermonde LWE (PV-LWE) problem and show the equality between PV-Knap and PV-LWE in Section 3.3. We proceed by presenting two variants of PV-Knap in Section 3.4.

The first variant of PV-Knap, called the PASS problem, is a *leaky* variant of PV-Knap, whose hardness is required in our security analysis of PASS Encrypt. Due to its leaky nature, it is clear that this problem is easier than the standard PV-Knap problem, but unfortunately, we do not know how to relate its hardness to known/standard assumptions; we have to introduce its hardness as an additional ad-hoc assumption needed for the security of PASS Encrypt. In Section 6.1 we provide some refined analysis on the concrete hardness of PASS against best known lattice attacks.

The second variant of PV-Knap, which we call the Hybrid-PV-P problem, can be viewed as a hybrid of PV-Knap with the standard P-LWE problem (see Definition 3). The hardness of Hybrid-PV-P is required in our security analysis of the PV Regev Encrypt scheme in Section 5. In contrast to PASS, however, as outlined in the Introduction and Figure 1, for Hybrid-PV-P we show a sequence of quantum hardness reductions (modulo a search-to-decision reduction for PNTT-PLWE that we leave as an open problem) from the hardness of the PV-Knap problem and the hardness of the decision NTRU problem. In this process, we also introduce the natural Partial- P-SIS problem (and its LWE dual counterpart), which may be of independent interest as a relaxation of the standard P-SIS problem.

3.1 Partial Vandermonde Knapsack and SIS

We now provide definitions of the search and decision versions of the partial Vandermonde knapsack problem (PV-Knap), which was first introduced by Hoffstein

et al. [HPS+14].[6] Throughout the section, we denote by $R = \mathbb{Z}[x]/\langle\phi(x)\rangle$ a cyclotomic ring of degree n. We use the same notation as in Section 2.1 and define the set $\mathcal{P}_t = \{\Omega \subseteq \{\omega_j\}_{j\in[n]} \colon |\Omega| = t\}$ of all subsets of primitive ν-th roots of unity in $\mathbb{Z}_q$ of size t. The corresponding set Ω for an instance of partial Vandermonde knapsack is chosen uniformly at random over $\mathcal{P}_t$.

Definition 9 (PV-Knap). *Let χ denote a distribution over $\mathbb{Z}_q$.*

The search partial Vandermonde knapsack problem PV-Knap$_\chi$ *is the following: let $\mathbf{a} \leftarrow \chi^n$ and $\Omega \leftarrow U(\mathcal{P}_t)$; given $\mathbf{b} = \mathbf{V}_\Omega\mathbf{a} \pmod q$, find the (unique) solution $\mathbf{a}$.*

The decision partial Vandermonde knapsack problem dec-PV-Knap$_\chi$ *asks to distinguish for a given tuple $(\mathbf{V}_\Omega, \mathbf{b}) \in \mathbb{Z}_q^{t\times n} \times \mathbb{Z}_q^t$, where $\Omega \leftarrow U(\mathcal{P}_t)$, if the vector $\mathbf{b}$ was sampled from the uniform distribution over $\mathbb{Z}_q^t$, or if the tuple is given as an* PV-Knap$_\chi$ *instance.*

The problem is assumed to be hard to solve if the distribution χ provides elements of $\mathbb{Z}_q$ with small norms, where smallness is with respect to q. The choice of χ regarding the other parameters also defines whether the solution to this problem is unique or not. We further define a homogeneous variant of PV-Knap with respect to the Euclidean norm.

Definition 10 (PV-SIS). *Let β be a positive integer. The* partial Vandermonde SIS problem PV-SIS$_\beta$ *consists in finding $\mathbf{a} \in R_q$ for a given $\Omega \leftarrow U(\mathcal{P}_t)$ such that $\mathbf{V}_\Omega\mathbf{a} = \mathbf{0}$ and $\|\mathbf{a}\|_2 \le \beta$.*

Similar to previous works [HPS+14,LZA18,DHSS20], we call this problem partial Vandermonde SIS in analogy to the standard short integer solution (SIS) problem, as introduced by Ajtai [Ajt96]. Here, the fully random matrix over $\mathbb{Z}_q$ is replaced by a random structured one, the partial Vandermonde transform $\mathbf{V}_\Omega$. Note that an instance of PV-SIS corresponds to an instance of Id-SVP (Sec. 2.2) in the special class of q-ary ideal lattices

$$\mathcal{I}_{\Omega,q} := \Lambda_q^\perp(\mathbf{V}_\Omega) = \{\mathbf{a} \in R \colon \mathbf{V}_\Omega\mathbf{a} = \mathbf{0} \bmod q\}.$$

We can see that an element $g \in R$ lies in $\mathcal{I}_{\Omega,q}$ if and only if all its evaluations at the roots $\omega_j \in \Omega$ equal 0 mod q. This is equivalent to lying in the ideal $\prod_{\omega_j\in\Omega}\langle q, x-\omega_j\rangle$ (as g must factorize in $g(x) = \prod_{\omega_j\in\Omega}(x-\omega_j)\cdot g'(x) \bmod q$, for some $g' \in R$). Hence, $I_{\Omega,q} = \prod_{\omega_j\in\Omega}\langle x-\omega_j, q\rangle$ (that is the ideal given by half of the factors of the ideal $\langle q\rangle$) and the language of ideal lattices allows us to look at the partial Vandermonde transform in another way. More precisely, for a given ring element $\mathbf{a}$ of small norm, the partial Vandermonde transform $\mathbf{V}_\Omega\mathbf{a} \bmod q$ is a way to specify the coset $\mathbf{a} + \mathcal{I}_{\Omega,q}$. The PV-SIS$_\beta$ problem for $\Omega \in \mathcal{P}_t$ consists in finding a short non-zero ring element $\mathbf{a} \in R$ in this ideal, i.e. satisfying $\mathbf{a} = \mathbf{0} \bmod \mathcal{I}_{\Omega,q}$. Further, the complement partial Vandermonde transform $\mathbf{V}_{\Omega^c}\mathbf{a} \bmod q$ specifies the coset $\mathbf{a}+\mathcal{I}'_{\Omega,q}$, where $\mathcal{I}'_{\Omega,q} := \langle q\rangle\mathcal{I}_{\Omega,q}^{-1}$. Given $\mathbf{a}+\mathcal{I}_{\Omega,q}$ and $\mathbf{a}+\mathcal{I}'_{\Omega,q}$ uniquely defines $\mathbf{a} + \langle q\rangle$.

[6] Even though they originally called it the partial Fourier recovery problem.

Remark 1. In an earlier version of this work, we didn't choose the Ω uniformly at random, but allowed for any fixed subset Ω, similar to earlier works, see for instance [HPS+14,LZA18,DHSS20]. We think that the latter may not be desirable in practice, as it fixes the underlying lattice and hence opens the door to pre-processing attacks, where the attacker focuses all their energy to find a good basis for this special lattice. Furthermore, there might exist bad choices of Ω for which the underlying lattice problems are not hard, and by sampling it at random, we minimize the probability of using such a set. In particular, an example of such a bad choice of Ω is one that contains a large subgroup of $\mathbb{Z}_\nu^\times$. It seems possible to generalize the attack of Pan et al. [PXWC21] against Id-SVP to non-primal ideal lattices to give a polynomial-time algorithm to solve PV-SIS for such choices of Ω [GPM21]; however, the probability of such 'bad' Ω for a uniformly random Ω is negligible for typical parameter choices.

In the case of power-of-2 cyclotomic rings, the equality of the minima of the ideal lattices $\mathcal{I}_{\Omega,q}$ gives the following lower bound on the minimum of such lattices.

Lemma 2 (Adapted from [SS11, Sec.2]). *Let n be a power of 2 and q a prime such that $q = 1 \bmod 2n$ and $\Omega \in \mathcal{P}_t$. The minimum of the ideal lattice $\Lambda_q^\perp(\mathbf{V}_\Omega)$ is lower bounded as*

$$\lambda_1(\Lambda_q^\perp(\mathbf{V}_\Omega)) \geq q^{t/n}.$$

In the following, we show that knowing $\mathbf{b} = \mathbf{V}_\Omega \mathbf{a} \bmod q$ (but not explicitly $\mathbf{a}$) suffices to compute $\mathbf{B} = \mathbf{V}_\Omega \mathrm{Rot}(\mathbf{a}) \bmod q$. In other words, knowing $\mathbf{B}$ does not reveal more information than knowing $\mathbf{b}$. The other direction is trivial, as the first column of $\mathbf{B}$ equals $\mathbf{b}$.

Lemma 3. *Given n, t, q, Ω as above, defining $\mathbf{V}_\Omega = (\omega_{ij}^k)_{j\in[t],k\in[n]}$. Let $\mathbf{b} = \mathbf{V}_\Omega \mathbf{a} \in \mathbb{Z}_q^t$ for some ring element $\mathbf{a} \in R_q$. Then, for $k \in [n]$, the k-th column of the matrix $\mathbf{B} = \mathbf{V}_\Omega \mathrm{Rot}(\mathbf{a})$ is given by $\mathbf{b} \circ (\omega_{i_0}^k, \ldots, \omega_{i_{t-1}}^k)^T$.*

Proof. Let $a(x)$ denote the polynomial that corresponds to $\mathbf{a} \in R_q = \mathbb{Z}_q[x]/\langle\phi(x)\rangle$. For $k \in [n]$, the k-th column of $\mathrm{Rot}(\mathbf{a})$ is given by the coefficient vector of $a(x) \cdot x^k \bmod \phi(x)$. For $\ell \in [t]$, the ℓ-th coefficient of $\mathbf{b}$ corresponds to the evaluation of $a(x)$ at $\omega_{i_\ell} \in \Omega$. We know that $\mathbf{B} = (b_{\ell k})_{\ell\in[t],k\in[n]}$ is given by $b_{\ell k} = (a(x) \cdot x^k)(\omega_{i_\ell})$. Using the homomorphic properties of $\mathbf{V}_\Omega$ completes the proof. □

3.2 Partial Vandermonde LWE

Similar to the standard knapsack problem, we can define a dual problem, called partial Vandermonde (decision) LWE (PV-LWE). Whereas the partial Vandermonde knapsack problem is defined with respect to the matrix $\mathbf{V}_\Omega$, its transpose $\mathbf{V}_\Omega^T$ is used for partial Vandermonde LWE. We use the same notation as in Section 3.1.

Definition 11 (PV-LWE). *Let χ be a distribution over $\mathbb{Z}$ and fix $\mathbf{s} \in \mathbb{Z}_q^t$. Let $B_{\mathbf{s},\chi}$ denote the* partial Vandermonde LWE distribution *over $\mathbb{Z}_q^{t\times n} \times \mathbb{Z}_q^n$, obtained by sampling $\Omega \leftarrow U(\mathcal{P}_t)$, $\mathbf{e} \leftarrow \chi^n$ and returning $(\mathbf{V}_\Omega, \mathbf{b} = \mathbf{V}_\Omega^T\mathbf{s} + \mathbf{e} \bmod q)$. The* partial Vandermonde LWE problem *comes in two variants:*

PV-SLWE$_\chi$**:** *Sample $\mathbf{s} \leftarrow U(\mathbb{Z}_q^t)$. Given a sample of $B_{\mathbf{s},\chi}$, find $\mathbf{s}$.*
PV-LWE$_\chi$**:** *Sample $\mathbf{s} \leftarrow U(\mathbb{Z}_q^t)$. Distinguish between a sample from $B_{\mathbf{s},\chi}$ and a sample of the form $(\mathbf{V}_\Omega, \mathbf{b})$, where $\mathbf{b} \leftarrow U(\mathbb{Z}_q^n)$ and $\Omega \leftarrow U(\mathcal{P}_t)$.*

An instance of PV-LWE defines an instance of BDD (Def. 2) in the q-ary ideal lattice

$$\Lambda_q(\mathbf{V}_\Omega) = \{\mathbf{a} \in R\colon \mathbf{a} = \mathbf{V}_\Omega^T\mathbf{s} \bmod q \text{ for some } \mathbf{s} \in \mathbb{Z}_q^t\}.$$

Lemma 4 below shows that it is not only closed with regard to addition, but also closed with regard to multiplication by any ring element, making it indeed an ideal lattice in R.

Lemma 4. *Let $\mathbf{a} \in \Lambda_q(\mathbf{V}_\Omega)$ and $\mathbf{r} \in R$. Then, $\mathbf{r} \cdot \mathbf{a} \in \Lambda_q(\mathbf{V}_\Omega)$.*

Proof. Let $\mathbf{a} \in \Lambda_q(\mathbf{V}_\Omega)$, i.e., $\mathbf{a} = \mathbf{V}_\Omega^T\mathbf{s} \bmod q$ for a vector $\mathbf{s} \in \mathbb{Z}_q^t$. Note that the multiplication $\mathbf{r} \cdot \mathbf{a}$ is done in $R = \mathbb{Z}[x]/\langle\phi(x)\rangle$. Let $\mathrm{Rot}(\mathbf{r}) = (\tilde{r}_{\ell k})_{\ell,k\in[n]} \in \mathbb{Z}^{n\times n}$ denote the matrix of multiplication of $\mathbf{r}$ in R with respect to the coefficient embedding. Then, it yields for $\ell \in [n]$ that $(\mathbf{r}\cdot\mathbf{a})_\ell = (\mathrm{Rot}(\mathbf{r})\cdot\mathbf{a})_\ell = \sum_{k\in[n]} \tilde{r}_{\ell k}\cdot a_k$. Using that $\mathbf{a} = \mathbf{V}_\Omega^T \cdot \mathbf{s} = (\sum_{j\in[t]} s_j \cdot \omega_{i_j}^k)_{k\in[n]}$ gives for $\ell \in [n]$ that

$$(\mathbf{r}\cdot\mathbf{a})_\ell = \sum_{k\in[n]} \tilde{r}_{\ell k}\left(\sum_{j\in[t]} s_j \cdot \omega_{i_j}^k\right) = \sum_{j\in[t]}\left[\sum_{k\in[n]} \tilde{r}_{\ell k}\omega_{i_j}^{k-\ell}\right]\omega_{i_j}^\ell = \sum_{j\in[t]} s'_j\omega_{i_j}^\ell,$$

where $s'_j = \sum_{k\in[n]} \tilde{r}_{\ell k}\omega_{i_j}^{k-\ell} \in \mathbb{Z}_q$. Thus, $\mathbf{r}\cdot\mathbf{a} = \mathbf{V}_\Omega^T \cdot \mathbf{s}'$, where $\mathbf{s}' = (s'_j)_{j\in[t]}$ and finally $\mathbf{r}\cdot\mathbf{a} \in \Lambda_q(\mathbf{V}_\Omega)$. □

Remark 2. We do *not* define a variant of the PV-LWE problem in the so-called Hermite normal form (HNF), where the secret follows the same distribution as the error, i.e., $\mathbf{s} \leftarrow \chi^t$, since that variant is easy to solve. Namely, let $\mathbf{b} = \mathbf{V}_\Omega^T\mathbf{s}+\mathbf{e}$. As the first column of $\mathbf{V}_\Omega$ is the vector that only contains 1's, the first coefficient of $\mathbf{b}$ is simply the sum of the coefficients of $\mathbf{s}$ plus the first coefficient of $\mathbf{e}$. If $\mathbf{s}$ and $\mathbf{e}$ are small, so is the first coefficient of $\mathbf{b}$, making it trivially distinguishable from uniform.

Remark 3. If we allow for more than one sample of PV-LWE (even with large uniform secret $\mathbf{s}$), the decision problem is easy to solve: Let $\mathbf{b} = \mathbf{V}_\Omega^T\mathbf{s} + \mathbf{e}$ be a first sample and $\mathbf{b}' = \mathbf{V}_{\Omega'}^T\mathbf{s} + \mathbf{e}'$ be a second sample, where $\Omega, \Omega' \leftarrow U(\mathcal{P}_t)$ and $\mathbf{e}, \mathbf{e}' \leftarrow \chi^n$. The first row of $\mathbf{V}_\Omega^T$ and $\mathbf{V}_{\Omega'}^T$ are always the same and thus the difference of the first coefficients of $\mathbf{b}$ and $\mathbf{b}'$ are the difference of the first coefficients of $\mathbf{e}$ and $\mathbf{e}'$ which are unusually small elements. If we delete the first row of $\mathbf{V}_\Omega^T$ (the all-1-row), then the duality of PV-LWE and PV-Knap (see the section below) does not hold anymore.

3.3 Equivalence of Partial Vandermonde Knapsack and LWE

We now show the equality between the PV-Knap and the PV-LWE matrices for the special case of power-of-2 cyclotomics. Let ν be a power of two, $n = \nu/2$ and $t = n/2$. As before, q is set to be a prime such that $q = 1 \bmod \nu$. Further, let ω be a primitive ν-th of unity modulo q and $\gamma = \omega^2$ be a primitive n-th root of unity modulo q. We sample $\Omega \leftarrow U(\mathcal{P}_t)$, and as ν is a power of two, the set Ω (and Ω^c) only contains odd powers of ω.

Additionally, we define $\mathrm{inv}(\mathbf{V}_\Omega) \in \mathbb{Z}_q^{t \times n}$ as the matrix whose (j, ℓ)-th element is the inverse in $\mathbb{Z}_q$ of the (j, ℓ)-th element of $\mathbf{V}_\Omega$ for $j \in [t]$ and $\ell \in [n]$. Since $\mathrm{inv}(\mathbf{V}_\Omega)$ is just $\mathbf{V}_{\mathrm{inv}(\Omega)}$, where $\mathrm{inv}(\Omega)$ is the set of inverses of the elements of Ω, we write $\mathbf{V}_{\mathrm{inv}(\Omega)}$ instead of $\mathrm{inv}(\mathbf{V}_\Omega)$ in the following.

