Partial Lattice Trapdoors
How to Split Lattice Trapdoors, Literally

Martin R. Albrecht[1], Russel W. F. Lai[2], Oleksandra Lapiha[3], and Ivy K. Y. Woo[2]

[1] King's College London and SandboxAQ
martin.albrecht@{kcl.ac.uk,sandboxaq.com}
[2] Aalto University
{russell.lai,ivy.woo}@aalto.fi
[3] Royal Holloway, University of London
sasha.lapiha.2021@live.rhul.ac.uk

Abstract. Lattice trapdoor algorithms allow us to sample hard random lattices together with their trapdoors, given which short lattice vectors can be sampled efficiently. This enables a wide range of advanced cryptographic primitives. In this work, we ask: can we distribute lattice trapdoor algorithms non-interactively?

We study a natural approach to sharing lattice trapdoors: splitting them into partial trapdoors for different lower-rank sublattices which allow the local sampling of short sublattice vectors. Given sufficiently many short sublattice vectors, these can then be combined to yield short vectors in the original lattice. Moreover, this process can be repeated an unbounded polynomial number of times without needing a party holding a full trapdoor to intervene. We further define one-wayness and indistinguishability properties for partial trapdoors.

We establish that such objects exist that have non-trivial performance under standard assumptions. Specifically, we prove these properties for a simple construction from the κ-SIS and κ-LWE assumptions, which were previously shown to be implied by the plain SIS and LWE assumptions, respectively. The security proofs extend naturally to the ring or module settings under the respective analogues of these assumptions, which have been conjectured to admit similar reductions.

Our partial trapdoors achieve non-trivial efficiency, with relevant parameters sublinear in the number of shareholders. Our construction is algebraic, without resorting to generic tools such as multiparty computation or fully homomorphic encryption. Consequently, a wide range of lattice-trapdoor-based primitives can be thresholdised non-interactively by simply substituting the trapdoor preimage sampling procedure with our partial analogue.

1 Introduction

Lattice-based cryptographic constructions commonly involve generating some public matrix $\mathbf{A}$ and a secret matrix $\mathbf{U}$ with short elements, the latter called a *trapdoor* of $\mathbf{A}$, which satisfy some non-trivial relation such as $\mathbf{A} \cdot \mathbf{U} \equiv \mathbf{0} \bmod q$. Many cryptographic tasks ranging from signatures to advanced encryption are then accomplished by using $\mathbf{U}$ to sample a short preimage $\mathbf{x}$ with respect to some target image vector $\mathbf{y}$, so that $\mathbf{A} \cdot \mathbf{x} \equiv \mathbf{y} \bmod q$ holds.

Since the first lattice trapdoor scheme [GPV08] almost two decades ago, lattice trapdoors have been all-or-nothing: Either finding short preimages w.r.t. any target is easy with a trapdoor, or such a problem w.r.t. any target is believed to be hard. In this work, we ask and answer the natural question:

can we partition a lattice trapdoor?

By this, we mean a non-interactive distribution of the preimage-sampling procedure using a trapdoor divided among k parties:

- With a *partial trapdoor* (possibly generated from a master (full) trapdoor), any party j can *locally* compute a partial preimage $\mathbf{x}_j$ of $\mathbf{y}$.
- After collecting enough partial preimages $(\mathbf{x}_j)_{j \in T}$ for some t-subset $T \subseteq_t [k]$, anyone can locally recover a full preimage $\mathbf{x}$ of $\mathbf{y}$.
- The distributed preimage sampling process should be secure even if run an unbounded polynomial number of times, without needing a party holding a full trapdoor to intervene at any time.

Such a partial trapdoor primitive is quite powerful, as it allows trivial and non-interactive thresholdisation of (the preimage sampling operations in) any trapdoor-based lattice-based primitives, such as (hash-and-sign) signatures [GPV08], homomorphic signatures [BF11, GVW15], identity-based encryption [GPV08, ABB10a], attribute-based encryption [BGG+14, Wee22], to name but a few.

Warm-up: A simple partial lattice trapdoor. Without size constraints, the above task may seem easy. Let us consider the trapdoor generation algorithm in [MP12] which outputs, say, two matrices $\mathbf{A}, \mathbf{U} = (\mathbf{U}_0 \ \mathbf{U}_1)$ satisfying

$$\begin{pmatrix} \mathbf{G} & \mathbf{0} \\ \mathbf{0} & \mathbf{G} \end{pmatrix} \equiv \mathbf{A} \cdot (\mathbf{U}_0 \ \mathbf{U}_1) \bmod q$$

where $\mathbf{G} \coloneqq \mathbf{I} \otimes \mathbf{g}^\intercal$, with $\mathbf{I}$ the identity matrix and $\mathbf{g}^\intercal = \left(1\ 2 \dots 2^{\lceil \log q \rceil - 1}\right)$ the "gadget vector". If we hand $\mathbf{U}_j$ to the j-th party, this immediately allows each party to compute a partial preimage $\mathbf{x}_j^*$ for some given target $\mathbf{y} \coloneqq (\mathbf{y}_0 \| \mathbf{y}_1)$,[4] i.e. $(\mathbf{y}_0 \| \mathbf{0}) \equiv \mathbf{A} \cdot \mathbf{x}_0 \bmod q$ and $(\mathbf{0} \| \mathbf{y}_1) \equiv \mathbf{A} \cdot \mathbf{x}_1 \bmod q$. We then have $(\mathbf{y}_0 \| \mathbf{y}_1) \equiv \mathbf{A} \cdot (\mathbf{x}_0 + \mathbf{x}_1) \bmod q$. This basic idea can be extended to more parties in the obvious way, and the "public parameters" $\mathbf{A}$ grow linearly in k. Moreover, with some effort, this can be generalised to the t-out-of-k setting by applying Shamir secret sharing on the image space, avoiding the usual issues surrounding secret-sharing short vectors. We expand on the latter in Section 2.1.

(In)security? We inspect security of the toy-example above. For a traditional full trapdoor, we require that the (inhomogeneous) Short Integer Solutions SIS and/or Learning With Errors (LWE) problems w.r.t. $\mathbf{A}$ remain hard so long as the trapdoor is not made available. Generalising to the partial trapdoor setting, for example, we may wish to argue that the inhomogeneous SIS problem remains hard even when given the partial trapdoors of up to $t-1$ corrupt parties and many partial preimages generated by honest parties.

Adapting the conventional proof strategy to our setting quickly reveals multiple challenges. To illustrate this, consider the original security proof of [GPV08]. There, a key step is to sample a short Gaussian vector $\mathbf{x}$ from a public lattice, e.g. $\mathbb{Z}^m$ if $\mathbf{A}$ has m columns, and to compute a challenge image $\mathbf{y} \equiv \mathbf{A} \cdot \mathbf{x} \bmod q$ for which the adversary must provide a short preimage. This is made possible by two convenient facts:

1. The tuples $(\mathbf{x}, \mathbf{y})$ sampled as above are statistically indistinguishable from those sampled in the real scheme, i.e. first sample a random $\mathbf{y}$, then sample $\mathbf{x}$ using a full trapdoor.
2. The preimages $\mathbf{x} \in \mathbb{Z}^m$ can be easily sampled without any secret information: they are just Gaussian vectors over the public lattice $\mathbb{Z}^m$ which admits a public short basis $\mathbf{I}_m$, the identity matrix. In particular, this implies that after seeing many preimages $\mathbf{x} \in \mathbb{Z}^m$ for random images, an adversary can (harmlessly) learn the trivial lattice where these preimages are from: $\mathbb{Z}^m$.

Neither of the above hold in the partial trapdoor setting. First, the distribution of a partial preimage $\mathbf{x}_j$ generated by a (secret) j-th partial trapdoor $\mathbf{U}_j$ is a priori unclear: it does not come from $\mathbb{Z}^m$, but some sublattice $\Lambda_j \subset \mathbb{Z}^m$ which is supposedly dependent on $\mathbf{U}_j$ and not fixed. Indeed, it is not even clear what the "distribution of the lattice Λ_j" is. Second and critically, it seems difficult to sample even a single vector from any such sublattice Λ_j without knowing any secrets about $\mathbf{A}$: if we sample a short vector $\mathbf{x}_j$ over $\mathbb{Z}^m$, it is highly unlikely that $\mathbf{x}_j$ falls into the desired sublattice Λ_j, e.g. it is unlikely that $\mathbf{A} \cdot \mathbf{x}_j \equiv (\mathbf{y}_0 \| \mathbf{0}) \bmod q$ for any $\mathbf{y}_0$. Third, adding to this list of challenges, it is not clear how to simulate the partial trapdoors for corrupt parties without already solving the inhomogeneous SIS problem. Finally, can we still meaningfully consider the hardness of an LWE problem, when any single party j is now able to distinguish an LWE sample $\mathbf{s}^\intercal \cdot \mathbf{A} + \mathbf{e}^\intercal \bmod q$ by sampling a preimage of $\mathbf{0}$ using its partial trapdoor $\mathbf{U}_j$?

1.1 Our Contributions

We resolve the questions above, conceptually and constructively.

[4] We write $(\mathbf{a} \| \mathbf{b})$ for $(\mathbf{a}^\intercal | \mathbf{b}^\intercal)^\intercal$, i.e. for stacking $\mathbf{a}$ on top of $\mathbf{b}$.

A new notion. We introduce the notion of "partial lattice trapdoors", trapdoors that allow sampling preimages to images belonging to designated subspaces. To enable plug-and-play thresholdisation of typical constructions based on lattice trapdoors, we endow partial trapdoors with two security properties: one-wayness and indistinguishability, which are natural analogues of the hardness of the (inhomogeneous) SIS and LWE problems for standard (full-)trapdoored lattices. Given their universal nature, partial trapdoors and plausible future improvements to them directly lead to better threshold variants of many lattice-based primitives. To demonstrate their utility, we give two example black-box applications of partial trapdoors in Section 6: threshold signatures and threshold identity-based encryption (IBE), which are straightforward adaptations of those from [GPV08].

A construction and new proof techniques. We provide a simple construction of partial lattice trapdoors and formally prove that our construction achieves both security properties, one-wayness and indistinguishability. Moreover, our construction achieves non-trivial efficiency, with public parameter and partial (and full) preimage sizes dependent only on the recovery threshold t, but otherwise independent of the total number k of shareholders.

Our security proofs rely on the varying-width κ-(M)SIS and κ-(M)LWE assumptions[5] respectively, which are natural generalisations of the respective assumptions to support hints with varying Gaussian widths. To gain confidence in this generalisation, we generalise an existing reduction [LPSS14] for the fixed-width case and show that the varying-width κ-SIS assumption is implied by the standard SIS assumption.

Our security proofs make use of a variety of techniques which may be of independent interest. In particular, the proofs require non-trivial analyses on lattice subspaces which are, to our knowledge, a setting not commonly seen in the literature. Moreover, in order to simulate partial preimages, we borrow ideas of the BASIS trapdoor sampling technique (but, critically, not the BASIS assumption) from [WW23] originally used for constructing functional commitments, demonstrating the usefulness of such techniques also in security proofs.

1.2 Related Work

This work presents an algebraic method for non-interactively distributing lattice preimage sampling. To the best of our knowledge, the only existing technique allowing the latter is the universal thresholdiser [BGG+18], which requires homomorphically evaluating the trapdoor preimage sampling algorithm. The work of [BKP13] considers sharing a lattice trapdoor, but either requires the trapdoor owner to pre-compute numerous one-time shared randomness to be consumed in the online phase or the parties to perform heavy interactive computation. A recent line of work [GKS24, DKM+24, EKT24, BKL+25] studied constructions of lattice-based threshold signatures without trapdoors, but all of them require at least two rounds of interaction. On the primitive level, [Wee21] proposed the notion of half trapdoors, which are only required to satisfy much weaker properties than partial trapdoors. We discuss the comparisons in more detail below.

Universal Thresholdiser. A general framework for thresholdising many primitives, called the universal thresholdiser, was introduced in [BGG+18], based on threshold fully-homomorphic encryption (FHE). Within this framework, [ASY22] represents the state of the art for round-optimal threshold lattice signatures. The idea is to use threshold FHE to evaluate the signing algorithm – a variant of Dilithium [LDK+22] in the case of [ASY22] – homomorphically. However, when it comes to concrete efficiency, this line of research is far from satisfying – the parameter size required are unrealistic in practice, let alone the expensive homomorphic evaluation of the signing circuit (which requires de-randomisation and additional pseudorandom function evaluations).

Homomorphic and Functional Secret Sharing. There also exists lines of work that consider homomorphic secret sharing schemes [BGI16a, COS+22] and functional secret sharing schemes [BGI15, BGI16b]. The former allows to evaluate arbitrary functions on secret-shared inputs and provides a reconstruction algorithm

[5] The (fixed-width) κ-MSIS and κ-MLWE are the generalisation of κ-SIS and κ-LWE [BF11, LPSS14] to modules over cyclotomic number rings. The κ-SIS and κ-LWE assumptions are known to be implied by the plain SIS and LWE assumptions respectively [BF11, LPSS14]. In most of this work, we consider modules of cyclotomic number rings for generality and all results apply to the integer setting as a special case.

for recovering the output value. The latter allows to evaluate secret-shared function on some input. Mapped to our setting, these works consider the t-out-of-t threshold scenario but not the t-out-of-k. The construction in [COS+22] is based on FHE, too, and in our setting relies on homomorphically evaluating the signing circuit. Thus, similar caveats as for the universal thresholdiser above apply.

Secure Multiparty Computation. Another generic approach of using secure multi-party computation (MPC) techniques and in particular linear secret sharing schemes (LSSS) has been explored in [BKP13]. This area exploits the linearity of trapdoor sampling in the framework of [MP12], which enables linear-secret sharing if certain perturbation-correction vectors are pre-computed (and consumed per signature).[6] However, the approach requires the secret dealer to pre-compute as many perturbation vectors as the number of preimages to be sampled distributedly, meaning that either the number of preimages to be given out is bounded or the secret dealer needs to participate periodically.

The work most similar to ours is [BKP13], so compare in more detail here. In [BKP13], a trusted setup, KeyGen, which generates a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and a trapdoor td which is additively secret shared between the parties. Then, to sample preimages in $\mathsf{SampPre}$ the trusted party generates some number B of perturbation vectors, which, too, are secret-shared to the parties. Moreover, $n \cdot q \cdot B$ additional pre-images are pre-computed by the trusted party and secret-shared to each party. In the online phase, the parties then linearly combine their local shares and the final signature the output of secret-sharing reconstruction. A second protocol removes the trusted setup, but requires generic MPC for some heavy offline interactive computation.

In comparison, our work does not require the pre-computation which limits signing to B queries. This also reduces the size of our shares in comparison to [BKP13] which needs to output $n \cdot q \cdot B$ preimages as part of the shares. As a corollary, as written, this limits [BKP13] to $q = \mathsf{poly}(\lambda)$ where we support larger q. On the other hand, partial pre-image generation in [BKP13] does not require knowledge of the signing set T, but our work requires this.

Tailored Constructions. A different family of threshold constructions directly incorporates a linear-secret-sharing scheme (LSSS) into a non-threshold primitive. Roughly speaking, any algorithm which computes a noisy linear function of the secret input can be non-interactively thresholdised by secret-sharing the secret with an LSSS and homomorphically evaluating the noisy linear function on the shares. Provided that the recovery algorithm of the LSSS computes linear functions with small coefficients of the shares, a noisy approximation of the function output can be recovered.

In the lattice setting, however, norm growth needs to be handled with care. For example, [DKM+24] explained that, in their threshold signatures template, naively using an LSSS with large recovery coefficients and without an additional masking would result in an insecure scheme. When interaction is possible, e.g. in typical application scenarios of threshold signatures, the issue of having large recovery coefficients can be dealt with via interactions that allow to introduce uniform masking terms to the shares, as done in e.g. [DKM+24, EKT24, BKL+25] for their Schnorr-style signatures.

In settings where interactions and/or masking is infeasible, there has been two major solutions: to use an LSSS with very small sharing- and/or reconstruction coefficients, e.g. $\{0,1\}$-LSSS [BGG+18], or to use the traditional Shamir's LSSS but with carefully chosen evaluation points over the ring, the latter also known as subtractive sets [AL21]. In both approaches, the norm growth can be controlled (in the sense that it can be upper-bounded), but each of them come with additional overheads: LSS with small sharing- and/or reconstruction coefficients mostly come with very large number of shares per party (e.g. $O(k^{4.3})$ for $\{0,1\}$-LSSS [Val84]), whereas for Shamir's LSSS over subtractive sets the existing norm-bound is exponential in the threshold t. These overheads make either approaches barely practical in many applications. An exception is LSSS for the majority function of [HMP06], which can emulate a threshold function with binary recovery coefficients with the number of shares per party asymptotically $O(k^{\sqrt{2}})$, and is used for example in [DLN+21] for constructing threshold public-key encryption.

Notably, all of these approaches do not thresholdise non-linear functions, e.g. trapdoor preimage sampling, non-interactively.

Half Trapdoors. On the level of techniques, [Wee21] appears most similar to ours. There, the author introduced a half trapdoor $\mathbf{A} \cdot \mathbf{T}_{1/2} \equiv [\mathbf{0} \| \mathbf{G}]$; thus, $\mathbf{T}_{1/2}$ matches $\mathbf{T}_1$ in our introduction. It was shown in [Wee21] that

[6] This is somewhat analogous to Beaver triples in generic MPC applications.

LWE w.r.t $\mathbf{A}$ remains hard in the presence of an oracle outputting certain preimages sampled using $\mathbf{T}_{1/2}$. In this work, we consider stronger properties: we hand out the partial trapdoor $\mathbf{T}_1$ and also simulate preimage queries for the image space of $\mathbf{T}_0$.

2 Technical Overview

For the purpose of this technical overview, we consider matrices and vectors over $\mathcal{R} = \mathbb{Z}$ and $\mathcal{R}_q = \mathbb{Z}_q$, where q is a prime, while keeping in mind that all results – except our reduction from SIS to κ-SIS – generalise to any cyclotomic ring $\mathcal{R} = \mathbb{Z}[\zeta]$ and $\mathcal{R}_q = \mathbb{Z}_q[\zeta]$ for any primitive root of unity $\zeta \in \mathbb{C}$ and $q \in \mathbb{N}$.

2.1 Partial Lattice Trapdoors

Lattice trapdoor algorithms [GPV08, MP12] allow to sample a matrix $\mathbf{A}$ together with a trapdoor td. Using a trapdoor, we can efficiently sample a short preimage $\mathbf{x}$ of $\mathbf{A}$ for any target vector $\mathbf{y}$. Without, such short preimages are hard to find based on the Short Integer Solution (SIS) assumption [Ajt96]. Lattice trapdoors are usually described as a tuple of algorithms $(\mathsf{TrapGen}, \mathsf{SampPre})$:

- $\mathsf{TrapGen}$: Generate a wide (pseudorandom) matrix $\mathbf{A}$ with a trapdoor td.
- $\mathsf{SampPre}$: Given the trapdoor td and any target vector $\mathbf{y}$, sample a short Gaussian vector $\mathbf{x}$ subject to $\mathbf{A} \cdot \mathbf{x} \equiv \mathbf{y} \bmod q$. The target vector $\mathbf{y}$ is also called an image and $\mathbf{x}$ is called the preimage.

We consider partial lattice trapdoors by extending ordinary lattice trapdoors with an additional triple of algorithms $(\mathsf{PTrapGen}, \mathsf{PSampPre}, \mathsf{Rec})$:

- $\mathsf{PTrapGen}$: Given a trapdoor td, generate k partial trapdoors $\mathsf{ptd}_0, \ldots, \mathsf{ptd}_{k-1}$.
- $\mathsf{PSampPre}$: Given a partial trapdoor ptd_j, a set of t parties $T \subseteq_t [k]$ that party j belongs to, and any target vector $\mathbf{y}$, sample a partial preimage $\mathbf{x}_j$.
- Rec: Recover a preimage $\mathbf{x}$ of $\mathbf{y}$ from partial preimages $(\mathbf{x}_j : j \in T)$.

Simple k-out-of-k Setting. We begin by recalling the simple k-out-of-k partial trapdoor example sketched in Section 1, where all parties need to contribute a partial preimage to recover a preimage. For a matrix $\bar{\mathbf{A}} \in \mathbb{Z}_q^{nk \times mk}$, its partial trapdoor for each party $j \in [k]$ is a short $\mathbf{U}_j \in \mathbb{Z}^{mk \times m}$ such that

$$\bar{\mathbf{A}} \cdot (\mathbf{U}_1, \ldots, \mathbf{U}_k) \equiv \mathbf{I}_k \otimes \mathbf{G}_n \bmod q, \tag{1}$$

where $\mathbf{G} = \mathbf{I}_n \otimes \mathbf{g}^\mathsf{T}$ is the gadget matrix. Notice that $(\mathbf{U}_1, \ldots, \mathbf{U}_k)$ jointly is a (full) gadget-trapdoor[7] of $\bar{\mathbf{A}}$. That is, the above is literally partitioning the full gadget-trapdoor of $\bar{\mathbf{A}}$ into k column chunks and handing one to each party j. Moreover, denoting the j-th canonical vector by $\iota_j \in \{0,1\}^k$, we have

$$\bar{\mathbf{A}} \cdot \mathbf{U}_j \equiv \iota_j \otimes \mathbf{G}_n \bmod q.$$

Therefore, what we have done is effectively partitioning the image space $\mathbb{Z}_q^{nk}$ into k subspaces, so that the j-th subspace, which a partial trapdoor $\mathbf{U}_j$ is able to sample preimages for, is spanned (over $\mathbb{Z}_q$) by $\iota_j \otimes \mathbf{G}_n$, equivalently by $\iota_j \otimes \mathbf{I}_n$.

To sample a preimage of a target vector $\mathbf{y}$, we partition it into k chunks $\mathbf{y} = (\mathbf{y}_0 \| \ldots \| \mathbf{y}_{k-1})$, and let the j-th party locally samples a partial preimage $\mathbf{x}_j$ satisfying $\bar{\mathbf{A}}_j \cdot \mathbf{x}_j \equiv \iota_j \otimes \mathbf{y}_j \bmod q$, which is possible since $\iota_j \otimes \mathbf{y}_j$ is spanned by $\iota_j \otimes \mathbf{G}_n$. Summing the partial preimages yields a short (full) preimage $\mathbf{x} = \sum_{j \in [k]} \mathbf{x}_j$ satisfying $\bar{\mathbf{A}} \cdot \mathbf{x} \equiv \sum_{j \in [k]} \bar{\mathbf{A}} \cdot \mathbf{x} \equiv \sum_{j \in [k]} \iota_j \otimes \mathbf{y}_j \equiv \mathbf{y} \bmod q$.[8]

[7] To recall, a gadget-trapdoor, or simply $\mathbf{G}$-trapdoor, is any short matrix $\mathbf{U}$ such that $\bar{\mathbf{A}} \cdot \mathbf{U} \equiv \mathbf{G} \bmod q$ [MP12].

[8] Another simple k-out-of-k partial trapdoor example is to let $\bar{\mathbf{A}}$ be a block-diagonal matrix with k blocks $(\mathbf{A}_j)_{j \in [k]}$, each block sampled with an individual full trapdoor $\mathbf{U}_j$. To sample preimage of $\mathbf{y}$, let it be partitioned into k chunks $\mathbf{y} = (\mathbf{y}_0 \| \ldots \| \mathbf{y}_{k-1})$, each party j sample a partial preimage s.t. $\mathbf{A}_j \cdot \mathbf{x}_j = \mathbf{y}_j \bmod q$, and recovery is the trivial concatenation.

Extending to t-out-of-k Setting. We can extend the above to the t-out-of-k setting for threshold $t \leq k$, where the rows of the public matrix $\bar{\mathbf{A}}$ are "compressed" by the share-generating matrix of a secret-sharing scheme. Specifically, now let the public matrix be

$$\mathbf{A} \coloneqq \left(\mathbf{V} \otimes \mathbf{I}_n\right) \cdot \bar{\mathbf{A}} \in \mathbb{Z}_q^{nt \times mk},$$

where $\bar{\mathbf{A}} \in \mathbb{Z}_q^{nk \times mk}$ and

$$\mathbf{V} \coloneqq \left(\mathbf{v}_0\ \mathbf{v}_1\ \ldots\ \mathbf{v}_{k-1}\right) \coloneqq \begin{pmatrix} 1 & 1 & \ldots & 1 \\ v_0 & v_1 & \ldots & v_{k-1} \\ \vdots & \vdots & \ddots & \vdots \\ v_0^{t-1} & v_1^{t-1} & \ldots & v_{k-1}^{t-1} \end{pmatrix} \in \mathbb{Z}_q^{t \times k} \tag{2}$$

is the (column style) Vandermonde matrix generated from them. This Vandermonde matrix corresponds to t-out-of-k Shamir secret-sharing with polynomial evaluation points $(v_j)_{j \in [k]}$. Each partial trapdoor $\mathbf{U}_j$ is again such that Eq. (1) holds. Notice that by construction it holds

$$\mathbf{A} \cdot (\mathbf{U}_1, \ldots, \mathbf{U}_k) \equiv (\mathbf{V} \otimes \mathbf{I}_n) \cdot (\mathbf{I}_k \otimes \mathbf{G}_n) \equiv (\mathbf{v}_0 \otimes \mathbf{G}_n, \ldots, \mathbf{v}_k \otimes \mathbf{G}_n) \bmod q.$$

That is, we again partition the image space $\mathbb{Z}_q^{nt}$ into k subspaces, but this time the j-th subspace corresponding to $\mathbf{U}_j$ is spanned (over $\mathbb{Z}_q$) by $\mathbf{v}_j \otimes \mathbf{G}_n$, equivalently by $\mathbf{v}_j \otimes \mathbf{I}_n$. In other words, we have secret-shared the image space $\mathbb{Z}_q^{nt}$ to the k parties via $\mathbf{V}$.

Next, for any set $T \subseteq_t [k]$ and any target image $\mathbf{y}$, we decompose the latter as $\mathbf{y} \equiv \sum_{j \in T} \mathbf{v}_j \otimes \mathbf{z}_j \bmod q$, which is possible since any subset of t columns of $\mathbf{V}$ form an invertible matrix over $\mathbb{Z}_q$.[9] Each partial preimage $\mathbf{x}_j$ is set to be a short solution to the equation $\mathbf{A} \cdot \mathbf{x}_j \equiv \mathbf{v}_j \otimes \mathbf{z}_j \bmod q$, which can be found given $\mathbf{U}_j$, since the image $\mathbf{v}_j \otimes \mathbf{z}_j$ lies in the $\mathbb{Z}_q$-subspace spanned by $\mathbf{v}_j \otimes \mathbf{G}_n$. Letting the full preimage $\mathbf{x}$ be the sum of $\mathbf{x}_j$ for $j \in T$, preimage recovery correctness follows from

$$\mathbf{A} \cdot \mathbf{x} \equiv \sum_{j \in T} \mathbf{A} \cdot \mathbf{x}_j \equiv \sum_{j \in T} \mathbf{v}_j \otimes \mathbf{z}_j \equiv \mathbf{y} \bmod q. \tag{3}$$

Note that, in this example, the public matrix $\mathbf{A}$ has size scaling with tk.

2.2 Our Partial Lattice Trapdoors Construction

Building upon the above ideas, we describe our partial trapdoor construction which achieves non-trivial efficiency, namely the size of the matrix $\mathbf{A}$ scales with t^2 but is otherwise independent of k.[10]

Recall that in the above, the tuple of all partial trapdoors $(\mathbf{U}_1, \ldots, \mathbf{U}_k)$ is itself a full (gadget-)trapdoor of $\mathbf{A}$. Our first observation is that, this condition can be lifted without altering functionality, so long as for each $\mathbf{U}_j$ the relation $\mathbf{A} \cdot \mathbf{U}_j = \mathbf{v}_j \otimes \mathbf{G}_n \bmod q$ continues to hold. In particular, we may adopt an $\mathbf{A}$ with fewer columns, which yields smaller parameters.

How narrow can we make $\mathbf{A}$? To answer this, we consider a reasonable security model, where an adversary may corrupt up to a set $\mathcal{C}$ of $t-1$ parties and receive their partial trapdoors $(\mathbf{U}_j)_{j \in \mathcal{C}} \in (\mathbb{Z}^{mt \times m})^{t-1}$, in which case security would solely be based on the only remaining party $j^* \notin \mathcal{C}$ in a recovery set T – it must be that, solving the ISIS problem of $\mathbf{A} \cdot \mathbf{x}_{j^*} \equiv \mathbf{v}_{j^*} \otimes \mathbf{y}_{j^*} \bmod q$ remains sufficiently hard even with knowledge of $(\mathbf{U}_j)_{j \in \mathcal{C}}$. Notice that this is asking for a solution $\mathbf{x}_{j^*}$ which falls into the m-dimensional $\mathbb{Z}$-sublattice generated by the hidden $\mathbf{U}_{j^*}$. From this, we may conclude that $\mathbf{A}$ must be of dimension at least $nt \times mt$, where m is large enough for the above ISIS problem to remain hard.

Putting these observations together, we obtain the following template, which will become the core of our final construction in Section 5.2.

[9] In the ring setting this is also possible, by pick the entries v_j's of $\mathbf{V}$ to be such that their differences are invertible over the ring [AL21, ACX21].

[10] We are ignoring polylogarithmic factors here.

Partial Trapdoors Generation. For each $j \in [k]$, the partial trapdoor is a short matrix $\mathbf{U}_j \in \mathbb{Z}^{mt \times m}$ satisfying[11]

$$\mathbf{A} \cdot \mathbf{U}_j \equiv \mathbf{v}_j \otimes \mathbf{G}_n \bmod q$$

where $\mathbf{A} \in \mathbb{Z}_q^{nt \times mt}$. To obtain these $\mathbf{U}_j$'s, we sample $\mathbf{A}$ together with a full trapdoor $\mathbf{T}_\mathbf{A}$, then sample each $\mathbf{U}_j$ subject to the above constraint using $\mathbf{T}_\mathbf{A}$.

Partial Preimage Sampling. Same as the t-out-of-k example in Section 2.1, for any recovery set $T \subseteq_t [k]$ and any target image $\mathbf{y}$, we decompose the latter as $\mathbf{y} \equiv \sum_{j \in T} \mathbf{v}_j \otimes \mathbf{z}_j \bmod q$, where $\mathbf{v}_j$ is the j-th column of the Vandermonde matrix in Eq. (2). A partial preimage $\mathbf{x}_j$ is such that

$$\mathbf{A} \cdot \mathbf{x}_j \equiv \mathbf{v}_j \otimes \mathbf{z}_j \bmod q, \tag{4}$$

and the full preimage $\mathbf{x}$ is the sum of $\mathbf{x}_j$ for $j \in T$. Full preimage recovery correctness follows from Eq. (3).

Given $\mathbf{U}_j$, a solution $\mathbf{x}_j$ to Eq. (4) can be obtained, for example, by simply outputting $\mathbf{x}_j \coloneqq \mathbf{U}_j \cdot \mathbf{G}^{-1}(\mathbf{z}_j)$ where $\mathbf{G}^{-1}$ denotes the binary decomposition operator, with which it holds $\mathbf{A} \cdot \mathbf{x}_j \equiv \mathbf{A} \cdot \mathbf{U}_j \cdot \mathbf{G}^{-1}(\mathbf{z}_j) \equiv (\mathbf{v}_j \otimes \mathbf{G}) \cdot \mathbf{G}^{-1}(\mathbf{z}_j) \equiv \mathbf{v}_j \otimes \mathbf{z}_j \bmod q$ and $\mathbf{x}_j$ is short. However, this simple approach is insecure, in the sense that the partial trapdoor $\mathbf{U}_j$ can be easily recovered by linear algebra after seeing a sufficient number of partial preimages. Instead, we use the public trapdoor of $\mathbf{G}$ to sample a Gaussian distributed $\mathbf{d}_j$ subject to $\mathbf{G} \cdot \mathbf{d}_j \equiv \mathbf{z}_j \bmod q$, where the covariance matrix of the Gaussian distribution is parametrised by $\mathbf{U}_j$ in such a way that $\mathbf{U}_j \cdot \mathbf{d}_j$ is distributed as a spherical Gaussian vector over (a suitable sublattice coset of) $\Lambda(\mathbf{U}_j) = \mathbf{U}_j \cdot \mathbb{Z}^m$.

Overall, this yields a partial preimage sampling procedure using $\mathbf{U}_j$ as follows:

1. For a target image vector $\mathbf{y}$, decompose to $\mathbf{y} \equiv \sum_{j \in T} \mathbf{v}_j \otimes \mathbf{z}_j \bmod q$ according to the collaborating set T.
2. Sample a Gaussian distributed $\mathbf{d}_j$ subject to $\mathbf{G} \cdot \mathbf{d}_j \equiv \mathbf{z}_j \bmod q$ and $\mathbf{U}_j \cdot \mathbf{d}_j$ being distributed as spherical Gaussian over $\Lambda(\mathbf{U}_j)$,
3. Output $\mathbf{U}_j \cdot \mathbf{d}_j$.

We verify that $\mathbf{A} \cdot \mathbf{x}_j \equiv \mathbf{A} \cdot \mathbf{U}_j \cdot \mathbf{d}_j \equiv (\mathbf{v}_j \otimes \mathbf{G}) \cdot \mathbf{d}_j \equiv \mathbf{v}_j \otimes \mathbf{z}_j \bmod q$, as desired. Moreover, since $\mathbf{U}_j \cdot \mathbf{d}_j$ is a spherical Gaussian vector, intuitively only information about $\Lambda(\mathbf{U}_j)$ but not $\mathbf{U}_j$ is leaked.

The more challenging task is to argue for security. We walk through how we address the two main concerns highlighted in Section 1 – on the distribution of partial preimages and their sampleability in security proofs – in the following Sections 2.3 and 2.4 respectively.

2.3 Distribution of Real and Ideal Partial Preimages

To understand the distribution of partial preimages induced from the scheme in Section 2.2, we borrow the proof strategy from [GPV08, MP12] for their full lattice trapdoors, which translates as follows: We want to argue that, for fixed $(\mathbf{A}, \mathbf{v}_j)$ the following joint distributions of partial preimage-image pairs $(\mathbf{x}_j, \mathbf{z}_j)$ are indistinguishable:

1. (Real preimages.) Sample $\mathbf{z}_j$ uniformly at random. Then, use $\mathsf{ptd}_j = \mathbf{U}_j$ to sample short $\mathbf{x}_j = \mathbf{U}_j \cdot \mathbf{d}_j$ such that $\mathbf{A} \cdot \mathbf{x}_j \equiv \mathbf{v}_j \otimes \mathbf{z}_j \bmod q$.
2. (Ideal preimages.) Sample a Gaussian $\mathbf{x}_j$ from a public "ambient lattice", independent of $\mathbf{U}_j$. Somehow compute $\mathbf{z}_j$ satisfying $\mathbf{A} \cdot \mathbf{x}_j \equiv \mathbf{v}_j \otimes \mathbf{z}_j \bmod q$.

In the classical full trapdoor setting, where there is only one party and the subscript j can be dropped, we can think of $\mathbf{v} = 1 \in \mathbb{Z}_q$ as being a single element, and the ambient lattice is $\mathbb{Z}^{mt} = \mathbb{Z}^m$. In this case, the ideal distribution of a preimage $\mathbf{x}$ is simply a wide enough Gaussian distribution over $\mathbb{Z}^m$, due to the well-known regularity lemma stating that $\mathbf{A} \cdot \mathbf{x} \bmod q$ is statistically close to uniform over $\mathbb{Z}_q^n$. In particular, it is likely that $\{\mathbf{A} \cdot \mathbf{x} \bmod q : \mathbf{x} \in \mathbb{Z}^m\}$ covers the entire image space $\mathbb{Z}_q^n$.

[11] Proving security with $\mathbf{G}$ as the image turns out to be tricky. In the main body, we instead sample some random $\mathbf{C}$ together with a trapdoor instead of using a fixed $\mathbf{G}$, which is functionally equivalent.

With the above analogy in mind, we examine the requirements for ambient lattices in our partial trapdoor scheme. We observe that, from sufficiently many real partial preimages $(\mathbf{x}_{j,0}, \ldots, \mathbf{x}_{j,L-1})$ where $\mathbf{x}_{j,\ell} = \mathbf{U}_j \cdot \mathbf{d}_{j,\ell}$, an adversary can learn that all partial preimages produced from $\mathbf{U}_j$ belong to a rank-m sublattice

$$\Lambda(\mathbf{U}_j) \coloneqq \left\{\mathbf{x} \in \mathbb{Z}^{mt} : \exists\, \mathbf{d} \in \mathbb{Z}^m,\ \mathbf{x} = \mathbf{U}_j \cdot \mathbf{d}\right\}$$

of the rank-mt lattice

$$\Lambda_j \coloneqq \left\{\mathbf{x} \in \mathbb{Z}^{mt} : \mathbf{A} \cdot \mathbf{x} \in \mathbb{Z}_q\text{-}\mathsf{span}(\mathbf{v}_j \otimes \mathbf{I}_n)\right\}. \tag{5}$$

A natural choice of the ambient lattice is thus a random rank-m sublattice of Λ_j. This yields the following concrete candidate ideal distribution:

2. (Ideal preimages.) Fix a random rank-m sublattice $\Lambda_j^* \subseteq \Lambda_j$ to be the ambient lattice. Sample a Gaussian $\mathbf{x}_j$ from Λ_j^*. Compute the unique $\mathbf{z}_j$ satisfying $\mathbf{A} \cdot \mathbf{x}_j \equiv \mathbf{v}_j \otimes \mathbf{z}_j \bmod q$.

Indeed, by suitably generalising existing regularity lemmas, invoking lattice smoothing lemmas, and picking appropriate Gaussian parameters, we are able to prove that the real and ideal partial preimages distributions stated above are statistically close. For more details on this we refer to Section 5.4 (D_0, D_1, D_2 therein).

2.4 Sampling Partial Preimages in Security Proofs

To prove security of any reasonable application of partial trapdoors, we aim to show that the (inhomogeneous) SIS and/or LWE problems w.r.t. $\mathbf{A}$ remain hard, even when given the partial trapdoors of (at most $t-1$) corrupt parties and many partial preimages generated by honest parties. We shall say that the partial lattice trapdoor has *one-wayness*, if (inhomogeneous) SIS w.r.t. $\mathbf{A}$ is hard; we say that it has *indistinguishability*, if LWE w.r.t. some random vector $\mathbf{y}^*$ in the appropriate space is hard even when given LWE samples w.r.t. $\mathbf{A}$.[12]

In the full trapdoor setting, these guarantees follow immediately from the standard SIS and LWE assumptions respectively, since preimages can be simulated without any secret information, as their ideal distribution is simply Gaussian over $\mathbb{Z}^m$. In contrast, as mentioned in Section 1, in our setting it is not immediately clear how to sample even a single vector from the lattice Λ_j in Eq. (5): If we sample a Gaussian vector $\mathbf{x}_j$ from $\mathbb{Z}^{mt}$, it is highly unlikely that $\mathbf{A} \cdot \mathbf{x}_j \bmod q$ falls into the span of $\mathbf{v}_j \otimes \mathbf{I}_n$, i.e. into $\mathbf{A} \cdot \Lambda_j \bmod q$. Furthermore, it is in fact impossible to have a reduction, from the SIS or LWE problem with respect to $\mathbf{A}$, to the same problem in presence of even one corrupt party, since it is easy to derive a short vector in the kernel of $\mathbf{A}$ from any partial trapdoor.

We resolve the above difficulties by instead relying on the κ-SIS and κ-LWE assumptions respectively, tentatively for $\kappa = m \cdot (t-1)$, where the reduction receives from the problem instance additionally as hints a set of Gaussian vectors $\mathbf{U}_{\neq 0}$ in the kernel of a public matrix $\tilde{\mathbf{A}}_0$. In essence, our reduction will, on the one hand, give away the $m \cdot (t-1)$ hints available from the κ-SIS/κ-LWE problem instance to the adversary as partial trapdoors for the set $\mathcal{C}$ of corrupt parties, and, on the other hand, use these hints to derive a BASIS trapdoor [WW23] for each honest party $j \notin \mathcal{C}$, using which the reduction can sample (ingenuine) partial trapdoors of honest parties. These allow to sample partial preimages following the ideal distribution stated in Section 2.3.[13]

Below we sketch our strategies for reducing the κ-SIS problem to the problem of breaking one-wayness of our partial trapdoor scheme sketched in Section 2.2. Most of the techniques are shared by the reduction from the κ-LWE problem to breaking indistinguishability, and we highlight the differences afterwards.

One-wayness. Our task is to design a reduction which solves the $(m \cdot (t-1))$-SIS problem, i.e. on input an $(m \cdot (t-1))$-SIS instance $(\tilde{\mathbf{A}}_0, \mathbf{U}_{\neq 0}) \in \mathbb{Z}_q^{n \times mt} \times \mathbb{Z}^{mt \times m(t-1)}$, where

$$\tilde{\mathbf{A}}_0 \cdot \mathbf{U}_{\neq 0} \equiv \mathbf{0}_{n \times m(t-1)} \bmod q \qquad \text{and} \qquad \mathbf{U}_{\neq 0} \approx \mathbf{0}_{mt \times m(t-1)},$$

[12] Requiring LWE samples w.r.t. $\mathbf{A}$ to be pseudorandom is not possible in presence of corruption, since any partial trapdoor holder can sample preimages of $\mathbf{0}$, right-multiply and distinguish by the resulting norm.

[13] Looking ahead, we will see that using the hints for both purposes would result in unfaithful simulation, i.e. the simulated partial preimages sketched in this subsection have too wide Gaussian widths. We solve this problem by doubling the number of hints, with two different sets of Gaussian widths. We expand on this in Section 2.5.

find a short vector $\tilde{\mathbf{x}}$ satisfying $\tilde{\mathbf{A}}_0 \cdot \tilde{\mathbf{x}} = \mathbf{0} \bmod q$ and $\tilde{\mathbf{x}}$ is not in the $\mathbb{Q}$-span of $\mathbf{U}_{\neq 0}$. The reduction interacts with the adversary in the following way:

- The adversary selects a set $\mathcal{C} \subset [k]$ of $|\mathcal{C}| = t-1$ corrupt parties.
- The reduction simulates a matrix $\mathbf{A} \in \mathbb{Z}_q^{nt \times mt}$, partial trapdoors $(\mathbf{U}_j)_{j \in \mathcal{C}}$ for all corrupt parties, as well as a partial preimage oracle which does the following: On input a set $T \subseteq_t [k]$ parties, outputs $((\mathbf{x}_j)_{j \in T}, \mathbf{y})$ where $\mathbf{y} \in \mathbb{Z}_q^{nt}$ is a seemingly uniform target vector and $(\mathbf{x}_j)_{j \in T}$ are partial preimages of $\mathbf{y}$.
- After interacting with the partial preimage oracle, the adversary requests to be challenged on (T^*, i^*), where $T^* \subseteq_t [k]$ and $i^* \in T^* \setminus \mathcal{C}$.
- The reduction returns $((\mathbf{x}_j^*)_{j \in T^* \setminus \{i^*\}}, \mathbf{y}^*)$ where $\mathbf{y}^* \in \mathbb{Z}_q^{nt}$ is a seemingly uniform target vector and $(\mathbf{x}_j^*)_{j \in T^* \setminus \{i^*\}}$ are partial preimages of $\mathbf{y}^*$. Note that one partial preimage, $\mathbf{x}_{i^*}$, is withheld from the adversary.
- The adversary returns a short vector $\mathbf{x}^*$ satisfying $\mathbf{A} \cdot \mathbf{x}^* \equiv \mathbf{y}^* \bmod q$.

We outline the key steps of our reduction. To begin, sample a uniformly random $\tilde{\mathbf{A}}_{\neq 0} \leftarrow \mathbb{Z}_q^{n(t-1) \times mt}$ subject to $\tilde{\mathbf{A}}_{\neq 0} \cdot \mathbf{U}_{\neq 0} \equiv \mathbf{I}_{t-1} \otimes \mathbf{G} \bmod q$. Define

$$\mathbf{A} \coloneqq \left(\mathbf{V}_{\{0\} \cup \mathcal{C}} \otimes \mathbf{I}_n\right) \cdot \begin{pmatrix} \tilde{\mathbf{A}}_0 \\ \tilde{\mathbf{A}}_{\neq 0} \end{pmatrix} \tag{6}$$

where $\mathbf{V}_{\{0\} \cup \mathcal{C}}$ is the Vandermonde matrix defined by $(v_0, v_{i_1}, \ldots, v_{i_{t-1}})$ where $\{i_1, \ldots, i_{t-1}\} \coloneqq \mathcal{C}$. To argue that $\mathbf{A}$ constructed in this way is indistinguishable from a random-looking one, we apply [BF11, Theorem 4.3], a result originally used in their κ-SIS-to-SIS reduction.[14]

Simulating corrupt partial trapdoors. To simulate partial trapdoors for corrupt parties, we partition $\mathbf{U}_{\neq 0} = (\mathbf{U}_{i_1} | \ldots | \mathbf{U}_{i_{t-1}})$ into $t-1$ chunks. Notice that $\mathbf{U}_j \in \mathbb{Z}_q^{mt \times m}$ satisfies $\mathbf{A} \cdot \mathbf{U}_j \equiv \mathbf{v}_j \otimes \mathbf{G}$ for all $j \in \mathcal{C}$ by construction. Set $\mathsf{ptd}_j \coloneqq \mathbf{U}_j$ for $j \in \mathcal{C}$.

Simulating ingenuine partial trapdoors for honest parties. The most interesting part is to simulate the partial trapdoors for honest parties $j \in [k] \setminus \mathcal{C}$. We make use of the BASIS trapdoor sampling technique of [WW23]. For each $j \in [k] \setminus \mathcal{C}$, parse $\mathbf{A}$ as $\mathbf{A} = (\mathbf{A}_0 \| \ldots \| \mathbf{A}_{t-1})$ and define

$$\mathbf{B}_j \coloneqq \left[\begin{array}{c|c} 1 \cdot \mathbf{A}_0 & -\mathbf{G} \\ v_j^{-1} \cdot \mathbf{A}_1 & -\mathbf{G} \\ \vdots & \vdots \\ v_j^{-(t-1)} \cdot \mathbf{A}_{t-1} & -\mathbf{G} \end{array}\right] \in \mathbb{Z}_q^{nt \times m(t+1)}.$$

Borrowing techniques from [WW23], we can turn $\mathbf{U}_{\neq 0}$ satisfying $\mathbf{A} \cdot \mathbf{U}_{\neq 0} = \mathbf{V}_{\{0\} \cup \mathcal{C}} \otimes \mathbf{G}$ into a short matrix $\mathbf{T}$ satisfying

$$\mathbf{B}_j \cdot \mathbf{T} \equiv \hat{\mathbf{H}}_j \otimes \mathbf{G} \bmod q$$

for some matrix $\hat{\mathbf{H}}_j$ invertible over $\mathbb{Z}_q$, i.e. $\mathbf{T}$ is a gadget trapdoor [MP12] of $\mathbf{B}_j$ with tag $\hat{\mathbf{H}}_j$. We refer to Section 5.4 (D_3, D_4 therein) for an overview of how this is achieved, and the proof of Lemma 20 for the details. Using such a trapdoor $\mathbf{T}$, the reduction samples Gaussian $(\mathbf{u}_j \| \mathbf{w}_j)$ such that

$$\mathbf{B}_j \cdot \begin{pmatrix} \mathbf{u}_j \\ \mathbf{w}_j \end{pmatrix} \equiv \mathbf{0} \bmod q \qquad \Longrightarrow \qquad \mathbf{A} \cdot \mathbf{u}_j \equiv \mathbf{v}_j \otimes \mathbf{G} \cdot \mathbf{w}_j \bmod q.$$

Repeating this m times, the reduction obtains a short matrix $\mathbf{U}_j$ and a statistically uniform matrix $\mathbf{C}_j$ such that $\mathbf{A} \cdot \mathbf{U}_j \equiv \mathbf{v}_j \otimes \mathbf{C}_j \bmod q$. The matrix $\mathbf{U}_j$ is thus a short basis of a rank-m sublattice $\Lambda(\mathbf{U}_j) \subseteq \Lambda_j$ of the lattice Λ_j defined in Eq. (5), as desired.

[14] We additionally generalise [BF11, Theorem 4.3] from the integers to the ring and module settings. The proofs are given in Appendix A and may be of independent interest.

Sampling partial preimage oracle. From the above two steps, the reduction now possesses short matrices $\mathbf{U}_j$ such that

$$\mathbf{A} \cdot \mathbf{U}_j = \begin{cases} \mathbf{v}_j \otimes \mathbf{G} \bmod q & j \in \mathcal{C} \\ \mathbf{v}_j \otimes \mathbf{C}_j \bmod q & j \in [k] \setminus \mathcal{C} \end{cases}$$

where $\mathbf{C}_j$ are statistically uniform. To respond to a partial preimage oracle query on any $T \subseteq_t [k]$, the reduction simply samples random Gaussian vectors $\mathbf{x}_j \leftarrow \Lambda(\mathbf{U}_j)$ for all $j \in T$, and then computes $\mathbf{y} \coloneqq \mathbf{A} \cdot \sum_{j \in T} \mathbf{x}_j \bmod q$. Similarly, to simulate the challenge $((\mathbf{x}^*_j)_{j \in T^* \setminus \{i^*\}}, \mathbf{y}^*)$, the reduction uses the same strategy as above, except that it does not return $\mathbf{x}^*_{i^*}$.

Breaking the $(m \cdot (t-1))$-SIS *problem.* Eventually, the adversary returns a short vector $\mathbf{x}^*$ satisfying $\mathbf{A} \cdot \mathbf{x}^* \equiv \mathbf{y}^* \bmod q$, implying $\mathbf{A} \cdot \left(\mathbf{x}^* - \sum_{j \in T^*} \mathbf{x}^*_j\right) \equiv \mathbf{0} \bmod q$. By recalling the construction of $\mathbf{A}$ in Eq. (6) and using the fact that $\mathbf{V}_{\{0\} \cup \mathcal{C}}$ is invertible, this implies

$$\tilde{\mathbf{A}}_0 \cdot \left(\mathbf{x}^* - \sum_{j \in T^*} \mathbf{x}^*_j\right) \equiv \mathbf{0} \bmod q.$$

By arguing about the min-entropy of $\mathbf{x}^*_{i^*}$ and about intersections of independent subspaces, we can conclude that $\tilde{\mathbf{x}} \coloneqq \mathbf{x}^* - \sum_{j \in T^*} \mathbf{x}^*_j$ is not in the span of $\mathbf{U}_{\neq 0}$ with high probability. We refer to Section 5.5 for an overview of this argument. The reduction therefore returns $\tilde{\mathbf{x}}$ as a solution to the $(m \cdot (t-1))$-SIS instance. Overall, for more details on one-wayness, we refer to Section 5.5.

Indistinguishability. We also sketch how to reduce the $(m \cdot (t-1))$-LWE problem to the indistinguishability of our partial trapdoor. Comparing to the previous reduction, now in addition to $(\tilde{\mathbf{A}}_0, \mathbf{U}_{\neq 0}) \in \mathbb{Z}_q^{n \times mt} \times \mathbb{Z}^{mt \times m(t-1)}$ where

$$\tilde{\mathbf{A}}_0 \cdot \mathbf{U}_{\neq 0} \equiv \mathbf{0}_{n \times m(t-1)} \bmod q \qquad \text{and} \qquad \mathbf{U}_{\neq 0} \approx \mathbf{0}_{mt \times m(t-1)},$$

the reduction is further given $\mathbf{b}$ and it is asked to distinguish:

- $\mathbf{b}^\intercal \approx \mathbf{s}_0^\intercal \cdot \tilde{\mathbf{A}}_0 \bmod q$ is an LWE sample, or
- $\mathbf{b}$ is uniform over the "noisy kernel" of $\mathbf{U}_{\neq 0}$: $\{\mathbf{x} : \mathbf{x}^\intercal \cdot \mathbf{U}_{\neq 0} \equiv \mathbf{0} \bmod q\} + \mathsf{noise}$.

It interacts with an adversary against the indistinguishability of the partial trapdoor. More precisely, the adversary distinguishes an LWE challenge $c_1 \approx \mathbf{s}^\intercal \cdot \mathbf{y}^* \bmod q$ for a random vector $\mathbf{y}^*$ from a uniform sample over $c_1 \leftarrow \mathbb{Z}_q^{mt}$, given the matrix $\mathbf{A}$,

(a) partial trapdoors of corrupt parties $(\mathbf{U}_j)_{j \in \mathcal{C}}$,
(b) a preimage oracle,
(c) a set of $t-1$ partial preimages $(\mathbf{x}^*_j)_{j \in T^* \setminus \{i^*\}}$ for the challenge image $\mathbf{y}^*$, where $T^* \subseteq_t [k], i \in T^* \notin \mathcal{C}$ are chosen by the adversary, and
(d) additional LWE samples $\mathbf{c}_0 \approx \mathbf{s}^\intercal \cdot \mathbf{A} \bmod q$.

The matrix $\mathbf{A}$ together with items (a), (b) can all be simulated by the reduction using $\tilde{\mathbf{A}}_0$ and $\mathbf{U}_{\neq 0}$ in the same way as in the previous reduction to one-wayness. We focus on the simulation of items (c), (d) and the challenge c_1.

Simulating $\mathbf{y}^*$ *and* $t-1$ *partial preimages.* For $j \in T^* \setminus \{i^*\}$, the reduction first samples a random preimage $\mathbf{x}_j \leftarrow \Lambda(\mathbf{U}_j)$ and computes $\mathbf{z}_j$ satisfying $\mathbf{A} \cdot \mathbf{x}_j \equiv \mathbf{v}_j \otimes \mathbf{z}_j \bmod q$. It then samples a short $\mathbf{r}$ and sets $\mathbf{y}^*_0 \equiv \tilde{\mathbf{A}}_0 \cdot \mathbf{r} \bmod q$, so that $\mathbf{y}^*_0$ is statistically close to uniform by a regularity lemma. Using these, it finds a solution $(\mathbf{y}^*_{\neq 0}, \mathbf{z}_{i^*})$ satisfying

$$\left(\mathbf{V}_{\{0\} \cup \mathcal{C}} \otimes \mathbf{I}_n\right) \cdot \begin{pmatrix} \mathbf{y}^*_0 \\ \mathbf{y}^*_{\neq 0} \end{pmatrix} \equiv \sum_{j \in T^* \setminus \{i^*\}} \mathbf{v}_j \otimes \mathbf{z}_j + \mathbf{v}_{i^*} \otimes \mathbf{z}_{i^*} \bmod q$$

and sets $\mathbf{y}^* \equiv \left(\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n\right) \cdot \begin{pmatrix} \mathbf{y}_0^* \\ \mathbf{y}_{\neq 0}^* \end{pmatrix} \bmod q$. Notice that the above linear equation is equivalent to $\mathbf{A} \cdot \sum_{j\in T^*} \mathbf{x}_j \equiv \mathbf{y}^* \bmod q$. In other words, the reduction solves for $(\mathbf{y}_{\neq 0}^*, \mathbf{z}_{i^*})$ such that $(\mathbf{z}_j)_{j\in T^*}$ (implicit from $\mathbf{x}_j$'s) all fall into the correct subspaces and jointly satisfy the linear relation with $\mathbf{y}^*$.[15]

Simulating $(\mathbf{c}_0, c_1)$ *from* $\mathbf{b}$. To simulate the LWE challenge $(\mathbf{c}_0, c_1)$, the reduction samples uniform $\mathbf{s}_{\neq 0} \leftarrow \mathbb{Z}_q^{n(t-1)}$ and sets

$$\mathbf{c}_0^\intercal \equiv \mathbf{b}^\intercal + \mathbf{s}_{\neq 0}^\intercal \cdot \tilde{\mathbf{A}}_{\neq 0} \bmod q, \qquad\qquad c_1 \approx \mathbf{b}^\intercal \cdot \mathbf{r} + \mathbf{s}_{\neq 0}^\intercal \cdot \mathbf{y}_{\neq 0}^* \bmod q.$$

Suppose $\mathbf{b}^\intercal \approx \mathbf{s}_0^\intercal \cdot \tilde{\mathbf{A}}_0$, then we observe that

$$\mathbf{c}_0^\intercal \approx \ (\mathbf{s}_0^\intercal, \mathbf{s}_{\neq 0}^\intercal) \cdot \begin{pmatrix} \tilde{\mathbf{A}}_0 \\ \tilde{\mathbf{A}}_{\neq 0} \end{pmatrix} \equiv \underbrace{(\mathbf{s}_0^\intercal, \mathbf{s}_{\neq 0}^\intercal) \cdot \left(\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n\right)^{-1}}_{=:\mathbf{s}^\intercal} \cdot \underbrace{\left(\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n\right) \cdot \begin{pmatrix} \tilde{\mathbf{A}}_0 \\ \tilde{\mathbf{A}}_{\neq 0} \end{pmatrix}}_{\mathbf{A}},$$

where $\mathbf{s}^\intercal$ is uniformly distributed, since $(\mathbf{s}_0^\intercal, \mathbf{s}_{\neq 0}^\intercal)$ is uniform and $\left(\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n\right)^{-1}$ is a bijective map. For c_1, we have

$$\begin{aligned} c_1 &\approx (\mathbf{s}_0^\intercal \cdot \tilde{\mathbf{A}}_0) \cdot \mathbf{r} + \mathbf{s}_{\neq 0}^\intercal \cdot \mathbf{y}_{\neq 0}^* + e \quad \equiv \quad \mathbf{s}_0^\intercal \cdot \mathbf{y}_0^* + \mathbf{s}_{\neq 0}^\intercal \cdot \mathbf{y}_{\neq 0}^* + e \bmod q \\ &\equiv (\mathbf{s}_0^\intercal, \mathbf{s}_{\neq 0}^\intercal) \cdot \left(\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n\right)^{-1} \cdot \left(\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n\right) \cdot \begin{pmatrix} \mathbf{y}_0^* \\ \mathbf{y}_{\neq 0}^* \end{pmatrix} + e \bmod q \\ &\equiv \mathbf{s}^\intercal \cdot \mathbf{y}^* + e \bmod q, \end{aligned}$$

where (implicit in the approximation) we use noise flooding via e, and the first equality is by construction that $\mathbf{y}_0^* \equiv \tilde{\mathbf{A}}_0 \cdot \mathbf{r} \bmod q$. Therefore, the simulated $(\mathbf{c}_0, c_1)$ is either distributed as

$$\text{(1a):} \quad \approx (\mathbf{s}^\intercal \cdot \mathbf{A},\ \mathbf{s}^\intercal \cdot \mathbf{y}^*) \qquad \text{or} \qquad \text{(1b):} \quad (\mathbf{b}^\intercal + \mathbf{s}_{\neq 0}^\intercal \cdot \tilde{\mathbf{A}}_{\neq 0},\ \mathbf{b}^\intercal \cdot \mathbf{r} + \mathbf{s}_{\neq 0}^\intercal \cdot \mathbf{y}_{\neq 0}^*)$$

for $\mathbf{b}$ a random vector in the "noisy kernel" of $\mathbf{A}$, and by the κ-LWE assumption the above are indistinguishable. Similarly, we have

$$\text{(2a):} \quad \approx (\mathbf{s}^\intercal \cdot \mathbf{A},\ \text{random}) \qquad \text{and} \qquad \text{(2b):} \quad (\mathbf{b}^\intercal + \mathbf{s}_{\neq 0}^\intercal \cdot \tilde{\mathbf{A}}_{\neq 0},\ \text{random})$$

are indistinguishable under κ-LWE, by another reduction that is identical except that it outputs a uniformly random c_1. To complete the chain, we show via a statistical argument that, when $\mathbf{b}$ is a sample from the noisy kernel of $\mathbf{U}_{\neq 0}$, then $\mathbf{b}^\intercal \cdot \mathbf{r} \bmod q$ is close to uniform over $\mathbb{Z}_q$, therefore (1b) is indistinguishable from (2b). Putting everything together, we conclude (1a) and (2a) are indistinguishable, implying indistinguishability of our partial trapdoor. We highlight that, the last statistically step requires a new leftover-hash lemma over subspaces, formally stated in Lemma 21 and may be of independent interest. For more details on indistinguishability, we refer to Section 5.6.