Lemma 5. *Given n, q, Ω as above, defining $\mathbf{V}_\Omega$ and $\mathbf{V}_{\Omega^c}$. It yields*

$$\mathbf{V}_\Omega \cdot \mathbf{V}_{\mathrm{inv}(\Omega^c)}^T = \mathbf{V}_{\Omega^c} \cdot \mathbf{V}_{\mathrm{inv}(\Omega)}^T = \mathbf{0} \in \mathbb{Z}_q^{t \times t}.$$

Proof. We only show the first part, as the second follows in an analog manner. For $j, \ell \in [t]$, the entry of the j-th row and the ℓ-th column of the matrix product $\mathbf{V}_\Omega \cdot \mathbf{V}_{\mathrm{inv}(\Omega^c)}^T$ is given by

$$\sum_{k \in [n]} \left(\omega_j \cdot {\omega_\ell}^{-1}\right)^k = \sum_{k \in [n]} \left(\omega^{i_j} \cdot \omega^{-i_\ell}\right)^k = \sum_{k \in [n]} \left(\omega^u\right)^k = \sum_{k \in [n]} \left(\gamma^v\right)^k,$$

where $\omega_j = \omega^{i_j} \in \Omega$ and $\omega_\ell = \omega^{i_\ell} \in \Omega^c$ and thus $0 \neq i_j - i_\ell = u$, with $u = 0 \bmod 2$. We can thus write $u = 2v$ for some non-zero integer v and deduce that $\omega^u = \gamma^v$. The last geometric sum $T := \sum_{k \in [n]} (\gamma^v)^k$ satisfies $(1 - \gamma^v) \cdot T = 1 - (\gamma^v)^n$. As γ is a primitive n-th root of unity and $0 < v < n$, we have that γ^v is an n-th root of unity and $\gamma^v \neq 1$, so the last sum $T = 0$. □

In the proof of Lemma 5, we use the orthogonality of the power bases in the power-of-2 cyclotomic field. It is easy to see that for other cyclotomic fields, the above statement is not true. We now show the equivalence of dec-PV-Knap and PV-LWE, both in their decision variant. The equivalence of their search versions follows in an analog manner. The proof is essentially the same as for standard knapsack and LWE, see for instance [MM11, Sec. 4.2], combined with the orthogonality Lemma 5. Essentially, the lemma above showed that $\mathbf{V}_{\mathrm{inv}(\Omega)}^T$ is the kernel of the ring homomorphism defined by $\mathbf{V}_\Omega$. We use this to explicate the ideal in which we instantiate the BDD problem given by PV-LWE. In Section 6.4, we show how to interpret their duality in terms of error-correcting codes.

Lemma 6. *Let n be a power of two, $t = n/2$ and $\Omega \leftarrow U(\mathcal{P}(t))$. Further, let χ denote a distribution over $\mathbb{Z}$. There is a* PPT *reduction from the problem* dec-PV-Knap$_\chi$ *to* PV-LWE$_\chi$.

Proof. Assume that we are given an efficient algorithm A for PV-LWE$_\chi$. We now build an efficient algorithm B for dec-PV-Knap$_\chi$ that is given an instance $(\mathbf{V}_\Omega, \mathbf{b})$ with $\mathbf{b} = \mathbf{V}_\Omega \mathbf{a} \bmod q$ or $\mathbf{b} \leftarrow U(\mathbb{Z}_q^t)$, where $\mathbf{a} \leftarrow \chi^n$ and $\Omega \leftarrow U(\mathcal{P}_t)$.

Set $\mathbf{V}_{\Omega'} = \mathbf{V}_{\mathrm{inv}(\Omega^c)}$ and note that if $\Omega \leftarrow U(\mathcal{P}_t)$, then $\mathrm{inv}(\Omega^c)$ also follows the uniform distribution over $\mathcal{P}_t$, where $\mathrm{inv}(\Omega^c) = \{\omega^{-j} = \omega^{\nu-j} : \omega^j \in \Omega^c\} \subseteq \{\omega^k : k \in \mathbb{Z}_\nu^\times\}$. Thus, $\mathbf{V}_{\Omega'}$ defines a valid matrix for PV-LWE. The algorithm B for dec-PV-Knap samples $\mathbf{s} \leftarrow U(\mathbb{Z}_q^t)$ and picks an arbitrary preimage $\mathbf{v} \in R$ of $\mathbf{b}$ under $\mathbf{V}_\Omega$, i.e., $\mathbf{V}_\Omega \mathbf{v} = \mathbf{b} \bmod q$ (such preimage exists since $\mathbf{V}_\Omega$ is of full rank t over $\mathbb{Z}_q$). The algorithm B then sets $\mathbf{b}' = \mathbf{V}_{\Omega'}^T \cdot \mathbf{s} + \mathbf{v}$ and runs A on input $(\mathbf{V}_{\Omega'}, \mathbf{b}')$, returning whatever A returns. We now argue that in the real case $\mathbf{b} = \mathbf{V}_\Omega \cdot \mathbf{a} \bmod q$, then $(\mathbf{V}_{\Omega'}, \mathbf{b}')$ is a valid real case instance of PV-LWE$_\chi$. Indeed, we know that $\mathbf{v} = \mathbf{v}' + \mathbf{a}$ for some $\mathbf{v}' \in R$ with $\mathbf{V}_\Omega \cdot \mathbf{v}' = \mathbf{0} \bmod q$. Using Lemma 5, we have $\mathbf{V}_\Omega \cdot \mathbf{V}_{\Omega'}^T = \mathbf{0} \bmod q$ and thus, $\mathbf{v}'$ has to be in the image of $\mathbf{V}_{\Omega'}^T$ and hence $\mathbf{v}' = \mathbf{V}_{\Omega'}^T \cdot \mathbf{s}'$, for some $\mathbf{s}' \in \mathbb{Z}_q^t$ and finally $\mathbf{b}' = \mathbf{V}_{\Omega'}^T \cdot (\mathbf{s}+\mathbf{s}') + \mathbf{a}$, so $\mathbf{b}'$ has the correct real case PV-LWE distribution with secret $\mathbf{s}'' := \mathbf{s}+\mathbf{s}'$ uniformly random in $\mathbb{Z}_q^t$ (thanks to the uniformly random choice of $\mathbf{s}$) and error $\mathbf{a}$ sampled from χ^n, as required. In the random case $\mathbf{b} \leftarrow U(\mathbb{Z}_q^t)$, $\mathbf{b}'$ is uniformly random in $\mathbb{Z}_q^n$ since $\mathbf{v}$ is in the uniformly random coset of $\Lambda_q^\perp(\mathbf{V}_\Omega)$ defined by $\mathbf{b}$, and $\mathbf{b}'$ is also uniformly random in this coset thanks to the uniformly random $\mathbf{s}'$. Therefore, the advantage of B is the same as that of A. □

Lemma 7. *Let n be a power of two, $t = n/2$ and $\Omega \leftarrow U(\mathcal{P}(t))$. Further, let χ denote a distribution over $\mathbb{Z}$. There is a* PPT *reduction from the problem* PV-LWE$_\chi$ *to* dec-PV-Knap$_\chi$.

Proof. Given an efficient algorithm A for dec-PV-Knap$_\chi$, we build an efficient algorithm B for PV-LWE$_\chi$, that is given an instance $(\mathbf{V}_\Omega, \mathbf{b})$ with either $\mathbf{b} = \mathbf{V}_\Omega^T \cdot \mathbf{s} + \mathbf{e} \bmod q$ or $\mathbf{b} \leftarrow U(\mathbb{Z}_q^n)$, where $\Omega \leftarrow U(\mathcal{P}_t)$, $\mathbf{s} \leftarrow U(\mathbb{Z}_q^t)$ and $\mathbf{e} \leftarrow \chi^n$. With the same argumentation as above, we define $\mathbf{V}_{\Omega'} = \mathbf{V}_{\mathrm{inv}(\Omega^c)}$ fulfilling $\mathbf{V}_{\Omega'} \cdot \mathbf{V}_\Omega^T = \mathbf{0} \bmod q$ and following the uniform distribution over $\mathcal{P}_t$. The algorithm B then computes $\mathbf{b}' = \mathbf{V}_{\Omega'} \cdot \mathbf{b}$ and runs A on input $(\mathbf{V}_{\Omega'}, \mathbf{b}')$, returning whatever A returns. We now argue that in the real case $\mathbf{b} = \mathbf{V}_\Omega^T \cdot \mathbf{s} + \mathbf{e} \bmod q$, then $(\mathbf{V}_{\Omega'}, \mathbf{b}')$ is a valid real case instance of dec-PV-Knap$_\chi$. Indeed, using Lemma 5, we have $\mathbf{b}' = \mathbf{V}_{\Omega'}(\mathbf{V}_\Omega^T \cdot \mathbf{s} + \mathbf{e}) = \mathbf{V}_{\Omega'} \cdot \mathbf{e}$, with $\mathbf{e} \leftarrow \chi^n$. In the random case $\mathbf{b} \leftarrow U(\mathbb{Z}_q^n)$, $\mathbf{b}'$ is uniformly random in $\mathbb{Z}_q^t$, as the matrix $\mathbf{V}_{\Omega'}$ has full rank t. Hence, B has the same advantage as A. □

3.4 Variants of Partial Vandermonde Knapsack

In this section, we introduce two additional variants of the PV-Knap problem needed for the security of our encryption schemes, and present our results on their hardness.

PASS Problem We call this special leaky variant of dec-PV-Knap the PASS problem, as it is used as the underlying hard problem of PASS Encrypt. As opposed to the problems before, it does not only make use of the partial Vandermonde transform $\mathbf{V}_\Omega$, but simultaneously also uses its complement $\mathbf{V}_{\Omega^c}$.

Definition 12 (PASS). *Given n, t, q as above defining the set $\mathcal{P}_t$. Let χ_f, χ_r and χ_s be distributions over $\mathbb{Z}$. The problem* $\mathrm{PASS}_{\chi_f,\chi_r,\chi_s}$ *asks to distinguish the following two cases, when given*

$$(\mathbf{V}_\Omega \mathbf{f}, \mathbf{b}, \mathbf{V}_{\Omega^c}\mathbf{r}, \mathbf{V}_{\Omega^c}\mathbf{s}) \in \left(\mathbb{Z}_q^t\right)^2 \times \left(\mathbb{Z}_q^{n-t}\right)^2.$$

In the first case $\Omega \leftarrow U(\mathcal{P}_t)$, $\mathbf{f} \leftarrow \chi_f^n$, $\mathbf{r} \leftarrow \chi_r^n, \mathbf{s} \leftarrow \chi_s^n$ and $\mathbf{b} = \mathbf{V}_\Omega(\mathbf{f} \cdot \mathbf{r} + \mathbf{s})$. In the second case, the only difference is that $\mathbf{b} \leftarrow U(\mathbb{Z}_q^t)$.

Intuitively, the vector $\mathbf{a} = \mathbf{f} \cdot \mathbf{s} + \mathbf{e}$ can be seen as the secret of an instance $(\mathbf{V}_\Omega, \mathbf{b})$ of the problem dec-PV-Knap, where we are given additional information in the form of $\mathbf{V}_\Omega \mathbf{f}, \mathbf{V}_{\Omega^c}\mathbf{s}$ and $\mathbf{V}_{\Omega^c}\mathbf{e}$, which we interpret as some leakage on the secret $\mathbf{a}$. It is clear that this problem is easier than the standard dec-PV-Knap. In Section 6.1 we provide some refined analysis on the concrete hardness of PASS.

Hybrid-PV-P Problem We call this combination of an instance of Polynomial LWE (P-LWE) and an instance of dec-PV-Knap the Hybrid-PV-P problem, where the underlying secrets of both instances are related to each other. It serves as the underlying hard problem of PV Regev Encrypt in Section 5 and is defined for power-of-two cyclotomics.

Definition 13 (Hybrid-PV-P**).** *Let n be a power two and let q be a prime such that $q = 1 \bmod 2n$, defining $R_q = \mathbb{Z}_q[x]/\langle x^n + 1\rangle$. Further, let χ_r and χ_e be two distributions over $\mathbb{Z}$. The* Hybrid-PV-$\mathrm{P}_{\chi_r,\chi_e}$ *problem asks to distinguish the following two cases, given a sample of the form*

$$(\mathbf{V}_\Omega, \mathbf{b}, \mathbf{u}, \mathbf{v}) \in \mathbb{Z}_q^t \times R_q.$$

In the first case $\Omega \leftarrow U(\mathcal{P}_t)$, $\mathbf{b} \leftarrow U(R_q)$, $\mathbf{u} = \mathbf{V}_\Omega \cdot \mathbf{r}$, where $\mathbf{r} \leftarrow \chi_r^n$, and $\mathbf{v} = \mathbf{b} \cdot \tilde{\mathbf{r}} + \mathbf{e}$, where $\mathbf{e} \leftarrow \chi_e^n$ and $\tilde{\mathbf{r}} = (r_0, -r_{n-1}, \ldots, -r_1)^T$ for $\mathbf{r} = (r_j)_{j\in[n]}$. In the second case, $\mathbf{V}_\Omega$ and $\mathbf{b}$ are as before, but $\mathbf{u} \leftarrow U(\mathbb{Z}_q^t)$ and $\mathbf{v} \leftarrow U(R_q)$.

Let $T\colon R \to R$ be the linear map that sends a ring element $\mathbf{r} = (r_j)_{j\in[n]}$ to the vector $T \cdot \mathbf{r} = (r_0, -r_{n-1}, \ldots, -r_1)^T$. Thus, in the first case $\mathbf{u}$ defines an instance of dec-PV-Knap_{χ_r} and $\mathbf{v}$ defines an instance of P-$\mathrm{LWE}_{x^n+1,\chi_e,\tilde{\chi_r}}$, where $\tilde{\chi_r} = T \cdot \chi_r$ and thus both secrets, $\mathbf{r}$ and $\tilde{\mathbf{r}} = T \cdot \mathbf{r}$, are related to each other. Note that this construction is based on the structure of R as the ring of integers of a power-of-two cyclotomic number field.

We now show a sequence of hardness reductions that implies the hardness of Hybrid-PV-P from PV-LWE and standard NTRU problems, via the natural Partial-P-SIS and P-LWE problems that we introduce below. Refer to Figure 1 for an illustration of these reductions and to Section 2.3 for a reminder of the definitions of NTRU, P-SIS and P-LWE. As a first step, we observe that the well-known 'dual lattice' attack (see, e.g. [MR10, Se.5.4]) on LWE, when specialized to PV-LWE, gives a reduction from PV-LWE to PV-SIS.

Lemma 8 (PV-LWE **to** PV-SIS**).** *Let $\chi_\alpha = D_{\mathbb{Z},\alpha q}$ be a discrete Gaussian noise distribution with parameter αq for some $\alpha \in (0, 1)$. There is a* PPT *reduction from the problem* PV-$\mathrm{LWE}_{\chi_\alpha}$ *to* PV-SIS_β*, with $\beta = \alpha^{-1}/2$.*

Proof. The dual lattice attack (see for instance [MR10]) makes use of a short vector in the dual lattice $\Lambda_q^\perp(\mathbf{V}_\Omega)$. For completeness we summarise the proof below. Namely, given an efficient algorithm A for PV-SIS$_\beta$, we build an efficient algorithm B for PV-LWE$_{\chi_\alpha}$, that works as follows. Given an instance $(\mathbf{V}_\Omega, \mathbf{b})$, the algorithm B runs A on $\mathbf{V}_\Omega$. If A succeeds to compute a short non-zero $\mathbf{v} \in \Lambda_q^\perp(\mathbf{V}_\Omega)$ with $\|\mathbf{v}\|_2 \leq \beta$, B computes $I := \mathbf{v}^T \cdot \mathbf{b} \bmod q \in \mathbb{Z}_q$ and returns 1 if $|I| \leq q/4$ and 0 else. If A fails, B returns a random bit. Note that in case $\mathbf{b}$ comes from the real PV-LWE distribution (i.e., $\mathbf{b} = \mathbf{V}_\Omega^T \cdot \mathbf{s} + \mathbf{e} \mod q$, where $\Omega \leftarrow U(\mathcal{P}_t)$, $\mathbf{s} \leftarrow U(\mathbb{Z}_q^t)$ and $\mathbf{e} \leftarrow \chi_\alpha^n$) then $I = \mathbf{v}^T\mathbf{e}$ has a Gaussian distribution with standard deviation $\alpha q/\sqrt{2\pi} \cdot \|\mathbf{v}\|_2 \leq \beta \cdot \alpha q/\sqrt{2\pi}$, so $I \leq q/4$ with probability $\geq 2/3$ using $\beta = \alpha^{-1}/2$, by a Gaussian tail bound. In the other case that $\mathbf{b}$ comes from the random PV-LWE distribution (i.e., $\mathbf{b}$ is uniformly random in $\mathbb{Z}_q^n$), I is uniformly random in $\mathbb{Z}_q$ so $|I| \leq q/4$ with probability $1/2$. Since $2/3 - 1/2 = \Omega(1)$, the advantage ϵ_B of B is lower bounded as $\Omega(\epsilon_A)$, where ϵ_A is the advantage of A, as required. □

Next, we define Partial-P-SIS, which is a relaxation version of the standard P-SIS problem (Definition 4) as discussed above; instead of requiring $\mathbf{z}_1 + \mathbf{b} \cdot \mathbf{z}_2$ to be $\mathbf{0}$ in R_q, we only require t out of n NTT coefficients to be zero. Note that we define it for two summands only and in the Hermite normal form, but that both restrictions are not necessary in general. In particular, it is straightforward to generalize the definition to $m \geq 2$ summands $\mathbf{b}_1, \ldots, \mathbf{b}_m$.

Definition 14 (PNTT-PSIS**).** *Given $n, t, q, \beta > 0$. The* Partial-NTT P-SIS problem (PNTT-PSIS$_{\mathbf{V}_\Omega,\beta}$) *is defined as follows: Given $\Omega \leftarrow U(\mathcal{P}_t)$ and $\mathbf{b} \leftarrow U(R_q)$, find $\mathbf{z}_1, \mathbf{z}_2 \in R$ such that $\mathbf{V}_\Omega(\mathbf{z}_1 + \mathbf{b} \cdot \mathbf{z}_2) = \mathbf{0} \bmod q$ (i.e. $\mathbf{z}_1 + \mathbf{b} \cdot \mathbf{z}_2 = \mathbf{0} \bmod \mathcal{I}_{\Omega,q}$, where $\mathcal{I}_{\Omega,q} = \prod_{\omega_j \in \Omega} \langle q, x - \omega_j \rangle$), and $0 < \|(\mathbf{z}_1, \mathbf{z}_2)\|_2 \leq \beta$.*

Partial-NTT P-SIS corresponds to solve SVP (Sec. 2.2) in the q-ary module lattice

$$\Lambda_q^\perp(\mathbf{V}_\Omega, \mathbf{b}) = \{\mathbf{a}_1, \mathbf{a}_2 \in R \colon \mathbf{V}_\Omega \cdot (\mathbf{a}_1 + \mathbf{b} \cdot \mathbf{a}_2) = \mathbf{0} \bmod q\}. \tag{1}$$

We now show that, for appropriate parameter choices, the hardness of PV-SIS, together with the hardness of the NTRU problems (from Section 2.3), implies the hardness of PNTT-PSIS.