Remark 1 (Removing noise flooding). In the indistinguishability proof we applied noise flooding, which requires setting the modulus q to be super-polynomial in the security parameter. This can be avoided by using a κ-LWE analogue of the error-leakage LWE (elLWE) assumption introduced in [DKL+23] to replace the noise flooding argument, as demonstrated in the same work in the context of a laconic encryption construction. The elLWE assumption is shown [DKL+23] to be implied by the standard LWE assumption with only slightly larger parameters, and as discussed in [DKL+23], the reduction is agnostic about the distributions of the LWE secrets and matrices, and any hints related to the LWE matrices. This means essentially the same reduction could show that the κ-LWE analogue of elLWE is implied by κ-LWE, which would then imply indistinguishability of our partial trapdoor under a polynomial modulus. Since this is not the focus of this work, we omit formalising.

[15] As we will shortly see, this form of sampling $\mathbf{y}^*$ allows us to invoke a leftover hash lemma involving $\mathbf{r}$.

2.5 Short ≠ Short, and Varying-Width Hints.

As hinted earlier, the reduction sketches above ignored an issue regarding Gaussian width, which we will now address. Recall that the reduction possesses

$$\mathbf{A} \cdot \mathbf{U}_j = \begin{cases} \mathbf{v}_j \otimes \mathbf{G} \bmod q & j \in \mathcal{C} \\ \mathbf{v}_j \otimes \mathbf{C}_j \bmod q & j \in [k] \setminus \mathcal{C} \end{cases}$$

where $\mathbf{U}_j$ is a subset of κ-MSIS hints if $j \in \mathcal{C}$ and otherwise constructed using the BASIS trapdoor sampling technique explained above. This implies that norms of $\mathbf{U}_j$, while still short, are longer for $j \notin \mathcal{C}$ than for $j \in \mathcal{C}$.

This is not immediately a problem since the adversary does not get to see $\mathbf{U}_j$ for $j \notin \mathcal{C}$. However, it gets to see many samples from $\Lambda(\mathbf{U}_j)$ and we expect this lattice to be distributed differently depending on the norms of the basis $\mathbf{U}_j$. For example, the volume of this lattice will differ. In other words, we expect it to be easy for the adversary to distinguish the above naive simulation from the real security experiment, forbidding us from provably using the adversary's output to solve our $(m \cdot (t-1))$-SIS/LWE problem instance.

To work around this issue, we consider a variant of the κ-SIS/LWE problem for $\kappa = 2m \cdot (t-1)$ with hints $\mathbf{U}_{\neq 0} = (\mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}})$, where half the hints $\mathbf{U}_{\mathtt{Hon}}$ are shorter than the other half $\mathbf{U}_{\mathtt{Cor}}$. We may then use these particularly short hints $\mathbf{U}_{\mathtt{Hon}}$ to generate the partial trapdoors $\mathbf{U}_j$ for the honest parties $j \notin \mathcal{C}$ of matching distributions – of identical width as that of the corrupt trapdoors $\mathbf{U}_{\mathtt{Cor}}$. To gain confidence in this approach, we generalise the SIS-to-k-SIS reduction of [LPSS14] over the integers setting, showing that the κ-SIS problem with hints of varying widths is as hard as the plain SIS problem. We expect a similar result to hold for the κ-LWE problem with hints of varying widths, namely that it is as hard as the plain LWE problem. Our final results are stated under the more general κ-MSIS/-κ-MLWE assumptions over rings, which we conjecture to be as hard as the MSIS/MLWE problems respectively.

2.6 Applications

To illustrate our threshold lattice trapdoor machinery, we give two example applications in Section 6: a GPV-style threshold signature scheme and a GPV-style threshold(-authority) IBE scheme [GPV08].

The former is a straight-forward adaptation of the GPV signature scheme to the threshold setting: For $\mathbf{A}$ being the verification key and a signature being a short preimage of the hash $\mathsf{H}(\mu)$ of a message μ, any t of our k signers, each of whom holding a partial trapdoor, can jointly generate a valid signature. Unforgeability relies on the one-wayness of the partial lattice trapdoor. Same as GPV, we prove security in the Random Oracle model, where we sample random partial preimages $\mathbf{x}_j$ satisfying $\mathbf{A} \cdot \mathbf{x}_j \equiv \mathbf{v}_j \otimes \mathbf{z}_j \bmod q$, compute the resulting image $\mathbf{y} \coloneqq \mathbf{A} \sum_{j \in T} \mathbf{x}_j \bmod q$, and programme $\mathsf{H}(\mu) = \mathbf{y}$ for any query μ.

For the GPV-style IBE, the trapdoor of $\mathbf{A}$ acts as the master secret key of an authority to generate identity keys to users, and the encryptor encrypts w.r.t. $\mathbf{A}$ and a user identity. Applying our technique, we obtain a threshold IBE where a user can decrypt upon collecting t keys from k authorities. To prove CPA security, we rely on the indistinguishability of our partial trapdoors. We emphasise that a threshold IBE from a naive approach (with block diagonal $\mathbf{A}_j$'s and mentioned in a footnote above) is not possible, since a block diagonal $\mathbf{A}$ would necessitate the encryptor to encrypt w.r.t. some specific set T of authorities which contradicts with the functionality. Alternatively, encrypting to all $T \subseteq_t [k]$ limits t to be constant.

We note that we opted to instantiate both of our example applications in rather restricted security games for ease of exposure and discuss extensions lifting these restrictions in Sections 2.7 and 6.

Looking ahead further, our technique can be applied generically to thresholdise other primitives where a lattice trapdoor is used, e.g. for key generation. For example, this is also the case for the classic attribute-based encryption (ABE) of [BGG+14], where a lattice trapdoor acts as the authority's master secret key to generate preimages corresponding to secret keys for specific function policies. Plugging in our partial lattice trapdoor, we generically obtain a threshold(-authority) ABE.

2.7 Open Problems

In this work, we introduce a new partial trapdoor notion which enables non-interactive thresholdisation of any lattice-trapdoor-based primitives and provide the first construction based on standard assumptions. Still,

there is a significant gap between what the construction in this work achieves and what is desirable from a generic building block for lattice-based threshold primitives. We list a few interesting research directions that we believe are avenues for further work. We discuss application-specific limitations and open problems in Section 6.

Asymptotics. Our preimages are linear in the threshold t and our public parameter $\mathbf{A}$ quadratic in t. While, as per our ambition, this means our construction is conceptually sublinear in k, in some applications we have $t = O(k)$. It would be interesting to explore if compression techniques such as those introduced in e.g. [ACL+22, CLM23, WW23, HLL23] can be applied to shrink our parameters.

Moreover, due to the use of noise-flooding, we require a super-polynomial modulus q to achieve indistinguishability (although a polynomial modulus suffices for one-wayness). As discussed in Remark 1, we believe it to be possible to use a variant of the elLWE-assumption [DKL+23] to avoid this.

Obliviousness to Collaborators. Our techniques necessitate that each party releasing a partial preimage must know their collaborators, i.e. the set T, which – while not uncommon in the lattice-based threshold setting [DKM+24, KRT24] – might be a difficult condition in some applications. Ultimately, this is due to the use of secret sharing techniques in the image space instead of the preimage space, leading to large recovery (i.e. Lagrange) coefficients. Techniques like those studied in [AL21] should allow to overcome this limitation.

Indistinguishability of Reconstructed Preimages. While the reconstruction of a preimage $\mathbf{x}$ is simply addition and thus does not need to know the set T of preimage generators, this does not guarantee that T cannot be recovered from $\mathbf{x}$.[16] This is because $\sum_{j\in T} \mathbf{x}_j^*$ is not guaranteed to follow a distribution independent of the lattices $(\Lambda(\mathbf{U}_j))_{j\in T}$. In particular, the lattice from which the preimages $\mathbf{x}$ are sampled may be a strict sublattice $\Lambda((\mathbf{U}_j)_{j\in T})$ of $\mathbb{Z}^{mt}$, allowing to identify T. We give two directions for overcoming this limitation:

1. Zero-Knowledge: Instead of $\mathbf{x}$, a zero-knowledge (ZK) argument of knowledge that one knows a short $\mathbf{x}$ s.t. $\mathbf{A}\cdot\mathbf{x} \equiv \mathbf{y} \bmod q$ is output. Note that the bit size of $\mathbf{x}$ is linear in t but independent of k, and linear-sized proofs for lattice statements have become highly efficient in practice [LNP22]. However, compiling with a random-oracle-based ZK-argument sacrifices proof-friendliness, since then proving knowledge of $\mathbf{x}$ requires proving statements involving random oracles.
2. Sampling: The idea is to exploit that each partial signer already knows T, and adapt the output distribution of $\mathbf{x}_j^*$ such that their sum does not live in a proper sublattice of $\mathbb{Z}^{mt}$, which may require adapting the way we sample $\mathbf{U}_j$. We consider formalising and realising this approach as the more interesting research direction.

General Access Structure. A natural extension of our work would be one for more expressive access structures. For example, one natural idea is to adopt techniques of using $\{0,1\}$-linear secret sharing schemes (LSSS) [BGG+18] in the lattice setting [DKW21]. Unfortunately, the combinatorial aspects of $\{0,1\}$-LSSS do not seem immediately compatible with our approach.

Trusted Setup. Our construction relies on a trusted setup which is undesirable. Depending on the concrete instantiation of the trapdoor for the public matrix $\mathbf{A}$, it might be possible to avoid this. For example, we could use the gadgets from [ENP24].

3 Preliminaries

We write $[k]$ for $\{0,\dots,k-1\}$ and start indexing at zero. For a set T and any $j \in T$, denote by $\hat{\mathbf{e}}_{T,j} \in \{0,1\}^T$ the unit-vector indexed by j. We denote a matrix $(1,\dots,1)^\intercal \in \mathcal{R}^t$ as $\mathbf{1}_t$. For $h \in \mathbb{C}$, write $\bar{h}$ for its complex conjugate. For a matrix $\mathbf{S}$ over $\mathbb{C}$, write $\mathbf{S}^\dagger \coloneqq \bar{\mathbf{S}}^\intercal$ for its conjugate transpose. A matrix $\mathbf{\Sigma}$ is said to be positive

[16] In the context of threshold signatures, this means our GPV-style construction does not guarantee anonymity – one might learn T given a (full) signature $\mathbf{x}$ – which may or may not matter depending on the application.

semi-definite if it can be written as $\mathbf{\Sigma} = \mathbf{S} \cdot \mathbf{S}^\dagger$. We write $\sqrt{\mathbf{\Sigma}}$ for any fixed matrix $\mathbf{S}$ satisfying $\mathbf{\Sigma} = \mathbf{S} \cdot \mathbf{S}^\dagger$, e.g. $\mathbf{S}$. We say that two positive semi-definitive matrices $\mathbf{\Sigma}_0$ and $\mathbf{\Sigma}_1$ satisfy $\mathbf{\Sigma}_0 > \mathbf{\Sigma}_1$ if $\mathbf{\Sigma}_0 - \mathbf{\Sigma}_1$ is positive semi-definite. We also say $\mathbf{\Sigma}_0 > \sigma^2$ for some $\sigma^2 \in \mathbb{R}$ if $\mathbf{\Sigma}_0 - \sigma^2 \cdot \mathbf{I}$ is positive semi-definite.

Let D_0 and D_1 two distributions. The statistical distance SD between D_0 and D_1 over a common domain X is defined as $\mathsf{SD}(D_0, D_1) = \frac{1}{2} \sum_{x \in X} |D_0(x) - D_1(x)|$. We write $D_0 \approx_s D_1$ when $\mathsf{SD}(D_0, D_1) = \mathsf{negl}(\lambda)$.

We write $\mathbf{V}_{\mathcal{C}}$ for the Vandermonde-style matrix where the j-th column is $(v_{c_j}^0, \ldots, v_{c_j}^{t-1})^\intercal$ and c_j is the j-th entry in $\mathcal{C}$ with $|\mathcal{C}| \leq t$. We write $(\mathbf{A} \| \mathbf{B})$ for the matrix obtained by stacking $\mathbf{A}, \mathbf{B}$ vertically. For matrices $\mathbf{C}_j \in \mathcal{R}_q^{n \times m}$ and $j \in [k]$ we write $\mathbf{C}_{\mathcal{C}} \in \mathcal{R}_q^{n|\mathcal{C}| \times m|\mathcal{C}|}$ to signify the block matrix with matrices $\mathbf{C}_{c_j}$ along the main diagonal, i.e. $\mathbf{C}_{\mathcal{C}} \coloneqq \mathsf{diag}\left(\{\mathbf{C}_j\}_{j \in \mathcal{C}}\right)$. If the matrix on the diagonal is the same we denote it as $\mathsf{diag}(\mathbf{C}, k) \in \mathcal{R}_q^{nk \times mk}$ where $k \geq 1$ is the number of repetitions.

3.1 Algebraic Number Theory and Lattices

Let $\mathcal{K} = \mathbb{Q}(\zeta)$ where $\zeta = \zeta_{\mathfrak{f}}$ be the cyclotomic field with conductor $\mathfrak{f}$ and degree $\varphi \coloneqq \varphi(\mathfrak{f})$, and $\mathcal{R} = \mathbb{Z}[\zeta]$ its ring of integers.

Splitting of primes and primitive matrices. For $q \in \mathbb{N}$, write $\mathcal{R}_q \coloneqq \mathcal{R}/q\mathcal{R}$. We generally write $\mathbf{A} \cdot \mathbf{x} \equiv \mathbf{y} \bmod q$ for $\mathbf{A} \in \mathcal{R}_q^{n \times m}, \mathbf{x} \in \mathcal{R}^m, \mathbf{y} \in \mathcal{R}_q^n$ to highlight an equality holds in $\mathcal{R}_q$, but we may drop the suffix "mod q" for space reasons. If $\mathcal{K} \neq \mathbb{Q}$, then we always assume q to be a rational prime unramified in $\mathcal{K}$, and which factors as $\langle q \rangle = \prod_{j=1}^{g} \mathfrak{q}_j$ with norm $\mathcal{N}(\mathfrak{q}_j) = q^{\varphi/g} > \mathsf{poly}(\lambda)$ in $\mathcal{R}$. In other words, $\mathcal{R}_q$ splits into fields $\mathcal{R}_{\mathfrak{q}_j} \coloneqq \mathcal{R}/\mathfrak{q}_j$ with $|\mathcal{R}_{\mathfrak{q}_j}| > \mathsf{poly}(\lambda)$. We note that this is a non-trivial assumption. We make use of this assumption, to apply (a) a regularity lemma for uniformly random matrices $\mathbf{A} \in \mathcal{R}_q^{n \times m}$ (Lemma 12) and to argue that (b) some $\mathbf{A}, \{\mathbf{C}_j\}$ sampled below are primitive with high probability. In both cases, our assumption can be waived by conditioning on such that $\mathbf{A}, \{\mathbf{C}_j\}$ contain invertible submatrices. Since we do not wish to "carry around" these conditional probabilities, we leave this for future work.

Canonical embedding and norms. We write $\sigma = (\sigma_i)_{i \in \mathbb{Z}_{\mathfrak{f}}^\times} : \mathcal{K} \to \mathbb{C}^\varphi$ for the canonical embedding of $\mathcal{K}$, and extend the embedding naturally to $\mathcal{K}$-vectors by concatenation. The norm $\|\cdot\| : \mathcal{K}^m \to \mathbb{R}^{\geq 0}$ is taken to be the ℓ_2-norm over the canonical embedding, i.e. $\|\mathbf{x}\| = \|\sigma(\mathbf{x})\|$. For any matrix $\mathbf{A} \in \mathcal{K}^{m \times n}$ of field elements we use the spectral norm $\|\mathbf{A}\| = s_{\max}(\mathbf{A}) \coloneqq \sup\{\|\mathbf{A} \cdot \mathbf{x}\| : \mathbf{x} \in \mathcal{K}^n, \|\mathbf{x}\| = 1\}$ which is the maximum singular value of $\mathbf{A}$. We also write $s_{\min}(\mathbf{A})$ for the minimum singular value of $\mathbf{A}$.

Module lattices. For $\mathbf{A} \in \mathcal{R}_q^{n \times m}$ and $\mathbf{U} \in \mathcal{R}^{m \times k}$, we write
– $\mathit{\Lambda}_q^\perp(\mathbf{A}) = \{\mathbf{u} \in \mathcal{R}^m : \mathbf{A} \cdot \mathbf{u} \equiv \mathbf{0} \bmod q\}$,
– $\mathit{\Lambda}_q^{\mathbf{v}}(\mathbf{A}) = \{\mathbf{u} \in \mathcal{R}^m : \mathbf{A} \cdot \mathbf{u} \equiv \mathbf{v} \bmod q\}$,
– $\mathit{\Lambda}_q(\mathbf{U}) = \left\{\mathbf{x} \in \mathcal{R}^m : \exists\ \mathbf{d} \in \mathcal{R}^k,\ \mathbf{U} \cdot \mathbf{d} = \mathbf{x} \bmod q\right\}$, and
– $\mathit{\Lambda}(\mathbf{U}) = \left\{\mathbf{x} \in \mathcal{R}^m : \exists\ \mathbf{d} \in \mathcal{R}^k,\ \mathbf{U} \cdot \mathbf{d} = \mathbf{x}\right\}$
viewed as lattice cosets via the canonical embedding σ. The above notation is generalised naturally to ideal moduli $\mathfrak{q}$. When discussing lattice quantities, e.g. the minimum distance $\lambda_1(\cdot)$ and the determinant $\det(\cdot)$, we treat $\mathcal{R}$-modules and $\mathcal{R}$-ideals as lattices without writing σ explicitly. For example, we write $\det(\mathcal{R}^m)$ instead of $\det(\sigma(\mathcal{R}^m))$.

For the definitions of the standard LWE and SIS assumptions we refer to [Reg05] and [Ajt96]. We denote their module counterparts as MLWE and MSIS.

3.2 Gaussians

Definition 1. *The n-dimensional Gaussian function is given by:*

$$\rho_{\mathbf{B},\mathbf{c}}(\mathbf{x}) \coloneqq \exp\left(-\pi \cdot (\mathbf{x} - \mathbf{c})^\intercal \cdot \mathbf{\Sigma}^{-1} \cdot (\mathbf{x} - \mathbf{c})\right)$$

for invertible $\mathbf{B} \in \mathbb{R}^{n \times n}$, $\mathbf{c} \in \mathbb{R}^n$ and $\mathbf{\Sigma} \coloneqq \mathbf{B} \cdot \mathbf{B}^\dagger$. Since $\rho_{\mathbf{B},\mathbf{c}}(\mathbf{x})$ is only distinguished up to $\mathbf{\Sigma}$, we write $\rho_{\sqrt{\mathbf{\Sigma}}}$. When $\mathbf{\Sigma} = \sigma^2 \cdot \mathbf{I}$ and $\mathbf{c} = \mathbf{0}$, we write ρ_σ.

The discrete Gaussian distribution over lattice Λ with parameters $\sqrt{\Sigma}$ and $\mathbf{c}$ is defined as follows: For any $\mathbf{x} \in \Lambda$, $\mathcal{D}_{\Lambda,\sqrt{\Sigma},\mathbf{c}}(\mathbf{x}) = \rho_{\sqrt{\Sigma},\mathbf{c}}(\mathbf{x})/\rho_{\sqrt{\Sigma},\mathbf{c}}(\Lambda)$, where $\rho_{\sqrt{\Sigma},\mathbf{c}}(\Lambda) \coloneqq \sum_{\mathbf{x}\in\Lambda} \rho_{\sqrt{\Sigma},\mathbf{c}}(\mathbf{x})$ is a finite normalisation factor. We treat Gaussians embedded in higher dimensional spaces by the following lemma.

Lemma 1 ([GMPW20, Lem. 1]). *Let $m \geq n$. For a set $A \subset \mathbb{R}^n$, a matrix $\sqrt{\Sigma} \in \mathbb{R}^{n\times n}$, and a matrix $\mathbf{T} \in \mathbb{R}^{m\times n}$ corresponding to an injective map, it holds that $\mathbf{T} \cdot \mathcal{D}_{A,\sqrt{\Sigma}} = \mathcal{D}_{\mathbf{T}\cdot A,\mathbf{T}\cdot\sqrt{\Sigma}}$.*

We write $\eta_\varepsilon(\Lambda)$ for the *smoothing parameter* of the lattice Λ which is the smallest value $\sigma > 0$ s.t. $\rho_{1/\sigma}(\Lambda^* \setminus \{\mathbf{0}\}) \leq \varepsilon$ where Λ^* is the dual lattice of Λ. Throughout, we will consider $\varepsilon = \mathsf{negl}(\lambda)$.

Lemma 2 ([BF11, Thm. 4.3]). *Suppose $m \geq 2n\log q, m > 2k$, and $\sigma > \omega(\sqrt{\log m})$. The following distributions are statistically close in n:*

$$\left\{(\mathbf{A},\mathbf{U}) \;\middle|\; \begin{array}{l}\mathbf{A} \leftarrow \mathbb{Z}_q^{n\times m}\\ \mathbf{U} \leftarrow \mathcal{D}_{\mathcal{L}_q^\perp(\mathbf{A}),\sigma}\end{array}\right\} \qquad \textit{and} \qquad \left\{(\mathbf{A},\mathbf{U}) \;\middle|\; \begin{array}{l}\mathbf{U} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}^{m\times k}\\ \mathbf{A} \leftarrow \mathbb{Z}_q^{n\times m} : \mathbf{AU} = \mathbf{0} \bmod q\end{array}\right\}.$$

In Appendix A we generalise Lemma 2 from the integers to the ring and module settings, and from the Gaussian distribution over the lattice $\mathcal{L}_q^\perp(\mathbf{A})$ to that over arbitrary cosets.

For uniformly random matrices and in general, we have respectively:

Lemma 3 (Generalised[17] from [LPR13, Thm. 4.1]). *Let n, m, q be positive integers with $n \leq m \leq \mathsf{poly}(\varphi)$. For $\mathbf{A} \leftarrow \mathcal{R}_q^{n\times m}$, with probability $2^{-\Omega(\varphi m)}$ we have*

$$\eta_{2^{-\Omega(\varphi m)}}(\Lambda_q^\perp(\mathbf{A})) \leq 8\,\varphi\sqrt{m}\cdot q^{n/m+2/(\varphi\cdot m)}.$$

The generalisation requires a stronger (e.g. $\varepsilon = 2^{-3\varphi m}$) smoothing condition for an ideal lattice in one of the substatements.

Lemma 4 ([MR04]). *For any n-dimensional lattice Λ and $\varepsilon > 0$,*

$$\eta_\varepsilon(\Lambda) \leq \sqrt{\frac{\ln(2n\cdot(1+1/\varepsilon))}{\pi}}\cdot\lambda_n(\Lambda)$$

In particular for any function in $\omega(\sqrt{\log n})$ there exist a negligible function $\varepsilon(n)$ such that $\eta_\varepsilon(\Lambda) \leq \omega(\sqrt{\log n})\cdot\lambda_n(\Lambda)$.

Lemma 5 (Adapted from [GPV08, Cor. 2.8]). *Let Λ, Λ' be n-dimensional lattices such that $\Lambda' \subseteq \Lambda$. Then for any $\varepsilon \in (0, 1/2)$, $s_{\min}(\sqrt{\Sigma}) \geq \eta_\varepsilon(\Lambda')$ and $\mathbf{c} \in \mathbb{R}^n$ the distribution of $(\mathcal{D}_{\Lambda,\sqrt{\Sigma},\mathbf{c}} \bmod \Lambda')$ has statistical distance at most $2\cdot\varepsilon$ to the uniform distribution over Λ/Λ'.*

Lemma 6 (Adapted from [MP12, Lem. 2.6]). *Let $\Lambda \subseteq \mathbb{R}^n$ be a lattice, $s_{\min}(\sqrt{\Sigma}) \geq \eta_\varepsilon(\Lambda)$ for some $\varepsilon > 0$, and $\mathbf{c} \in \mathsf{span}(\Lambda)$. We have*

$$\Pr\left[\left\|\mathcal{D}_{\Lambda+\mathbf{c},\sqrt{\Sigma}}\right\| \geq \sigma\sqrt{n}\right] \leq 2^{-n}\cdot\frac{1+\varepsilon}{1-\varepsilon}.$$

Furthermore, if $\mathbf{c} = \mathbf{0}$ then the bound holds for all $\sqrt{\Sigma} > 0$, with $\varepsilon = 0$.

Lemma 7 (Adapted from [PR06, Lem. 2.10]). *Let Λ be a rank k lattice in $\mathbb{R}^n$. Let $\varepsilon > 0$ and $\sqrt{\Sigma} \geq 2\cdot\eta_\varepsilon(\Lambda)$. Let $\mathbf{y} \in \Lambda$ then*

$$\Pr_{\mathbf{y}'\leftarrow\mathcal{D}_{\Lambda,\sqrt{\Sigma}}}[\mathbf{y}' = \mathbf{y}] \leq 2^{-k}\cdot\frac{1+\varepsilon}{1-\varepsilon}\ .$$

[17] We generalise [LPR13] from probabilities negligible in φ to negligible in $\varphi\cdot m$ so that the results of this paper apply for the typical parameters in both integer and ring setting.

3.3 Probabilities, Distances

Lemma 8 (Primitive (Lem. 2.6 of [BJRW23])). *Let $n \leq m$ be positive integers and q be an unramified prime that factors as $\langle q \rangle = \prod_{j=1}^{g} \mathfrak{q}_j$ with norm $\mathcal{N}(\mathfrak{q}_j) = q^{\varphi/g}$ in $\mathcal{R}$. Then*

$$\Pr_{\mathbf{A} \leftarrow \mathcal{R}_q^{n \times m}} \left[\mathbf{A} \cdot \mathcal{R}_q^m = \mathcal{R}_q^n \right] = \prod_{\ell=0}^{n-1} \prod_{j=1}^{g} \left(1 - q^{-\frac{\varphi \cdot (m-\ell)}{g}} \right).$$

The above probability is at least $(1 - \frac{ng}{q^{\varphi \cdot (m-n+1)/g}})$, overwhelming in λ when $q^{\varphi/g} > \mathsf{poly}(\lambda)$.

Lemma 9 (Noise drowning (adapted from [GKPV10])). *For any $\mathbf{c} \in \mathcal{R}^m$, it holds that*

$$\mathsf{SD}\left(\mathcal{D}_{\mathcal{R}^m, \sqrt{\mathbf{\Sigma}}},\ \mathcal{D}_{\mathcal{R}^m, \sqrt{\mathbf{\Sigma}}} + \mathbf{c} \right) \leq O\left(\|\mathbf{c}\| / s_m(\sqrt{\mathbf{\Sigma}}) \right).$$

In particular, if $\sqrt{\mathbf{\Sigma}} \geq \lambda^{\omega(1)} \cdot \|\mathbf{c}\|$, then $\mathcal{D}_{\mathcal{R}^m, \sqrt{\mathbf{\Sigma}}} \approx_s \mathcal{D}_{\mathcal{R}^m, \sqrt{\mathbf{\Sigma}}} + \mathbf{c}$.