Lemma 9 (PV-SIS+NTRU **to** PNTT-PSIS**).** *Let $\chi_\alpha = D_{\mathbb{Z},\alpha q}$ be a discrete Gaussian distribution with parameter αq for some $\alpha \in (0, 1)$. There is a* PPT *reduction from either* PV-SIS$_{\beta'}$, dec-NTRU$_{\chi_\alpha}$ *or* s-NTRU$_{\chi_\alpha,\beta}$ *to* PNTT-PSIS$_\beta$ *with $\beta' = \beta \cdot (\alpha q) \cdot \sqrt{2n}$.*

Proof. The proof strategy from a high level perspective is that the NTRU instance acts as a lossy argument in order to transform a successful adversary A against PNTT-PSIS into a successful adversary C against PV-SIS or against a successful adversary D against s-NTRU. Given an efficient algorithm A for PNTT-PSIS$_\beta$ that, given $\Omega \leftarrow U(\mathcal{P}_t)$ and $\mathbf{b} \leftarrow U(R_q)$, returns with non-negligible probability ϵ_A a valid PNTT-PSIS solution $(\mathbf{z}_1, \mathbf{z}_2)$ satisfying $\mathbf{z}_1 + \mathbf{b} \cdot \mathbf{z}_2 = \mathbf{0} \mod \mathcal{I}_{\Omega,q}$ and $0 < \|(\mathbf{z}_1, \mathbf{z}_2)\|_2 \leq \beta$.

First, let PNTT-PSIS$'_{\alpha,\beta}$ denote a modification of PNTT-PSIS$_\beta$, where the distribution of $\mathbf{b}$ is changed to the NTRU key distribution $\mathbf{b} \leftarrow N_{\chi_\alpha}$ (i.e. obtained by sampling $\mathbf{g} \leftarrow \chi_\alpha^n$ and $\mathbf{f} \leftarrow (\chi_\alpha^n \cap R_q^\times)$, and returning $\mathbf{b} = \mathbf{g}/\mathbf{f} \in R_q$), instead of the original uniform distribution on R_q. We claim that the hardness of dec-NTRU$_{\chi_\alpha}$ implies that A has a non-negligible advantage ϵ'_A against PNTT-PSIS$'_{\alpha,\beta}$. Indeed, if ϵ'_A differs non-negligibly from ϵ_A then, since the validity of the PNTT-PSIS solution returned by A is efficiently verifiable, we can construct an efficient distinguisher B against dec-NTRU$_{\chi_\alpha}$ with advantage $\epsilon'_A - \epsilon_A$. Namely, given $\mathbf{b}$, B runs A on $\mathbf{b}$ to get $(\mathbf{z}_1, \mathbf{z}_2)$ and B returns 1 iff $(\mathbf{z}_1, \mathbf{z}_2)$ is a valid PNTT-PSIS$_\beta$ solution.

Second, we construct algorithms C and D against PV-SIS$_{\beta'}$ and s-NTRU$_{\chi_\alpha,\beta}$ respectively. Algorithm C, given $\Omega \leftarrow U(\mathcal{P}_t)$, works as follows: C samples $\mathbf{b} = \mathbf{g}/\mathbf{f} \in R_q$ from N_{χ_α} by sampling $\mathbf{f}, \mathbf{g}$ as above, and runs A on input $(\mathbf{V}_\Omega, \mathbf{b})$ to get $(\mathbf{z}_1, \mathbf{z}_2)$. C then computes and returns $\mathbf{z} = \mathbf{f} \cdot \mathbf{z}_1 + \mathbf{g} \cdot \mathbf{z}_2 \in R$. Algorithm D, given $\mathbf{b} = \mathbf{g}/\mathbf{f} \in R_q$ sampled from N_{χ_α}, runs A on input $(\mathbf{V}_\Omega, \mathbf{b})$ to get $(\mathbf{z}_1, \mathbf{z}_2)$ and returns $(\mathbf{z}_1, \mathbf{z}_2)$.

We claim that $\epsilon_C + \epsilon_D \geq \epsilon'_A - 2^{-n+1}$ (where ϵ_C and ϵ_D denote the advantages of C and D respectively), so that if ϵ'_A is non-negligible, then one of ϵ_C or ϵ_D is non-negligible. Indeed, in the PV-SIS$_{\beta'}$ game of algorithm C, let S denote the event that A outputs a valid PNTT-PSIS$_\beta$ solution, i.e. $(\mathbf{z}_1, \mathbf{z}_2)$ satisfying $\mathbf{z}_1 + \mathbf{b} \cdot \mathbf{z}_2 = \mathbf{0} \mod \mathcal{I}_{\Omega,q}$ and $0 < \|(\mathbf{z}_1, \mathbf{z}_2)\|_2 \leq \beta$. We have $\epsilon'_A = \Pr[S]$. Now let Z denote the event that $\mathbf{z} = \mathbf{f} \cdot \mathbf{z}_1 + \mathbf{g} \cdot \mathbf{z}_2 = \mathbf{0} \in R$. Note that if event $S \cap Z$ occurs, then we have (since $\mathbf{f} \in R_q^\times$) $\mathbf{f}^{-1} \cdot \mathbf{z} \bmod q = \mathbf{z}_1 + \mathbf{b} \cdot \mathbf{z}_2 \bmod q = \mathbf{0}$, so $(\mathbf{z}_1, \mathbf{z}_2)$ is a valid solution to s-NTRU$_{\chi_\alpha,\beta}$, i.e. $\epsilon_D \geq \Pr[S \cap Z]$. On the other hand, if event $S \cap \neg Z$ occurs, then $\mathbf{z} = \mathbf{f} \cdot \mathbf{z}_1 + \mathbf{g} \cdot \mathbf{z}_2$ is a valid solution to PV-SIS$_{\beta'}$ because $\mathbf{z}_1 + \mathbf{b} \cdot \mathbf{z}_2 = \mathbf{0} \mod \mathcal{I}_{\Omega,q}$ so that $\mathbf{z} = \mathbf{f} \cdot (\mathbf{z}_1 + \mathbf{b} \cdot \mathbf{z}_2) = \mathbf{0} \mod \mathcal{I}_{\Omega,q}$, $\mathbf{z} \neq 0$ and by the Schwartz inequality, $\|\mathbf{z}\|_2 \leq \sqrt{n} \cdot \|(\mathbf{f}, \mathbf{g})\|_2 \cdot \|(\mathbf{z}_1, \mathbf{z}_2)\|_2 \leq \sqrt{n} \cdot (\alpha q \sqrt{2n}) \cdot \beta := \beta'$, where we have used the discrete Gaussian tail bound $\|(\mathbf{f}, \mathbf{g})\|_2 \leq (\alpha q \sqrt{2n})$ which holds except with negligible probability $\leq 2^{-n+1}$. It follows that $\epsilon_C + \epsilon_D \geq \epsilon'_A - 2^{-n+1}$, as required. □

We now define a dual problem to the problem above by transposing the matrix $\mathbf{V}_\Omega$, or more precisely the matrix $(\mathbf{V}_\Omega | \mathbf{V}_\Omega \cdot \mathrm{Rot}(\mathbf{b}))$. We call it Partial-NTT P-LWE.

Definition 15 (Partial-NTT P-LWE). *Let χ denote a distribution over $\mathbb{R}$ and fix $\mathbf{s} \in \mathbb{Z}_q^t$. Let $C_{\mathbf{s},\chi}$ denote the* Partial-NTT P-LWE *distribution over $\mathbb{Z}_q^{t \times n} \times (R_q)^3$, obtained by sampling $\Omega \leftarrow U(\mathcal{P}_t)$, $\mathbf{b} \leftarrow U(R_q)$, $\mathbf{e}_1, \mathbf{e}_2 \leftarrow \chi^n$ and returning $(\mathbf{V}_\Omega, \mathbf{b}, \mathbf{y}_1 = \mathbf{V}_\Omega^T \cdot \mathbf{s} + \mathbf{e}_1, \mathbf{y}_2 = \mathrm{Rot}(\mathbf{b}) \cdot \mathbf{V}_\Omega^T \cdot \mathbf{s} + \mathbf{e}_2)$. The* Partial-NTT P-LWE problem *comes in two variants:*

s-PNTT-PLWE$_\chi$*: Let $\mathbf{s} \leftarrow U(\mathbb{Z}_q^t)$. Given a sample of $C_{\mathbf{s},\chi}$, find $\mathbf{s}$.*

d-PNTT-PLWE$_\chi$*: Let $\mathbf{s} \leftarrow U(\mathbb{Z}_q^t)$. Distinguish between a sample from $C_{\mathbf{s},\chi}$ and a sample of the form $(\mathbf{V}_\Omega, \mathbf{b}, \mathbf{y}_1, \mathbf{y}_2)$, where $\mathbf{V}_\Omega$ and $\mathbf{b}$ as before, but now $\mathbf{y}_1, \mathbf{y}_2 \leftarrow U(R_q)$.*

Partial-NTT P-LWE corresponds to solve BDD (Def. 2) in the q-ary module lattice

$$\Lambda_q(\mathbf{V}_\Omega, \mathbf{b}) = \{\mathbf{a}_1, \mathbf{a}_2 \in R\colon \exists \mathbf{s} \in \mathbb{Z}_q^t \text{ s.t. } \mathbf{V}_\Omega^T \mathbf{s} = \mathbf{a}_1 \text{ and } \mathrm{Rot}(\mathbf{b})\mathbf{V}_\Omega^T \mathbf{s} = \mathbf{a}_2 \bmod q\}.$$

As for the PNTT-PSIS problem introduced above, it is possible to define a more general version of PNTT-PLWE, where not only one $\mathbf{b}$ is used, but several $\mathbf{b}_1, \ldots, \mathbf{b}_m$, defining multiple $\mathbf{y}_1, \ldots, \mathbf{y}_{m+1}$. We now show that the quantum reduction of [SSTX09] from P-SIS to search P-LWE can be specialized to the Partial-NTT case.

Lemma 10 (PNTT-PSIS **to** s-PNTT-PLWE**).** *Let* $\chi_\alpha := D_{\alpha q}$ *be a* continuous *Gaussian noise distribution with parameter* $\alpha q < \sqrt{\pi} q^{1-t/n}/(4\sqrt{n \ln(20n)})$. *There is a quantum* PPT *reduction from* PNTT-PSIS$_\beta$ *with* $\beta = \sqrt{n/2} \cdot \alpha^{-1}$ *to* s-PNTT-PLWE$_{\chi_\alpha}$.

Proof. In the following, we need a slightly modified version of PNTT-PLWE, that we call PNTT-PLWE′. The only difference between both variants is the way how $\mathbf{y}_2$ is defined in the plain distribution: Whereas in PNTT-PLWE it is set to $\mathbf{y}_2 = \mathrm{Rot}(\mathbf{b}) \cdot \mathbf{V}_\Omega^T \cdot \mathbf{s} + \mathbf{e}_2$, it is defined as $\mathbf{y}_2' = \mathrm{Rot}(\mathbf{b})^T \cdot \mathbf{V}_\Omega^T \cdot \mathbf{s} + \mathbf{e}_2$ in the latter. Hence, the only difference is that the rotation matrix of $\mathbf{b}$ is transposed in PNTT-PLWE′. By mapping $\mathbf{y}_2$ to $T \cdot \mathbf{y}_2$, where T is as in the definition of Hybrid-PV-P (Def. 13) and assuming that χ is a balanced distribution over $\mathbb{R}$ (i.e., $\chi(-x) = \chi(x)$ for all $x \in \mathbb{R}$), it is easy to see that both variants are equivalent. Here, we use that $\mathrm{Rot}(\mathbf{b})^T = \mathrm{Rot}(T \cdot \mathbf{b})$ for power-of-two cyclotomics. Note that PNTT-PLWE′ corresponds to BDD in $\Lambda_q(\mathbf{V}_\Omega, T \cdot \mathbf{b})$ and that its dual lattice is given by $\Lambda_q^\perp(\mathbf{V}_\Omega, \mathbf{b})$ (see Equation 1).

The proof is almost the same as the proof of [SSTX09, Thm. 3], so we only summarize the main differences here. Given an efficient algorithm A for s-PNTT-PLWE$_{\chi_\alpha}$, we can construct an efficient algorithm A' for s-PNTT-PLWE$'_{\chi_\alpha}$ by using the map T as elaborated above. Note that χ_α is a balanced distribution. This is essentially a BDD algorithm on the lattice $\Lambda_q(\mathbf{V}_\Omega, T \cdot \mathbf{b})$ and the general quantum PPT reduction in [SSTX09, Lem. 9] transforms it into a quantum algorithm B that samples a short vector $\mathbf{v}$ from a distribution with bounded statistical distance from the discrete Gaussian distribution $D_{\Lambda_q^\perp(\mathbf{V}_\Omega, \mathbf{b}), \frac{1}{2\alpha}}$ on the corresponding dual SIS lattice, which in our case is $\Lambda_q^\perp(\mathbf{V}_\Omega, \mathbf{b})$. To apply [SSTX09, Lem. 9], the hypothesis $\alpha q < \lambda/(4\sqrt{n})$ must be satisfied, where $\lambda := \lambda_1(\Lambda_q(\mathbf{V}_\Omega, T \cdot \mathbf{b}))$ is the first minimum of the lattice $\Lambda_q(\mathbf{V}_\Omega, T \cdot \mathbf{b})$. We observe that for any lattice vector $\mathbf{v} \in \Lambda_q(\mathbf{V}_\Omega, T \cdot \mathbf{b})$, the top n coordinates of $\mathbf{v}$ are a vector $\mathbf{v}_1$ in the ideal lattice $\Lambda_q(\mathbf{V}_\Omega)$. Therefore, $\lambda \geq \lambda_1(\Lambda_q(\mathbf{V}_\Omega)) = \lambda_1(\Lambda_q^\perp(\mathbf{V}_{\Omega'})) \geq q^{1-t/n}$ with $\Omega' := \mathrm{inv}(\Omega^c)$, using Lemma 5, and the fact that the minimum of the ideal lattice $\Lambda_q^\perp(\mathbf{V}_{\Omega'})$ is lower bounded by $q^{1-t/n}$ using Lemma 2, since $|\mathrm{inv}(\Omega^c)| = n - t$. Therefore, the hypothesis of [SSTX09, Lem. 9] is satisfied if $\alpha q < q^{1-t/n}/(4\sqrt{n})$. The other condition needed following the proof of [SSTX09, Thm. 3] is that the probability that $\mathbf{v} = 0$ is exponentially small in n; using [SSTX09, Lem. 1] this is true

if $1/(2\alpha) \geq \frac{2\sqrt{\ln(20n)/\pi}}{\lambda^\infty}$, where λ^∞ is the infinity-norm minimum of $\Lambda_q^\perp(\mathbf{V}_\Omega, \mathbf{b})$. The above lower bound on λ gives the lower bound $\lambda^\infty \geq q^{-t/n}/\sqrt{n}$ and therefore, the desired condition $1/(2\alpha) \geq \frac{2\sqrt{\ln(20n)/\pi}}{\lambda^\infty}$ is implied by the condition $(\alpha q) < \sqrt{\pi} q^{1-t/n}/(4\sqrt{n\ln(20n)})$, as required. □

As explained in the introduction, we conjecture that PNTT-PLWE enjoys a search-to-decision reduction from s-PNTT-PLWE to d-PNTT-PLWE. One way to show such a reduction would be to adapt the corresponding reduction shown for P-LWE by Lyubashevsky et al. [LPR13, Sec. 5]. Unfortunately, we currently don't see how to perform this adaption as we meet several issues. First, the secret $\mathbf{s} \in \mathbb{Z}_q^t$ of the underlying problem in our case is, in contrary to P-LWE, not a ring element. This makes it difficult to perform the same trick as in [LPR13, Lem. 5.5], where the full secret can be recovered by learning its values modulo all the different prime ideals appearing in the factorization of $\langle q \rangle$. Second, when amplifying the success probability while performing the worst-case to average-case reduction [LPR13, Lem. 5.12], one has to have access to several samples of the problem. As explained in Remark 3, this is impossible for PV-LWE and hence also for PNTT-PLWE. We thus leave it as an open problem to investigate other search-to-decision techniques, ideally sample-preserving as for instance [MM11].

Using such a search-to-decision reduction, we could complete our sequence of reductions by showing that d-PNTT-PLWE reduces to the Hybrid-PV-P problem, which is needed for the security of our PV Regev Encrypt.

Lemma 11 (d-PNTT-PLWE **to** Hybrid-PV-P**).** *Let χ be a balanced distribution over $\mathbb{Z}$. There is a* PPT *reduction from* d-PNTT-PLWE$_\chi$ *to* Hybrid-PV-P$_{\chi,\chi}$.

Proof. The proof uses the syndrome reduction [MM11], similar to the proof of Lemma 7. Let A be an efficient algorithm for Hybrid-PV-P$_{\chi,\chi}$. In the following, we transform A to an efficient algorithm B for d-PNTT-PLWE$_\chi$ via two intermediate algorithms A' and B'.

We recall that real case instances of Hybrid-PV-P$_{\chi,\chi}$ are of the following form $(\mathbf{V}_{\Omega'}, \mathbf{b}, \mathbf{u}, \mathbf{v})$, where $\Omega' \leftarrow U(\mathcal{P}_t)$, $\mathbf{b} \leftarrow U(R_q)$ and where $\mathbf{u} = \mathbf{V}_{\Omega'} \cdot \mathbf{r}$ and $\mathbf{v} = \mathbf{b} \cdot \tilde{\mathbf{r}} + \mathbf{e} = \mathrm{Rot}(\mathbf{b}) \cdot T \cdot \mathbf{r} + \mathbf{e} = T \cdot \mathrm{Rot}(\mathbf{b})^T \cdot \mathbf{r} + \mathbf{e}$, with $\mathbf{r}, \mathbf{e} \leftarrow \chi^n$, where we used the fact that $\mathrm{Rot}(\mathbf{b}) \cdot T = T \cdot \mathrm{Rot}(\mathbf{b})^T$ and $T^2 = I$, the identity function.

We first build an efficient algorithm A' for a variant of Hybrid-PV-P$_{\chi,\chi}$, that we call Hybrid-PV-P$'_{\chi,\chi}$, where in the latter, real case instances have the form $(\mathbf{V}_{\Omega'}, \mathbf{b}, \mathbf{u}, \mathbf{v})$, where $\mathbf{u} = \mathbf{V}_{\Omega'} \cdot \mathbf{r}$ and $\mathbf{v} = -\mathrm{Rot}(-\mathbf{b})^T \cdot \mathbf{r} + \mathbf{e}$, with $\Omega' \leftarrow U(\mathcal{P}_t)$, $\mathbf{b} \leftarrow U(R_q)$, $\mathbf{e}, \mathbf{r} \leftarrow \chi^n$. The only difference is that in the latter $\mathbf{v}$, there is no T. The algorithm A' maps an input instance $(\mathbf{V}_{\Omega'}, \mathbf{b}, \mathbf{u}, \mathbf{v})$ to $(\mathbf{V}_{\Omega'}, \mathbf{b}, \mathbf{u}, \mathbf{v}')$, where $\mathbf{v}' = T \cdot \mathbf{v}$, runs A and returns whatever A does. Since χ is balanced, T maps χ^n to itself and hence maps the real (resp. random) case instance distribution of Hybrid-PV-P$'_{\chi,\chi}$ to real (resp. random) case instance distribution of Hybrid-PV-P$_{\chi,\chi}$.

Recall that real case instances of d-PNTT-PLWE$_\chi$ are given by some tuple $(\mathbf{V}_\Omega, \mathbf{b}, \mathbf{y}_1, \mathbf{y}_2)$, with $(\mathbf{y}_1, \mathbf{y}_2) = \mathbf{A} \cdot \mathbf{s} + (\mathbf{e}_1, \mathbf{e}_2) \bmod q$, where $\mathbf{A}$ has $\mathbf{V}_\Omega^T$ as

its first n rows and $\mathrm{Rot}(\mathbf{b}) \cdot \mathbf{V}_\Omega^T$ as its last n rows, with $\Omega \leftarrow U(\mathcal{P}_t)$, $\mathbf{b} \leftarrow U(R_q)$ and $\mathbf{e}_1, \mathbf{e}_2 \leftarrow \chi^n$. We now use the efficient algorithm A' to build an efficient algorithm B' for a variant of d-PNTT-PLWE$_\chi$, that we call d-PNTT-PLWE$'_\chi$, where in the latter, real case instances have the form $(\mathbf{V}_\Omega, \mathbf{b}, \mathbf{y}_1, \mathbf{y}_2)$, where $\mathbf{y}_1 = \mathbf{V}_\Omega^T \cdot \mathbf{s} + \mathbf{e}_1$ and $\mathbf{y}_2 = \mathrm{Rot}(\mathbf{b})^T \cdot \mathbf{V}_\Omega^T \cdot \mathbf{s} + \mathbf{e}_2$. Thus, the only difference is that the rotation matrix is now transposed (similar to the proof of Lemma 10).

The reduction from d-PNTT-PLWE$'_\chi$ to Hybrid-PV-P$'_{\chi,\chi}$ runs as follows: given input $(\mathbf{V}_\Omega, \mathbf{b}, \mathbf{y}_1, \mathbf{y}_2)$, B' transforms it to $(\mathbf{V}_{\Omega'}, \mathbf{b}, \mathbf{u}, \mathbf{v})$, where $\Omega' := \mathrm{inv}(\Omega^c)$, with $(\mathbf{u}, \mathbf{v}) = \mathbf{M} \cdot (\mathbf{y}_1, \mathbf{y}_2) \bmod q$, where $\mathbf{M}$ is the $(2n-t) \times 2n$ matrix having as its first $n-t$ rows the submatrix $\mathbf{M}_1 := (\mathbf{V}_{\Omega'} | \mathbf{0}_{(n-t)\times n})$ and as its last n rows the submatrix $\mathbf{M}_2 := (-\mathrm{Rot}(\mathbf{b})^T | \mathbf{I}_n)$. The algorithm B' runs A' on input $(\mathbf{V}_{\Omega'}, \mathbf{b}, \mathbf{u}, \mathbf{v})$ and returns whatever A' returns.

For real case d-PNTT-PLWE$'_\chi$ instances, we have $(\mathbf{y}_1, \mathbf{y}_2) = \mathbf{A} \cdot \mathbf{s} + (\mathbf{e}_1, \mathbf{e}_2) \bmod q$, where $\mathbf{A}$ has $\mathbf{V}_\Omega^T$ as its first n rows and $\mathrm{Rot}(\mathbf{b})^T \cdot \mathbf{V}_\Omega^T$ as its last n rows, with $\mathbf{e}_1, \mathbf{e}_2 \leftarrow \chi^n$. Observe that $\mathbf{M}$ is the check matrix for $\mathbf{A}$, i.e $\mathbf{M} \cdot \mathbf{A} = \mathbf{0} \bmod q$ and $\mathbf{M}$ has full rank $2n-t$ over $\mathbb{Z}_q$. It follows that $(\mathbf{u}, \mathbf{v}) = \mathbf{M} \cdot (\mathbf{y}_1, \mathbf{y}_2) = \mathbf{M} \cdot (\mathbf{e}_1, \mathbf{e}_2) \bmod q$ has the correct real case distribution required by Hybrid-PV-P$'_{\chi,\chi}$. For random case d-PNTT-PLWE$_\chi$ instances, we have $(\mathbf{y}_1, \mathbf{y}_2) \leftarrow U(\mathbb{Z}_q^{2n})$ and so $(\mathbf{u}, \mathbf{v}) = M \cdot (\mathbf{y}_1, \mathbf{y}_2)$ is uniformly random in $\mathbb{Z}_q^{2n-t}$ thanks to the full-rank of M, as required.

Finally we build an efficient algorithm B for d-PNTT-PLWE$_\chi$. The algorithm B maps an input instance $(\mathbf{V}_\Omega, \mathbf{b}, \mathbf{y}_1, \mathbf{y}_2)$ to an instance $(\mathbf{V}_\Omega, \mathbf{b}, \mathbf{y}_1, \mathbf{y}'_2)$, where $\mathbf{y}'_2 = T \cdot \mathbf{y}_2$, runs B' and does whatever B does. Again, since χ is balanced, T maps χ^n to itself and thus maps the real (resp. random) case instance distribution of d-PNTT-PLWE$_\chi$ to the real (resp. random) case instance distribution of d-PNTT-PLWE$'_\chi$. This concludes our proof. □

4 Provable Secure PASS Encrypt

PASS Encrypt is a public key encryption (PKE) scheme introduced by Hoffstein and Silverman in 2015. In its original version, PASS Encrypt comes without a security proof with respect to the hardness of explicit computational problems, and the scheme is deterministic and thus cannot satisfy the standard notion of IND-CPA security.[7] Further, partial Vandermonde SIS over the ring $\mathbb{Z}_q[x]/\langle x^n - 1\rangle$ is easy to solve as we explain in Section 6.3.

In the following, we propose a modified version of PASS Encrypt, see Figure 3. First, we move to the ring $\mathbb{Z}_q[x]/\langle x^n + 1\rangle$. It is the ring of integers of the ν-th cyclotomic number field, where ν is a power of 2 (and thus $n = \nu/2$), and does not allow the same trivial solution to PV-SIS as before. Second, we make the scheme probabilistic by adding the random terms $\mathbf{V}_\Omega(p \cdot \mathbf{s})$ to $\mathbf{e}$ and $\mathbf{V}_{\Omega^c}(p \cdot \mathbf{s})$ to $\mathbf{e}''$, where $\mathbf{s}$ is a random ring element of small norm. We then give a proof of correctness (Lemma 12) and a proof of security (Lemma 13).

[7] A deterministic PKE scheme cannot be IND-CPA secure as an adversary can simply encrypt both messages using the public key and decide which one is used in the challenge ciphertext.

Let λ denote the security parameter. We use the same notation as in Section 3.1, where we assume that ν is a power of 2 and q is a prime such that $q = 1 \bmod \nu$. Further, we set $n = \nu/2$ and $t = n/2$. There are exactly n primitive ν-th roots of unity over $\mathbb{Z}_q$ and for the key generation of our scheme we need to choose at random t (i.e., half) of them. To this end, we denote by $\mathcal{P}_t$ the set of all subsets Ω of size t of all primitive ν-th roots of unity over $\mathbb{Z}_q$. Every Ω defines the corresponding partial Vandermonde matrix $\mathbf{V}_\Omega$ and its complement $\mathbf{V}_{\Omega^c}$. Let χ_f, χ_r, χ_s be distributions over $\mathbb{Z}$. Recall that we denote by $+$ and $\circ$ component-wise addition and multiplication of vectors over $\mathbb{Z}_q$. The message space M is given by $\{0,1\}^n$, and we select a message $\mathbf{m} \in \mathrm{M}$. Finally, let p be a small prime which is coprime to q.

$\mathsf{KGen}(1^\lambda)$.	Sample $\mathbf{f} \leftarrow \chi_f^n$ and $\Omega \leftarrow U(\mathcal{P}_t)$,
	return $\mathsf{sk} = \mathbf{f} \in \mathbb{Z}^n$, $\mathsf{pk} = (\mathsf{pk}_0, \mathsf{pk}_1) = (\Omega, \mathbf{V}_\Omega \cdot \mathbf{f}) \in \mathbb{Z}_q^t \times \mathbb{Z}_q^t$.
$\mathsf{Enc}(\mathsf{pk}, \mathbf{m})$.	Sample $\mathbf{r} \leftarrow \chi_r^n, \mathbf{s} \leftarrow \chi_s^n$,
	set $\mathbf{r}' = p \cdot \mathbf{r}$ and $\mathbf{m}' = p \cdot \mathbf{s} + \mathbf{m}$,
	set $\mathbf{e} = (\mathbf{V}_\Omega \cdot \mathbf{r}' \circ \mathsf{pk}_1) + \mathbf{V}_\Omega \cdot \mathbf{m}'$
	set $\mathbf{e}' = \mathbf{V}_{\Omega^c} \cdot \mathbf{r}'$,
	set $\mathbf{e}'' = \mathbf{V}_{\Omega^c} \cdot \mathbf{m}'$,
	return $\mathbf{c} = (\mathbf{e}, \mathbf{e}', \mathbf{e}'') \in \mathbb{Z}_q^t \times \mathbb{Z}_q^{n-t} \times \mathbb{Z}_q^{n-t}$.
$\mathsf{Dec}(\mathsf{sk}, \mathbf{c})$.	Compute $\mathbf{c}' = (\mathbf{e}' \circ \mathbf{V}_{\Omega^c} \cdot \mathsf{sk}) + \mathbf{e}'' \in \mathbb{Z}_q^{n-t}$,
	combine $\mathbf{e}$ and $\mathbf{c}'$ to obtain $\mathbf{c}''$ as a vector over $\mathbb{Z}_q^n$,
	return $\mathbf{V}^{-1} \cdot \mathbf{c}'' \bmod p$.

Fig. 3: Our slightly modified PASS Encrypt.

We now describe our slightly modified version of PASS Encrypt, as summarized in Figure 3. During key generation, we sample $\mathbf{f}$ from the distribution χ_f^n, defining the secret key $\mathsf{sk} = \mathbf{f} \in \mathbb{Z}^n$. Then, we sample $\Omega \in \mathbb{Z}_q^t$ uniformly at random over the set $\mathcal{P}_t$, which determines the public key $\mathsf{pk} = (\mathsf{pk}_0, \mathsf{pk}_1) = (\Omega, \mathbf{V}_\Omega \mathbf{f}) \in \mathbb{Z}_q^t \times \mathbb{Z}_q^t$, where the second part is the partial Vandermonde transform of $\mathbf{f}$ evaluated at the roots given by Ω. In order to encrypt a message $\mathbf{m} \in \mathrm{M}$, we sample two random vectors $\mathbf{r} \leftarrow \chi_r^n$ and $\mathbf{s} \leftarrow \chi_s^n$, which define $\mathbf{r}' = p\mathbf{r}$ and $\mathbf{m}' = p\mathbf{s} + \mathbf{m}$. This randomizes the message vector and thus converts PASS Encrypt from a deterministic in a randomized scheme. The ciphertext $\mathbf{c}$ is then given by three elements. The first is $\mathbf{e} = (\mathbf{V}_\Omega \mathbf{r}' \circ \mathsf{pk}_1) + \mathbf{V}_\Omega \mathbf{m}' \in \mathbb{Z}_q^t$, using the partial Vandermonde matrix $\mathbf{V}_\Omega$. And the other two are given by $\mathbf{e}' = \mathbf{V}_{\Omega^c} \mathbf{r}' \in \mathbb{Z}_q^{n-t}$ and $\mathbf{e}'' = \mathbf{V}_{\Omega^c} \mathbf{m}' \in \mathbb{Z}_q^{n-t}$, using the complementary partial Vandermonde matrix $\mathbf{V}_{\Omega^c}$. In order to decrypt a ciphertext $\mathbf{c}$, we use the secret key sk to first compute $\mathbf{c}' = (\mathbf{e}' \circ \mathbf{V}_{\Omega^c} \mathsf{sk}) + \mathbf{e}'' \in \mathbb{Z}_q^{n-t}$. Now, using the knowledge of Ω and Ω^c, we can combine $\mathbf{e}$ and $\mathbf{c}'$ to obtain a full vector $\mathbf{c}''$ over $\mathbb{Z}_q^n$. The decryption algorithm then returns $\mathbf{V}^{-1}\mathbf{c}'' \bmod p$. For completeness, we give some sample parameters in Section 7 in Figure 8.

Our version of PASS Encrypt differs in two aspects from the original version as presented in [HS15, Sec. 4]. First, they use the partial Fourier transform (that they denote by $\mathcal{F}_S$) instead of the partial Vandermonde transform $\mathbf{V}_\Omega$, and second, in their case it always yields $\mathbf{s} = \mathbf{0}$ and thus $\mathbf{m}' = \mathbf{m}$. This makes the third part $\mathbf{e}''$ of their ciphertext only dependent on $\mathbf{m}$ and hence the scheme deterministic. Additionally, the modifications also apply to the second version of the original proposed scheme, see [HS15, Sec. 6].

We would like to emphasize the following connection of PASS Encrypt to ideal lattices. As elaborated in Section 3.1, the first part of the public key of PASS Encrypt given by the partial Vandermonde transform $\mathsf{pk}_1 = \mathbf{V}_\Omega \mathbf{f}$ can be seen as a way to specify the coset $\mathbf{f} + \mathcal{I}_{\Omega,q}$, where $\mathcal{I}_{\Omega,q} = \prod_{\omega_j \in \Omega} \langle q, x - \omega_j \rangle$ is an ideal lattice. Simultaneously, the complement partial Vandermonde transforms $\mathbf{V}_{\Omega^c}\mathbf{r}'$ and $\mathbf{V}_{\Omega^c}\mathbf{m}'$ (i.e., the second and third part of the ciphertext of PASS Encrypt) can be seen as a way to specify the cosets $\mathbf{r}' + \mathcal{I}_{\Omega^c,q}$ and $\mathbf{m}' + \mathcal{I}_{\Omega^c,q}$, where $\mathcal{I}_{\Omega,q} \cdot \mathcal{I}_{\Omega^c,q} = \langle q \rangle$. In other words, PASS Encrypt allows a formulation directly in the language of ideal lattices, as we summarize in Figure 4.

$\mathsf{KGen}(1^\lambda)$.	Sample $\mathbf{f} \leftarrow \chi_f^n$ and $\Omega \leftarrow U(\mathcal{P}_t)$, return $\mathsf{sk} = \mathbf{f}$, $\mathsf{pk} = (\mathsf{pk}_0, \mathsf{pk}_1) = (\Omega, \mathbf{f} + I_{\Omega,q}) \in \mathbb{Z}_q^t \times R/\mathcal{I}_{\Omega,q}$.
$\mathsf{Enc}(\mathsf{pk}, \mathbf{m})$.	Sample $\mathbf{r} \leftarrow \chi_r^n, \mathbf{s} \leftarrow \chi_s^n$, set $\mathbf{r}' = p\mathbf{r}$ and $\mathbf{m}' = p\mathbf{s} + \mathbf{m}$, set $\mathbf{e} = ((\mathbf{r}' + I_{\Omega,q}) \cdot \mathsf{pk}_1) + (\mathbf{m}' + I_{\Omega,q})$ set $\mathbf{e}' = \mathbf{r}' + I_{\Omega^c,q}$, set $\mathbf{e}'' = \mathbf{m}' + I_{\Omega^c,q}$, return $\mathbf{c} = (\mathbf{e}, \mathbf{e}', \mathbf{e}'')$.
$\mathsf{Dec}(\mathsf{sk}, \mathbf{c})$.	Compute $\mathbf{c}' = (\mathbf{e}' \cdot (\mathsf{sk} + I_{\Omega^c,q})) + \mathbf{e}''$, combine $\mathbf{e} \in R/\mathcal{I}_{\Omega,q}$ and $\mathbf{c}' \in R/\mathcal{I}_{\Omega^c,q}$ to obtain $\mathbf{c}'' + \langle q \rangle$, return $\mathbf{c}''$.

Fig. 4: Our PASS Encrypt formulated over ideals.

Correctness We now show that the PKE scheme defined above is perfectly correct under a proper choice of parameters. See Definition 7 for the formal statement of the correctness property of a PKE scheme.