Lemma 10 ([KNSW20, Thm. 4]). *Let $n > 60$, $\varepsilon > 0$, $\sigma > 20\sqrt{n}$ and $m > 1355n \log(\sigma)$. Then*

$$\Pr_{\mathbf{X} \leftarrow (\mathcal{D}_{\mathbb{Z}^n, \sigma})^m} (\eta_\varepsilon(\Lambda^\perp(\mathbf{X}) \leq 77\sqrt{(n + \log(m)) \cdot \log(2m/\varepsilon)})) \leq 1 - 2^{\Omega(n)}$$

3.4 Norm

Lemma 11. *Let $\mathbf{U} = (\mathbf{u}_j)_{j \in [k]} \in \mathcal{K}^{m \times k}$. Then $\|\mathbf{U}\| \leq \sqrt{\sum_{j \in [k]} \|\mathbf{u}_j\|^2}$.*

Proof. Fix any non-zero $\mathbf{x} \in \mathcal{K}^k$. We have $\left\| \sum_j \mathbf{u}_j x_j \right\|^2 \leq \left(\sum_j \|\mathbf{u}_j x_j\| \right)^2 \leq \left(\sum_j \|\mathbf{u}_j\| \|x_j\| \right)^2 \leq \left(\sum_j \|\mathbf{u}_j\|^2 \right) \cdot \left(\sum_j \|x_j\|^2 \right)$ due to the triangle inequality, the sub-multiplicativity of $\|\cdot\|$, and the Cauchy-Schwartz inequality respectively. Also note that $\|\mathbf{x}\|^2 = \sum_j \|x_j\|^2$. Therefore $\|\mathbf{U} \cdot \mathbf{x}\| / \|\mathbf{x}\| \leq \sqrt{\sum_j \|\mathbf{u}_j\|^2}$. □

3.5 Trapdoor Sampling and Other Algorithms

We recall the trapdoor generation from [MP12] adapted to the module setting.

Definition 2 (Lattice Trapdoors). *Fix $\mathcal{R}, q$ parametrised by λ. A lattice trapdoor scheme over $\mathcal{R}_q$ is a tuple of PPT algorithms* $(\mathsf{TrapGen}, \mathsf{SampPre})$*:*

- $(\mathbf{A}, \mathsf{td}) \leftarrow \mathsf{TrapGen}(1^\lambda, 1^n, 1^m)$*: On input dimensions $n, m \in \mathbb{N}$, generate a matrix $\mathbf{A} \in \mathcal{R}_q^{n \times m}$ together with a trapdoor* td*.*
- $\mathbf{u} \leftarrow \mathsf{SampPre}(\mathsf{td}, \mathbf{v}, \mathbf{S})$*: On input a trapdoor* td*, a target image $\mathbf{v} \in \mathcal{R}_q$-$\mathsf{span}(\mathbf{A})$, and a Gaussian parameter $\mathbf{S}^{\varphi m \times \varphi m} > 0$, output a vector $\mathbf{u} \in \mathcal{R}^m$. The syntax is extended naturally to matrix images.*

Let $n, m \in \mathbb{N}$ and $\tilde{s} > 0$ be parametrised by λ. The scheme is said to be $(n, m, \tilde{s})$-correct if, for $s_{\min}(\mathbf{S}) \geq \tilde{s}$, the following hold:

- *The distribution $\left\{ \mathbf{A} : (\mathbf{A}, \mathsf{td}) \leftarrow \mathsf{TrapGen}(1^\lambda, 1^n, 1^m) \right\}$ is statistically close to the uniform distribution over $\mathcal{R}_q^{n \times m}$.*
- *For any $\mathbf{v} \in \mathcal{R}_q^n$ and $(\mathbf{A}, \mathsf{td}) \leftarrow \mathsf{TrapGen}(1^\lambda, 1^n, 1^m)$, the distributions $\mathcal{D}_{\Lambda_q^{\mathbf{v}}(\mathbf{A}), \mathbf{S}}$ and $\mathsf{SampPre}(\mathsf{td}, \mathbf{v}, \mathbf{S})$ are statistically close except with negligible probability.*

Gadget matrix. Let $\delta \geq 2$. We set $\tilde{q} \coloneqq \lceil \log_\delta q \rceil$, $\mathbf{g}^\intercal = [1, \delta, \ldots, \delta^{\tilde{q}-1}] \in \mathcal{R}_q^{1\times\tilde{q}}$ and $\mathbf{G}_n \coloneqq \mathbf{I}_n \otimes \mathbf{g}^\intercal \in \mathcal{R}_q^{n\times(n\cdot\tilde{q})}$. When the dimensions are clear from context we simply write $\mathbf{G}$. Write $\mathbf{G}_n^{-1} : \mathcal{R}_q^{n\times t} \to \mathcal{R}_q^{(n\cdot\tilde{q})\times t}$ for the inverse function that takes a matrix of entries in $\mathcal{R}_q$, and decomposes each entry w.r.t. the base δ. We also write $\mathbf{g}^{-1}$ for $\mathbf{G}_1^{-1}$. As an immediate corollary of Theorem 4.1 in [MP12] and Lemma 4 there exists a negligible function $\varepsilon(\varphi \cdot n \cdot \tilde{q})$ such that $\eta_\varepsilon(\Lambda_q^\perp(\mathbf{G}_n)) \leq \sqrt{5} \cdot \omega(\sqrt{\log(\varphi \cdot n \cdot \tilde{q})})$.

Lemma 12 (Regularity Lemma (adapted from [LPR13, Cor. 4.2])). *Let $\mathcal{R}_q$ have degree φ and split into fields of super-polynomial size. Let $\mathbf{A} \leftarrow \mathcal{R}_q^{n\times m}$. Then with probability $1 - 2^{-\Omega(\varphi m)}$ over the choice of $\mathbf{A}$, the distribution of $\mathbf{A} \cdot \mathbf{x} \in \mathcal{R}_q^n$ where each coordinate of $\mathbf{x} \in \mathcal{R}^m$ is chosen from a discrete Gaussian of width $\sigma > 8\,\varphi\sqrt{m} \cdot q^{n/m+2/(\varphi\cdot m)}$ is within $\mathsf{negl}(\lambda)$ statistical distance to uniform.*

Lemma 13 (Gadget Trapdoors (adapted from [MP12])). *Let $\mathcal{R}_q$ have degree φ and split into fields of super-polynomial size, $m \geq n \cdot (\tilde{q}+1)$ and $\tilde{s} = \delta \cdot \tilde{r} \cdot \omega(\sqrt{n \cdot \tilde{q} \cdot \log(m \cdot \varphi)})$ where $\tilde{r} = \sigma \cdot \sqrt{2\,(m - n \cdot \tilde{q}) \cdot \varphi}$ and $\sigma > 2\,\varphi \cdot q^{n/m+2/(\varphi\cdot m)}$. There exists a $(n, m, \tilde{s})$-correct lattice trapdoor scheme over $\mathcal{R}_q$.*

Furthermore, the following advanced properties are satisfied:

- *A trapdoor td takes the form of a "$\mathbf{G}_n$-trapdoor", i.e. $\mathsf{td} = (\mathbf{R}, \mathbf{H}) \in \mathcal{R}_q^{m\times n\cdot\tilde{q}} \times \mathcal{R}_q^{n\times n}$ satisfying $\mathbf{A} \cdot \mathbf{R} \equiv \mathbf{H} \cdot \mathbf{G}_n \bmod q$, $\mathbf{H}$ being invertible over $\mathcal{R}_q$, and each column $\mathbf{r} \in \mathcal{R}_q^m$ of $\mathbf{R}$ satisfies $\|\mathbf{r}\| \leq \tilde{r}$, implying $\|\mathbf{R}\| \leq \sqrt{n \cdot \tilde{q}} \cdot \tilde{r}$. The tag is taken to be $\mathbf{H} = \mathbf{I}_n$ when omitted.*
- *The statistical closeness between $\mathcal{D}_{\Lambda_q^\mathbf{y}(\mathbf{A}),\mathbf{S}}$ and $\mathsf{SampPre}(\mathsf{td}, \mathbf{v}, \mathbf{S})$ holds for any $\mathsf{td} = (\mathbf{R}, \mathbf{H})$ with $\mathbf{A} \cdot \mathbf{R} \equiv \mathbf{H} \cdot \mathbf{G}_n \bmod q$, $\mathbf{H}$ invertible over $\mathcal{R}_q$, and $s_{\min}(\mathbf{S}) \geq \delta \cdot \|\mathbf{R}\| \cdot \omega(\sqrt{\log(m \cdot \varphi)})$.*

The following result is immediate from linear algebra for fields, but requires a bit more care for more general $\mathcal{R}_q$.

Proposition 1. *Let $\mathcal{R}_q \coloneqq \mathcal{R}/(q\mathcal{R})$. Let $\mathbf{A} \in \mathcal{R}_q^{m\times k}$ and $\mathbf{B} \in \mathcal{R}_q^{n\times k}$. There exists a PPT algorithm sampling $\mathbf{X} \in \mathcal{R}_q^{m\times n}$ uniformly at random subject to the condition that $\mathbf{A} \equiv \mathbf{X} \cdot \mathbf{B} \bmod q$.*

Proof. We first note that $\mathcal{R}_q$ is a principal ideal ring (PIR) because – by the Chinese Remainder Theorem – it is the direct product of fields which are PIRs.[18] In [Sto00] an algorithm is given for computing the Howell Form of matrices over PIRs. Using the Howell Form of $\mathbf{B}^\intercal$, we can compute an arbitrary solution $\tilde{\mathbf{X}}$ satisfying $\mathbf{A} \equiv \tilde{\mathbf{X}} \cdot \mathbf{B} \bmod q$ [Sto00, p.27]. Moreover, using the Howell Form of $\mathbf{B}^\intercal$, we can compute a basis for the kernel of $\mathbf{B}^\intercal$ [Sto00, p.70]. Picking a random element from the kernel (which is well-defined as a random linear combination of its basis), we can sample $\mathbf{X}'$ s.t. $\mathbf{X}' \cdot \mathbf{B} \equiv \mathbf{0} \bmod q$. The final solution is $\tilde{\mathbf{X}} + \mathbf{X}'$. □

4 Varying-Width κ-MSIS and κ-MLWE Assumptions

The security of the trapdoor sharing and its applications relies on versions of κ-SIS defined in [BF11] and κ-LWE from [LPSS14]. We generalise these assumptions to structured lattices and to varying hint distributions.

The κ-SIS assumption was generalised from $\mathbb{Z}_q$ to the $\mathcal{R}_q$ setting before in [ACL+22], but not proven. In this work we also define a ring-generalisation of κ-LWE. These definitions can be found in Appendix B.

The security arguments in this work require two sets of hints that have different covariance matrices. Intuitively, one set of the hints has to be more 'narrow' than the other to cover for the subsequent norm blow-up. Therefore, we further generalise both assumptions to allow varied hints.

Definition 3 (κ-MSIS Assumption With Varied Hints.). *Let $\varphi, n, m, \log q, \kappa, \{s_{\max}(\mathbf{S}_i)\}_i, \beta \in \mathsf{poly}(\lambda)$ with $n, \kappa \leq m$. An instance of κ-$\mathsf{MSIS}_{\mathcal{R}_q,n,m,\{\mathbf{S}_i\}_i,\beta}$ problem is a uniform matrix $\mathbf{A} \leftarrow \mathcal{R}_q^{n\times m}$ and κ vectors $\mathbf{e}_i \leftarrow \mathcal{D}_{\Lambda_q^\perp(\mathbf{A}),\mathbf{S}_i}$ for $i \in [\kappa]$ and arbitrary full-rank matrix $\mathbf{S}_i \in \mathbb{R}^{\varphi m\times\varphi m}$. A solution to the problem is a vector $\mathbf{v} \in \mathcal{R}_q^m$ such that: (1) $\mathbf{A} \cdot \mathbf{v} \equiv \mathbf{0} \bmod q$, (2) $\|\mathbf{v}\| \leq \beta$, (3) $\mathbf{v} \notin \mathcal{K}\text{-}\mathsf{span}(\mathbf{e}_0, \ldots, \mathbf{e}_{\kappa-1})$. The κ-$\mathsf{MSIS}_{\mathcal{R}_q,n,m,\{\mathbf{S}_i\}_i,\beta}$ assumption states that any PPT $\mathcal{A}$ finds a solution with probability $\leq \mathsf{negl}(\lambda)$.*

[18] Following, [Sto00], we do not require PIRs to not have zero divisors, in contrast to principal ideal domains (PIDs).

Definition 4 (κ-Uniform Distribution[19]). *Let $\varphi, \kappa, m, \log q \in \mathsf{poly}(\lambda)$ with $\kappa \leq m$. For any $\mathbf{U} \in \mathcal{R}_q^{m\times\kappa}$, denote $\mathbf{U}^\perp = \left\{\mathbf{x} \in \mathcal{R}_q^m : \mathbf{x}^\mathsf{T} \cdot \mathbf{U} = \mathbf{0}^\mathsf{T} \bmod q\right\}$. The κ-uniform distribution with parameters $\mathbf{U}, \mathcal{R}, q, \chi$ is*

$$\mathcal{U}(\mathbf{U}^\perp) + \mathcal{D}^m_{\mathcal{R},\chi} \coloneqq \left\{\mathbf{b} \;\middle|\; \mathbf{x} \leftarrow \mathbf{U}^\perp,\ \mathbf{e} \leftarrow \mathcal{D}^m_{\mathcal{R},\chi}; \mathbf{b} \coloneqq \mathbf{x} + \mathbf{e} \quad \right\}.$$

Definition 5 (κ-MLWE Assumption With Varied Hints.). *Let $n, m, \log q, \kappa, \chi, \{s_{\max}(\mathbf{S}_i)\}_i \in \mathsf{poly}(\lambda)$ with $n, \kappa \leq m$. Let $\{\mathbf{S}_i\}_i$ be a set of arbitrary full-rank matrices $\mathbf{S}_i \in \mathbb{R}^{\varphi m\times\varphi m}$. The κ-$\mathsf{MLWE}_{\mathcal{R}_q,n,m,\{\mathbf{S}_i\}_i,\chi}$ assumption states that for any PPT $\mathcal{A}$, the following distributions over $(\mathbf{A}, \mathbf{U}, \mathbf{b})$ are indistinguishable:*

$$\left\{\begin{array}{l}\mathbf{A} \leftarrow \mathcal{R}_q^{n\times m} \\ \mathbf{U} \leftarrow \prod_{i=0}^{\kappa-1} \mathcal{D}_{\Lambda_q^\perp(\mathbf{A}),\mathbf{S}_i} \\ \mathbf{s} \leftarrow \mathcal{R}_q^n,\ \mathbf{e} \leftarrow \mathcal{D}^m_{\mathcal{R},\chi} \\ \mathbf{b}^\mathsf{T} \coloneqq \mathbf{s}^\mathsf{T} \cdot \mathbf{A} + \mathbf{e}^\mathsf{T} \bmod q\end{array}\right\} \qquad \textit{and} \qquad \left\{\begin{array}{l}\mathbf{A} \leftarrow \mathcal{R}_q^{n\times m} \\ \mathbf{U} \leftarrow \prod_{i=0}^{\kappa-1} \mathcal{D}_{\Lambda_q^\perp(\mathbf{A}),\mathbf{S}_i} \\ \mathbf{b} \leftarrow \mathcal{U}(\mathbf{U}^\perp) + \mathcal{D}^m_{\mathcal{R},\chi}\end{array}\right\}.$$

Since $\mathbb{Z}_q$ is a cyclotomic ring of degree 1 some instances of κ-$\mathsf{MSIS}_{\mathcal{R}_q,n,m,\{\mathbf{S}_i\}_i,\beta}$ and κ-$\mathsf{LWE}_{\mathcal{R}_q,n,m,\{\mathbf{S}_i\}_i,\chi}$ problems are as hard as SIS and LWE via reductions in [BF11, LPSS14]. In Theorem 1 we adapt the reduction of [LPSS14] over $\mathbb{Z}_q$ to hints with different covariance matrices.

Previous work only covered the case where $\mathbf{S}_i = \mathbf{S} = \mathsf{diag}\,(\sigma\mathbf{I}, \sigma'\mathbf{I})$. Now the reduction allows for $\mathbf{S}_i = \mathsf{diag}\,(\sigma_i\mathbf{I}, \sigma'_i\mathbf{I})$ that is slightly more general than what we need for the security proof. The difference in the top and bottom standard deviations of the hints σ_i and σ'_i is a proof artefact that we inherit from [LPSS14].

Theorem 1 (Generalised from [LPSS14, Thm. 18 (eprint)]). *Let $n, m, \kappa, q, \beta, \beta', \chi, \chi', \{\sigma_i, \sigma'_i\}_{i\in[\kappa]}$ satisfy the following:*

- $\kappa \geq 100$, $\forall i \in [\kappa] : \sigma_i > 0, \sigma'_i > 0$,
- $m \geq \max\{\Omega(\kappa \log(\kappa \max_i(\sigma_i))), \Omega(n \log q)\}$,
- $q > \max_i(\sigma'_i)\sqrt{\log m}$ *is prime,*
- $\min_i(\sigma_i) > \max(\Omega(\sqrt{\kappa m \log m}), \max(\sigma'_i)^{\kappa/(m+\kappa)})$, *and*
- $\min_i(\sigma'_i) \geq \Omega(\kappa\sqrt{m}\max(\sigma_i)^2 \log^{3/2}(\kappa m \max(\sigma_i)))$.

Let $\mathbf{S}_i = \mathsf{diag}\,(\sigma_i\mathbf{I}, \sigma'_i\mathbf{I})$. If

$$\beta > \Omega(m^{3/2}\max(\sigma'_i) \cdot \beta') \qquad \textit{and} \qquad \chi' > \Omega(m^{3/2}\max(\sigma'_i) \cdot \chi)$$

then there exists a PPT reduction from $\mathsf{SIS}_{\mathbb{Z}_q,n,m,\beta}$ to κ-$\mathsf{SIS}_{\mathbb{Z}_q,n,m+\kappa,\{\mathbf{S}_i\}_i,\beta'}$ and a PPT reduction from $\mathsf{LWE}_{\mathbb{Z}_q,n,m,\chi}$ to κ-$\mathsf{LWE}_{\mathbb{Z}_q,n,m+\kappa,\{\mathbf{S}_i\}_i,\chi'}$.

Proof. (Sketch) We obtain the proof by following exactly the strategy of [LPSS14]. The only adaptation we require is proving Lemma 29 that corresponds to Lemma 16 in the original paper (eprint numbering). For that we also adapt Lemma 27 and Lemma 28 that correspond to Lemmas 6 and 7 of the original work accordingly. See Appendix B for these adapted lemmas.

Remark 2. We conjecture that for more general rings κ-$\mathsf{MSIS}_{\mathcal{R}_q,n,m+\kappa,\{\mathbf{S}_i\}_i,\beta'}$ is as hard as $\mathsf{MSIS}_{\mathcal{R}_q,n,m,\beta}$ and κ-$\mathsf{MLWE}_{\mathcal{R}_q,n,m+\kappa,\{\mathbf{S}_i\}_i,\chi'}$ is as hard as $\mathsf{MLWE}_{\mathcal{R}_q,n,m,\chi}$ for adapted parameters.

5 Partial Lattice Trapdoors

In Section 5.1 we formally define the syntax and security of a threshold partial lattice trapdoor, in Section 5.2 we provide our full construction and its correctness analysis. Sections 5.3 and 5.4 are devoted to a handful of lemmas and their proofs, which will be useful for proving security (one-wayness and indistinguishability) of our construction in Sections 5.5 and 5.6.

[19] In the sense of [LPSS14].

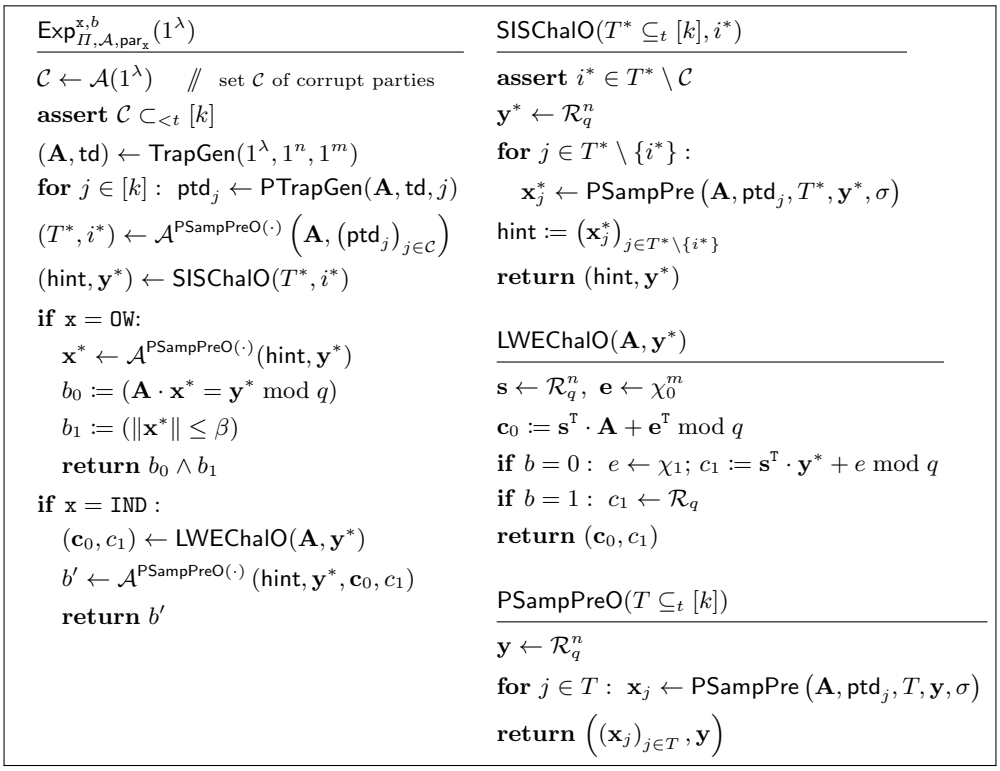

Fig. 1. Security and indistinguishability experiments for partial lattice trapdoor.

5.1 Definitions

Definition 6 (Partial Lattice Trapdoors). *A (t,k)-threshold partial lattice trapdoor scheme over $\mathcal{R}_q$ is an extension of a (full) lattice trapdoor scheme* $(\mathsf{TrapGen}, \mathsf{SampPre})$ *over $\mathcal{R}_q$ with the PPT algorithms* $(\mathsf{PTrapGen}, \mathsf{PSampPre}, \mathsf{Rec})$*:*

- $\mathsf{ptd}_j \leftarrow \mathsf{PTrapGen}(\mathbf{A}, \mathsf{td}, j)$*: The partial trapdoor generation algorithm inputs a matrix $\mathbf{A}$, its (full) trapdoor* td *and an index $j \in [k]$, and generates a partial trapdoor* ptd_j *of $\mathbf{A}$ for j.*
- $\mathbf{x}_j \leftarrow \mathsf{PSampPre}\left(\mathbf{A}, \mathsf{ptd}_j, T, \mathbf{y}, \sigma\right)$*: The partial preimage sampling algorithm, given a partial trapdoor* ptd_j *for index j, a set $T \subseteq_t [k]$, a target image $\mathbf{y}$, and a Gaussian parameter σ, samples a partial preimage $\mathbf{x}_j$ for j.*
- $\mathbf{x} \leftarrow \mathsf{Rec}\left((\mathbf{x}_j)_{j\in T}\right)$*: The reconstruction algorithm, given a tuple of partial preimages $(\mathbf{x}_j)_{j\in T}$, reconstructs a preimage $\mathbf{x}$.*

Definition 7 (Correctness). *Let $k, n, m, \tilde{m}, t \in \mathbb{N}$ with $t \le k$ and $n \le m$, and $\tilde{s}, \beta, \sigma > 0$. A (t,k)-threshold partial lattice trapdoor scheme is said to be $(n, m, \tilde{s}, \beta, \sigma)$-correct if the underlying lattice trapdoor scheme* $(\mathsf{TrapGen}, \mathsf{SampPre})$ *is $(n, m, \tilde{s})$-correct, and for any t-subset $T \subseteq_t [k]$ and vector $\mathbf{y} \in \mathcal{R}_q^n$, it holds that*

$$\Pr\left[\begin{array}{l|l} \mathbf{A}\cdot\mathbf{x} \equiv \mathbf{y} & (\mathbf{A},\mathsf{td}) \leftarrow \mathsf{TrapGen}(1^\lambda, 1^n, 1^m) \\ \wedge\ \|\mathbf{x}\| \le \beta & \quad \mathsf{ptd}_j \leftarrow \mathsf{PTrapGen}(\mathbf{A}, \mathsf{td}, j)\ \forall j \in [k] \\ & \quad\ \mathbf{x}_j \leftarrow \mathsf{PSampPre}\left(\mathbf{A}, \mathsf{ptd}_j, T, \mathbf{y}, \sigma\right)\ \forall j \in T \\ & \quad\ \mathbf{x} \leftarrow \mathsf{Rec}\left((\mathbf{x}_j)_{j\in T}\right) \end{array}\right] \ge 1 - \mathsf{negl}(\lambda).$$

Definition 8 (One-wayness and Indistinguishability). *Let $k, n, m, t \in \mathbb{N}$ with $t \le k$ and $n \le m$, and $\beta, \chi_0, \chi_1, \sigma > 0$. Denote $\mathsf{par}_{\mathtt{OW}} = (n, m, \beta, \sigma)$ and $\mathsf{par}_{\mathtt{IND}} = (n, m, \chi_0, \chi_1, \sigma)$. A (t,k)-threshold partial lattice*

$\mathsf{PTrapGen}(\mathbf{A} \in \mathcal{R}_q^{nt \times 2mt}, \mathsf{td}, j)$	$\mathsf{PSampPre}(\mathbf{A}, \mathsf{ptd}_j, T, \mathbf{y} \in \mathcal{R}_q^{nt}, \sigma)$
$(\mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}) \leftarrow \mathsf{TrapGen}(1^\lambda, 1^n, 1^m)$	$\mathbf{V}_T \coloneqq (\mathbf{v}_j)_{j \in T} \in \mathcal{R}_q^{t \times t}$
$\mathbf{U}_j \leftarrow \mathsf{SampPre}(\mathbf{A}, \mathsf{td}, \mathbf{v}_j \otimes \mathbf{C}_j, \mathbf{S})$	$\mathbf{z}_j \coloneqq \left((\hat{\mathbf{e}}_{T,j}^\intercal \cdot \mathbf{V}_T^{-1}) \otimes \mathbf{I}_n\right) \cdot \mathbf{y} \bmod q \in \mathcal{R}_q^n$
if $\mathsf{Rank}_{\mathcal{K}}(\mathbf{U}_j) \neq m$: **return** $\perp$	// Decomp. $\mathbf{y} = \sum_{j \in T} \mathbf{v}_j \otimes \mathbf{z}_j \bmod q$ is unique
$\mathsf{ptd}_j \coloneqq (\mathbf{U}_j \in \mathcal{R}^{2mt \times m}, \mathbf{C}_j \in \mathcal{R}_q^{n \times m}, \mathsf{td}_{\mathbf{C}_j})$	$\mathbf{\Sigma}_j \coloneqq \sigma^2 \cdot (\mathbf{U}_j^\dagger \cdot \mathbf{U}_j)^{-1}$
return ptd_j	$\mathbf{d}_j \leftarrow \mathsf{SampPre}(\mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}, \mathbf{z}_j, \sqrt{\mathbf{\Sigma}_j})$
	return $\mathbf{x}_j \coloneqq \mathbf{U}_j \cdot \mathbf{d}_j \in \mathcal{R}^{2mt}$

Fig. 2. Partial lattice trapdoor construction with $\mathsf{Rec}((\mathbf{x}_j)_{j \in T}) = \sum_{j \in T} \mathbf{x}_j$.

trapdoor scheme is said to be $\mathsf{par}_{\mathtt{OW}}$*-one-way, if for any PPT* $\mathcal{A}$ *it holds that*

$$\Pr\left[\mathsf{Exp}^{\mathtt{OW}}_{\Pi,\mathcal{A},\mathsf{par}_{\mathtt{OW}}}(1^\lambda) = 1\right] \leq \mathsf{negl}(\lambda),$$

and it is said to have $\mathsf{par}_{\mathtt{IND}}$*-indistinguishability, if for any PPT* $\mathcal{A}$

$$\left|\Pr\left[\mathsf{Exp}^{\mathtt{IND},0}_{\Pi,\mathcal{A},\mathsf{par}_{\mathtt{IND}}}(1^\lambda) = 1\right] - \Pr\left[\mathsf{Exp}^{\mathtt{IND},1}_{\Pi,\mathcal{A},\mathsf{par}_{\mathtt{IND}}}(1^\lambda) = 1\right]\right| \leq \mathsf{negl}(\lambda),$$

where the experiments $\mathsf{Exp}^{\mathtt{x},b}_{\Pi,\mathcal{A},\mathsf{par}_{\mathtt{x}}}$ *for* $\mathtt{x} \in \{\mathtt{OW}, \mathtt{IND}\}$ *and* $b \in \{0,1\}$ *are defined as in Fig. 1.*

Remark 3. In the above, we define $\mathsf{PSampPre}$ to take as input the reconstruction set T and Rec does not. For our proof of feasibility, this is a natural choice since it simplifies the construction and its analysis while already being useful in some contexts. Indeed, we note that such a setting is common among lattice-based threshold signature schemes (e.g. [DKM+24, KRT24, BKL+25]). This choice, however, is not inherent to the notion of partial trapdoors: we could consider a variant where $\mathsf{PSampPre}$ is oblivious to T while Rec takes T as input, and define correctness, one-wayness and indistinguishability analogously.