Lemma 12 (Correctness). *Let $\mathcal{P}_t, p$ and χ_f, χ_r, χ_s be the public parameters of the scheme* PASS Encrypt. *Assume that there exists $\alpha, \beta > 0$ such that for $\mathbf{f} \leftarrow \chi_f^n, \mathbf{r} \leftarrow \chi_r^n$ and $\mathbf{s} \leftarrow \chi_s^n$ it yields with probability 1 that $\|\mathbf{f}\|_\infty \leq 1, \|\mathbf{r}\|_1 \leq \alpha$ and $\|\mathbf{s}\|_\infty \leq \beta$. Further, we require $p(\alpha + \beta) + 1 < q/2$. For every key pair $(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{KGen}(1^\lambda)$ and message $\mathbf{m} \in \mathrm{M}$, it holds*

$$\Pr\left[\mathsf{Dec}(\mathsf{sk}, \mathsf{Enc}(\mathsf{pk}, \mathbf{m})) = \mathbf{m}\right] = 1.$$

Proof. The decryption oracle first computes $\mathbf{c}' = (\mathbf{e}' \circ \mathbf{V}_{\Omega^c}\mathsf{sk}) + \mathbf{e}'' \in \mathbb{Z}_q^{n-t}$ and then combines it with $\mathbf{e} \in \mathbb{Z}_q^t$ in order to obtain a full vector $\mathbf{c}'' \in \mathbb{Z}_q^n$. To guarantee correctness, we need to make sure that $\mathbf{V}^{-1}\mathbf{c}'' = \mathbf{m} \bmod p$. Using the definition of sk, $\mathbf{e}'$ and $\mathbf{e}''$ it yields

$$\mathbf{c}' = (\mathbf{V}_{\Omega^c}\mathbf{r}' \circ \mathbf{V}_{\Omega^c}\mathbf{f}) + \mathbf{V}_{\Omega^c}\mathbf{m}' = \mathbf{V}_{\Omega^c}(\mathbf{r}' \cdot \mathbf{f} + \mathbf{m}').$$

Simultaneously, using the definition of pk and $\mathbf{e}$ it yields

$$\mathbf{e} = (\mathbf{V}_{\Omega}\mathbf{r}' \circ \mathbf{V}_{\Omega}\mathbf{f}) + \mathbf{V}_{\Omega}\mathbf{m}' = \mathbf{V}_{\Omega}(\mathbf{r}' \cdot \mathbf{f} + \mathbf{m}').$$

In both cases we use the homomorphic properties of $\mathbf{V}_\Omega$ and $\mathbf{V}_{\Omega^c}$. Thus, combining both $\mathbf{c}'$ and $\mathbf{e}$ provides $\mathbf{c}'' = \mathbf{V}(\mathbf{r}' \cdot \mathbf{f} + \mathbf{m}')$ and thus $\mathbf{V}^{-1}\mathbf{c}'' = \mathbf{r}' \cdot \mathbf{f} + \mathbf{m}' = p\mathbf{r} \cdot \mathbf{f} + p\mathbf{s} + \mathbf{m} \bmod q$. Hence, $\mathbf{V}^{-1}\mathbf{c}'' \bmod p = \mathbf{m} \bmod p$ if $\|p\mathbf{r} \cdot \mathbf{f} + p\mathbf{s} + \mathbf{m}\|_\infty < q/2$. Using the properties of power-of-2 cyclotomics to bound the infinity norm of the product of two elements as presented in Lemma 1 and that $\mathbf{m} \in \mathrm{M} = \{0,1\}^n$, it yields

$$\begin{aligned}\|p\mathbf{r} \cdot \mathbf{f} + p\mathbf{s} + \mathbf{m}\|_\infty &\leq p\|\mathbf{r} \cdot \mathbf{f}\|_\infty + p\|\mathbf{s}\|_\infty + \|\mathbf{m}\|_\infty \\ &\leq p \cdot \|\mathbf{f}\|_\infty \cdot \|\mathbf{r}\|_1 + p\|\mathbf{s}\|_\infty + \|\mathbf{m}\|_\infty \\ &\leq p\alpha + p\beta + 1.\end{aligned}$$

As we require $p(\alpha + \beta) + 1 < q/2$, the decryption algorithm decrypts correctly with probability 1. □

Security We now prove the security of PASS Encrypt as defined above based on the hardness of PASS and dec-PV-Knap, both problems are defined in Section 3. We use the standard notion of IND-CPA security whose proper definition we recall in Definition 8.

In order to show the IND-CPA security of PASS Encrypt, we use a common game-hopping argument, as summarized in Figure 5. Game 1 corresponds to the our version of PASS Encrypt. In Game 2 we change the definition of $\mathbf{e}$, in Game 3 the one of $\mathbf{e}'$ and last in Game 4 the one of $\mathbf{e}''$. Note that in Game 2, 3 and 4 the decryption algorithm does in general not succeed as the ciphertext parts $\mathbf{e}, \mathbf{e}'$ or/and $\mathbf{e}''$, when chosen uniformly at random, do in general not possess a small preimage under $\mathbf{V}_\Omega$ or $\mathbf{V}_{\Omega^c}$, respectively. For the proof of IND-CPA security, however, this does not pose any problem.

Lemma 13 (Security). *Let $\mathcal{P}_t, p$ and χ_f, χ_r, χ_s be the public parameters of* PASS Encrypt *and the message $\mathbf{m} \in \mathrm{M}$. Assuming the hardness* dec-PV-Knap$_{\chi_1}$, dec-PV-Knap$_{\chi_2}$ *and* PASS$_{\chi_f,\chi_1,\chi_2}$, *where $\chi_1 = p \cdot \chi_r$ and $\chi_2 = p \cdot \chi_s + \mathbf{m}$, the encryption scheme as summarized in Figure 3 is* IND-CPA *secure.*

Proof. Note that Game 1 corresponds to the proposed PASS Encrypt. The only difference between Game 1 and Game 2 is the way how $\mathbf{e}$ is defined. In the first game, it is a partial Vandermonde transform, given by $(\mathbf{V}_\Omega\mathbf{r}' \circ \mathsf{pk}) + \mathbf{V}_\Omega\mathbf{m}' =$

	Game 1	Game 2
KGen:	$\mathsf{sk} = \mathbf{f} \leftarrow \chi_f^n,\ \Omega \leftarrow U(\mathcal{P}_t)$ $\mathsf{pk} = (\mathsf{pk}_0, \mathsf{pk}_1) = (\Omega, \mathbf{V}_\Omega \mathbf{f} \bmod q)$	$\mathsf{sk} = \mathbf{f} \leftarrow \chi_f^n,\ \Omega \leftarrow U(\mathcal{P}_t)$ $\mathsf{pk} = (\Omega, \mathbf{V}_\Omega \mathbf{f} \bmod q)$
Enc:	$\mathbf{r} \leftarrow \chi_r^n, \mathbf{s} \leftarrow \chi_s^n$ $\mathbf{r}' = p\mathbf{r},\ \mathbf{m}' = p\mathbf{s} + \mathbf{m}$ $\mathbf{e} = (\mathbf{V}_\Omega \mathbf{r}' \circ \mathsf{pk}_1) + \mathbf{V}_\Omega \mathbf{m}'$ $\mathbf{e}' = \mathbf{V}_{\Omega^c} \mathbf{r}'$ $\mathbf{e}'' = \mathbf{V}_{\Omega^c} \mathbf{m}'$	$\mathbf{r} \leftarrow \chi_r^n, \mathbf{s} \leftarrow \chi_s^n$ $\mathbf{r}' = p\mathbf{r},\ \mathbf{m}' = p\mathbf{s} + \mathbf{m}$ $\mathbf{e} \leftarrow U(\mathbb{Z}_q^t)$ $\mathbf{e}' = \mathbf{V}_{\Omega^c} \mathbf{r}'$ $\mathbf{e}'' = \mathbf{V}_{\Omega^c} \mathbf{m}'$
	Game 3	Game 4
KGen:	$\mathsf{sk} = \mathbf{f} \leftarrow \chi_f^n,\ \Omega \leftarrow U(\mathcal{P}_t)$ $\mathsf{pk} = (\Omega, \mathbf{V}_\Omega \mathbf{f} \bmod q)$	$\mathsf{sk} = \mathbf{f} \leftarrow \chi_f^n,\ \Omega \leftarrow U(\mathcal{P}_t)$ $\mathsf{pk} = (\Omega, \mathbf{V}_\Omega \mathbf{f} \bmod q)$
Enc:	$\mathbf{r} \leftarrow \chi_r^n, \mathbf{s} \leftarrow \chi_s^n$ $\mathbf{r}' = p\mathbf{r},\ \mathbf{m}' = p\mathbf{s} + \mathbf{m}$ $\mathbf{e} \leftarrow U(\mathbb{Z}_q^t)$ $\mathbf{e}' \leftarrow U(\mathbb{Z}_q^{n-t})$ $\mathbf{e}'' = \mathbf{V}_{\Omega^c} \mathbf{m}'$	$\mathbf{r} \leftarrow \chi_r^n, \mathbf{s} \leftarrow \chi_s^n$ $\mathbf{r}' = p\mathbf{r},\ \mathbf{m}' = p\mathbf{s} + \mathbf{m}$ $\mathbf{e} \leftarrow U(\mathbb{Z}_q^t)$ $\mathbf{e}' \leftarrow U(\mathbb{Z}_q^{n-t})$ $\mathbf{e}'' \leftarrow U(\mathbb{Z}_q^{n-t})$

Fig. 5: Game hopping for IND-CPA security of PASS Encrypt.

$\mathbf{V}_\Omega(\mathbf{r}' \cdot \mathbf{f} + \mathbf{m}')$, and in the second game it is sampled uniformly at random over $\mathbb{Z}_q^t$. Notice that $\mathsf{pk}, \mathbf{e}'$ and $\mathbf{e}''$ are *not* independent from $\mathbf{e}$, but assuming the hardness of $\text{PASS}_{\chi_f, \chi_1, \chi_2}$, with $\chi_1 = p \cdot \chi_r$ and $\chi_2 = p \cdot \chi_s + \mathbf{m}$, an adversary cannot distinguish between the two games.

Now, we are studying the difference between Game 2 and Game 3. Here, the second ciphertext part $\mathbf{e}'$ is replaced by a uniform element over $\mathbb{Z}_q^{n-t}$. We remark, that $\mathbf{e}'$ is independent from the other two ciphertext parts $\mathbf{e}$ and $\mathbf{e}''$ and also independent from the secret key. Thus, assuming the hardness of dec-PV-Knap_{χ_1}, an adversary cannot distinguish Game 2 from Game 3.

The only difference between Game 3 and Game 4 is the definition of $\mathbf{e}''$. With the same argument, they cannot distinguish Game 3 from Game 4, assuming the hardness of dec-PV-Knap_{χ_2}.

In the last Game 4, the ciphertext $\mathbf{c} = (\mathbf{e}, \mathbf{e}', \mathbf{e}'')$ is independent of the message $\mathbf{m}$ and the key pair $(\mathsf{sk}, \mathsf{pk})$. Thus, the adversary has no chance to distinguish two ciphertexts in the IND-CPA security game better than to guess. □

Homomorphic Properties In the following, we show that our slight modifications on PASS Encrypt preserve its additive and multiplicative homomorphic properties, as originally demonstrated by Hoffstein and Silverman [HS15, Sec. 5].

Additive Homomorphic For addition, we can simply sum the different parts of two given ciphertexts to obtain the encryption of the sum of the original messages. To decrypt the sum, we can use the same decryption algorithm as for a single ciphertext. More precisely, given for a fixed key pair $(\mathsf{sk}, \mathsf{pk})$ two ciphertexts $\mathbf{c}_1 = (\mathbf{e}_1, \mathbf{e}_1', \mathbf{e}_1'')$ and $\mathbf{c}_2 = (\mathbf{e}_2, \mathbf{e}_2', \mathbf{e}_2'')$ on two messages $\mathbf{m}_1$ and $\mathbf{m}_2$, where during encryption the random ring elements $\mathbf{r}_1, \mathbf{s}_1$ and $\mathbf{r}_2, \mathbf{s}_2$ were used.

Then, the element $\mathbf{c} = (\mathbf{e}_1 + \mathbf{e}_2, \mathbf{e}'_1 + \mathbf{e}'_2, \mathbf{e}''_1 + \mathbf{e}''_2)$ defines the ciphertext of the message $\mathbf{m} = \mathbf{m}_1 + \mathbf{m}_2$ with encryption randomness $\mathbf{r} = \mathbf{r}_1 + \mathbf{r}_2$ and $\mathbf{s} = \mathbf{s}_1 + \mathbf{s}_2$. Here, we only use the linearity of matrix-vector products.

Multiplicative Homomorphic The situation is slightly more complex for multiplication, where an additional cross term has to be provided in the ciphertext in order to enable the decryption of the product of two ciphertexts. In more details, assume that we are given for a fixed key pair $(\mathsf{sk}, \mathsf{pk})$ two ciphertexts $\mathbf{c}_1 = (\mathbf{e}_1, \mathbf{e}'_1, \mathbf{e}''_1)$ and $\mathbf{c}_2 = (\mathbf{e}_2, \mathbf{e}'_2, \mathbf{e}''_2)$ on two messages $\mathbf{m}_1$ and $\mathbf{m}_2$, where during encryption the random elements $\mathbf{r}_1, \mathbf{s}_1$ and $\mathbf{r}_2, \mathbf{s}_2$ were used. In order to provide enough information to recover the product message $\mathbf{m}_1 \cdot \mathbf{m}_2 \bmod p$, we need to transmit in the ciphertext the respective products $\mathbf{e} = \mathbf{e}_1 \circ \mathbf{e}_2$, $\mathbf{e}' = \mathbf{e}'_1 \circ \mathbf{e}'_2$ and $\mathbf{e}'' = \mathbf{e}''_1 \circ \mathbf{e}''_2$, and additionally a cross term $\mathbf{E} = \mathbf{e}'_1 \circ \mathbf{e}''_2 + \mathbf{e}'_2 \circ \mathbf{e}''_1$. To decrypt, we use $\mathbf{e}, \mathbf{e}', \mathbf{e}'', \mathbf{E}$ and the secret key sk to compute

$$\begin{aligned}
\mathbf{c}' &= \left(\mathbf{e}' \circ (\mathbf{V}_{\Omega^c}\mathsf{sk})^2\right) + (\mathbf{E} \circ \mathbf{V}_{\Omega^c}\mathsf{sk}) + \mathbf{e}'' \\
&= \left(\mathbf{V}_{\Omega^c}\mathbf{r}'_1 \circ \mathbf{V}_{\Omega^c}\mathbf{r}'_2 \circ (\mathbf{V}_{\Omega^c}\mathbf{f})^2\right) + (\mathbf{V}_{\Omega^c}\mathbf{r}'_1 \circ \mathbf{V}_{\Omega^c}\mathbf{m}'_2 \circ \mathbf{V}_{\Omega^c}\mathbf{f}) \\
&+ (\mathbf{V}_{\Omega^c}\mathbf{r}'_2 \circ \mathbf{V}_{\Omega^c}\mathbf{m}'_1 \circ \mathbf{V}_{\Omega^c}\mathbf{f}) + (\mathbf{V}_{\Omega^c}\mathbf{m}'_1 \circ \mathbf{V}_{\Omega^c}\mathbf{m}'_2) \\
&= \mathbf{V}_{\Omega^c}\left((\mathbf{r}'_1\mathbf{r}'_2\mathbf{f}^2) + (\mathbf{r}'_1\mathbf{m}'_2\mathbf{f}) + (\mathbf{r}'_2\mathbf{m}'_1\mathbf{f}) + (\mathbf{m}'_1\mathbf{m}'_2)\right).
\end{aligned}$$

On the other hand, it yields

$$\begin{aligned}
\mathbf{e} &= \mathbf{e}_1 \circ \mathbf{e}_2 \\
&= ((\mathbf{V}_\Omega\mathbf{r}'_1 \circ \mathsf{pk}) + \mathbf{V}_\Omega\mathbf{m}'_1) \circ ((\mathbf{V}_\Omega\mathbf{r}'_2 \circ \mathsf{pk}) + \mathbf{V}_\Omega\mathbf{m}'_2) \\
&= (\mathbf{V}_\Omega(\mathbf{r}'_1\mathbf{f} + \mathbf{m}'_1)) \circ (\mathbf{V}_\Omega(\mathbf{r}'_2\mathbf{f} + \mathbf{m}'_2)) \\
&= \mathbf{V}_\Omega\left((\mathbf{r}'_1\mathbf{r}'_2\mathbf{f}^2) + (\mathbf{r}'_1\mathbf{m}'_2\mathbf{f}) + (\mathbf{r}'_2\mathbf{m}'_1\mathbf{f}) + (\mathbf{m}'_1\mathbf{m}'_2)\right).
\end{aligned}$$

Combining $\mathbf{c}'$ and $\mathbf{e}$ gives the full Vandermonde transform and by applying $\mathbf{V}^{-1}$, we obtain $\mathbf{c}'' = (\mathbf{r}'_1\mathbf{r}'_2\mathbf{f}^2) + (\mathbf{r}'_1\mathbf{m}'_2\mathbf{f}) + (\mathbf{r}'_2\mathbf{m}'_1\mathbf{f}) + (\mathbf{m}'_1\mathbf{m}'_2) \bmod q$. If $\|\mathbf{c}''\|_\infty < q/2$, we can compute $\mathbf{c}'' \bmod p = \mathbf{m}'_1\mathbf{m}'_2 \bmod p = \mathbf{m}_1\mathbf{m}_2 \bmod p$. Here, we use that $\mathbf{r}'_j = p\mathbf{r}_j$ and $\mathbf{m}'_j = p\mathbf{s}_j + \mathbf{m}_j$ for $j \in \{1, 2\}$.

5 PV Regev Encrypt

In this section, we propose PV Regev Encrypt, a PKE scheme following Regev's approach for LWE-based encryption [Reg05], adapted to the partial Vandermonde setting. We first present the scheme, illustrated in Figure 6, then give a proof of correctness (Lemma 14) and a proof of security (Lemma 15).

Let λ denote the security parameter. We use the same notation as in Section 3.1, where we assume that ν is a power of two and q is a prime such that $q = 1 \bmod \nu$, defining $R_q = \mathbb{Z}_q[x]/\langle x^n + 1\rangle$. Further, let $\mathcal{P}_t$ denote the set of all subsets Ω of size t of all primitive ν-th roots of unity over $\mathbb{Z}_q$. The message space M is given by $\{0,1\}^n$, and we select a message $\mathbf{m} \in \mathrm{M}$. Let χ_e and χ_r denote two distributions over $\mathbb{Z}$.

$\mathsf{KGen}(1^\lambda)$.	Sample $\Omega \leftarrow U(\mathcal{P}_t)$, $\mathbf{s} \leftarrow U\left(\mathbb{Z}_q^t\right)$ and $\mathbf{e} \leftarrow \chi_e^n$,
	return $\mathsf{sk} = \mathbf{s} \in \mathbb{Z}_q^t$,
	and $\mathsf{pk} = (\Omega, \mathbf{b} = \mathbf{V}_\Omega^T \cdot \mathbf{s} + \mathbf{e} \bmod q) \in \mathbb{Z}_q^t \times \mathbb{Z}_q^n$.
$\mathsf{Enc}(\mathsf{pk}, \mathbf{m})$.	Sample $\mathbf{r} \leftarrow \chi_r^n$ and $\mathbf{e}' \leftarrow \chi_e^n$,
	set $\mathbf{u} = \mathbf{V}_\Omega \cdot \mathbf{r}$
	set $\mathbf{v} = \mathrm{Rot}(\mathbf{r})^T \cdot \mathbf{b} + \mathbf{e}' + \lfloor q/2 \rfloor \cdot \mathbf{m} \bmod q$,
	return $\mathbf{c} = (\mathbf{u}, \mathbf{v}) \in \mathbb{Z}_q^t \times \mathbb{Z}_q^n$.
$\mathsf{Dec}(\mathsf{sk}, \mathbf{c})$.	Construct $\mathbf{U} = \mathbf{V}_\Omega \cdot \mathrm{Rot}(\mathbf{r})$ out of $\mathbf{u}$,
	compute $\mathbf{c}' = \mathbf{v} - \mathbf{U}^T \cdot \mathbf{s}$.
	For each coefficient c'_k of $\mathbf{c}'$:
	If c'_k is closer to 0 than to $\lfloor q/2 \rfloor$, set $\hat{m}_k = 0$, else $\hat{m}_k = 1$.
	Return $\hat{\mathbf{m}}$.

Fig. 6: PV Regev Encrypt.

In analog to the standard Regev-like PKE, the public key pk is an instance of the problem PV-LWE and the secret key sk contains the underlying secret $\mathbf{s}$ and noise $\mathbf{e}$. To encrypt a message $\mathbf{m}$, a random vector $\mathbf{r}$ of small norm is chosen. The first part $\mathbf{u}$ of a ciphertext $\mathbf{c}$ is given as an instance of PV-Knap in order to mask the vector $\mathbf{r}$. The second part $\mathbf{v}$ uses the public key $\mathbf{b}$ and the randomness $\mathbf{r}$ to hide a message $\mathbf{m}$ via a sample of P-LWE. Using the structure of the partial Vandermonde transform $\mathbf{V}_\Omega$, we can encrypt an n-dimensional binary message vector. The security of this scheme is simultaneously based on the hardness of PV-LWE and Hybrid-PV-P, a hybrid variant of dec-PV-Knap together with an instance of P-LWE, where the underlying secrets are related to each other. For a more detailed discussion on this variant, we refer to Section 3.4. For completeness, we give some sample parameters in Section 7 in Figure 9.

Correctness We now show that the PKE scheme defined above is perfectly correct under a proper choice of parameters. See Definition 7 for the formal statement of the correctness property of a PKE scheme.

Lemma 14 (Correctness). *Let $\mathcal{P}_t$ and χ_e, χ_r be the public parameters of* PV Regev Encrypt *as presented in Figure 6. Assume that there exists $\alpha, \beta > 0$, such that for $e \leftarrow \chi_e$ and $r \leftarrow \chi_r$ it yields with probability 1 that $|e| \leq \alpha$ and that $|r| \leq \beta$. Further, we require $\alpha(n\beta + 1) < q/4$. For every key pair $(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{KGen}(\lambda)$ and message $\mathbf{m} \in \mathrm{M}$ it holds*

$$\Pr[\mathsf{Dec}(\mathsf{sk}, \mathsf{Enc}(\mathsf{pk}, \mathbf{m})) = \mathbf{m}] = 1.$$

Proof. As shown in Lemma 3, in order to construct the k-th column of $\mathbf{U} = \mathbf{V}_\Omega \cdot \mathrm{Rot}(\mathbf{r})$ out of $\mathbf{u}$ for $k \in [n]$, we multiply component-wise $\mathbf{u}$ with the vec-

tor $(\omega_{i_0}^k, \ldots, \omega_{i_{t-1}}^k)^T$. It yields

$$\begin{aligned}\mathbf{c}' &= \mathbf{v} - \mathbf{U}^T \cdot \mathbf{s} \\ &= \mathrm{Rot}(\mathbf{r})^T \cdot \mathbf{b} + \mathbf{e}' + \lfloor q/2 \rfloor \cdot \mathbf{m} - \mathrm{Rot}(\mathbf{r})^T \cdot \mathbf{V}_\Omega^T \cdot \mathbf{s} \\ &= \mathrm{Rot}(\mathbf{r})^T \cdot (\mathbf{V}_\Omega^T \cdot \mathbf{s} + \mathbf{e}) + \mathbf{e}' + \lfloor q/2 \rfloor \cdot \mathbf{m} - \mathrm{Rot}(\mathbf{r})^T \cdot \mathbf{V}_\Omega^T \cdot \mathbf{s} \\ &= \mathrm{Rot}(\mathbf{r})^T \cdot \mathbf{e} + \mathbf{e}' + \lfloor q/2 \rfloor \cdot \mathbf{m}.\end{aligned}$$

For $k \in [n]$, the k-th row of $\mathrm{Rot}(\mathbf{r})^T$ and the k-th coefficient of $\mathbf{e}'$ satisfy

$$|(\mathrm{Rot}(\mathbf{r})^T)_k \cdot \mathbf{e} + e'_k| \leq \|(\mathrm{Rot}(\mathbf{r})^T)_k\|_2 \cdot \|\mathbf{e}\|_2 + |e'_k| \leq \sqrt{n}\beta \cdot \sqrt{n}\alpha + \alpha,$$

where we use the special form of the rotation matrix in power-of-two cyclotomics. As we require $\alpha(n\beta + 1) < q/4$, the decryption algorithm always succeeds. □

Security We now prove the security of PV Regev Encrypt as defined above based on the hardness of PV-LWE and Hybrid-PV-P, both problems are defined in Section 3. We use the standard notion of IND-CPA security whose proper definition we recall in Definition 8.

Again, we use a standard game-hopping argument, as summarized in Figure 7. Game 1 corresponds to the proposed PKE scheme. Note that in Game 2 and 3 the decryption algorithm does in general not succeed anymore. For the proof of IND-CPA security, however, this does not pose any problem.

	Game 1	Game 2	Game 3
KGen:	$\mathsf{sk} = (\mathbf{s}, \mathbf{e})$	$\mathsf{sk} = (\mathbf{s}, \mathbf{e})$	$\mathsf{sk} = (\mathbf{s}, \mathbf{e})$
	$\mathbf{b} = \mathbf{V}_\Omega^T \cdot \mathbf{s} + \mathbf{e}$	$\mathbf{b} \leftarrow U(\mathbb{Z}_q^n)$	$\mathbf{b} \leftarrow U(\mathbb{Z}_q^n)$
	$\mathsf{pk} = (\mathbf{V}_\Omega, \mathbf{b})$	$\mathsf{pk} = (\mathbf{V}_\Omega, \mathbf{b})$	$\mathsf{pk} = (\mathbf{V}_\Omega, \mathbf{b})$
Enc:	$\mathbf{u} = \mathbf{V}_\Omega \cdot \mathbf{r}$	$\mathbf{u} = \mathbf{V}_\Omega \cdot \mathbf{r}$	$\mathbf{u} \leftarrow U(\mathbb{Z}_q^t)$
	$\mathbf{v} = \mathrm{Rot}(\mathbf{r})^T\mathbf{b} + \mathbf{e}' + \lfloor q/2 \rfloor \mathbf{m}$	$\mathbf{v}' = \mathrm{Rot}(\mathbf{r})^T\mathbf{b} + \mathbf{e}'$	$\mathbf{v}' \leftarrow U(R_q)$
		$\mathbf{v} = \mathbf{v}' + \lfloor q/2 \rfloor \mathbf{m}$	$\mathbf{v} = \mathbf{v}' + \lfloor q/2 \rfloor \mathbf{m}$

Fig. 7: Game hopping for IND-CPA security of PV Regev Encrypt.

Lemma 15 (Security). *Let $\mathcal{P}_t$ and χ_e, χ_r be the public parameters of* PV Regev Encrypt *with message space $\mathrm{M} = \{0,1\}^n$. Assuming the hardness of* PV-LWE$_{\chi_e}$ *and* Hybrid-PV-P$_{\chi_r,\chi_e}$*, the encryption scheme as summarized in Figure 6 is* IND-CPA *secure.*

Proof. Note that Game 1 corresponds to the proposed PV Regev Encrypt. The only difference between Game 1 and Game 2 is how the vector $\mathbf{b}$ is defined. Assuming the hardness of PV-LWE$_{\chi_e}$, the second is computationally indistinguishable from the first one.

In Game 2, we are given $\mathbf{u} = \mathbf{V}_\Omega \cdot \mathbf{r}$ and $\mathbf{v} = \mathrm{Rot}(\mathbf{r})^T \cdot \mathbf{b} + \mathbf{e}' + \lfloor q/2 \rfloor \cdot \mathbf{m}$, where $\mathbf{b}$ is a uniform random ring element. We observe that $\mathrm{Rot}(\mathbf{r})^T \cdot \mathbf{b} = \mathrm{Rot}(\tilde{\mathbf{r}}) \cdot \mathbf{b} = \tilde{\mathbf{r}} \cdot \mathbf{b}$, where $\tilde{\mathbf{r}} = T \cdot \mathbf{r} = (r_0, -r_{n-1}, \ldots, -r_1)^T$ for $\mathbf{r} = (r_j)_{j \in [n]}$. Hence, $\mathbf{v}' = \mathbf{b} \cdot \tilde{\mathbf{r}} + \mathbf{e}'$, which defines an instance of P-LWE with secret $\tilde{\mathbf{r}}$ and error $\mathbf{e}'$. Note that $\mathbf{r}$ is used in $\mathbf{u}$ and $\mathbf{v}$, so we cannot argue with dec-PV-Knap$_{\chi_r}$ and P-LWE$_{x^n+1,\chi_e,\tilde{\chi_r}}$ with $\tilde{\chi_r} = T\chi_r$, independently. We need a hybrid assumption, that one dec-PV-Knap instance *together* with one P-LWE instance, where the underlying secrets depend on each other, are computationally indistinguishable from uniform. This is exactly the Hybrid-PV-P problem that we introduced in Section 3.4. Assuming the hardness of Hybrid-PV-P$_{\chi_r,\chi_e}$, Game 3 and Game 2 are computationally indistinguishable. In the last game, the adversary has no chance to distinguish two ciphertexts better than guessing. □

Remark 4. It is also possible to build a PKE scheme following the dual-Regev framework, in the spirit of [GPV08]. In this setting, the public key consists of an instance of PV-Knap and the ciphertext consists of an instance of PV-LWE together with an instance of P-LWE, where the corresponding secrets are related to each other.

6 Concrete Security against Best Known Attacks

We start this section by investigating the concrete security against best known attacks of both schemes PASS Encrypt and PV Regev Encrypt. We then argue why we move from the Fourier to the Vandermonde transform in Section 6.3. Furthermore, we show how to interpret the partial Vandermonde problems in terms of error-correcting codes in Section 6.4.

6.1 Concrete Security of PASS Encrypt

In this section, we analyze the concrete security of PASS Encrypt, as proposed in Section 4, by presenting three different attacks.

The first two attacks, that we call the *key recovery* and *randomness recovery* attack in the following, were already studied in the original PASS Encrypt proposal [HS15]. We restate them for completeness and rephrase them in the primal attack framework of LWE as done by Alkim et al. [ADPS16]. Further, [HS15] use the less common notion of MIPS-years, where one MIPS-year equals the number of instructions executed during one year of computing at one million instructions per second. Note that in the parallel line of work concerning PASS Sign, the same type of attacks is studied as well ([HPS+14,LZA18,DHSS20]).

Essentially, recovering the secret key sk (resp. the randomness $\mathbf{r}'$) used in the encryption algorithm (see Figure 3) corresponds to solve an instance of PV-Knap with regard to the partial Vandermonde matrix $\mathbf{V}_\Omega$ (resp. the complement partial Vandermonde matrix $\mathbf{V}_{\Omega^c}$). However, no attack that aims at recovering the secret vector of the PASS instance given by a ciphertext has been studied so far. This leads us to the third attack, that we call the *plaintext recovery using hints*

attack in the following. This novel attack takes the design of PASS Encrypt into account and thus improves our understanding of its security.

We give concrete sample parameters and values for all three attacks in Section 7 in Figure 8. We also compare PASS Encrypt with PV Regev Encrypt and two other efficient lattice-based PKE schemes in Section 7.4.

Key Recovery Attack We now describe the first attack against PASS Encrypt, as already considered in the original proposal [HS15, Sec. 7]. We restate it for completeness and rephrase it in the attack framework of LWE as done by Alkim et al. [ADPS16], using the BKZ algorithm with quantum sieving to solve the associated unique shortest vector problem (u-SVP). The second component of the public key pk of PASS Encrypt is a vector $\mathbf{g} \in \mathbb{Z}_q^t$ defined as $\mathbf{g} = \mathbf{V}_\Omega \mathbf{f} \bmod q$, where $\mathbf{f} \leftarrow \chi_f^n$. We can write $\mathbf{V}_\Omega = [\mathbf{A}|\mathbf{B}] \in \mathbb{Z}_q^{t \times n}$ with $\mathbf{A} \in \mathbb{Z}_q^{t \times (n-t)}$ and $\mathbf{B} \in \mathbb{Z}_q^{t \times t}$, where $\mathbf{B}$ has full rank t and thus by multiplying $\mathbf{V}_\Omega$ by the inverse of $\mathbf{B}$, it takes the form $[\tilde{\mathbf{A}}|\mathbf{I}_t] \in \mathbb{Z}_q^{t \times n}$, for some matrix $\tilde{\mathbf{A}} \in \mathbb{Z}_q^{t \times (n-t)}$ and $\mathbf{I}_t$ the identity matrix of order t. Hence, we can transform the equation above to

$$\tilde{\mathbf{g}} = \tilde{\mathbf{A}} \cdot \mathbf{f}_1 + \mathbf{f}_2 \bmod q, \tag{2}$$

where $\tilde{\mathbf{g}} = \mathbf{B}^{-1}\mathbf{g}$ and $\mathbf{f} = (\mathbf{f}_1, \mathbf{f}_2)^T$ with $\mathbf{f}_1 \in \mathbb{Z}^{n-t}$ and $\mathbf{f}_2 \in \mathbb{Z}^t$. Equation 2 can be seen as an instance of LWE in its Hermite normal form (HNF) with public matrix $\tilde{\mathbf{A}}$ of LWE dimension $n - t$ and with t denoting the number of given samples. In doing so, we ignore the known structure of the matrix $\tilde{\mathbf{A}}$ and treat it as a uniform random matrix. This is a common approach used in structured lattice-based cryptography as no cryptanalytic technique making use of the algebraic structure is known, for a more elaborated discussion see [ACD+18]. We proceed as we usually do for LWE (see [ADPS16] for more details) and rewrite the equation above as $\tilde{\mathbf{A}} \cdot \mathbf{f}_1 + \mathbf{f}_2 - \tilde{\mathbf{g}} = \mathbf{0} \bmod q$. This defines an instance of u-SVP in the lattice $\Lambda = \Lambda(\tilde{\mathbf{A}}, \tilde{\mathbf{g}})$ given by

$$\Lambda = \{(\mathbf{x}, \mathbf{y}, w)^T \in \mathbb{Z}^t \times \mathbb{Z}^{n-t} \times \mathbb{Z} \colon \mathbf{x} + \tilde{\mathbf{A}} \cdot \mathbf{y} - w\tilde{\mathbf{g}} = \mathbf{0} \bmod q\}.$$

A basis of this lattice is given by the column vectors of

$$\mathbf{C} = \begin{bmatrix} q\mathbf{I}_t & -\tilde{\mathbf{A}} & \tilde{\mathbf{g}} \\ \mathbf{0}_{(n-t)\times t} & \mathbf{I}_{n-t} & \mathbf{0}_{(n-t)\times 1} \\ \mathbf{0}_{1\times t} & \mathbf{0}_{1\times(n-t)} & 1 \end{bmatrix} \in \mathbb{Z}^{(n+1)\times(n+1)},$$

where for $n, m \in \mathbb{N}$, we denote by $\mathbf{0}_{n\times m}$ the $n \times m$ matrix composed of zeros. The lattice Λ has full rank $n + 1$ as it is has an upper triangular form. Further, its determinant is q^t, corresponding to the determinant of $\mathbf{C}$. It is easy to see that the vector $(\mathbf{f}_2, \mathbf{f}_1, 1)^T \in \mathbb{Z}^{n+1}$ lies in Λ, where its norm depends on the distribution χ_f. Assuming that $\chi_f = U(\{-1, 0, 1\})$, we expect its Euclidean norm to be $\approx \sqrt{\frac{2n}{3}}$.

When estimating the expected length of a shortest vector, one can use the Gaussian heuristic. More precisely, the Gaussian heuristic for the given lattice Λ

with determinant q^t and of dimension $n+1$ states that the shortest vector in Λ has Euclidean norm approximately $\sqrt{\frac{n+1}{2\pi e}} \cdot q^{\frac{t}{(n+1)}} \approx \sqrt{q \cdot n}$, for $t = n/2$. This is much larger than the norm of $(\mathbf{f}_2, \mathbf{f}_1, 1)^T$. In other words, we assume that $(\mathbf{f}_2, \mathbf{f}_1, 1)^T$ is the unique shortest vector and the second shortest vector has the norm following the Gaussian heuristic. This is an instance of u-SVP. If we assume that q is linear in n, then the ratio of the shortest and the second shortest vector (also called the SVP gap) is approximately $\frac{1}{\sqrt{n}}$.

The u-SVP instance can be solved by the BKZ algorithm and its running time depends on the used blocksize within BKZ. We use the publicly accessible[8] Leaky LWE Estimator [DDGR20] to estimate the necessary blocksize for the BKZ algorithm, denoted as *bikz*. The algorithm BKZ itself uses an SVP oracle in dimension *bikz*. As in [ADPS16], we evaluate the running time of BKZ using the core SVP hardness, thus considering only the cost of one call to an SVP oracle in dimension *bikz*.

As the best known heuristic quantum algorithm to solve SVP in dimension *bikz* runs in time $2^{0.265 \cdot bikz}$ [Laa15], we give the number of quantum security bits by $0.265 \cdot bikz$.