5.2 Construction

Let $t, n, \log q \in \mathsf{poly}(\lambda)$ and $m = n \cdot (\lceil \log q \rceil + 1)$. Let $\{v_0, v_1, \ldots, v_{k-1}\} \subseteq \mathcal{R}_q^\times$ be an arbitrary set such that $v_i - v_j \in \mathcal{R}_q^\times$ for all $i \neq j$. In our setting this holds with overwhelming probability for random v_i and v_j. More generally, this can be achieved, for example, by choosing v_j with small $\|v_j\|$ for all j [ACX21]. For $j \in [k]$, let $\mathbf{v}_j \coloneqq (1, v_j, \ldots, v_j^{t-1})$. Let $\hat{\mathbf{e}}_{T,j} \in \{0,1\}^T$ the unit-vector indexed by j. Let $(\mathsf{TrapGen}, \mathsf{SampPre})$ be a lattice trapdoor scheme, e.g. that specified in Lemma 13, which is $(nt, 2mt, \tilde{s})$-correct with some $\tilde{s} > 0$. In Fig. 2 we construct a (t,k)-threshold partial lattice trapdoor scheme, where we let $\mathbf{S} \in \mathbb{R}^{2\varphi mt \times 2\varphi mt}$ be Gaussian parameters hardwired in the $\mathsf{PTrapGen}$ algorithm.

Theorem 2 (Correctness). *Let parameters be as in Table 2, in particular* $\beta \geq t \cdot \sigma \cdot \sqrt{2mt\varphi}$*. The partial lattice trapdoor in Fig. 2 is* $(nt, 2mt, \tilde{s}, \beta, \sigma)$*-correct.*

Proof. With overwhelming probability, $\mathbf{U}_j$ sampled in $\mathsf{PTrapGen}$ have full $\mathcal{K}$-rank by Theorem 9, so that it does not abort. Assume this is the case in the rest.

For each $j \in T$, the output $\mathbf{x}_j$ of $\mathsf{PSampPre}(\mathbf{A}, \mathsf{ptd}_j, T, \mathbf{y}, \sigma)$ satisfies

$$\mathbf{A} \cdot \mathbf{x}_j = \mathbf{A} \cdot \mathbf{U}_j \cdot \mathbf{d}_j \equiv (\mathbf{v}_j \otimes \mathbf{C}_j) \cdot \mathbf{d}_j = \mathbf{v}_j \otimes (\mathbf{C}_j \cdot \mathbf{d}_j) \equiv \mathbf{v}_j \otimes \mathbf{z}_j \bmod q.$$

Observe that the vertical concatenation of $\mathbf{z}_j$'s (ordered by T) yields

$$\begin{bmatrix} \vdots \\ \mathbf{z}_j \\ \vdots \end{bmatrix} \equiv \begin{bmatrix} \vdots \\ ((\hat{\mathbf{e}}_{T,j}^\intercal \cdot \mathbf{V}_T^{-1}) \otimes \mathbf{I}_n) \cdot \mathbf{y} \\ \vdots \end{bmatrix} = (\mathbf{V}_T^{-1} \otimes \mathbf{I}_n) \cdot \mathbf{y} = (\mathbf{V}_T \otimes \mathbf{I}_n)^{-1} \cdot \mathbf{y} \bmod q.$$

Hence the reconstructed vector $\mathbf{x} = \sum_{j \in T} \mathbf{x}_j$ satisfies

$$\begin{aligned}\mathbf{A} \cdot \mathbf{x} = &\sum_{j \in T} \mathbf{A} \cdot \mathbf{x}_j \equiv \sum_{j \in T} \mathbf{v}_j \otimes \mathbf{z}_j = \sum_{j \in T} (\mathbf{V}_T \otimes \mathbf{I}) \cdot (\hat{\mathbf{e}}_{T,j} \otimes \mathbf{z}_j) \\ = &\ (\mathbf{V}_T \otimes \mathbf{I}) \cdot \begin{bmatrix} \vdots \\ \mathbf{z}_j \\ \vdots \end{bmatrix} = (\mathbf{V}_T \otimes \mathbf{I}) \cdot (\mathbf{V}_T \otimes \mathbf{I})^{-1} \cdot \mathbf{y} = \mathbf{y} \bmod q.\end{aligned}$$

For any $j \in [k]$, $\mathbf{C}_j$ generated by $\mathsf{TrapGen}$ is primitive with overwhelming probability by Lemmas 8 and 13. Conditioned on this, the distribution of $\mathbf{x}_j^*$ is identical to $\mathcal{D}_{\Lambda(\mathbf{U}_j) \cap \Lambda_q^{\mathbf{v} \otimes \mathbf{z}_j}(\mathbf{A}), \sigma}$ by Lemma 14, and $\Pr[\|\mathbf{x}_j\| \geq \sigma \cdot \sqrt{2mt\varphi}] \leq 2^{-2mt\varphi}$ by Lemma 6. Finally, $\Pr[\|\mathbf{x}\| \geq t \cdot \sigma \cdot \sqrt{2mt\varphi}] \leq t \cdot 2^{-2mt\varphi}$. □

5.3 Useful Lemmas for Security Proofs

In this subsection, we provide several lemmas which will be useful in arguing security of our partial trapdoor scheme. Using these, we establish in Section 5.4 a sequence of distributions and their statistical closeness. Building upon these, in Sections 5.5 and 5.6 we state and prove our main theorems on the one-wayness and indistinguishability of our partial lattice trapdoor.

Lemma 14 states that, for $\mathbf{U}$ satisfying $\mathbf{A} \cdot \mathbf{U} \equiv \mathbf{v} \otimes \mathbf{C}$, the distribution of $\mathbf{U} \cdot \mathbf{d}$ for an appropriately distributed $\mathbf{d}$ follows Gaussian over the intersection of the lattice $\Lambda(\mathbf{U})$ and the coset $\Lambda_q^{\mathbf{v} \otimes \mathbf{z}}(\mathbf{A})$.

Lemma 14 (Sampling from the intersection). *Let $\mathbf{A} \in \mathcal{R}_q^{nt \times 2mt}$, $\mathbf{v} \in \mathcal{R}_q^t$, and $\sigma \in \mathbb{R}^+$. Let $\mathbf{C} \in \mathcal{R}_q^{n \times m}$ be primitive. Let $\mathbf{U} \in \mathcal{R}^{mt \times m}$ be any matrix satisfying $\mathbf{A} \cdot \mathbf{U} \equiv \mathbf{v} \otimes \mathbf{C} \bmod q$. Let $\mathbf{\Sigma} \coloneqq \sigma^2 \cdot (\mathbf{U}^\dagger \cdot \mathbf{U})^{-1}$. Then for any $\mathbf{z} \in \mathcal{R}_q^n$, the following distributions are identical:*

$$\left\{ \mathbf{x} \;\middle|\; \mathbf{d} \leftarrow \mathcal{D}_{\mathcal{R}^m, \sqrt{\mathbf{\Sigma}}} : \mathbf{C} \cdot \mathbf{d} \equiv \mathbf{z} \bmod q;\ \mathbf{x} \coloneqq \mathbf{U} \cdot \mathbf{d} \right\} \qquad \text{and} \qquad \mathcal{D}_{\Lambda(\mathbf{U}) \cap \Lambda_q^{\mathbf{v} \otimes \mathbf{z}}(\mathbf{A}), \sigma}.$$

Proof. We show that $\mathbf{U} \cdot \mathcal{D}_{\mathcal{R}^m, \sqrt{\mathbf{\Sigma}}} = \mathcal{D}_{\Lambda(\mathbf{U}), \mathbf{U}\sqrt{\mathbf{\Sigma}}} = \mathcal{D}_{\Lambda(\mathbf{U}), \sigma}$. The first equality is by definition. We show the second equality by direct calculation. Fix any $\mathbf{x} \in \Lambda(\mathbf{U})$ and write $\mathbf{x} = \mathbf{U} \cdot \mathbf{d}$ for some $\mathbf{d} \in \mathcal{R}^m$. Notice that

$$\begin{aligned}\rho_{\mathbf{U} \cdot \sqrt{\mathbf{\Sigma}}}(\mathbf{x}) &= \rho_{\mathbf{U} \cdot \sqrt{\mathbf{\Sigma}}}(\mathbf{U} \cdot \mathbf{d}) = \rho_{\sqrt{\mathbf{\Sigma}}}(\mathbf{d}) \\ &= \exp(-\pi \cdot \mathbf{d}^\dagger \cdot \mathbf{\Sigma}^{-1} \cdot \mathbf{d}) = \exp(-\pi \cdot \mathbf{d}^\dagger \cdot \mathbf{U}^\dagger \cdot \mathbf{U} \cdot \mathbf{d} / \sigma^2) \\ &= \rho_\sigma(\mathbf{x}).\end{aligned}$$

Therefore,

$$\rho_{\mathbf{U}\sqrt{\mathbf{\Sigma}}}(\Lambda(\mathbf{U})) = \sum_{\mathbf{x} \in \Lambda(\mathbf{U})} \rho_{\mathbf{U}\sqrt{\mathbf{\Sigma}}}(\mathbf{x}) = \sum_{\mathbf{x} \in \Lambda(\mathbf{U})} \rho_\sigma(\mathbf{x}) = \rho_\sigma(\Lambda(\mathbf{U}))$$

and

$$\mathcal{D}_{\Lambda(\mathbf{U}), \mathbf{U}\sqrt{\mathbf{\Sigma}}}(\mathbf{x}) = \frac{\rho_{\mathbf{U}\sqrt{\mathbf{\Sigma}}}(\mathbf{x})}{\rho_{\mathbf{U}\sqrt{\mathbf{\Sigma}}}(\Lambda(\mathbf{U}))} = \frac{\rho_\sigma(\mathbf{x})}{\rho_\sigma(\Lambda(\mathbf{U}))} = \mathcal{D}_{\Lambda(\mathbf{U}), \sigma}(\mathbf{x}).$$

The claim then follows by conditioning on $\mathbf{C} \cdot \mathbf{d} \equiv \mathbf{z} \bmod q$. □

Proposition 2 is a regularity lemma analogous to the classic result of [GPV08].

Proposition 2 (Image distribution). *Fix any primitive $\mathbf{A} \in \mathcal{R}_q^{nt \times 2mt}$, any $v \in \mathcal{R}_q^\times$, and $\sqrt{\mathbf{\Sigma}} \in \mathbb{R}^{2mt \times 2mt}$. Define $\mathbf{v} = (1, v, \ldots, v^{t-1})^\intercal$. If $\sqrt{\mathbf{\Sigma}} \geq \eta_\varepsilon(\Lambda_q^\perp(\mathbf{A}))$ for some $\varepsilon \leq \mathsf{negl}(\lambda)$ then the following distributions (parametrised by $(\mathbf{A}, v)$) are statistically close in λ.*

$$\left\{ \mathbf{U} \;\middle|\; \mathbf{U} \leftarrow (\mathcal{D}_{\mathcal{R}^{mt}, \sqrt{\mathbf{\Sigma}}})^m : \mathbf{A} \cdot \mathbf{U} \in (\mathcal{R}_q\text{-}\mathsf{span}(\mathbf{v} \otimes \mathbf{I}))^m \right\},$$
$$\left\{ \mathbf{U} \;\middle|\; \mathbf{C} \leftarrow \mathcal{R}_q^{n \times m};\ \mathbf{U} \leftarrow (\mathcal{D}_{\mathcal{R}^{mt}, \sqrt{\mathbf{\Sigma}}})^m : \mathbf{A} \cdot \mathbf{U} = \mathbf{v} \otimes \mathbf{C} \bmod q \right\}.$$

Proof. Define the lattice Λ and consider:

$$\begin{aligned}\Lambda &\coloneqq \left\{\mathbf{u} \in \mathcal{R}^{mt} : \exists\, \mathbf{z} \in \mathcal{R}_q^n,\ \mathbf{A}\cdot\mathbf{u} \equiv \mathbf{v}\otimes\mathbf{z} \bmod q\right\},\\ \Lambda_q^\perp(\mathbf{A}) &\coloneqq \left\{\mathbf{u} \in \mathcal{R}^{mt} : \mathbf{A}\cdot\mathbf{u} \equiv \mathbf{0} \bmod q\right\},\\ \Lambda_q^{\mathbf{v}\otimes\mathbf{c}}(\mathbf{A}) &\coloneqq \left\{\mathbf{u} \in \mathcal{R}^{mt} : \mathbf{A}\cdot\mathbf{u} \equiv \mathbf{v}\otimes\mathbf{c} \bmod q\right\}.\end{aligned}$$

Since $\mathbf{A}$ is primitive, the image space satisfies

$$\begin{aligned}\left\{\mathbf{A}\cdot\mathbf{u} \bmod q : \mathbf{u} \in \mathcal{R}^{mt}\right\} &= \mathcal{R}_q^{nt},\\ \left\{\mathbf{A}\cdot\mathbf{u} \bmod q : \mathbf{u} \in \Lambda\right\} &= \left\{\mathbf{v}\otimes\mathbf{z} \bmod q : \mathbf{z} \in \mathcal{R}_q^n\right\}.\end{aligned}$$

Since $s_{\min}(\sqrt{\mathbf{\Sigma}}) \geq \eta_\varepsilon(\Lambda_q^\perp(\mathbf{A}))$, by Lemma 5, $\mathcal{D}_{\Lambda,\sqrt{\mathbf{\Sigma}}} \bmod \Lambda_q^\perp(\mathbf{A})$ is statistically close to $\mathcal{U}(\Lambda/\Lambda_q^\perp(\mathbf{A}))$. Since (multiplication by) $\mathbf{A}$ is an isomorphism from $\Lambda/\Lambda_q^\perp(\mathbf{A})$ to the image space $\left\{\mathbf{v}\otimes\mathbf{z} \bmod q : \mathbf{z} \in \mathcal{R}_q^n\right\}$, we have $\mathbf{A}\cdot(\mathcal{D}_{\Lambda,\sqrt{\mathbf{\Sigma}}} \bmod \Lambda_q^\perp(\mathbf{A})) \bmod q$ being statistically close to $\mathcal{U}(\left\{\mathbf{v}\otimes\mathbf{z} \bmod q : \mathbf{z} \in \mathcal{R}_q^n\right\})$. The claim follows by repeating the above argument for each column of $\mathbf{C}$. □

For $\mathbf{U}$ satisfying $\mathbf{A}\cdot\mathbf{U} \equiv \mathbf{v}\otimes\mathbf{C}$, Proposition 3 gives a bound on the smoothing parameter of the intersection of the lattices $\Lambda(\mathbf{U})$ and $\Lambda_q^\perp(\mathbf{A})$.

Proposition 3 (Smoothing of intersection). *Let $\varepsilon > 0$, $\mathbf{v} \in \mathcal{R}_q^t$ be a vector with invertible coefficients, $\mathbf{A} \in \mathcal{R}_q^{nt\times 2mt}$, $\mathbf{C} \in \mathcal{R}_q^{n\times m}$ and $\mathbf{U} \in \mathcal{R}^{mt\times m}$ with $\mathbf{A}\cdot\mathbf{U} \equiv \mathbf{v}\otimes\mathbf{C}$. Then*

$$\eta_\varepsilon(\Lambda(\mathbf{U}) \cap \Lambda_q^\perp(\mathbf{A})) \leq s_{\max}(\mathbf{U}) \cdot \eta_\varepsilon(\Lambda_q^\perp(\mathbf{C})).$$

Proof. Observe

$$\begin{aligned}\Lambda(\mathbf{U}) \cap \Lambda_q^\perp(\mathbf{A}) &= \{\mathbf{U}\cdot\mathbf{d} \mid \mathbf{d} \in \mathcal{R}^m : \mathbf{A}\cdot\mathbf{U}\cdot\mathbf{d} = \mathbf{0} \bmod q\}\\ &= \{\mathbf{U}\cdot\mathbf{d} \mid \mathbf{d} \in \mathcal{R}^m : \mathbf{C}\cdot\mathbf{d} = \mathbf{0} \bmod q\} \ =\ \mathbf{U}\cdot\Lambda_q^\perp(\mathbf{C})\end{aligned}$$

where the first equality holds since the coefficients of $\mathbf{v}$ are invertible elements. The claim follows from a direct inspection of the smoothing parameters of $\mathbf{U}\cdot\Lambda_q^\perp(\mathbf{C})$ and from the inequality $\|\mathbf{U}\cdot\mathbf{x}\| \leq s_{\max}(\mathbf{U})\cdot\|\mathbf{x}\|$. □

Lemma 15 below will be useful for arguing the partial trapdoor sampling in later subsections.

Lemma 15. *Let $\mathbf{A} \in \mathcal{R}_q^{nt\times 2mt}$ be primitive, $\mathbf{v} \in \mathcal{R}_q^t$, $\sqrt{\mathbf{\Sigma}_1} \in \mathbb{R}^{2mt\times 2mt}$ and $\sqrt{\mathbf{\Sigma}_0} \in \mathbb{R}^{\tilde{m}\times\tilde{m}}$. If $\sqrt{\mathbf{\Sigma}_1} \geq \eta_\varepsilon(\Lambda_q^\perp(\mathbf{A}))$, $\sqrt{\mathbf{\Sigma}_0} \geq \eta_\varepsilon(\Lambda_q^\perp(\mathbf{G}))$, and $\mathbf{\Sigma}$ be such that $\mathbf{\Sigma}^{-1} = \begin{pmatrix}\mathbf{\Sigma}_1^{-1} & \\ & \mathbf{\Sigma}_0^{-1}\end{pmatrix}$, then following distributions are statistically close in λ:*

$$\begin{aligned}&\left\{\mathbf{U} \;\middle|\; \mathbf{U} \leftarrow (\mathcal{D}_{\mathcal{R}^{mt},\sqrt{\mathbf{\Sigma}_1}})^m : \mathbf{A}\cdot\mathbf{U} \in (\mathcal{R}_q\text{-}\mathsf{span}(\mathbf{v}\otimes\mathbf{I}_n))^m\right\},\\ &\left\{\mathbf{U} \;\middle|\; [\mathbf{U}\|\mathbf{W}] \leftarrow (\mathcal{D}_{\mathcal{R}^{mt+\tilde{m}},\sqrt{\mathbf{\Sigma}}})^m : \mathbf{A}\cdot\mathbf{U} = \mathbf{v}\otimes\mathbf{G}\cdot\mathbf{W} \bmod q\right\}.\end{aligned}$$

Proof. Applying Proposition 2, the first distribution is statistically close to

$$\left\{\mathbf{U} \;\middle|\; \begin{array}{l}\mathbf{C} \leftarrow \mathcal{R}_q^{n\times m}\\ \mathbf{U} \leftarrow (\mathcal{D}_{\mathcal{R}^{mt},\sqrt{\mathbf{\Sigma}_1}})^m : \mathbf{A}\cdot\mathbf{U} = \mathbf{v}\otimes\mathbf{C} \bmod q\end{array}\right\}.$$

For $\mathbf{W} \leftarrow (\mathcal{D}_{\mathcal{R}^{\tilde{m}},\sqrt{\mathbf{\Sigma}_0}})^m$, since $\sqrt{\mathbf{\Sigma}_0} \geq \eta_\varepsilon(\Lambda_q^\perp(\mathbf{G}))$, $\mathbf{G}\mathbf{W} \bmod q$ is statistically close to uniform over $\mathcal{R}_q^{n\times m}$ by Lemma 5, hence the above is statistically close to

$$\left\{\mathbf{U} \;\middle|\; \begin{array}{l}\mathbf{W} \leftarrow (\mathcal{D}_{\mathcal{R}^{\tilde{m}},\sqrt{\mathbf{\Sigma}_0}})^m\\ \mathbf{U} \leftarrow (\mathcal{D}_{\mathcal{R}^{mt},\sqrt{\mathbf{\Sigma}_1}})^m : \mathbf{A}\cdot\mathbf{U} = \mathbf{v}\otimes\mathbf{G}\cdot\mathbf{W} \bmod q\end{array}\right\}.$$

Below we show that

$$D \coloneqq \left\{ (\mathbf{u}, \mathbf{w}) \;\middle|\; \begin{array}{l} \mathbf{w} \leftarrow \mathcal{D}_{\mathcal{R}^{\tilde{m}}, \sqrt{\mathbf{\Sigma}_0}} \\ \mathbf{u} \leftarrow \mathcal{D}_{\mathcal{R}^{mt}, \sqrt{\mathbf{\Sigma}_1}} : \mathbf{A} \cdot \mathbf{u} = \mathbf{v} \otimes \mathbf{G} \cdot \mathbf{w} \bmod q \end{array} \right\},$$
$$D' \coloneqq \left\{ (\mathbf{u}, \mathbf{w}) \;\middle|\; \begin{bmatrix} \mathbf{u} \\ \mathbf{w} \end{bmatrix} \leftarrow \mathcal{D}_{\mathcal{R}^{mt+\tilde{m}}, \sqrt{\mathbf{\Sigma}}} : \mathbf{A} \cdot \mathbf{u} = \mathbf{v} \otimes \mathbf{G} \cdot \mathbf{w} \bmod q \right\}$$

are statistically close, then the claim follows by repeating over all columns.

The probability mass function of D is

$$\begin{aligned} D(\mathbf{u}, \mathbf{w}) = & \frac{\exp(-\pi \mathbf{w}^{\mathsf{T}} \mathbf{\Sigma}_0^{-1} \mathbf{w})}{\rho_{\sqrt{\mathbf{\Sigma}_0}}(\mathcal{R}^m)} \cdot \frac{\exp(-\pi \mathbf{u}^{\mathsf{T}} \mathbf{\Sigma}_1^{-1} \mathbf{u})}{\rho_{\sqrt{\mathbf{\Sigma}_1}}(\Lambda_q^{\mathbf{v} \otimes \mathbf{G} \mathbf{w}}(\mathbf{A}))} \\ = & (1 + \mathsf{negl}(\lambda)) \frac{\exp(-\pi \mathbf{w}^{\mathsf{T}} \mathbf{\Sigma}_0^{-1} \mathbf{w})}{\rho_{\sqrt{\mathbf{\Sigma}_0}}(\mathcal{R}^m)} \cdot \frac{\exp(-\pi \mathbf{u}^{\mathsf{T}} \mathbf{\Sigma}_1^{-1} \mathbf{u})}{\rho_{\sqrt{\mathbf{\Sigma}_1}}(\Lambda_q^{\perp}(\mathbf{A}))} \end{aligned}$$

where the second equality is due to $\mathbf{A}$ primitive, $\sqrt{\mathbf{\Sigma}_1} \geq \eta_\varepsilon(\Lambda_q^{\perp}(\mathbf{A}))$ and Corollary 2.

Denote by $\Lambda \coloneqq \left\{ \mathbf{x} = \begin{bmatrix} \mathbf{u} \\ \mathbf{w} \end{bmatrix} \in \mathcal{R}^{mt+\tilde{m}} : \mathbf{A} \cdot \mathbf{u} = \mathbf{v} \otimes \mathbf{G} \cdot \mathbf{w} \bmod q \right\}$. The probability mass function of D' is

$$D'(\mathbf{u}, \mathbf{w}) = \frac{\exp(-\pi\, \mathbf{x}^{\mathsf{T}} \cdot \mathbf{\Sigma}^{-1} \cdot \mathbf{x})}{\rho_{\sqrt{\mathbf{\Sigma}}}(\Lambda)} = \frac{\exp(-\pi\, \mathbf{w}^{\mathsf{T}} \cdot \mathbf{\Sigma}_0^{-1} \cdot \mathbf{w}) \cdot \exp(-\pi\, \mathbf{u}^{\mathsf{T}} \cdot \mathbf{\Sigma}_1^{-1} \cdot \mathbf{u})}{\rho_{\sqrt{\mathbf{\Sigma}}}(\Lambda)},$$

where

$$\begin{aligned} \rho_{\sqrt{\mathbf{\Sigma}}}(\Lambda) = & \sum_{\mathbf{u}, \mathbf{w} : \mathbf{A} \cdot \mathbf{u} = \mathbf{v} \otimes \mathbf{G} \cdot \mathbf{w} \bmod q} \exp(-\pi\, \mathbf{w}^{\mathsf{T}} \mathbf{\Sigma}_0^{-1} \cdot \mathbf{w}) \cdot \exp(-\pi\, \mathbf{u}^{\mathsf{T}} \cdot \mathbf{\Sigma}_1^{-1} \cdot \mathbf{u}) \\ = & \sum_{\mathbf{w} \in \mathcal{R}^{\tilde{m}}} \exp(-\pi\, \mathbf{w}^{\mathsf{T}} \cdot \mathbf{\Sigma}_0^{-1} \cdot \mathbf{w}) \cdot \sum_{\mathbf{u} \in \Lambda_q^{\mathbf{v} \otimes \mathbf{G} \cdot \mathbf{w}}(\mathbf{A})} \exp(-\pi\, \mathbf{u}^{\mathsf{T}} \cdot \mathbf{\Sigma}_1^{-1} \cdot \mathbf{u}) \\ = & (1 - \mathsf{negl}(\lambda)) \cdot \sum_{\mathbf{w} \in \mathcal{R}^{\tilde{m}}} \exp(-\pi\, \mathbf{w}^{\mathsf{T}} \cdot \mathbf{\Sigma}_0^{-1} \cdot \mathbf{w}) \cdot \sum_{\mathbf{u} \in \Lambda_q^{\perp}(\mathbf{A})} \exp(-\pi \mathbf{u}^{\mathsf{T}} \cdot \mathbf{\Sigma}_1^{-1} \cdot \mathbf{u}) \\ = & (1 - \mathsf{negl}(\lambda)) \cdot \rho_{\sqrt{\mathbf{\Sigma}_0}}(\mathcal{R}^{\tilde{m}}) \cdot \rho_{\sqrt{\mathbf{\Sigma}_1}}(\Lambda_q^{\perp}(\mathbf{A})), \end{aligned}$$

where in the third equality we used again that $\mathbf{A}$ is primitive, $\sqrt{\mathbf{\Sigma}_1} \geq \eta_\varepsilon(\Lambda_q^{\perp}(\mathbf{A}))$ and Corollary 2. Therefore,

$$D'(\mathbf{u}, \mathbf{w}) = (1 + \mathsf{negl}(\lambda)) \cdot \frac{\exp(-\pi \mathbf{w}^{\mathsf{T}} \cdot \mathbf{\Sigma}_0^{-1} \cdot \mathbf{w}) \cdot \exp(-\pi \mathbf{u}^{\mathsf{T}} \cdot \mathbf{\Sigma}_1^{-1} \cdot \mathbf{u})}{\rho_{\sqrt{\mathbf{\Sigma}_0}}(\mathcal{R}^{\tilde{m}}) \rho_{\sqrt{\mathbf{\Sigma}_1}}(\Lambda_q^{\perp}(\mathbf{A}))}.$$

The statistical distance between D and D' is hence

$$\frac{1}{2} \sum_{\mathbf{u}, \mathbf{w} :\ \mathbf{A} \cdot \mathbf{u} \equiv \mathbf{v} \otimes \mathbf{G} \cdot \mathbf{w}} |D(\mathbf{u}, \mathbf{w}) - D'(\mathbf{u}, \mathbf{w})| \leq \mathsf{negl}(\lambda). \qquad \square$$

5.4 Simulation of Partial Trapdoors and Preimages

Looking ahead, in Sections 5.5 and 5.6, we will prove the one-wayness and indistinguishability of our partial trapdoor scheme from the κ-SIS and κ-LWE assumptions respectively for $\kappa = 2m(t-1)$. Both proofs consist of two main steps:

1. Given a κ-SIS/LWE instance, simulate the partial trapdoors of corrupt parties and responses to partial preimage queries to PSampPreO.
2. Use the adversary's response to solve the κ-SIS/LWE instance.

Table 1. Summary of distributions in Section 5.4.

	D_0	D_1	D_2	D_3	D_4
$\mathbf{A}$	TrapGen	\$	$\to$	Reverse-sampling from $2m(t\text{-}1)$-MSIS $\tilde{\mathbf{A}}_0, \mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}}$	$\to$
$\mathbf{C}_j, j \in \mathcal{C}$	TrapGen	$\to$	$\to$	$\to$	$\to$
$\mathbf{U}_j, j \in \mathcal{C}$	SampPre	$\mathcal{D}_{\Lambda_q^{\mathbf{v}_j \otimes \mathbf{C}_j}(\mathbf{A}), \mathbf{S}_{\mathtt{Cor}}}$	$\to$	From $2m(t\text{-}1)$-MSIS $\mathbf{U}_{\mathtt{Cor}}$	$\to$
$\mathbf{C}_j, j \notin \mathcal{C}$	TrapGen	\$	Random image from $\mathbf{U}_j$	$\to$	$\to$
$\mathbf{U}_j, j \notin \mathcal{C}$	SampPre	$\mathcal{D}_{\Lambda_q^{\mathbf{v}_j \otimes \mathbf{C}_j}(\mathbf{A}), \mathbf{S}_{\mathtt{Cor}}}$	Basis of random subspace	$\to$	BASIS-style sampling from $2m(t\text{-}1)$-MSIS $\mathbf{U}_{\mathtt{Hon}}$
$\mathbf{x}_j, j \in T$	SampPre	$\mathcal{D}_{\Lambda(\mathbf{U}_j) \cap \Lambda_q^{\mathbf{v}_j \otimes \mathbf{z}_j}(\mathbf{A}), \sigma}$	Random image from $\mathbf{U}_j$	$\to$	$\to$

Left-most column are main components in the construction (Fig. 2), where the view of an adversary is highlighted in gray. D_0 refers to the distribution induced by the construction, D_4 is efficiently sampleable given an $2m(t-1)$-MSIS(/-MLWE) instance. "$\to$" means same as previous distribution; "\$" means sample uniformly at random.

This subsection is dedicated to Step 1, i.e. simulating partial trapdoors and partial preimages, which is common in both proofs.

In more detail, for any fixed set $\mathcal{C} \subset_{t-1} [k]$ of corrupt parties, $L \leq \mathsf{poly}(\lambda)$ number of partial preimage queries, and reconstruction sets $T_\ell \subseteq [k]$ for $\ell \in [L]$, we define in Figs. 3 and 4 a sequence of distributions $D_i, i \in \{0,1,2,3,4\}$ on the tuple

$$\left(\mathbf{A}, \left(\mathsf{ptd}_j = (\mathbf{U}_j, \mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j})\right)_{j \in \mathcal{C}}, (\mathbf{x}_{\ell,j}, \mathbf{y}_\ell)_{\ell \in [L], j \in T_\ell}\right)$$

which are statistically close to each other. The first distribution D_0 captures the distribution of (a core part of) the view of the adversary in the one-wayness and indistinguishability experiments[20] (c.f. Fig. 1), where in particular partial preimages $\mathbf{x}_{\ell,j}$ are generated using genuinely distributed partial trapdoors ptd_j. After a number of modifications, in the last distribution D_4, the partial preimages and partial trapdoors are instead efficiently simulated from a varying-width κ-SIS/LWE instance. In particular, the matrices $\mathbf{A}$ and $\mathbf{C}_j$ are no longer generated with trapdoors. For an overview of the modifications between subsequent distributions, see Table 1.

In this subsection, in particular for distributions D_3 and D_4, we will assume for notational convenience that party 0 is not corrupt, i.e. $\mathcal{C} \subseteq_{t-1} [k] \setminus \{0\}$. Under such an assumption we introduce the shorthand $\mathbf{V}_{\{0\} \cup \mathcal{C}} \in \mathcal{R}_q^{t \times t}$ which denotes the Vandermonde matrix with (ordered) columns $(\mathbf{v}_j)_{j \in \{0\} \cup \mathcal{C}}$. For distribution D_4, we further introduce the shorthands $\mathbf{H}_j, \hat{\mathbf{H}}_j \in \mathcal{R}_q^{nt \times nt}$ for $j \notin \mathcal{C}$,

$$\mathbf{H}_j \coloneqq \begin{bmatrix} 1 & & & \\ & v_j^{-1} & & \\ & & \ddots & \\ & & & v_j^{-(t-1)} \end{bmatrix} \otimes \mathbf{I}_n, \qquad \hat{\mathbf{H}}_j \coloneqq \begin{bmatrix} 1 & 1 & \dots & 1 \\ 1 & \frac{v_1}{v_j} & \dots & \frac{v_{t-1}}{v_j} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & \left(\frac{v_1}{v_j}\right)^{t-1} & \dots & \left(\frac{v_{t-1}}{v_j}\right)^{t-1} \end{bmatrix} \otimes \mathbf{I}_n, \tag{7}$$

where we abuse notation by implicitly relabelling the set $\{0\} \cup \mathcal{C}$ as $[t]$, to avoid additional layers of subscripts. In this case, since $v_j \notin \{v_1, v_2, \dots, v_{t-1}\}$ for any $j \notin \mathcal{C}$, the matrix $\hat{\mathbf{H}}_j$ is guaranteed to be invertible over $\mathcal{R}_q$.

Lemma 16 ($D_0 \approx_s D_4$). *Let parameters be as in Table 2. The distributions D_0 in Fig. 3 and D_4 in Fig. 4 are statistically close in λ.*

Proof. The lemma follows from Lemmas 17 to 20 below.

We given an overview of the main ideas of the proof of Lemma 16.

[20] More precisely, in $\mathsf{Exp}^{\mathtt{OW}}_{\Pi,\mathcal{A},\mathsf{par}_{\mathtt{OW}}}$, $\mathcal{A}$ is given these together with the challenge image $\mathbf{y}^*$ and all but one partial preimages $\left(\mathbf{x}_j^*\right)_{j \in T^* \setminus \{i^*\}}$, whereas in $\mathsf{Exp}^{\mathtt{IND}}_{\Pi,\mathcal{A},\mathsf{par}_{\mathtt{IND}}}$, $\mathcal{A}$ is further given the challenge LWE sample $(\mathbf{c}_0, \mathbf{c}_1)$. We address simulating the challenges in the main security proofs, i.e. Theorems 3 and 4.

$$
\begin{array}{l}
D_0 \to \left(\mathbf{A}, \left(\mathbf{U}_j, \mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}\right)_{j\in\mathcal{C}}, (\mathbf{x}_{\ell,j}, \mathbf{y}_\ell)_{\ell\in[L], j\in T_\ell}\right) \\
\hline
\qquad\qquad (\mathbf{A}, \mathsf{td}) \leftarrow \mathsf{TrapGen}(1^\lambda, 1^{nt}, 1^{2mt}) \\
\forall j \in [k], \qquad (\mathbf{U}_j, \mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}) \leftarrow \mathsf{PTrapGen}(\mathbf{A}, \mathsf{td}, j) \\
\forall \ell \in [L], \qquad \mathbf{y}_\ell \leftarrow \mathcal{R}_q^{nt} \\
\forall \ell \in [L], j \in T_\ell, \quad \mathbf{x}_{\ell,j} \leftarrow \mathsf{PSampPre}(\mathbf{A}, (\mathbf{U}_j, \mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}), T, \mathbf{y}_\ell)
\end{array}
$$

$$
\begin{array}{l}
D_1 \to \left(\mathbf{A}, \left(\mathbf{U}_j, \mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}\right)_{j\in\mathcal{C}}, (\mathbf{x}_{\ell,j}, \mathbf{y}_\ell)_{\ell\in[L], j\in T_\ell}\right) \\
\hline
\qquad\qquad \mathbf{A} \leftarrow \mathcal{R}_q^{nt\times 2mt} \\
\forall j \in [k], \qquad \begin{cases} (\mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}) \leftarrow \mathsf{TrapGen}(1^\lambda, 1^n, 1^m) & j \in \mathcal{C} \\ \mathbf{C}_j \leftarrow \mathcal{R}_q^{n\times m} & j \notin \mathcal{C} \end{cases} \\
\forall j \in [k], \qquad \mathbf{U}_j \leftarrow (\mathcal{D}_{\mathcal{R}^{2mt}, \mathbf{S}_{\mathsf{Cor}}})^m \text{ s.t. } \mathbf{A}\cdot\mathbf{U}_j \equiv \mathbf{v}_j \otimes \mathbf{C}_j \\
\forall \ell \in [L], j \in T_\ell, \quad \mathbf{z}_{\ell,j} \leftarrow \mathcal{R}_q^n; \quad \Lambda_j \coloneqq \Lambda(\mathbf{U}_j) \cap \Lambda_q^{\mathbf{v}_j \otimes \mathbf{z}_{\ell,j}}(\mathbf{A}); \quad \mathbf{x}_{\ell,j} \leftarrow \mathcal{D}_{\Lambda_j,\sigma} \\
\forall \ell \in [L], \qquad \mathbf{y}_\ell \coloneqq \sum_{j\in T_\ell} \mathbf{v}_j \otimes \mathbf{z}_{\ell,j} \bmod q
\end{array}
$$

$$
\begin{array}{l}
D_2 \to \left(\mathbf{A}, \left(\mathbf{U}_j, \mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}\right)_{j\in\mathcal{C}}, (\mathbf{x}_{\ell,j}, \mathbf{y}_\ell)_{\ell\in[L], j\in T_\ell}\right) \\
\hline
\qquad\qquad \mathbf{A} \leftarrow \mathcal{R}_q^{nt\times 2mt} \\
\forall j \in \mathcal{C}, \qquad (\mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}) \leftarrow \mathsf{TrapGen}(1^\lambda, 1^n, 1^m) \\
\forall j \in [k], \qquad \begin{cases} \mathbf{U}_j \leftarrow (\mathcal{D}_{\mathcal{R}^{2mt}, \mathbf{S}_{\mathsf{Cor}}})^m \text{ s.t. } \mathbf{A}\cdot\mathbf{U}_j \equiv \mathbf{v}_j \otimes \mathbf{C}_j & j \in \mathcal{C} \\ \mathbf{U}_j \leftarrow (\mathcal{D}_{\mathcal{R}^{2mt}, \mathbf{S}_{\mathsf{Cor}}})^m \text{ s.t. } \mathbf{A}\cdot\mathbf{U}_j \in (\mathcal{R}_q\text{-}\mathsf{span}(\mathbf{v}_j \otimes \mathbf{I}_n))^m & j \notin \mathcal{C} \end{cases} \\
\forall \ell \in [L], j \in T_\ell, \quad \mathbf{x}_{\ell,j} \leftarrow \mathcal{D}_{\Lambda(\mathbf{U}_j),\sigma} \\
\forall \ell \in [L], \qquad \mathbf{y}_\ell \coloneqq \mathbf{A}\cdot\sum_{j\in T_\ell} \mathbf{x}_{\ell,j} \bmod q
\end{array}
$$

Fig. 3. Distributions D_0, D_1, D_2, with respect to some fixed set $\mathcal{C} \subset_{t-1} [k]$ of corrupt parties, $L \le \mathsf{poly}(\lambda)$ number of partial preimage queries, and reconstruction sets $T_\ell \subseteq [k]$ for $\ell \in [L]$, satisfying $D_0 \overset{(Lem.\ 17)}{\approx} D_1 \overset{(Lem.\ 18)}{\approx} D_2 \overset{(Lem.\ 19)}{\approx} D_3 \overset{(Lem.\ 20)}{\approx} D_4$.

$\mathsf{ExtendA}(\tilde{\mathbf{A}}_0, \mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}}) \to (\mathbf{A}, \left(\mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}\right)_{j\in\mathcal{C}})$

$\forall j \in \mathcal{C},\ (\mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}) \leftarrow \mathsf{TrapGen}(1^\lambda, 1^n, 1^m)$

$\tilde{\mathbf{A}}_{\neq 0} \leftarrow \mathcal{R}_q^{n(t-1)\times 2mt}$ s.t. $\tilde{\mathbf{A}}_{\neq 0} \cdot (\mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}}) \equiv \left(\mathsf{diag}\left(\{\mathbf{C}_j\}_{j\in\mathcal{C}}\right), \mathbf{G}_{n(t-1)}\right)$

$\mathbf{A} \coloneqq \left(\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n\right) \cdot \mathsf{diag}\left(\tilde{\mathbf{A}}_0, \tilde{\mathbf{A}}_{\neq 0}\right) \bmod q$

$D_3 \to (\mathbf{A}, \left(\mathbf{U}_j, \mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}\right)_{j\in\mathcal{C}}, (\mathbf{x}_{\ell,j}, \mathbf{y}_\ell)_{\ell\in[L], j\in T_\ell})$

$\tilde{\mathbf{A}}_0 \leftarrow \mathcal{R}_q^{n\times 2mt}$

$\forall \mathtt{Typ} \in \{\mathtt{Cor}, \mathtt{Hon}\},\ \mathbf{U}_{\neq 0,\mathtt{Typ}} \leftarrow (\mathcal{D}_{\Lambda_q^\perp(\tilde{\mathbf{A}}_0), \mathbf{S}_{\mathtt{Typ}}})^{m(t-1)}$

$(\mathbf{A}, \left(\mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}\right)_{j\in\mathcal{C}}) \leftarrow \mathsf{ExtendA}(\tilde{\mathbf{A}}_0, \mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}})$

$\forall j \in [k],$
$\begin{cases} \mathbf{U}_j \coloneqq \mathbf{U}_{\mathtt{Cor}} \cdot (\hat{\mathbf{e}}_{\mathcal{C},j} \otimes \mathbf{I}_n) \quad /\!/\ (\mathbf{U}_j)_{j\in\mathcal{C}} = \mathbf{U}_{\mathtt{Cor}} & j \in \mathcal{C} \\ \mathbf{U}_j \leftarrow (\mathcal{D}_{\mathcal{R}^{2mt}, \mathbf{S}_{\mathtt{Cor}}})^m \text{ s.t. } \mathbf{A}\cdot\mathbf{U}_j \in (\mathcal{R}_q\text{-}\mathsf{span}(\mathbf{v}_j \otimes \mathbf{I}_n))^m & j \notin \mathcal{C} \end{cases}$

$\forall \ell \in [L], j \in T_\ell,\quad \mathbf{x}_{\ell,j} \leftarrow \mathcal{D}_{\Lambda(\mathbf{U}_j),\sigma}$

$\forall \ell \in [L],\quad \mathbf{y}_\ell \coloneqq \mathbf{A} \sum_{j\in T_\ell} \mathbf{x}_{\ell,j} \bmod q$

$D_4 \to (\mathbf{A}, \left(\mathbf{U}_j, \mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}\right)_{j\in\mathcal{C}}, (\mathbf{x}_{\ell,j}, \mathbf{y}_\ell)_{\ell\in[L], j\in T_\ell})$

$\tilde{\mathbf{A}}_0 \leftarrow \mathcal{R}_q^{n\times 2mt}$

$\forall \mathtt{Typ} \in \{\mathtt{Cor}, \mathtt{Hon}\},\ \mathbf{U}_{\neq 0,\mathtt{Typ}} \leftarrow (\mathcal{D}_{\Lambda_q^\perp(\tilde{\mathbf{A}}_0), \mathbf{S}_{\mathtt{Typ}}})^{m(t-1)}$

$(\mathbf{A}, \left(\mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}\right)_{j\in\mathcal{C}}) \leftarrow \mathsf{ExtendA}(\tilde{\mathbf{A}}_0, \mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}})$

$\mathbf{T} \coloneqq \begin{bmatrix} \mathbf{0}_{2mt\times\tilde{m}} & \mathbf{U}_{\mathtt{Hon}} \\ -\mathbf{I}_{\tilde{m}} & \mathbf{0}_{\tilde{m}\times m(t-1)} \end{bmatrix}$

$\forall j \in [k],\quad \mathbf{B}_j \coloneqq \left[\mathbf{H}_j\mathbf{A} \middle| \mathbf{1}_t \otimes -\mathbf{G}_n\right] \quad /\!/\ \mathbf{B}_j\,\mathbf{T} \equiv \hat{\mathbf{H}}_j\,\mathbf{G}_{nt}$

$\forall j \in [k],$
$\begin{cases} \mathbf{U}_j \coloneqq \mathbf{U}_{\mathtt{Cor}} \cdot (\hat{\mathbf{e}}_{\mathcal{C},j} \otimes \mathbf{I}_n) \quad /\!/\ (\mathbf{U}_j)_{j\in\mathcal{C}} = \mathbf{U}_{\mathtt{Cor}} & j \in \mathcal{C} \\ \begin{bmatrix} \mathbf{U}_j \in \mathcal{R}_q^{2mt\times m} \\ \mathbf{W}_j \in \mathcal{R}_q^{\tilde{m}\times m} \end{bmatrix} \leftarrow (\mathsf{SampPre}(\mathbf{B}_j, \mathbf{T}, \hat{\mathbf{H}}_j, \mathbf{0}, \mathbf{S}_{\mathtt{BASIS}}))^m & j \notin \mathcal{C} \end{cases}$

$\forall \ell \in [L], j \in T_\ell,\quad \mathbf{x}_{\ell,j} \leftarrow \mathcal{D}_{\Lambda(\mathbf{U}_j),\sigma}$

$\forall \ell \in [L],\quad \mathbf{y}_\ell \coloneqq \mathbf{A} \sum_{j\in T_\ell} \mathbf{x}_{\ell,j} \bmod q$

Fig. 4. Definitions of a subroutine $\mathsf{ExtendA}$ and distributions D_3, D_4, with respect to some fixed set $\mathcal{C} \subset_{t-1} [k]$ of corrupt parties, $L \le \mathsf{poly}(\lambda)$ number of partial preimage queries, and reconstruction sets $T_\ell \subseteq [k]$ for $\ell \in [L]$, satisfying $D_0 \overset{(Lem.\ 17)}{\approx} D_1 \overset{(Lem.\ 18)}{\approx} D_2 \overset{(Lem.\ 19)}{\approx} D_3 \overset{(Lem.\ 20)}{\approx} D_4$.

$D_0 \approx_s D_1$. The modifications from D_0 to D_1 are mostly about switching between running $\mathsf{TrapGen}$ and sampling uniformly random matrices, and between running $\mathsf{SampPre}$ and sampling perfect Gaussian. The statistical closeness of these changes can be argued by standard arguments. The most interesting difference is how the $\mathbf{x}_{\ell,j}$ are sampled in D_0 and D_1, which despite the differing representations are in fact identical. The latter fact is shown in Lemma 14, which in turn is proven by a direct calculation of the Gaussian weights.

Lemma 17 ($D_0 \approx_s D_1$). *Let parameters be as in Table 2. The distributions D_0 and D_1 in Fig. 3 are statistically close in λ.*

Proof. For any $\mathbf{y}_\ell \in \mathcal{R}_q^{nt}$, note that its decomposition into $\mathbf{y}_\ell = \sum_{j\in T} \mathbf{v}_j \otimes \mathbf{z}_{\ell,j} \bmod q$ with respect to $(\mathbf{v}_j)_{j\in T}$ is unique. Consequently, a uniformly random $\mathbf{y}_\ell \in \mathcal{R}_q^{nt}$ induces uniformly random $(\mathbf{z}_{\ell,j})_{j\in T}$ and vice versa. It therefore suffices to consider the statement for any fixed $(\mathbf{y}_\ell, (\mathbf{z}_{\ell,j})_{j\in T})_{\ell\in[L]}$ satisfying $\mathbf{y}_\ell = \sum_{j\in T} \mathbf{v}_j \otimes \mathbf{z}_{\ell,j} \bmod q$ and the result follows from averaging. In the below, we denote the variants of D_0 and D_1 with fixed $\mathbf{Y} \coloneqq (\mathbf{y}_\ell)_{\ell\in[L]}$ (and hence $(\mathbf{z}_{\ell,j})_{j\in T}$) by $D_{0,\mathbf{Y}}$ and $D_{1,\mathbf{Y}}$ respectively. We will show that $D_{0,\mathbf{y}}$ is statistically close to $D_{1,\mathbf{y}}$.

In $D_{0,\mathbf{Y}}$, the tuple $(\mathbf{A}, \left(\mathbf{U}_j, \mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}\right)_{j\in[k]}, (\mathbf{x}_{\ell,j})_{\ell\in[L],j\in T})$ is sampled as follows: Run $\mathsf{TrapGen}$ to generate $\mathbf{A}$ and its trapdoor. Run $\mathsf{TrapGen}$ to generate all the $\mathbf{C}_j$ with their trapdoors. For $j \in [k]$, sample $\mathbf{U}_j$ from $\mathsf{SampPre}(\mathbf{A}, \mathsf{td}, \mathbf{v}_j \otimes \mathbf{C}_j, \mathbf{S}_{\mathsf{Cor}})$. For $\ell \in [L]$ and $j \in T_\ell$, sample $\mathbf{d}_{\ell,j}$ from $\mathsf{SampPre}(\mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}, \mathbf{z}_{\ell,j}, \sqrt{\mathbf{\Sigma}_j})$, and set $\mathbf{x}_{\ell,j} = \mathbf{U}_j\, \mathbf{d}_{\ell,j}$.

In $D_{1,\mathbf{Y}}$, the tuple is sampled as follows: Sample $\mathbf{A}$ uniformly at random. Run $\mathsf{TrapGen}$ to generate $\mathbf{C}_j$ with their trapdoors for $j \in \mathcal{C}$, and sample the remaining $\mathbf{C}_j$ uniformly at random. For $j \in [k]$, sample $\mathbf{U}_j$ from $(\mathcal{D}_{\mathcal{R}^{2mt},\mathbf{S}_{\mathsf{Cor}}})^m$ subject to the constraint $\mathbf{A}\,\mathbf{U}_j \equiv \mathbf{v}_j \otimes \mathbf{C}_j$. For $\ell \in [L]$ and $j \in T$, sample $\mathbf{x}_{\ell,j}$ from $\mathcal{D}_{\Lambda_j,\sigma}$ where $\Lambda_j \coloneqq \Lambda(\mathbf{U}_j) \cap \Lambda_q^{\mathbf{v}_j\otimes\mathbf{z}_{\ell,j}}(\mathbf{A})$.

To see that the above are statistically close, we first invoke Lemma 13 to replace calls to $\mathsf{SampPre}$ in $D_{0,\mathbf{Y}}$ by perfect Gaussian sampling. We can then invoke Lemma 14, which states that sampling $\mathbf{x}_{\ell,j}$ from $\Lambda(\mathbf{U}_j) \cap \Lambda_q^{\mathbf{v}_j\otimes\mathbf{z}_{\ell,j}}(\mathbf{A})$ (as in $D_{1,\mathbf{Y}}$), where $\mathbf{A}\,\mathbf{U}_j \equiv \mathbf{v}_j \otimes \mathbf{C}_j$, is identical to first sampling Gaussians from $\Lambda_q^{\mathbf{z}_{\ell,j}}(\mathbf{C}_j)$ and then multiplying the result by $\mathbf{U}_j$ (as in $D_{1,\mathbf{Y}}$). Finally, we invoke Lemma 13 to replace the sampling of $\mathbf{A}$ from $\mathsf{TrapGen}(1^\lambda, 1^{nt}, 1^{2mt})$ and of all $(\mathbf{C}_j)_{j\notin\mathcal{C}}$ from $\mathsf{TrapGen}(1^\lambda, 1^n, 1^m)$ in $D_{0,\mathbf{Y}}$to sampling them uniformly at random. □

$D_1 \approx_s D_2$. Recall that in distribution D_1 the values $\mathbf{U}_j, \mathbf{C}_j, \mathbf{x}_{\ell,j}$ are sampled subject to various constraints. In distribution D_2, the constraints for $(\mathbf{U}_j, \mathbf{C}_j)_{j\notin\mathcal{C}}$ and $(\mathbf{x}_{\ell,j})_{\ell\in[L],j\in T_\ell}$ are removed. To prove that $D_1 \approx_s D_2$, we rely on a generalised regularity lemma (Proposition 2) which states that first sampling a uniformly random $\mathbf{C}$ and then a Gaussian $\mathbf{U}$ subject to $\mathbf{A}\cdot\mathbf{U} \equiv \mathbf{v}\otimes\mathbf{C}$ is statistically close to first sampling a Gaussian $\mathbf{U}$ subject to $\mathbf{A}\cdot\mathbf{U} \in (\mathcal{R}_q\text{-}\mathsf{span}(\mathbf{v}\otimes\mathbf{I}))^m$ and then computing the corresponding $\mathbf{C}$. This regularity lemma generalises that of [GPV08] as the latter can be recovered by setting $\mathbf{v} = 1$. After that, the only difference between D_1 and D_2 is in how the partial preimages $\mathbf{x}_{\ell,j}$ (and hence $\mathbf{y}_\ell$) are sampled. In more detail, in D_1 they are sampled from $\Lambda(\mathbf{U}_j) \cap \Lambda_q^{\mathbf{v}_j\otimes\mathbf{z}_{\ell,j}}(\mathbf{A})$ for some random $\mathbf{z}_{\ell,j}$, while in D_2 they are simply sampled from $\Lambda(\mathbf{U}_j)$. Note that the relation between this pair of distributions is almost identical to that between the two ways of sampling $(\mathbf{U}_j, \mathbf{C}_j)$ handled above, and thus the same technique as in the proof of Proposition 2 can be applied. The only complication is that we have to upper bound the smoothing parameter of the intersection $\Lambda(\mathbf{U}_j) \cap \Lambda_q^\perp(\mathbf{A})$, which can be done using Proposition 3.

Lemma 18 ($D_1 \approx_s D_2$). *Let parameters be as in Table 2. The distributions D_1 and D_2 in Fig. 3 are statistically close in λ.*

Proof. With overwhelming probability, $\mathbf{A} \leftarrow \mathcal{R}_q^{nt\times 2mt}$ is primitive since $\mathcal{R}_q$ splits into fields of super-polynomial size. Our analysis below is conditioning on $\mathbf{A}$ being primitive. Since $\mathbf{S}_{\mathsf{Cor}} \geq \eta_\varepsilon(\Lambda_q^\perp(\mathbf{A}))$, by Proposition 2, the distribution D_2 is statistically close in λ to a similar distribution D_2' where, for $j \notin \mathcal{C}$, we sample $\mathbf{C}_j$ uniformly at random and then sample $\mathbf{U}_j$ subject to the constraint $\mathbf{A}\,\mathbf{U}_j \equiv \mathbf{v}_j \otimes \mathbf{C}_j$, i.e. as in D_1. Now, the only difference between D_1 and D_2' is in how $\mathbf{x}_{\ell,j}$ (and hence $\mathbf{y}_\ell$) are sampled, and we can consider $(\mathbf{A}, \left(\mathbf{U}_j, \mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}\right)_{j\in\mathcal{C}})$ as being fixed. Note that it suffices to consider only primitive $\mathbf{C}_j$, since they are sampled uniformly at random and are thus primitive with overwhelming probability, given that $\mathcal{R}_q$ splits into fields of super-polynomial size. Since in both D_1 and D_2' the distributions of $\mathbf{x}_{\ell,j}$ for different (ℓ, j) are independent, it suffices to analyse the distribution of $\mathbf{x}_{\ell,j}$ for each fixed (ℓ, j).

It remains to show that, for any fixed primitive $\mathbf{A} \in \mathcal{R}_q^{nt \times 2mt}$, primitive $\mathbf{C} \in \mathcal{R}_q^{n \times m}$, $\mathbf{v} = (1, v, \ldots, v^{t-1})$ with $v \in \mathcal{R}_q^\times$, and $\mathbf{U} \in \mathcal{R}^{2mt \times m}$ with $\mathbf{A} \cdot \mathbf{U} \equiv \mathbf{v} \otimes \mathbf{C}$, the following distributions are statistically close in λ:

$$\left\{\mathbf{x} \;\middle|\; \mathbf{z} \leftarrow \mathcal{R}_q^n;\ \mathbf{x} \leftarrow \mathcal{D}_{\Lambda(\mathbf{U}) \cap \Lambda_q^{\mathbf{v} \otimes \mathbf{z}}(\mathbf{A}), \sigma}\right\} \qquad \text{and} \qquad \left\{\mathbf{x} \;\middle|\; \mathbf{x} \leftarrow \mathcal{D}_{\Lambda(\mathbf{U}), \sigma}\right\}.$$

For this, we will use the same technique as in the proof of Proposition 2.