Randomness Recovery Attack We now describe the second attack against PASS Encrypt, which aims at recovering the underlying randomness $\mathbf{r}$ used during encryption. The second component of the ciphertext $\mathbf{c}$ of PASS Encrypt is given by a vector $\mathbf{e}' \in \mathbb{Z}_q^{n-t}$ satisfying $\mathbf{e}' = \mathbf{V}_{\Omega^c}\mathbf{r}' \bmod q$, with $\mathbf{r}' = p\mathbf{r}$ for $\mathbf{r} \leftarrow \chi_r$. As p is publicly known and coprime to q, we can divide the above by p to obtain $\mathbf{g} := \mathbf{e}'/p = \mathbf{V}_{\Omega^c}\mathbf{r} \bmod q$. As for the key recovery attack from above, we can interpret this as an instance of LWE. More precisely, we write $\mathbf{V}_{\Omega^c} = [\mathbf{C}|\mathbf{D}] \in \mathbb{Z}_q^{(n-t)\times n}$ with $\mathbf{C} \in \mathbb{Z}_q^{(n-t)\times t}$ and $\mathbf{D} \in \mathbb{Z}_q^{(n-t)\times(n-t)}$, where $\mathbf{D}$ has full rank $n-t$. We multiply the equation above by the inverse of $\mathbf{D}$ to obtain

$$\tilde{\mathbf{g}} = \tilde{\mathbf{C}} \cdot \mathbf{r}_1 + \mathbf{r}_2 \bmod q, \tag{3}$$

where $[\tilde{\mathbf{C}}|\mathbf{I}_{n-t}] = [\mathbf{C}|\mathbf{D}] \cdot \mathbf{D}^{-1}$, $\tilde{\mathbf{g}} = \mathbf{D}^{-1}\mathbf{g}$ and $\mathbf{r} = (\mathbf{r}_1, \mathbf{r}_2)^T$ with $\mathbf{r}_1 \in \mathbb{Z}^t$ and $\mathbf{r}_2 \in \mathbb{Z}^{n-t}$. In other words, Equation 3 describes an instance of LWE in Hermite normal form with public matrix $\tilde{\mathbf{C}}$ of dimension t, where $n-t$ samples are given. Note that the roles of the dimension and the number of samples are exactly the reversed roles as in the key recovery attack. However, for $t = n/2$, the dimension and number of samples of the LWE instances are in both attacks the same. Further, for $\chi_r = \chi_f$, as we do in the sample parameters in Figure 8, both attacks are equally hard.

In order to recover the plaintext $\mathbf{m}$ from the ciphertext $\mathbf{c} = (\mathbf{e}, \mathbf{e}', \mathbf{e}'')$, an attacker can use the same approach as in the key recovery attack to solve the associated LWE instance (obtaining $\mathbf{r}$) and to compute $\mathbf{r}' = p\mathbf{r}$. They can then use $\mathbf{r}'$ to compute $\mathbf{V}_\Omega\mathbf{m}' = \mathbf{e} - (\mathbf{V}_\Omega\mathbf{r}' \circ \mathsf{pk})$. By combining $\mathbf{V}_\Omega\mathbf{m}'$ and $\mathbf{e}'' = \mathbf{V}_{\Omega^c}\mathbf{m}'$ to the full discrete Vandermonde transform $\mathbf{V}\mathbf{m}'$, one can multiply it

[8] https://github.com/lducas/leaky-LWE-Estimator

by $\mathbf{V}^{-1}$ to obtain $\mathbf{m}' = p\mathbf{s} + \mathbf{m} \bmod q$. Finally, the plaintext message $\mathbf{m} \bmod p$ can be recovered by computing $\mathbf{m}' \bmod p$.

Plaintext Recovery Using Hints Attack In the following we present a new attack against PASS Encrypt, which is inspired by the recent work of Dachman-Soled et al. [DDGR20] on exploiting hints that are given on the LWE secret or noise. The first part of the ciphertext $\mathbf{c}$ of PASS Encrypt is given by $\mathbf{e} = (\mathbf{V}_\Omega \mathbf{r}' \circ \mathsf{pk}) + \mathbf{V}_\Omega \mathbf{m}'$, where $\mathbf{r}' = p\mathbf{r}$, $\mathsf{pk} = \mathbf{V}_\Omega \mathbf{f}$ and $\mathbf{m}' = p\mathbf{s} + \mathbf{m}$ with $\mathbf{f} \leftarrow \chi_f^n$, $\mathbf{r} \leftarrow \chi_r^n$ and $\mathbf{s} \leftarrow \chi_s^n$. Using the homomorphic properties of $\mathbf{V}_\Omega$ we can rewrite $\mathbf{e}$ as $\mathbf{e} = \mathbf{V}_\Omega \cdot (\mathbf{f} \cdot \mathbf{r}' + \mathbf{m}')$. Recall that $\mathrm{Rot}(\mathbf{f})$ denotes the matrix describing the multiplication by $\mathbf{f}$ in the coefficient embedding. In matrix form this gives

$$\mathbf{e} = \mathbf{V}_\Omega \cdot (\mathrm{Rot}(\mathbf{f}) \cdot \mathbf{r}' + \mathbf{I}_n \cdot \mathbf{m}') = [\mathbf{A} | \mathbf{V}_\Omega] \cdot \begin{pmatrix} \mathbf{r}' \\ \mathbf{m}' \end{pmatrix} \bmod q,$$

where $\mathbf{A} = \mathbf{V}_\Omega \cdot \mathrm{Rot}(\mathbf{f}) \in \mathbb{Z}_q^{t \times n}$. Note that we can compute $\mathbf{A}$ by knowing the roots $\omega_{i_\ell} \in \Omega$ for $\ell \in [t]$ and the public key pk, and not necessarily the secret key $\mathsf{sk} = \mathbf{f}$, as explained in Lemma 3. As $\mathbf{V}_\Omega$ has full rank t, we can use Gauss elimination to transform this equation into

$$\tilde{\mathbf{e}} = [\mathbf{B} | \mathbf{I}_t] \cdot \begin{pmatrix} \mathbf{r}' \\ \mathbf{m}' \end{pmatrix} = \mathbf{B} \cdot \tilde{\mathbf{r}}' + \tilde{\mathbf{m}}' \bmod q,$$

with $\mathbf{B} \in \mathbb{Z}_q^{t \times (2n-t)}$, $\tilde{\mathbf{r}}'$ containing $\mathbf{r}'$ and the first $n-t$ coefficients of $\mathbf{m}'$ and $\tilde{\mathbf{m}}'$ containing the last t coefficients of $\mathbf{m}'$. This corresponds to an instance of LWE of dimension $2n-t$ and t the number of given samples, with $\mathbf{B}$ the public matrix.

At first sight, one can see that the LWE instance is of much larger dimension than in the two previous attacks, and thus one may wonder why this attack should provide tighter security estimates. As we will see now, this is because of the additional information provided by the rest of the ciphertext. The second and third part of the ciphertext $\mathbf{e}'$ and $\mathbf{e}''$ can be viewed as hints on $\mathbf{r}'$ and $\mathbf{m}'$. To be more precise, $\mathbf{e}' = \mathbf{V}_{\Omega^c} \cdot \mathbf{r}'$ and $\mathbf{e}'' = \mathbf{V}_{\Omega^c} \cdot \mathbf{m}'$. This can be rewritten as

$$\mathbf{e}' = [\mathbf{V}_{\Omega^c} | \mathbf{0}_{(n-t)\times n}] \cdot \begin{pmatrix} \mathbf{r}' \\ \mathbf{m}' \end{pmatrix} \bmod q, \quad \mathbf{e}'' = [\mathbf{0}_{(n-t)\times n} | \mathbf{V}_{\Omega^c}] \cdot \begin{pmatrix} \mathbf{r}' \\ \mathbf{m}' \end{pmatrix} \bmod q.$$

Note that the vector $(\tilde{\mathbf{r}}', \tilde{\mathbf{m}}')^T$ is simply a re-labeling of the vector $(\mathbf{r}', \mathbf{m}')^T$. In the language of Dachman-Soled et al. [DDGR20] this corresponds to $2(n-t)$ modular hints.

For simplicity, we assume that $\mathbf{m} = \mathbf{0}$ and thus $\mathbf{m}' = p\mathbf{s}$. As p is a public parameter we can assume that the secret $\mathbf{r}'$ and the noise $\mathbf{m}'$ of the corresponding LWE sample are drawn from the distributions χ_r^n and χ_s^n, respectively.

As in the key recovery attack, the number of security bits claims that a quantum algorithm would need at least a running time of $2^{0.265 \cdot bikz}$, where $bikz$ is the blocksize resulting from the Leaky LWE Estimator [DDGR20].

6.2 Concrete Security of PV Regev Encrypt

In the following we analyze the concrete security of PV Regev Encrypt, as proposed in Section 5, by presenting three different attacks. We call the first attack the *key recovery* attack, where the attacker aims at solving the search variant of the PV-LWE instance given by the public key. The second attack is called the *randomness recovery* attack, where the attacker aims at solving the PV-Knap instance given by the first component of the ciphertext. The third attack, called *plaintext recovery using hints* attack, is an application of the theory of LWE with side information, as recently studied by Dachman-Soled et al. [DDGR20].

We give concrete sample parameters and values for all three attacks in Section 7 in Figure 9. We also compare PV Regev Encrypt with PASS Encrypt and two other efficient lattice-based PKE schemes in Section 7.4.

Key Recovery Attack The public key pk of PV Regev Encrypt as presented in Figure 6 is given by a tuple $(\mathbf{V}_\Omega, \mathbf{b}) \in \mathbb{Z}_q^{t\times n} \times \mathbb{Z}_q^n$, satisfying $\mathbf{b} = \mathbf{V}_\Omega^T \cdot \mathbf{s} + \mathbf{e} \bmod q$, where $\mathbf{s} \leftarrow U(\mathbb{Z}_q^t)$ and $\mathbf{e} \leftarrow \chi_e^n$. This can be seen as a sample of the learning with errors (LWE) problem with public matrix $\mathbf{V}_\Omega^T$ of LWE dimension t and with n denoting the number of given samples. In doing so, we ignore the known structure of the partial Vandermonde transformation matrix $\mathbf{V}_\Omega$ and treat it as a uniform random matrix. As already mentioned before, this is a common technique used in structured lattice-based cryptography. As for PASS Encrypt in Section 6.1, we use the Leaky LWE Estimator to estimate the necessary blocksize for the BKZ algorithm, denoted as *bikz*. Note that for the given LWE instance, the secret vector $\mathbf{s}$ is uniformly sampled over $\mathbb{Z}_q^t$, but the instance we feed to the Leaky LWE Estimator is in HNF, i.e., the secret follows the same distribution χ_e as the noise $\mathbf{e}$. From a cryptanalytic point of view, the LWE instance defines an easier problem when it is in its HNF as the norm of the shortest vector becomes smaller and thus the SVP gap, that we mentioned in Section 6.1, becomes larger and finally the given u-SVP instance becomes easier.

Randomness Recovery Attack The first component of the ciphertext $\mathbf{c}$ of PV Regev Encrypt as presented in Figure 6 is given by a vector $\mathbf{u} \in \mathbb{Z}_q^t$, satisfying $\mathbf{u} = \mathbf{V}_\Omega \cdot \mathbf{r} \bmod q$. As we explained in detail for the *key recovery* attack of PASS Encrypt in Section 6.1, this defines an instance of LWE of dimension $n - t$, where t samples are given. In order to recover the plaintext $\mathbf{m}$ from the ciphertext $\mathbf{c} = (\mathbf{u}, \mathbf{v})$, an attacker can first solve the associated LWE instance to recover $\mathbf{r}$. The attacker can then use $\mathbf{r}$ to compute $\mathbf{v} - \mathrm{Rot}(\mathbf{r})^T \cdot \mathbf{b}$ and finally coefficient-wise recover $\mathbf{m}$.

Plaintext Recovery Using Hints Attack We now present a third attack against PV Regev Encrypt scheme from Section 5. Similar to the *plaintext recovery using hints* attack against PASS Encrypt, it is inspired by the recent work of Dachman-Soled et al. [DDGR20] on investigating the hardness of LWE in the presence of side information. In order to facilitate notations, we set

the linear transformation $T\colon R \to R$ as the map that sends any ring element $\mathbf{r} = (r_0, \ldots, r_{n-1})^T$ to the vector $T \cdot \mathbf{r} = (r_0, -r_{n-1}, \ldots, -r_1)^T$, as we did in Section 3.4. The second component of the ciphertext $\mathbf{c}$ of PV Regev Encrypt is given by $\mathbf{v} = \mathrm{Rot}(\mathbf{r})^T \cdot \mathbf{b} + \mathbf{e}' + \lfloor q/2 \rfloor \mathbf{m} = \mathrm{Rot}(\mathbf{b}) \cdot T \cdot \mathbf{r} + \mathbf{e}' + \lfloor q/2 \rfloor \mathbf{m}$. This corresponds to an instance of LWE of dimension n and n the number of given samples, with $\mathrm{Rot}(\mathbf{b}) \cdot T$ the public matrix. In an analog manner as for the *plaintext recovery using hints* attack against PASS Encrypt, the first component $\mathbf{u} = \mathbf{V}_\Omega \cdot \mathbf{r}$ of the ciphertext $\mathbf{c}$ can be viewed as t different modular hints on the secret $\mathbf{r}$. For simplicity, we assume that $\mathbf{m} = \mathbf{0}$, and thus that the secret $\mathbf{r}$ and the error $\mathbf{e}'$ of the given LWE instance are drawn from the same distribution.

6.3 Choice of Ring

In the original description of PASS Encrypt [HS15], the partial Fourier transform is used, and not as we propose the partial Vandermonde transform. The main difference between the original and our version is that the latter works over the ring of integers of some cyclotomic number field, whereas the first one works over the cyclic ring $\mathbb{Z}[x]/\langle x^n - 1\rangle$, for some prime n. The setting in [HS15] is the following. Let n and q be primes satisfying $q = 1 \bmod n$ and let ω be a primitive n-th root of unity in $\mathbb{Z}_q$. Further, let S be a subset of $[n]$ of size t and let $\Omega = \{\omega^{k-1} \colon k \in S\}$ and $\Omega^c = \{\omega^{k-1} \colon k \in [n] \setminus S\}$. The partial Fourier transformation matrix $\mathbf{F}_\Omega$ is defined as $\mathbf{F}_\Omega := (\omega_j^{k-1})_{j\in[t],k\in[n]}$, where $\omega_j \in \Omega$ for $j \in [t]$. In an analog manner to Section 3.1, where we define partial Vandermonde SIS (Definition 10), we can define partial Fourier SIS, denoted by PF-SIS. More concretely, for a given parameter $\beta > 0$, PF-SIS$_\beta$ asks to find an element $\mathbf{a} \in \mathbb{Z}[x]/\langle x^n - 1\rangle$ of norm $\|\mathbf{a}\|_2 \leq \beta$ satisfying $\mathbf{F}_\Omega \cdot \mathbf{a} = \mathbf{0} \bmod q$.

Lemma 16 (Solution to Partial Fourier SIS). *The problem* PF-SIS$_\beta$ *is easy to solve for any* $\beta \geq \sqrt{n}$.

Proof. Let $\mathbf{1}$ denote the element in $\mathbb{Z}[x]/\langle x^n - 1\rangle$ whose coefficient vector is given by $(1, \ldots, 1)^T \in \mathbb{Z}^n$. The polynomial $x^n - 1$ can be factorized in the product $x^n - 1 = (x-1)(x^{n-1} + \cdots + x + 1)$. As for any $j \in [t]$ the element ω_j is a solution to the equation $x^n - 1 \bmod q$, we also know that $\sum_{k\in[n]} \omega_j^k = 0 \bmod q$ and thus $\mathbf{F}_\Omega \mathbf{1} = \mathbf{0} \bmod q$ with $\|\mathbf{1}\|_2 = \sqrt{n} \leq \beta$. □

This generic solution only holds for the *homogeneous* problem PF-SIS, and not for the *inhomogeneous* Knapsack counterpart. However, we prefer to move to the ring $R = \mathbb{Z}[x]/\langle x^n + 1\rangle$, where n is a power of two. For this ring, the so-called evaluation-at-1 attack does not work. Note that the evaluating-at-1 approach already led to successful attacks against NTRU.

6.4 Partial Vandermonde Problems as Error Correcting Codes

We now explain how the partial Vandermonde transform can be interpreted in terms of error-correcting codes. For a gentle introduction to coding theory, we refer to the book of Roth [Rot06].

More formally, the partial Vandermonde transformation matrix $\mathbf{V}_\Omega \in \mathbb{Z}_q^{t\times n}$, as defined in Section 3.1, describes the check matrix of an $[n, t, d]$ Reed-Solomon code, with $d = n - t + 1$ its minimum distance. Using the duality connection from Section 3.3, the PV-LWE matrix $\mathbf{V}^T_{\mathrm{inv}(\Omega^c)}$ corresponds to the generating matrix of the same code. More concretely, it is a punctured Reed-Solomon code fulfilling the optimal singleton-bound. Thus, it is a maximum distance separable code with correction capability $\lfloor (d-1)/2 \rfloor = \lfloor (n-t)/2 \rfloor$. As we set $t = \lfloor n/2 \rfloor$, the correction capability is bounded above by $\lfloor n/4 \rfloor$. It is further a unit-derived code and the theory of error-decoding pairs provides efficient decoder, see for instance the work by Hurley and Hurley [HH18].

In order to prevent coding-based attacks, we need to choose distributions for the secret in PV-Knap and for the noise in PV-LWE such that the expected Hamming weight is not near the correction capability bound. For example, setting the distribution χ as the uniform distribution over $\{-1, 0, 1\}^n$ and sampling $\mathbf{e} \leftarrow \chi$, we expect the Hamming weight of $\mathbf{e}$ to be $2n/3$, which is much larger than the error-correcting capacity, which is at most $\lfloor n/4 \rfloor$. So this would be a choice that does not allow for code-based attacks. If, however, we set the distribution χ as a sparse distribution over $\{-1, 0, 1\}^n$, where an element $\mathbf{e} \leftarrow \chi$ has only very few non-zero coefficients, let's say about $n/4$ non-zero coefficients, then the efficient decoder from [HH18] would apply.

7 Concrete Parameters

In the following we present sample parameters for both encryption schemes together with the estimated quantum security. We then discuss the influence of the parameter t on their security. We conclude with a comparison of our two schemes with two other efficient lattice-based PKE schemes.

7.1 Parameters for PASS Encrypt

We propose the following sample parameters for PASS Encrypt, under which the scheme as presented in Figure 3 is correct (Lemma 12).

We consider the case where K is the ν-th cyclotomic number field with ν a power of 2. By $n = \nu/2$ we denote its degree and the number of rows of the partial Vandermonde matrix t is given by $n/2$. In Section 7.3 below we argue why this is the optimal choice for t. The parameter q denotes the modulus over which the partial Vandermonde transformation matrix $\mathbf{V}_\Omega$ is taken. We require $q = 1 \bmod \nu$ such that the defining polynomial of K, given by $x^n + 1$, fully splits modulo q. Concretely, we provide two parameter sets, as summarized in Figure 8. In the first, we choose $\nu = 2048$ and $q = 12289$, and in the second, we keep the same q and set $\nu = 4096$. Note that the relevant parameter for security is t.