Fix any $\mathbf{p} \in \mathcal{R}^{2mt}$ satisfying $\mathbf{A} \cdot \mathbf{p} = \mathbf{v} \otimes \mathbf{z} \bmod q$. Define the lattice $\Psi \subset \Lambda(\mathbf{U})$ and lattice coset $\Psi + \mathbf{p}$:

$$\begin{aligned} \Psi &\coloneqq \{\mathbf{x} \in \Lambda(\mathbf{U}) : \mathbf{A} \cdot \mathbf{x} \equiv \mathbf{0} \bmod q\}, \\ \Psi + \mathbf{p} &\coloneqq \{\mathbf{x} \in \Lambda(\mathbf{U}) : \mathbf{A} \cdot \mathbf{x} \equiv \mathbf{v} \otimes \mathbf{z} \bmod q\}. \end{aligned}$$

Since $\mathbf{C}$ is primitive, the set of all images is

$$\begin{aligned} \{\mathbf{A} \cdot \mathbf{x} \bmod q : \mathbf{x} \in \Lambda(\mathbf{U})\} &= \{\mathbf{A} \cdot \mathbf{U} \cdot \mathbf{d} \bmod q : \mathbf{d} \in \mathcal{R}^m\} \\ &= \{\mathbf{v} \otimes \mathbf{C} \cdot \mathbf{d} \bmod q : \mathbf{d} \in \mathcal{R}^m\} \\ &= \left\{\mathbf{v} \otimes \mathbf{z} \bmod q : \mathbf{z} \in \mathcal{R}_q^n\right\}. \end{aligned}$$

By Proposition 3, $\eta_\varepsilon(\Lambda(\Psi)) \le \|\mathbf{U}\| \cdot \eta_\varepsilon(\Lambda_q^\perp(\mathbf{C}))$, and $\|\mathbf{U}\| \le s_{\max}(\mathbf{S}_{\mathtt{Cor}}) \cdot \sqrt{\varphi m^2 t}$ with overwhelming probability by standard tail bound. Therefore we have $\sigma \ge \eta_\varepsilon(\Lambda(\Psi))$ with overwhelming probability by the choice of parameter, and by Lemma 5, $\mathcal{D}_{\Lambda(\mathbf{U}), \sigma} \bmod \Psi$ is statistically close to $\mathcal{U}(\Lambda(\mathbf{U})/\Psi)$. Since $\Lambda(\mathbf{U})/\Psi$ is isomorphic to the image space $\left\{\mathbf{v} \otimes \mathbf{z} \bmod q : \mathbf{z} \in \mathcal{R}_q^n\right\}$, we have that $\mathbf{A} \cdot (\mathcal{D}_{\Lambda(\mathbf{U}), \sigma} \bmod \Psi) \bmod q$ is statistically close to $\mathcal{U}(\left\{\mathbf{v} \otimes \mathbf{z} \bmod q : \mathbf{z} \in \mathcal{R}_q^n\right\})$. □

$D_2 \approx_s D_3$. By hopping from D_0 to D_2, what we have accomplished so far is to let partial preimages be sampled as $\mathbf{x}_{\ell,j} \leftarrow \mathcal{D}_{\Lambda(\mathbf{U}_j), \sigma}$. However, the matrices $\mathbf{U}_j$ are sampled inefficiently from Gaussian distributions subject to some hard constraints. Through defining D_3 and D_4, our goal is to simulate the matrices $\mathbf{U}_j$ (along with $\mathbf{A}$) given a varying-width κ-SIS/LWE instance $(\tilde{\mathbf{A}}_0, \mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}})$ satisfying $\tilde{\mathbf{A}}_0 \cdot (\mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}}) \equiv \mathbf{0}$, where $\mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}}$ are Gaussian matrices with varying Gaussian parameters $\mathbf{S}_{\mathtt{Cor}}$ and $\mathbf{S}_{\mathtt{Hon}}$ respectively. It is crucial that the Gaussian parameter $\mathbf{S}_{\mathtt{Cor}}$ associated with the corrupt parties is "larger" than that $\mathbf{S}_{\mathtt{Hon}}$ associated to the honest parties. This is because, while we can directly split $\mathbf{U}_{\mathtt{Cor}}$ into $(\mathbf{U}_j)_{j \in \mathcal{C}}$ for the corrupt parties, simulating partial trapdoors and partial preimages for honest parties requires rerandomising $\mathbf{U}_{\mathtt{Hon}}$ which makes the Gaussian width of the resulting partial preimages to be larger than $\mathbf{S}_{\mathtt{Hon}}$. Therefore, there must be a large enough gap between $\mathbf{S}_{\mathtt{Cor}}$ and $\mathbf{S}_{\mathtt{Hon}}$ so that partial preimages for both corrupt and honest parties have identical Gaussian widths.

Although being visually quite different, the closeness of D_2 and D_3 can be proven by some simple arithmetic and two invocations of "reverse sampling" (Theorem 9, generalisation of Lemma 2), which states that for fixed $\mathbf{V}$ first sampling uniformly random $\mathbf{A}$ and then Gaussian $\mathbf{U}$ subject to $\mathbf{A} \cdot \mathbf{U} \equiv \mathbf{V}$ is statistically close to first sampling a Gaussian $\mathbf{U}$ and then a uniformly random $\mathbf{A}$ subject to the same constraint.

Lemma 19 ($D_2 \approx_s D_3$). *Let parameters be as in Table 2. The distributions D_2 in Fig. 3 and D_3 in Fig. 4 are statistically close in λ.*

Proof. We argue by gradually modifying D_3 into D_2 through statistically close hybrids. Note that the only difference between D_2 and D_3 lies in how $(\mathbf{A}, (\mathbf{U}_j)_{j \in \mathcal{C}})$ are sampled. It therefore suffices to argue that the distributions of $(\mathbf{A}, (\mathbf{U}_j)_{j \in \mathcal{C}})$ in D_2 and D_3 are statistically close. We will denote $\mathbf{U}_{\mathtt{Cor}} \coloneqq (\mathbf{U}_j)_{j \in \mathcal{C}} \in \mathcal{R}_q^{2mt \times m(t-1)}$ the horizontal concatenation of $\mathbf{U}_j$ for all $j \in \mathcal{C}$, equivalently $\mathbf{U}_j = \mathbf{U}_{\mathtt{Cor}}(\hat{\mathbf{e}}_{\mathcal{C},j} \otimes \mathbf{I}_n)$. Similarly, denote $\mathbf{U}_{\mathtt{Hon}} \coloneqq (\overline{\mathbf{U}}_j)_{j \in \mathcal{C}} \in \mathcal{R}_q^{2mt \times m(t-1)}$.

By Theorem 9, $(\tilde{\mathbf{A}}_0, \mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}})$ sampled from

$$\left\{\begin{aligned} &\tilde{\mathbf{A}}_0 \leftarrow \mathcal{R}_q^{n \times 2mt}, \\ &\mathbf{U}_{\mathtt{Cor}} \leftarrow (\mathcal{D}_{\Lambda_q^\perp(\tilde{\mathbf{A}}_0), \mathbf{S}_{\mathtt{Cor}}})^{m(t-1)} \\ &\mathbf{U}_{\mathtt{Hon}} \leftarrow (\mathcal{D}_{\Lambda_q^\perp(\tilde{\mathbf{A}}_0), \mathbf{S}_{\mathtt{Hon}}})^{m(t-1)} \end{aligned}\right\} \qquad \text{and} \qquad \left\{\begin{aligned} &\mathbf{U}_{\mathtt{Cor}} \leftarrow (\mathcal{D}_{\mathcal{R}^{2mt}, \mathbf{S}_{\mathtt{Cor}}})^{m(t-1)}, \\ &\mathbf{U}_{\mathtt{Hon}} \leftarrow (\mathcal{D}_{\mathcal{R}^{2mt}, \mathbf{S}_{\mathtt{Hon}}})^{m(t-1)}, \\ &\tilde{\mathbf{A}}_0 \leftarrow \mathcal{R}_q^{n \times 2mt} : \tilde{\mathbf{A}}_0 \cdot (\mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}}) = \mathbf{0} \bmod q \end{aligned}\right\}$$

are statistically close. Hence, the joint distribution of $(\tilde{\mathbf{A}}_0, \tilde{\mathbf{A}}_{\neq 0}, \mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}})$ from D_3, conditioned on $\mathbf{C}_j$'s, is statistically close to that sampled from

$$\left\{ \begin{array}{l} \mathbf{U}_{\mathtt{Cor}} \leftarrow (\mathcal{D}_{\mathcal{R}^{2mt}, \mathbf{S}_{\mathtt{Cor}}})^{m(t-1)}, \mathbf{U}_{\mathtt{Hon}} \leftarrow (\mathcal{D}_{\mathcal{R}^{2mt}, \mathbf{S}_{\mathtt{Hon}}})^{m(t-1)}, \\ \\ \begin{bmatrix} \tilde{\mathbf{A}}_0 \\ \tilde{\mathbf{A}}_{\neq 0} \end{bmatrix} \leftarrow \mathcal{R}_q^{nt \times 2mt} \text{ s.t. } \begin{bmatrix} \tilde{\mathbf{A}}_0 \\ \tilde{\mathbf{A}}_{\neq 0} \end{bmatrix} \cdot (\mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}}) = \left[\begin{array}{ccc} & \mathbf{0}_{n \times 2m(t-1)} & \\ \hline \ddots & & \ddots \\ & \mathsf{diag}\left((\mathbf{C}_j)_{j \in \mathcal{C}}\right) & \mathbf{G}_{n(t-1)} \\ & \ddots & \ddots \end{array} \right] =: \hat{\mathbf{C}} \bmod q \end{array} \right\}.$$

Since multiplication by $\mathbf{V}_{\{0\} \cup \mathcal{C}} \otimes \mathbf{I}_n$ is bijective over $\mathcal{R}_q$, $\begin{bmatrix} \tilde{\mathbf{A}}_0 \\ \tilde{\mathbf{A}}_{\neq 0} \end{bmatrix}$ from above is identically distributed as $\mathbf{A}$ sampled from

$$\begin{aligned} \mathbf{A} \leftarrow \mathcal{R}_q^{nt \times 2mt} : \mathbf{A} \cdot (\mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}}) = & \left(\mathbf{V}_{\{0\} \cup \mathcal{C}} \otimes \mathbf{I}_n\right) \cdot \hat{\mathbf{C}} \bmod q \\ = & \left((\ldots, \mathbf{v}_j \otimes \mathbf{C}_j, \ldots)_{j \in \mathcal{C}}, (\ldots, \mathbf{v}_j \otimes \mathbf{G}_n, \ldots)_{j \in \mathcal{C}}\right) \bmod q. \end{aligned}$$

Notice that the above sampling constraint is equivalent to $\mathbf{A}\mathbf{U}_j = \mathbf{v}_j \otimes \mathbf{C}_j \bmod q$ for all $j \in \mathcal{C}$, similarly $\mathbf{A}\overline{\mathbf{U}}_j = \mathbf{v}_j \otimes \mathbf{G}_n \bmod q$ for all $j \in \mathcal{C}$. Applying Theorem 9 again, $(\mathbf{A}, (\mathbf{U}_j, \overline{\mathbf{U}}_j)_{j \in \mathcal{C}})$ from above is statistically close to that from

$$\left\{ \begin{array}{l} \mathbf{A} \leftarrow \mathcal{R}_q^{nt \times 2mt} \\ \mathbf{U}_j \leftarrow (\mathcal{D}_{\mathcal{R}^{2mt}, \mathbf{S}_{\mathtt{Cor}}})^m : \mathbf{A} \cdot \mathbf{U}_j \equiv \mathbf{v}_j \otimes \mathbf{C}_j \bmod q \ \ \forall j \in \mathcal{C} \\ \overline{\mathbf{U}}_j \leftarrow (\mathcal{D}_{\mathcal{R}^{2mt}, \mathbf{S}_{\mathtt{Hon}}})^m : \mathbf{A} \cdot \overline{\mathbf{U}}_j \equiv \mathbf{v}_j \otimes \mathbf{G}_n \bmod q \ \ \forall j \in \mathcal{C} \end{array} \right\}.$$

At this point, $(\mathbf{A}, (\mathbf{U}_j)_{j \in \mathcal{C}})$ is sampled exactly as in D_2. □

$D_3 \approx_s D_4$. Finally, the only difference between D_3 and D_4 is that, in the latter, $\mathbf{U}_j$ for honest parties $j \notin \mathcal{C}$ are simulated from $\mathbf{U}_{\mathtt{Hon}}$. Specifically, the step $\mathbf{U}_j \leftarrow (\mathcal{D}_{\mathcal{R}^{2mt}, \mathbf{S}_{\mathtt{Cor}}})^m$ s.t. $\mathbf{A} \cdot \mathbf{U}_j \in (\mathcal{R}_q\text{-}\mathsf{span}(\mathbf{v}_j \otimes \mathbf{I}_n))^m$ is replaced by the following procedure: Let $\mathbf{T} \coloneqq \begin{bmatrix} \mathbf{0}_{2mt \times \tilde{m}} & \mathbf{U}_{\mathtt{Hon}} \\ -\mathbf{I}_{\tilde{m}} & \mathbf{0}_{\tilde{m} \times m(t-1)} \end{bmatrix}$ be a trapdoor of the BASIS-like matrices [WW23] $\mathbf{B}_j \coloneqq \left[\mathbf{H}_j \cdot \mathbf{A} \middle| \mathbf{1}_t \otimes -\mathbf{G}_n\right]$ for all $j \notin \mathcal{C}$. Note that $\mathbf{B}_j \cdot \mathbf{T} \equiv \hat{\mathbf{H}}_j \cdot \mathbf{G}_{nt}$. Using such a trapdoor, we can sample $\mathbf{U}_j$ by $\begin{bmatrix} \mathbf{U}_j \in \mathcal{R}_q^{2mt \times m} \\ \mathbf{W}_j \in \mathcal{R}_q^{\tilde{m} \times m} \end{bmatrix} \leftarrow (\mathsf{SampPre}(\mathbf{B}_j, \mathbf{T}, \hat{\mathbf{H}}_j, \mathbf{0}, \mathbf{S}_{\mathtt{Cor}}))^m$, where $\mathbf{W}_j$ can be discarded afterwards.

The proof of statistical closeness between D_3 and D_4 is the most involved among all hops. It consists of two main steps. The first is to show that, assuming that the trapdoor $\mathbf{T}$ is a valid $\mathbf{G}$-trapdoor of $\mathbf{B}_j$ with tag matrix $\hat{\mathbf{H}}_j$ for all $j \notin \mathcal{C}$, then $\mathbf{U}_j$ sampled in D_3 and D_4 are statistically close. For this, we use Lemma 15 which in turn follows from the regularity lemma (Proposition 2) and direct inspection of Gaussian weights. The second step is to show that $\mathbf{T}$ is indeed a trapdoor and with the desirable quality. This follows by bounding the norm of $\mathbf{T}$ and proving that $\mathbf{B}_j \cdot \mathbf{T} \equiv \hat{\mathbf{H}}_j \cdot \mathbf{G}_{nt}$ by direct inspection.

Lemma 20 ($D_3 \approx_s D_4$). *Let parameters be as in Table 2. Let* $(\mathsf{TrapGen}, \mathsf{SampPre})$ *be that in Lemma 13. Then D_3 and D_4 in Fig. 4 are statistically close in λ.*

Proof. It suffices to consider the distribution of $\mathbf{U}_j$ for $j \notin \mathcal{C}$, since in both D_3 and D_4 the conditional distributions of all other variables conditioning on $(\mathbf{U}_j)_{j \notin \mathcal{C}}$ are identical.

Assume for a moment that, $\mathbf{T}$ constructed in D_4 is a valid $\mathbf{G}$-trapdoor of $\mathbf{B}_j$ with $\hat{\mathbf{H}}_j$ the appropriate tag matrix (given by Eq. (7)), for all $j \notin \mathcal{C}$, so that sampling $\begin{bmatrix} \mathbf{U}_j \\ \mathbf{W}_j \end{bmatrix}$ from $\mathsf{SampPre}(\mathbf{B}_j, \mathbf{T}, \hat{\mathbf{H}}_j, \mathbf{0}, \mathbf{S}_{\mathtt{Cor}})$ is possible. For any $j \notin \mathcal{C}$, we inspect the relation between $\mathbf{U}_j, \mathbf{W}_j$.

Write $\mathbf{A} = \begin{bmatrix} \mathbf{A}_0 \\ \vdots \\ \mathbf{A}_{t-1} \end{bmatrix}$, i.e. each $\mathbf{A}_i \in \mathcal{R}_q^{n \times 2mt}, i \in [t]$, is a row chunk of $\mathbf{A}$. Recall the definition of $\mathbf{H}_j$ in Eq. (7), and observe that the constructed $\mathbf{B}_j$ is of the form

$$\mathbf{B}_j = \left[\mathbf{H}_j \cdot \mathbf{A} \middle| \mathbf{1}_t \otimes -\mathbf{G}_n\right]$$

$$= \left[\begin{array}{c|c} \begin{bmatrix} v_j^{-0}\mathbf{I}_n & & \\ & \ddots & \\ & & v_j^{-(t-1)}\mathbf{I}_n \end{bmatrix} \cdot \mathbf{A} & \mathbf{1}_t \otimes -\mathbf{G}_n \end{array}\right] = \left[\begin{array}{c|c} 1 \cdot \mathbf{A}_0 & -\mathbf{G} \\ v_j^{-1} \cdot \mathbf{A}_1 & -\mathbf{G} \\ \vdots & \vdots \\ v_j^{-(t-1)} \cdot \mathbf{A}_{t-1} & -\mathbf{G} \end{array}\right] \in \mathcal{R}_q^{nt \times (mt+\tilde{m})}.$$

Therefore the sampled $\mathbf{U}_j, \mathbf{W}_j$ satisfy

$$\left[\begin{array}{c|c} 1 \cdot \mathbf{A}_0 & -\mathbf{G} \\ v_j^{-1} \cdot \mathbf{A}_1 & -\mathbf{G} \\ \vdots & \vdots \\ v_j^{-(t-1)} \cdot \mathbf{A}_{t-1} & -\mathbf{G} \end{array}\right] \cdot \begin{bmatrix} \mathbf{U}_j \\ \mathbf{W}_j \end{bmatrix} \equiv \mathbf{0} \bmod q.$$

The above is equivalent to that, for each $i \in [t]$,

$$\mathbf{A}_i \cdot \mathbf{U}_j \equiv v_j^i \cdot \mathbf{G} \cdot \mathbf{W}_j \bmod q.$$

Concatenating vertically for all $i \in [t]$, we obtain

$$\mathbf{A} \cdot \mathbf{U}_j \equiv \mathbf{v}_j \otimes \mathbf{G} \cdot \mathbf{W}_j \bmod q.$$

We will see that by the choice of parameters, we can invoke Lemma 13 so that $\mathbf{U}_j$ obtained from $\mathsf{SampPre}$ is statistically close to

$$\left\{ \mathbf{U}_j \;\middle|\; \begin{bmatrix} \mathbf{U}_j \\ \mathbf{W}_j \end{bmatrix} \leftarrow (\mathcal{D}_{\mathcal{R}^{2mt+\tilde{m}}, \mathbf{S}_{\mathtt{Cor}}})^m : \mathbf{A} \cdot \mathbf{U}_j \equiv \mathbf{v}_j \otimes \mathbf{G} \cdot \mathbf{W}_j \bmod q \right\}.$$

Further, with overwhelming probability, $\mathbf{A}$ is primitive by Lemma 8 and $\mathbf{S}_{\mathtt{Cor}} \geq \eta_\varepsilon(\Lambda_q^\perp(\mathbf{A}))$ by Lemma 4. Conditioned on these and applying Lemma 15 (where we set $\sqrt{\mathbf{\Sigma}_0} = \sigma_G \mathbf{I}$ and $\sqrt{\mathbf{\Sigma}_1} = \mathbf{S}_{\mathtt{Cor}}$), we have $\mathbf{U}_j$ from above is statistically close to that in D_3.

It remains to show that, with overwhelming probability, $\mathbf{T}$ is indeed a $\mathbf{G}$-trapdoor of $\mathbf{B}_j$ with tag $\hat{\mathbf{H}}_j$ (cf. Lemma 13) under the specified parameters. By Lemmas 6 and 11, $\|\mathbf{U}_{\mathtt{Hon}}\| \leq s_{\max}(\mathbf{S}_{\mathtt{Hon}})\sqrt{\varphi 2mt \cdot m(t-1)}$ with overwhelming probability. Also, by Lemma 13, for any $j \in \mathcal{C}$, each column of $\mathsf{td}_{\mathbf{C}_j}$ has norm upper-bounded by $\eta_C\sqrt{2(m-\tilde{m})\cdot\varphi}$ with overwhelming probability, where η_C is defined in Table 2. Conditioned on these, it holds $\|\mathbf{T}\| \leq 2 \cdot \eta_C \cdot s_{\max}(\mathbf{S}_{\mathtt{Hon}}) \cdot \varphi \cdot m\sqrt{(m-\tilde{m})t(t-1)}$, therefore $\mathbf{T}$ satisfies the desired norm bound of $\delta \cdot \|\mathbf{T}\| \cdot \omega(\sqrt{\log((2mt+\tilde{m})\cdot\varphi)}) \leq s_{\min}(\mathbf{S}_{\mathtt{Cor}})$ required by Lemma 13. Finally, to see that $\mathbf{B}_j \cdot \mathbf{T} = \hat{\mathbf{H}}_j \cdot \mathbf{G}_{nt} \bmod q$ holds, the following abbreviations will be handy:

$$\mathbf{H}_j \cdot \left(\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n\right) = \left[\begin{array}{c|ccc} 1 & 1 & \dots & 1 \\ \hline \frac{v_0}{v_j} & \frac{v_1}{v_j} & \dots & \frac{v_{t-1}}{v_j} \\ \vdots & \vdots & \ddots & \vdots \\ \left(\frac{v_0}{v_j}\right)^{t-1} & \left(\frac{v_1}{v_j}\right)^{t-1} & \dots & \left(\frac{v_{t-1}}{v_j}\right)^{t-1} \end{array}\right] \otimes \mathbf{I}_n$$

$$= \left[\begin{array}{c|ccc} \mathbf{I}_n & \mathbf{I}_n & \dots & \mathbf{I}_n \\ \hline \frac{v_0}{v_j}\mathbf{I}_n & \frac{v_1}{v_j}\mathbf{I}_n & \dots & \frac{v_{t-1}}{v_j}\mathbf{I}_n \\ \vdots & \vdots & \ddots & \vdots \\ \left(\frac{v_0}{v_j}\right)^{t-1}\mathbf{I}_n & \left(\frac{v_1}{v_j}\right)^{t-1}\mathbf{I}_n & \dots & \left(\frac{v_{t-1}}{v_j}\right)^{t-1}\mathbf{I}_n \end{array}\right] =: \left[\begin{array}{c|c} \mathbf{H}_{00} & \mathbf{H}_{01} \\ \hline \mathbf{H}_{10} & \mathbf{H}_{11} \end{array}\right],$$

where $\mathbf{H}_{00} \in \mathcal{R}_q^{n\times n}$, $\mathbf{H}_{01} \in \mathcal{R}_q^{n\times n(t-1)}$, $\mathbf{H}_{10} \in \mathcal{R}_q^{n(t-1)\times n}$, $\mathbf{H}_{11} \in \mathcal{R}_q^{n(t-1)\times n(t-1)}$.

Now recall that by construction, $\mathbf{A} = \left(\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n\right) \cdot \begin{bmatrix} \tilde{\mathbf{A}}_0 \\ \tilde{\mathbf{A}}_{\neq 0} \end{bmatrix} \bmod q$. Using the above shorthand, we can write

$$\mathbf{H}_j \cdot \mathbf{A} = \mathbf{H}_j \cdot \left(\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n\right) \cdot \begin{bmatrix} \tilde{\mathbf{A}}_0 \\ \tilde{\mathbf{A}}_{\neq 0} \end{bmatrix} = \begin{bmatrix} \mathbf{H}_{00} \cdot \tilde{\mathbf{A}}_0 + \mathbf{H}_{01} \cdot \tilde{\mathbf{A}}_{\neq 0} \\ \mathbf{H}_{10} \cdot \tilde{\mathbf{A}}_0 + \mathbf{H}_{11} \cdot \tilde{\mathbf{A}}_{\neq 0} \end{bmatrix}.$$

Consequently,

$$\begin{aligned}
\mathbf{B}_j \cdot \mathbf{T} &= \left[\begin{array}{c|c} \mathbf{H}_{00} \cdot \tilde{\mathbf{A}}_0 + \mathbf{H}_{01} \cdot \tilde{\mathbf{A}}_{\neq 0} & -\mathbf{G}_n \\ \mathbf{H}_{10} \cdot \tilde{\mathbf{A}}_0 + \mathbf{H}_{11} \cdot \tilde{\mathbf{A}}_{\neq 0} & \mathbf{1}_{t-1} \otimes -\mathbf{G}_n \end{array}\right] \cdot \begin{bmatrix} \mathbf{0}_{2mt \times \tilde{m}} & \mathbf{U}_{\mathtt{Hon}} \\ -\mathbf{I}_{\tilde{m}} & \mathbf{0}_{\tilde{m} \times m(t-1)} \end{bmatrix} \\
&= \left[\begin{array}{c|c} \mathbf{G}_n & (\mathbf{H}_{00} \cdot \tilde{\mathbf{A}}_0 + \mathbf{H}_{01} \cdot \tilde{\mathbf{A}}_{\neq 0}) \cdot \mathbf{U}_{\mathtt{Hon}} \\ \mathbf{1}_{t-1} \otimes \mathbf{G}_n & (\mathbf{H}_{10} \cdot \tilde{\mathbf{A}}_0 + \mathbf{H}_{11} \cdot \tilde{\mathbf{A}}_{\neq 0}) \cdot \mathbf{U}_{\mathtt{Hon}} \end{array}\right] \\
&= \left[\begin{array}{c|c} \mathbf{G}_n & \mathbf{H}_{01} \cdot (\mathbf{I}_{t-1} \otimes \mathbf{G}_n) \\ \mathbf{1}_{t-1} \otimes \mathbf{G}_n & \mathbf{H}_{11} \cdot (\mathbf{I}_{t-1} \otimes \mathbf{G}_n) \end{array}\right] \\
&= \begin{bmatrix} \mathbf{I}_n & \mathbf{H}_{01} \\ \mathbf{1}_{t-1} \otimes \mathbf{I}_n & \mathbf{H}_{11} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{G}_n & \\ & \mathbf{I}_{t-1} \otimes \mathbf{G}_n \end{bmatrix} \\
&= \hat{\mathbf{H}}_j \cdot \mathbf{G}_{nt} \bmod q,
\end{aligned}$$

where the third equality follows from $\tilde{\mathbf{A}}_0 \cdot \mathbf{U}_{\mathtt{Hon}} = \mathbf{0} \bmod q$ and $\tilde{\mathbf{A}}_{\neq 0} \cdot \mathbf{U}_{\mathtt{Hon}} = \mathbf{I}_{t-1} \otimes \mathbf{G}_n \bmod q$ by construction. The proof is completed by noting that $\hat{\mathbf{H}}_j$ is invertible over $\mathcal{R}_q$, since for $j \notin \mathcal{C}$ we have $v_j \neq v_1, \ldots, v_{t-1}$. □

5.5 Security: One-Wayness

Towards showing our partial lattice trapdoor is one-way, we need one more lemma concerning the $\mathcal{K}$-span of matrices $\mathbf{U}$ satisfying relation of the form $\mathbf{AU} \equiv \mathbf{v} \otimes \mathbf{C}$.

Proposition 4. *Let $\mathbf{A} \in \mathcal{R}_q^{nt \times 2mt}$ be primitive. For $i \in [t]$, let $v_i \in \mathcal{R}_q^\times$, $v_i \neq v_j$ for $i \neq j$, $\mathbf{v}_i \coloneqq (1, v_i, \ldots, v_i^{t-1})^\intercal$, and let $\mathbf{U}_i \in \mathcal{R}_q^{2mt \times m_i}, \mathbf{C}_i \in \mathcal{R}_q^{n \times m_i}$ be arbitrary satisfying $\mathbf{A} \cdot \mathbf{U}_i \equiv \mathbf{v}_i \otimes \mathbf{C}_i$. Fix any $i^* \in [t]$ and write $\tilde{\mathbf{U}} \coloneqq (\ldots, \mathbf{U}_j, \ldots)_{j \neq i^*}$. Then $\mathcal{K}\text{-}\mathsf{span}\left(\mathbf{U}_{i^*}, \tilde{\mathbf{U}}\right) \neq \mathcal{K}\text{-}\mathsf{span}\left(\tilde{\mathbf{U}}\right)$.*

The idea of the proof is basically to transform the linear independence claim about the preimages $\left(\mathbf{U}_{i^*}, \tilde{\mathbf{U}}\right)$ into one about the images $\left(\mathbf{v}_{i^*}, \tilde{\mathbf{V}}\right)$, where $\tilde{\mathbf{V}} \coloneqq (\mathbf{v}_j)_{j \neq i^*}$.

Proof. Denote $\tilde{\mathbf{V}} \boxtimes \tilde{\mathbf{C}} \coloneqq (\ldots, \mathbf{v}_j \otimes \mathbf{C}_j, \ldots)_{j \neq i^*}$. Write $\tilde{m} \coloneqq \sum_{j \neq i^*} m_j$. Assume towards contradiction that $\mathcal{K}\text{-}\mathsf{span}\left(\mathbf{U}_{i^*}\right) \subset \mathcal{K}\text{-}\mathsf{span}\left(\tilde{\mathbf{U}}\right)$, meaning that for any $\mathbf{x} \in \mathcal{K}\text{-}\mathsf{span}\left(\mathbf{U}_{i^*}\right)$, we can write $\mathbf{x} = \tilde{\mathbf{U}} \cdot \mathbf{z}$ for some $\mathbf{z} \in \mathcal{K}^{\tilde{m}}$. Pick an arbitrary $\mathbf{y} \in (\mathcal{R}_q^\times)^n$. Observe that $\mathbf{v}_{i^*} \otimes \mathbf{y} \notin \mathcal{R}_q\text{-}\mathsf{span}(\tilde{\mathbf{V}} \boxtimes \tilde{\mathbf{C}})$ since $\mathbf{V}_{[t]} \coloneqq (\mathbf{v}_i)_{i \in [t]}$ is invertible modulo q. Pick an arbitrary $\mathbf{x} \in \Lambda(\mathbf{U}_{i^*}) \subset \mathcal{K}\text{-}\mathsf{span}\left(\mathbf{U}_{i^*}\right)$ s.t. $\mathbf{A} \cdot \mathbf{x} \equiv \mathbf{v}_{i^*} \otimes \mathbf{y}$. Such an $\mathbf{x}$ must exist because $\mathbf{A}$ is primitive. By assumption, there exists $\mathbf{z} \in \mathcal{K}^{\tilde{m}}$ s.t. $\mathbf{x} = \tilde{\mathbf{U}} \cdot \mathbf{z}$. We write $\mathbf{z} = \mathbf{b}/a$ with $\mathbf{b} \in \mathcal{R}^{\tilde{m}}$ and $0 \neq a \in \mathcal{R}$ and where q does not divide a and $\mathbf{b}$ simultaneously (otherwise we can clear the common factor). It holds that $\mathbf{x} \cdot a = \tilde{\mathbf{U}} \cdot \mathbf{b}$. We have $(\mathbf{v}_{i^*} \otimes \mathbf{y}) \cdot a \equiv \mathbf{A} \cdot \mathbf{x} \cdot a = \mathbf{A} \cdot \tilde{\mathbf{U}} \cdot \mathbf{b} \equiv (\tilde{\mathbf{V}} \boxtimes \tilde{\mathbf{C}}) \cdot \mathbf{b}$. Since $\mathbf{v}_{i^*} \otimes \mathbf{y} \notin \mathcal{R}_q\text{-}\mathsf{span}(\tilde{\mathbf{V}} \boxtimes \tilde{\mathbf{C}})$, we have $a \equiv 0$ and $\mathbf{b} \equiv \mathbf{0}$. This contradicts that q does not divide a and $\mathbf{b}$ simultaneously. □

We are ready to prove that our construction in Fig. 2 satisfies one-wayness. Given the results in Section 5.4, the proof is fairly standard. We highlight some key points below.