We set the distributions χ_f, χ_r and χ_s of the secret key and the encryption randomness as $U(T_n(d))$, the uniform distribution over $T_n(d)$, the set of ternary polynomials with exactly d coefficients that equal 1, and d coefficients

that equal -1, and $n - 2d$ coefficients that equal 0, where $d = \lfloor n/3 \rfloor$. Thus, for every element $\mathbf{f} \leftarrow \chi_f$ (resp. $\mathbf{r} \leftarrow \chi_r$ and $\mathbf{s} \leftarrow \chi_s$) it yields $\|\mathbf{f}\|_\infty \leq 1$ (resp. $\|\mathbf{r}\|_1 \leq 2n/3$ and $\|\mathbf{s}\|_\infty \leq 1$) with probability 1. Hence, we can set the parameter α to $2n/3$ and β to 1. Fixing the number of coefficients that equal -1 and 1 makes it possible to set $\alpha = 2n/3$ (in order to keep *perfect* correctness for the given q), but adds a structural hint, as exploited by Dachman-Soled et al. [DDGR20, Sec. 6.3]. This structural hint roughly decreases the estimated $bikz$ by 1. Further, we set p as 2.

We then provide the needed block sizes of the BKZ algorithm in order to perform the three attacks on PASS Encrypt for both parameter sets, as presented in Section 6. All estimations are computed with SageMath using the Leaky LWE Estimator [DDGR20].

Parameter	Set 1	Set 2
ν	2048	4096
n	1024	2048
t	512	1024
q	12289	12289
p	2	2
α	$\lfloor 2n/3 \rfloor$	$\lfloor 2n/3 \rfloor$
β	1	1
$\chi_f = \chi_r = \chi_s$	$U(T_n(\lfloor n/3 \rfloor))$	$U(T_n(\lfloor n/3 \rfloor))$
key recovery ($bikz_0$)	298.87	710.11
randomness recovery ($bikz_1$)	298.87	710.11
plaintext recovery using hints ($bikz_2$)	298.14	712.95
quantum security (bits)	79	188

Fig. 8: Sample parameters and security estimations for PASS Encrypt. The number of quantum security bits is computed as $0.265 \cdot \min_{j \in [3]}(bikz_j)$.

With the first sample set, we achieve a quantum bit security of 79 and with the second one, we achieve a quantum bit security of 188. We made the SageMath code of our experiments publicly available.[9]

7.2 Parameter for PV Regev Encrypt

We propose the following sample parameters for PV Regev Encrypt, under which the scheme as presented in Figure 6 is correct (Lemma 14).

Again, we consider the case where K is the ν-th cyclotomic number field with ν a power of two with R its ring of integers. By n we denote its degree and t is given by $n/2$. The parameter q denotes the modulus over which the matrix $\mathbf{V}_\Omega$ is taken. We require $q = 1 \bmod \nu$ to ensure that the defining polynomial of K, given by $x^n + 1$, fully splits modulo q. Concretely, we provide two parameter

[9] https://github.com/KatinkaBou/SecurityAnalysisPASSEncrypt

sets, as summarized in Figure 9. In the first, we choose $\nu = 2048$ and $q = 12289$, and in the second, we keep the same q and set $\nu = 4096$. Note that the relevant parameter for security is t.

We set the distributions χ_e and χ_r of the LWE noise and the encryption randomness as the uniform distribution over ternary elements, i.e., $U(\{-1, 0, 1\})$.[10] Thus, for every element $\mathbf{e} \leftarrow \chi_e^n$ (resp. $\mathbf{r} \leftarrow \chi_r^n$) it yields $\|\mathbf{e}\|_\infty \leq 1$ (resp. $\|\mathbf{r}\|_\infty \leq 1$) with probability 1. Hence, we can set the parameter α and β to 1.

We then provide the needed block sizes of the BKZ algorithm in order to perform the three attacks on PV Regev Encrypt for both parameter sets, as presented in Section 6. All estimations are computed with SageMath using the Leaky LWE Estimator [DDGR20].

Parameter	Set 1	Set 2
ν	2048	4096
n	1024	2048
t	512	1024
q	12289	12289
α	1	1
β	1	1
$\chi_e = \chi_r$	$U(\{-1, 0, 1\})$	$U(\{-1, 0, 1\})$
key recovery ($bikz_0$)	299.64	711.06
randomness recovery ($bikz_1$)	299.64	711.06
plaintext recovery using hints ($bikz_2$)	299.64	711.06
quantum security (bits)	79	188

Fig. 9: Sample parameters and security estimations for PV Regev Encrypt. The quantum security bits are computed as $0.265 \cdot \min_{j\in[3]}(bikz_j)$.

Using the first sample set, we achieve a quantum bit security of 79 and using the second sample set, we achieve a quantum bit security of 188. We made the SageMath code of our experiments publicly available.[11]

7.3 Choice of the Number of Rows

We now discuss the influence of the parameters t, i.e., the number of rows of the Vandermonde matrix chosen to construct $\mathbf{V}_\Omega$, on the security of our scheme. This observation also applies to the the original proposal in [HS15].

Increasing t leads to an easier key recovery attack against PASS Encrypt and an easier randomness recovery attack against PV Regev Encrypt, as the underlying LWE dimension $n - t$ decreases. On the other hand, decreasing t leads to

[10] In contrast to PASS Encrypt we don't need to bound the ℓ_1-norm for the correctness of PV Regev Encrypt and thus there is no motivation to use the uniform distribution over $T_n(d)$ as before.

[11] https://github.com/KatinkaBou/SecurityAnalysisPVRegevEncrypt

an easier randomness recovery attack against PASS Encrypt and an easier key recovery attack against PV Regev Encrypt, as the underlying LWE dimension t decreases. Hence, choosing $t = \lfloor n/2 \rfloor$ is the optimal choice, as it balances the hardness of all attacks. Our experiments with $t = \lfloor n/3 \rfloor$ (Set A), $t = \lfloor n/2 \rfloor$ (Set B) and $t = \lfloor 2n/3 \rfloor$ (Set C) validate those observations and are summarized in Figure 10. In both variations (Set A and Set C) the quantum security of PASS Encrypt and PV Regev Encrypt decreases from 79 to 45.

We emphasize that the observations made above do not apply to the sequence of works on PASS Sign. In the recent publication on the aggregate variant of PASS Sign by Doröz et al. [DHSS20], the parameter t is set to $\lfloor n/3 \rfloor$. Note that there is only the partial Fourier SIS problem with the matrix $\mathbf{F}_\Omega \in \mathbb{Z}_q^{t \times n}$ arising in the design of the signature scheme, and not the complement matrix $\mathbf{F}_{\Omega^c}$. Hence, decreasing t only makes the corresponding dimension of the LWE instance, that is defined by an instance of partial Fourier SIS, larger and thus the problem harder.

Parameter	Set A	Set B	Set C
ν	2048	2048	2048
n	1024	1024	1024
t	341	512	682
q	12289	12289	12289
PASS Encrypt			
key recovery ($bikz_0$)	474.89	298.87	171.82
randomness recovery ($bikz_1$)	171.09	298.87	473.45
plaintext recovery using hints ($bikz_2$)	202.87	298.14	430.49
quantum security (bits)	45	79	45
PV Regev Encrypt			
key recovery ($bikz_0$)	171.86	299.64	432.39
randomness recovery ($bikz_1$)	476.45	299.64	172.59
plaintext recovery using hints ($bikz_2$)	433.18	299.64	172.59
quantum security (bits)	45	79	45

Fig. 10: Security estimations with different number of rows t for PASS Encrypt and PV Regev Encrypt. The number of quantum security bits is computed as $0.265 \cdot \min_{j \in [3]}(bikz_j)$.

7.4 Comparison

Finally, we provide a comparison between the asymptotic parameters of PASS Encrypt and PV Regev Encrypt with two other efficient lattice-based PKE schemes. We therefore compute the asymptotic parameters in bits for the secret key sk, the public key pk and the ciphertext $\mathbf{c}$. As the variable t is the important parameter defining the asymptotic security of PASS Encrypt and PV Regev Encrypt, we state everything with regard to t.

In PASS Encrypt the secret key is a ring element sampled from the uniform distribution over $T_n(\lfloor n/3 \rfloor)$. Assuming $n = 2t$, the bit size of the secret key is $2t \cdot \log_2 3$. The public key is given by $\mathsf{pk} = (\Omega, \mathbf{V}_\Omega \mathsf{sk})$ and lies in $\mathbb{Z}_q^t \times \mathbb{Z}_q^t$, requiring $2t \cdot \log_2 q$ bits to transmit it.[12] Finally, the ciphertext is an element of $\mathbb{Z}_q^t \times \mathbb{Z}_q^{n-t} \times \mathbb{Z}_q^{n-t}$, with $n - t = t$, requiring $3t \cdot \log_2 q$ bits.

In PV Regev Encrypt the secret key is given by $\mathbf{s} \in \mathbb{Z}_q^t$, i.e., of bit size $t \cdot \log_2 q$. The public key is given by $\mathsf{pk} = (\Omega, \mathbf{b})$ and lies in $\mathbb{Z}_q^t \times \mathbb{Z}_q^n$, requiring $3t \cdot \log_2 q$ bits to transmit it. The ciphertext is an element of $\mathbb{Z}_q^t \times \mathbb{Z}_q^n$, and thus of bit size $3t \cdot \log_2 q$.

We now compare our schemes with two other efficient lattice-based PKE schemes, as illustrated in Figure 11. The first is the Regev-like PKE scheme based on P-LWE, as presented in [LP11], and the second is the NTRU scheme, as presented in [HPS98]. In [LP11], the secret key and the public key are both ring elements of the ring $R = \mathbb{Z}[x]/\langle x^t + 1 \rangle$, where t is a power of two and the parameter that is determining the asymptotic security. For a better comparison, we assume that the secret is, as in PASS Encrypt, sampled uniformly over $T_t(\lfloor n/3 \rfloor)$. The ciphertext is composed of two ring elements, allowing to encrypt a t-bit message. In [HPS98], the ring $R = \mathbb{Z}[x]/\langle x^t - 1 \rangle$ is used. The secret key is a ring element of small norm. Again, for better comparison, we use the same distribution as in PASS Encrypt. The public key and the ciphertext are elements of R and the schemes allows to encrypt a t-bit message. We note that for simplicity we consider non-optimized versions of the four schemes.

An important characteristic of an PKE scheme is the ratio between the sum of the bit size of its public key and ciphertext and the bit size of the encrypted message. Figure 11 shows that this ratio is $2.5 \log_2 q$ for PASS Encrypt and $3 \log_2 q$ for PV Regev Encrypt and thus placing them in the same range as NTRU [HPS98] and the P-LWE-based Regev scheme [LP11].

Param.	PASS Encrypt	PV Regev Encrypt	[LP11]	[HPS98]
sk	$2t \cdot \log_2 3$	$t \cdot \log_2 q$	$t \cdot \log_2 3$	$t \cdot \log_2 3$
pk	$2t \cdot \log_2 q$	$3t \cdot \log_2 q$	$t \cdot \log_2 q$	$t \cdot \log_2 q$
$\mathbf{c}$	$3t \cdot \log_2 q$	$3t \cdot \log_2 q$	$2t \cdot \log_2 q$	$t \cdot \log_2 q$
$\mathbf{m}$	$2t$	$2t$	t	t
$(\mathsf{pk}, \mathbf{c})/\mathbf{m}$	$2.5 \cdot \log_2 q$	$3 \cdot \log_2 q$	$3 \cdot \log_2 q$	$2 \cdot \log_2$

Fig. 11: Asymptotic parameters in bits for PASS Encrypt, PV Regev Encrypt, the Regev-like PKE over P-LWE [LP11] and NTRU [HPS98].

[12] We could further save in storage and bandwidth by only transmitting an index vector in $\{0,1\}^n$ (instead of the full vector Ω) indicating which row of $\mathbf{V}$ is used for the public key.

8 Acknowledgments

Katharina Boudgoust was funded by the Direction Générale de l'Armement (Pôle de Recherche CYBER). This work was supported in part by Australian Research Council Discovery Grant DP180102199. We thank our anonymous PKC'2021 and DCC referees for their helpful and constructive feedback and we also thank Alice Pellet–Mary for making us aware that there are unsafe choices of the partial Vandermonde matrix.

References

ACD+18. Martin R. Albrecht, Benjamin R. Curtis, Amit Deo, Alex Davidson, Rachel Player, Eamonn W. Postlethwaite, Fernando Virdia, and Thomas Wunderer. Estimate all the {LWE, NTRU} schemes! In *SCN*, volume 11035 of *Lecture Notes in Computer Science*, pages 351–367. Springer, 2018.

AD97. Miklós Ajtai and Cynthia Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *STOC*, pages 284–293. ACM, 1997.

ADPS16. Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange - A new hope. In *USENIX Security Symposium*, pages 327–343. USENIX Association, 2016.

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

CDPR16. Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. In *EUROCRYPT (2)*, volume 9666 of *Lecture Notes in Computer Science*, pages 559–585. Springer, 2016.

CDW17. Ronald Cramer, Léo Ducas, and Benjamin Wesolowski. Short stickelberger class relations and application to ideal-svp. In *EUROCRYPT (1)*, volume 10210 of *Lecture Notes in Computer Science*, pages 324–348, 2017.

DDGR20. Dana Dachman-Soled, Léo Ducas, Huijing Gong, and Mélissa Rossi. LWE with side information: Attacks and concrete security estimation. In *CRYPTO (2)*, volume 12171 of *Lecture Notes in Computer Science*, pages 329–358. Springer, 2020.

DHSS20. Yarkin Doröz, Jeffrey Hoffstein, Joseph H. Silverman, and Berk Sunar. MMSAT: A scheme for multimessage multiuser signature aggregation. *IACR Cryptol. ePrint Arch.*, page 520, 2020.

DvW21. Léo Ducas and Wessel P. J. van Woerden. NTRU fatigue: How stretched is overstretched? In *ASIACRYPT (4)*, volume 13093 of *Lecture Notes in Computer Science*, pages 3–32. Springer, 2021.

GPM21. E. Gachon and Alice Pellet-Mary. Private communication, 2021.

GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206. ACM, 2008.

GSW13. Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *CRYPTO (1)*, volume 8042 of *Lecture Notes in Computer Science*, pages 75–92. Springer, 2013.

HH18. Ted Hurley and Donny Hurley. Coding theory: the unit-derived methodology. *Int. J. Inf. Coding Theory*, 5(1):55–80, 2018.

HPS98. Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In *ANTS*, volume 1423 of *Lecture Notes in Computer Science*, pages 267–288. Springer, 1998.

HPS+14. Jeffrey Hoffstein, Jill Pipher, John M. Schanck, Joseph H. Silverman, and William Whyte. Practical signatures from the partial fourier recovery problem. In *ACNS*, volume 8479 of *Lecture Notes in Computer Science*, pages 476–493. Springer, 2014.

HS15. Jeffrey Hoffstein and Joseph H. Silverman. Pass-encrypt: a public key cryptosystem based on partial evaluation of polynomials. *Des. Codes Cryptogr.*, 77(2-3):541–552, 2015.

Laa15. Thijs Laarhoven. Search problems in cryptography, 2015. http://www.thijs.com/docs/phd-final.pdf, last accessed on 08.07.2021.

LM06. Vadim Lyubashevsky and Daniele Micciancio. Generalized compact knapsacks are collision resistant. In *ICALP (2)*, volume 4052 of *Lecture Notes in Computer Science*, pages 144–155. Springer, 2006.

LP11. Richard Lindner and Chris Peikert. Better key sizes (and attacks) for lwe-based encryption. In *CT-RSA*, volume 6558 of *Lecture Notes in Computer Science*, pages 319–339. Springer, 2011.

LPR10. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In *EUROCRYPT*, volume 6110 of *Lecture Notes in Computer Science*, pages 1–23. Springer, 2010.

LPR13. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. *J. ACM*, 60(6):43:1–43:35, 2013.

LTV12. Adriana López-Alt, Eran Tromer, and Vinod Vaikuntanathan. On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In *STOC*, pages 1219–1234. ACM, 2012.

LZA18. Xingye Lu, Zhenfei Zhang, and Man Ho Au. Practical signatures from the partial fourier recovery problem revisited: A provably-secure and gaussian-distributed construction. In *ACISP*, volume 10946 of *Lecture Notes in Computer Science*, pages 813–820. Springer, 2018.

MM11. Daniele Micciancio and Petros Mol. Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In *CRYPTO*, volume 6841 of *Lecture Notes in Computer Science*, pages 465–484. Springer, 2011.

MR10. Daniele Micciancio, , and Oded Regev. Lattice-based cryptography. In *Post-Quantum Cryptography*, pages 147–191. Springer, 2010.

Pei16. Chris Peikert. A decade of lattice cryptography. *Found. Trends Theor. Comput. Sci.*, 10(4):283–424, 2016.

PHS19. Alice Pellet-Mary, Guillaume Hanrot, and Damien Stehlé. Approx-svp in ideal lattices with pre-processing. In *EUROCRYPT (2)*, volume 11477 of *Lecture Notes in Computer Science*, pages 685–716. Springer, 2019.

PR06. Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, volume 3876 of *Lecture Notes in Computer Science*, pages 145–166. Springer, 2006.

PXWC21. Yanbin Pan, Jun Xu, Nick Wadleigh, and Qi Cheng. On the ideal shortest vector problem over random rational primes. In *EUROCRYPT (1)*, volume 12696 of *Lecture Notes in Computer Science*, pages 559–583. Springer, 2021.

Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, pages 84–93. ACM, 2005.

Rot06. Ron M. Roth. *Introduction to coding theory*. Cambridge University Press, 2006.

SS11. Damien Stehlé and Ron Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In *EUROCRYPT*, volume 6632 of *Lecture Notes in Computer Science*, pages 27–47. Springer, 2011.

SSTX09. Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In *ASIACRYPT*, volume 5912 of *Lecture Notes in Computer Science*, pages 617–635. Springer, 2009.

Ste14. Ron Steinfeld. Ntru cryptosystem: Recent developments and emerging mathematical problems in finite polynomial rings. *Algebraic Curves and Finite Fields: Cryptography and Other Applications*, 16:179, 2014.