We aim to show that, given an adversary $\mathcal{A}$ against one-wayness, we can construct an algorithm $\mathcal{B}^{\mathcal{A}}$ against the κ-MSIS problem for $\kappa = 2m \cdot (t-1)$. First, we observe that the view of the adversary $\mathcal{A}$ in the one-wayness experiment essentially follows the distribution D_0 in Fig. 3, except that a challenge partial preimage $\mathbf{x}_{i^*}$ is withheld from the adversary. Therefore, we could define a hybrid experiment where the adversary's view is simulated as in distribution D_4 in Fig. 4, and by Lemma 16, we can show that the two experiments are statistically close.

It remains to show that if $\mathcal{A}$ has non-negligible advantage in winning the hybrid experiment, then $\mathcal{B}^{\mathcal{A}}$ can extract a solution for a given κ-MSIS instance $\left(\tilde{\mathbf{A}}_0, \mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}}\right)$. Conveniently, distribution D_4 and hence the hybrid experiment is by design efficiently simulatable given $\left(\tilde{\mathbf{A}}_0, \mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}}\right)$. Our algorithm $\mathcal{B}^{\mathcal{A}}$ therefore simulates the hybrid experiment for $\mathcal{A}$, and upon receiving an output $\mathbf{x}^*$ from $\mathcal{A}$ outputs $\mathbf{x}^\dagger \coloneqq \mathbf{x}^* - \sum_{j \in T} \mathbf{x}_j$. By direct inspection, one could show that $\mathbf{x}^\dagger$ is indeed a vector satisfying $\tilde{\mathbf{A}}_0 \cdot \mathbf{x}^\dagger = \mathbf{0} \bmod q$, and its shortness

Table 2. Table of parameters for the partial lattice trapdoor scheme in Fig. 2, when $(\mathsf{TrapGen}, \mathsf{SampPre})$ is instantiated with the (full) lattice trapdoor scheme specified in Lemma 13. Lemma 16 uses Lemmas 17 to 20.

φ	$\in \mathsf{poly}(\lambda)$, $\varphi(m-k) \geq \Omega(\lambda)$	Ring degree
q	unramified prime with $q\mathcal{R} = \prod_{j\in[g]} \mathfrak{q}_j$ where $ng/q^{\varphi\cdot(m-n+1)/g} \leq \mathsf{negl}(\lambda)$ (Thm. 2, Lem. 19)	Modulus
δ	≥ 2	Base of $\mathbf{G}$
n	≥ 1	Matrix row parameter
$\tilde{m}$	$= n \cdot \lceil \log_\delta q \rceil$	Width of $\mathbf{G}_n$
m	$= \tilde{m} + n$	Matrix column parameter
t, k	$t \leq k$	Threshold, number of parties
η_A	$\geq 8\,\varphi\sqrt{2mt} \cdot q^{n/2m+1/(\varphi\cdot mt)}$	Upper bound of $\eta_\varepsilon(\Lambda_q^\perp(\mathbf{A}))$
η_C	$\geq 8\,\varphi\sqrt{m} \cdot q^{n/m+2/(\varphi\cdot m)}$	Upper bound of $\eta_\varepsilon(\Lambda_q^\perp(\mathbf{C}_j))$
η_G	$\geq \sqrt{5} \cdot \omega(\sqrt{\log(\varphi\cdot\tilde{m})})$	Upper bound of $\eta_\varepsilon(\Lambda_q^\perp(\mathbf{G}))$
B_A	$\geq \eta_A \cdot \sqrt{2\,\tilde{m}\cdot(m-\tilde{m})t\cdot\varphi}$	Upper bound of $\|\mathsf{td}_\mathbf{A}\|$
B_C	$\geq \eta_C \cdot \sqrt{2\,\tilde{m}\cdot(m-\tilde{m})\cdot\varphi}$	Upper bound of $\left\|\mathsf{td}_{\mathbf{C}_j}\right\|$
B_U	$\geq s_{\max}(\mathbf{S}_{\mathtt{Cor}}) \cdot \sqrt{m^2 t\varphi}$	Upper bound of $\|\mathbf{U}_j\|$
$\mathbf{S}$	$= \mathbf{S}_{\mathtt{Cor}}$	Gaussian width of $\mathsf{PTrapGen}$
$\tilde{\mathbf{S}}_i$	$\tilde{\mathbf{S}}_i = \mathbf{S}_{\mathtt{Cor}}\ \ \forall i \in [m(t-1)]$, $\tilde{\mathbf{S}}_i = \mathbf{S}_{\mathtt{Hon}}\ \ \forall i \in [2m(t-1)] \setminus [m(t-1)]$	Assumption parameter
a	≥ 1	Parameter regulating $\tilde{\mathbf{S}}_i$, $i = 0, 1$
$\mathbf{S}_{\mathtt{Hon}}$	$(2m(t-1)\text{-}\mathsf{MSIS}_{\mathcal{R}_q,n,2mt,\{\tilde{\mathbf{S}}_i\}_i,\alpha}$ is plausibly hard), $(2m(t-1)\text{-}\mathsf{MLWE}_{\mathcal{R}_q,n,2mt,\{\tilde{\mathbf{S}}_i\}_i,\chi_0}$ is plausibly hard), $\max\{\eta_A, 2\sqrt{\varphi}\cdot(a^k q^n)^{1/(m-k)}\} < s_{\min}(\mathbf{S}_{\mathtt{Hon}})$ (Lem. 19), $s_{\max}(\mathbf{S}_{\mathtt{Hon}}) \leq \min\{q^{1/g}/\sqrt{m}, a\cdot s_{\min}(\mathbf{S}_{\mathtt{Hon}})\}$ (Lem. 19)	Assumption parameter
$\mathbf{S}_{\mathtt{Cor}}$	$(2m(t-1)\text{-}\mathsf{MSIS}_{\mathcal{R}_q,n,2mt,\{\tilde{\mathbf{S}}_i\}_i,\alpha}$ is plausibly hard), $(2m(t-1)\text{-}\mathsf{MLWE}_{\mathcal{R}_q,n,2mt,\{\tilde{\mathbf{S}}_i\}_i,\chi_0}$ is plausibly hard), $s_{\min}(\mathbf{S}_{\mathtt{Cor}}) \geq \delta\cdot B_A\cdot\omega(\sqrt{\log(2mt\cdot\varphi)})$ (Lem. 17), $s_{\min}(\mathbf{S}_{\mathtt{Cor}}) \geq \eta_A$ (Lems. 18 and 20), $\max\{\eta_A, 2\sqrt{\varphi}\cdot(a^k q^n)^{1/(m-k)}\} < s_{\min}(\mathbf{S}_{\mathtt{Cor}})$ (Lem. 19), $s_{\max}(\mathbf{S}_{\mathtt{Cor}}) \leq \min\{q^{1/g}/\sqrt{m}, a\cdot s_{\min}(\mathbf{S}_{\mathtt{Cor}})\}$ (Lem. 19) $s_{\min}(\mathbf{S}_{\mathtt{Cor}}) \geq 2\cdot\delta\cdot\eta_C\cdot s_{\max}(\mathbf{S}_{\mathtt{Hon}})\cdot\varphi\cdot m\sqrt{(m-\tilde{m})t(t-1)}\cdot \omega(\sqrt{\log((2mt+\tilde{m})\,\varphi)})$ (Lem. 20)	assumption parameter
$\mathbf{S}_{\mathtt{BASIS}}$	$= \begin{pmatrix}\mathbf{S}_{\mathtt{Cor}} \\ & \sigma_G\mathbf{I}_{\tilde{m}}\end{pmatrix}, \sigma_G \geq \eta_G$	Parameter in Lemma 20
σ	$\geq B_U\cdot\delta\cdot B_C\cdot\omega(\sqrt{\log(m\cdot\varphi)})$ (Lem. 17), $\geq B_U\cdot\eta_C$ (Lem. 18), $\geq 2\cdot B_U\cdot\omega(\sqrt{\log(\varphi\cdot m\cdot t)})$ (Thm. 3)	Gaussian width of $\mathsf{PSampPre}$
β	$\geq t\cdot\sigma\sqrt{2mt\varphi}$	Correctness parameter
β'	$\leq \alpha - t\cdot\sigma\cdot\sqrt{2mt\varphi}$ (Thm. 3)	One-wayness parameter
α	$(2m(t-1)\text{-}\mathsf{MSIS}_{\mathcal{R}_q,n,2mt,\{\tilde{\mathbf{S}}_i\}_i,\alpha}$ is plausibly hard)	Assumption parameter
χ_0	$(2m(t-1)\text{-}\mathsf{MLWE}_{\mathcal{R}_q,n,2mt,\{\tilde{\mathbf{S}}_i\}_i,\chi_0}$ is plausibly hard)	Indistinguishability/ assumption parameter
χ_1	$\geq \lambda^{\omega(1)}\cdot\chi_0\cdot\chi_2 2gmt$ (Thm. 4)	Indistinguishability parameter
χ_2	$\geq 8\varphi\sqrt{2gmt}\cdot q^{n/2gm+2/(\varphi\cdot 2gmt)}$ (Thm. 4), $> 8\varphi\sqrt{2m}\cdot q^{1/2m+1/\varphi m}$ (Thm. 4)	Parameter in proof of Thm. 4

follows from standard norm bounds. We could also show that $\mathbf{x}^\dagger \neq \mathbf{0}$ with overwhelming probability basically due to the entropy of the withheld preimage $\mathbf{x}_{i^*}$. Finally, we argue that $\mathbf{x}^\dagger$ is not in the span of $(\mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}})$ by Proposition 4.

Theorem 3. *Let parameters be as in Table 2. The partial trapdoor scheme in Fig. 2 is $(nt, 2mt, \beta', \sigma)$-one-way, if the κ-$\mathsf{MSIS}_{\mathcal{R}_q,n,2mt,\{\tilde{\mathbf{S}}_i\}_i,\alpha}$ assumption holds for $\kappa = 2m\cdot(t-1)$.*

Proof. Let $\mathcal{A}$ be a PPT adversary winning the one-wayness experiment defined in Fig. 1 with a non-negligible probability. We build a PPT algorithm $\mathcal{B}^\mathcal{A}$ solving the κ-$\mathsf{MSIS}_{\mathcal{R}_q,n,2mt,\{\tilde{\mathbf{S}}_i\}_i,\alpha}$ problem with a non-negligible probability, where $\kappa = 2m\cdot(t-1)$. As before, we suppose that $\mathcal{A}$ has declared $\mathcal{C} \subset_{t-1} [k]$ to be the set of

$t-1$ corrupt parties in the security experiment and w.l.o.g. assume that $0 \notin \mathcal{C}$. Consider the following hybrid experiments:[21]

- Hyb_0: The real one-wayness experiment as defined in Fig. 1. Note that for any possible sequence of queries $(T_1, \ldots, T_L, (T^*, i^*))$ the view

$$\left(\mathbf{A}, \left(\mathsf{ptd}_j = (\mathbf{U}_j, \mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j})\right)_{j\in\mathcal{C}}, (\mathbf{x}_{\ell,j}, \mathbf{y}_\ell)_{\ell\in[L], j\in T_\ell}, \left(\mathbf{x}_j^*\right)_{j\in T^*\setminus\{i^*\}}, \mathbf{y}^*\right) \tag{8}$$

  of the adversary $\mathcal{A}$ almost follows the distribution D_0 defined in Fig. 3, except that it additionally obtains the values $\left(\mathbf{x}_j^*\right)_{j\in T^*\setminus\{i^*\}}, \mathbf{y}^*$. The latter can be generated in the same way that $(\mathbf{x}_{\ell,j}, \mathbf{y}_\ell)_{\ell\in[L], j\in T_\ell}$ are generated, except that $\mathbf{x}_{i^*}$ is withheld from the adversary $\mathcal{A}$.
- Hyb_4: A modified one-wayness experiment where the view Eq. (9) of $\mathcal{A}$ is instead generated as in distribution D_4 defined in Fig. 4, with the same treatment to $\left(\mathbf{x}_j^*, \mathbf{y}^*\right)_{j\in T^*\setminus\{i^*\}}$ mentioned above in Hyb_0.

By Lemma 16, we can conclude that Hyb_0 is statistically close in λ to Hyb_4.

We next show that the adversary $\mathcal{A}$'s response in Hyb_4 can be converted into a solution for the κ-MSIS instance $\left(\tilde{\mathbf{A}}_0, \mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}}\right)$ with non-negligible probability. First, we note that $\mathcal{B}^{\mathcal{A}}$ can simulate Hyb_4 when given a κ-MSIS instance $\left(\tilde{\mathbf{A}}_0, \mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}}\right)$ by running the subroutine $(\mathbf{A}, \left(\mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}\right)_{j\in\mathcal{C}}) \leftarrow \mathsf{ExtendA}(\mathbf{A}, \tilde{\mathbf{A}}_0, \mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}})$ defined in Fig. 4 and then following the procedures in Hyb_4 (i.e. distribution D_4). To extract a κ-MSIS solution, we recall that in the simulation $\mathcal{B}^{\mathcal{A}}$ samples $\mathbf{x}_j \leftarrow \mathcal{D}_{\Lambda(\mathbf{U}_j),\sigma}$ for $j \in T^*$, and let $\mathbf{y}^* = \mathbf{A} \cdot \sum_{j\in T} \mathbf{x}_j \bmod q$. $\mathcal{A}$ is given $\mathbf{y}^*$ and $\mathbf{x}_j$ for $j \neq i^*$ and returns $\mathbf{x}^*$ such that $\mathbf{A} \cdot \mathbf{x}^* = \mathbf{y}^* \bmod q$. Since $(\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n)$ is invertible over $\mathcal{R}_q$, we have

$$\begin{aligned}
(\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n) \begin{pmatrix} \tilde{\mathbf{A}}_0 \\ \tilde{\mathbf{A}}_{\neq 0} \end{pmatrix} &= \mathbf{A} \bmod q, \\
(\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n) \begin{pmatrix} \tilde{\mathbf{A}}_0 \\ \tilde{\mathbf{A}}_{\neq 0} \end{pmatrix} (\mathbf{x}^* - \sum_{j\in T^*} \mathbf{x}_j) &= \mathbf{0} \bmod q, \\
\tilde{\mathbf{A}}_0 (\mathbf{x}^* - \sum_{j\in T^*} \mathbf{x}_j) &= \mathbf{0} \bmod q
\end{aligned}$$

i.e. we have a candidate SIS solution $\mathbf{x}^\dagger \coloneqq \mathbf{x}^* - \sum_{j\in T^*} \mathbf{x}_j \bmod q$ for $\tilde{\mathbf{A}}_0$. We show that $\mathbf{x}^\dagger$ is a valid κ-MSIS solution with overwhelming probability. That is, it satisfies the following with overwhelming probability:

$$(1)\ \ \mathbf{x}^\dagger \neq \mathbf{0}, \qquad (2)\ \ \left\|\mathbf{x}^\dagger\right\| \leq \alpha, \qquad (3)\ \ \mathbf{x}^\dagger \notin \mathcal{K}\text{-}\mathsf{span}(\mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}}).$$

The vector $\mathbf{x}_{i^*}$ is sampled from the lattice $\Lambda(\mathbf{U}_{i^*})$. By Lemma 4 and Lemma 6:

$$\eta_\varepsilon(\Lambda(\mathbf{U}_{i^*})) \leq s_{\max}(\mathbf{U}_{i^*}) \cdot \omega(\sqrt{\log(2mt\varphi)}) \leq B_U \cdot \omega(\sqrt{\log(2mt\varphi)})$$

with overwhelming probability, where B_U is specified in Table 2. Consequently, $\sigma > 2\eta_\varepsilon(\Lambda(\mathbf{U}_{i^*}))$ with overwhelming probability. Since $\mathbf{x}_{i^*}$ is not revealed to the adversary the distributions of $\mathbf{x}^*$ and $\mathbf{x}_{i^*}$ are independent. Then the probability of $\mathbf{x}_{i^*} = \mathbf{x}^* - \sum_{j\in T^*\setminus\{i^*\}} \mathbf{x}_j$ is bounded by Lemma 7:

$$\max_{\mathbf{x}^*} \Pr_{\mathbf{x}_{i^*}\leftarrow\mathcal{D}_{\Lambda(\mathbf{U}_{i^*}),\sigma}} \left(\mathbf{x}_{i^*} = \mathbf{x}^* - \sum_{j\in T^*\setminus\{i^*\}} \mathbf{x}_j\right) \leq 2^{-m\varphi} \cdot \frac{1+\varepsilon}{1-\varepsilon} \leq \mathsf{negl}(\lambda).$$

To bound the norm of $\mathbf{x}^\dagger$ we simply use triangle inequality and the Gaussian tail bound from Lemma 6, so

$$\left\|\mathbf{x}^\dagger\right\| \leq \|\mathbf{x}^*\| + \sum_{j\in T^*} \|\mathbf{x}_j\| \leq \beta' + t \cdot \sigma \cdot \sqrt{2mt\varphi} \leq \alpha$$

[21] We only define Hyb_0 and Hyb_4 since they naturally correspond to distributions D_0 and D_4 in Figs. 3 and 4 respectively, which are statistically close in λ by Lemma 16.

holds with overwhelming probability. Finally, we show that

$$\Pr\left[\mathbf{x}^\dagger \notin \mathcal{K}\text{-}\mathsf{span}(\mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}})\right] \geq 1 - \mathsf{negl}(\lambda).$$

Define the shorthand $\mathbf{U}_{\neq 0} \coloneqq (\mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}})$. Suppose $\mathbf{x}^\dagger \notin \mathcal{K}\text{-}\mathsf{span}\left(\mathbf{U}_{i^*}, \mathbf{U}_{\neq 0}\right)$, then $\mathbf{x}^\dagger \notin \mathcal{K}\text{-}\mathsf{span}\left(\mathbf{U}_{\neq 0}\right)$ and we are done. In the rest consider $\mathbf{x}^\dagger \in \mathcal{K}\text{-}\mathsf{span}\left(\mathbf{U}_{i^*}, \mathbf{U}_{\neq 0}\right)$.

Any $\mathbf{x}_{i^*} \in \mathcal{K}\text{-}\mathsf{span}\left(\mathbf{U}_{i^*}\right)$ can be written as $\mathbf{x}_{i^*}^{(0)} + \mathbf{x}_{i^*}^{(\neq 0)}$, where $\mathbf{x}_{i^*}^{(0)}$ is orthogonal to $\mathcal{K}\text{-}\mathsf{span}\left(\mathbf{U}_{\neq 0}\right)$ and $\mathbf{x}_{i^*}^{(\neq 0)} \in \mathcal{K}\text{-}\mathsf{span}\left(\mathbf{U}_{\neq 0}\right)$. By Proposition 4[22], we know that $\mathbf{x}_{i^*}^{(0)}$ is not trivial, i.e. lives in a space of dimension at least one. Since $\mathcal{D}_{\Lambda(\mathbf{U}_{i^*}),\sigma}$ is spherical, projecting orthogonally to $\mathcal{K}\text{-}\mathsf{span}\left(\mathbf{U}_{\neq 0}\right)$ produces a Gaussian of the same width but in dimension $\dim\left(\mathcal{K}\text{-}\mathsf{span}\left(\mathbf{U}_{i^*}\right) \setminus \mathcal{K}\text{-}\mathsf{span}\left(\mathbf{U}_{\neq 0}\right)\right) \geq 1$ over the ring, or $\geq \varphi$ over the reals. Since $\sigma > 2\,\eta_\varepsilon(\Lambda(\mathbf{U}_{i^*}))$ and the smoothing parameter of a projected sublattice can only shrink, we apply Lemma 7 and conclude that $\mathbf{x}_{i^*}^{(0)} = \mathbf{x}_{i^*} \bmod \mathbf{U}_{\neq 0}$ has min-entropy $\Omega(\varphi)$. Suppose $\mathbf{x}^\dagger = \mathbf{x}^* - \sum_{j\in T} \mathbf{x}_j \in \mathcal{K}\text{-}\mathsf{span}\left(\mathbf{U}_{\neq 0}\right)$, we have $\mathbf{x}_{i^*}^{(0)} = \mathbf{x}^* - \sum_{j\in T\setminus\{i^*\}} \mathbf{x}_j \bmod \mathbf{U}_{\neq 0}$. This happens with probability at most $2^{-\Omega(\varphi)} \leq \mathsf{negl}(\lambda)$.

Therefore, $\mathbf{x}^\dagger$ is a valid solution with overwhelming probability, which means that $\mathcal{B}^{\mathcal{A}}$ solves the $2m\cdot(t-1)$-MSIS problem with overwhelming probability. □

5.6 Security: Indistinguishability

Towards showing indistinguishability of our partial lattice trapdoor, we prove a version of leftover hash lemma concerning the κ-uniform (see Definition 4) and discrete Gaussian distributions. This result may be of independent interest.

Lemma 21 (Leftover hash lemma for κ-uniform distribution). *Let $\mathcal{R}_q$ be a ring of degree φ splitting into g fields. Let $m, \kappa, \log q, \chi, \chi' \in \mathsf{poly}(\lambda)$, where $\kappa < m$, and $\varepsilon', \varepsilon > 0$, $\chi > 8\varphi\sqrt{m-\kappa}\cdot q^{1/(m-\kappa)+2/\varphi\cdot(m-\kappa)}$. For any matrix $\mathbf{U} = (\mathbf{U}_0 || \ldots || \mathbf{U}_{g-1})$ where $\mathbf{U}_i \in \mathcal{R}_q^{m\times\kappa}$ are primitive, it holds that*

$$\left\{ \left(\mathbf{b}, v\right) \;\middle|\; \begin{array}{l} \mathbf{b} \leftarrow \mathcal{U}(\mathbf{U}^\perp) + \mathcal{D}_{\mathcal{R},\chi'}^{g\cdot m}, \mathbf{r} \leftarrow \mathcal{D}_{\mathcal{R},\chi}^{g\cdot m} \\ v \coloneqq \mathbf{b}^{\mathsf{T}} \cdot \mathbf{r} \bmod q \end{array} \right\} \approx_s \left\{ \left(\mathbf{b}, v\right) \;\middle|\; \begin{array}{l} \mathbf{b} \leftarrow \mathcal{U}(\mathbf{U}^\perp) + \mathcal{D}_{\mathcal{R},\chi'}^{g\cdot m} \\ v \leftarrow \mathcal{R}_q \end{array} \right\},$$

where $\mathbf{U}^\perp = \left\{\mathbf{x} \in \mathcal{R}_q^{g\cdot m} \mid \mathbf{x}^{\mathsf{T}} \cdot \mathbf{U} \equiv \mathbf{0} \bmod q\right\}$.

The key difference between Lemma 21 and standard versions of the leftover hash lemma for lattices is that the "uniform part", $\mathbf{b}$, is not uniformly random but random in the kernel of $\mathbf{U}$. The proof handles this by effectively considering a basis of this kernel in normal form. Since $\mathbf{U}$ is full rank, over a field this implies the presence of a identity matrix in this normal form. Thus, the coefficients corresponding to this identity matrix in the basis are uniformly random as required. Then, since the entries of $\mathbf{r}$ are index-wise independent, the dependence of the non-pivot indices on these pivots does not affect the uniformity of the result. Finally, the proof handles that $\mathcal{R}_q$ may not be a field by running this argument in each "CRT slot" of $\mathcal{R}_q$, which are fields. The price we pay for this is an additional factor g in the dimension. Note that we place no restriction on the distribution χ' since it is not used in the proof.

Proof. For $\mathbf{x} \in \mathbf{U}^\perp$ we write $\mathbf{x} = (\mathbf{x}_0, \ldots, \mathbf{x}_{g-1})$, where $\mathbf{x}_i \in \mathcal{R}^m$. Then it holds that $\sum_{i=0}^{g-1} \mathbf{x}_i \cdot \mathbf{U}_i = \mathbf{0} \bmod q$. For each set $\{\alpha_i\}_i$ s.t. $\sum_{i=0}^{g-1} \alpha_i^{\mathsf{T}} = \mathbf{0} \bmod q$ we prove the statement for uniform $\mathbf{x}_i^{\mathsf{T}}$ satisfying $\forall i : \mathbf{x}_i^{\mathsf{T}} \cdot \mathbf{U}_i = \alpha_i^{\mathsf{T}} \bmod q$. Let $\langle q\rangle = \prod_{i=0}^{g-1} \mathfrak{q}_i$. W.l.o.g pick $i = 0$ then $\mathcal{R}_q/\mathfrak{q}_0 = \mathbb{F}_{(0)}$ is a field. The matrix $\mathbf{U}_0$ is primitive in $\mathcal{R}_q$ which implies $\mathbf{U}_0 \bmod \mathfrak{q}_0$ contains a $\kappa\times\kappa$ invertible submatrix over $\mathbb{F}_{(0)}$. There exists a permutation matrix $\mathbf{T}_0 \in \mathcal{R}_q^{m\times m}$ such that $\mathbf{T}_0 \cdot \mathbf{U}_0 = [\mathbf{U}_{00} \| \mathbf{U}_{01}] \bmod \mathfrak{q}_0$ where $\mathbf{U}_{00} \in \mathbb{F}_{(0)}^{\kappa\times\kappa}$ is invertible and $\mathbf{U}_{01} \in \mathbb{F}_{(0)}^{(m-k)\times\kappa}$. Write $\mathbf{x}_0^{\mathsf{T}} \cdot \mathbf{T}_0^{-1} \coloneqq (\mathbf{x}_{00}^{\mathsf{T}}, \mathbf{x}_{01}^{\mathsf{T}}) \in \mathcal{R}_q^{\kappa} \times \mathcal{R}_q^{m-\kappa}$. Then for any $\mathbf{x} \in \mathbf{U}^\perp$, we can compute $\mathbf{x}_{00} \bmod \mathfrak{q}_0$ from the other variables

$$\mathbf{x}_0^{\mathsf{T}} \cdot \mathbf{T}_0^{-1} \cdot \mathbf{T}_0 \cdot \mathbf{U}_0 \equiv\ \alpha_0 \bmod \mathfrak{q}_0$$

[22] Recall $\mathbf{A}\mathbf{U}_{\neq 0} = (\ldots, \mathbf{v}_j \otimes (\mathbf{C}_j, \mathbf{G}), \ldots)_{j\in\mathcal{C}} \bmod q$ up to permutation of columns, $\mathbf{A}\mathbf{U}_{i^*} = \mathbf{v}_{i^*} \otimes \mathbf{C}_{i^*} \bmod q$ for some $\mathbf{C}_{i^*} \in \mathcal{R}_q^{n\times m}$, and $[t] \setminus \mathcal{C} = \{i^*\}$.

$$(\mathbf{x}_{00}^{\intercal}, \mathbf{x}_{01}^{\intercal}) \cdot \begin{bmatrix}\mathbf{U}_{00}\\ \mathbf{U}_{01}\end{bmatrix} \equiv \alpha_0 \bmod \mathfrak{q}_0$$
$$\mathbf{x}_{00}^{\intercal} \equiv \alpha_0 - \mathbf{x}_{01}^{\intercal} \cdot \mathbf{U}_{01} \cdot \mathbf{U}_{00}^{-1} \bmod \mathfrak{q}_0.$$

Since $\mathbf{x}$ is uniform over $\mathbf{U}^{\perp}$ and the mapping $\mathbf{T}_0^{-1}$ is a bijection, we have that $\mathbf{x}_{01}$ is uniform over $\mathbb{F}_{(0)}^{m-\kappa}$.

Write $\mathbf{b} = \mathbf{x} + \mathbf{e}$ where $\mathbf{x} \leftarrow \mathcal{U}(\mathbf{U}^{\perp})$, $\mathbf{e} = (\mathbf{e}_0, \ldots, \mathbf{e}_{g-1}) \leftarrow \mathcal{D}_{\mathcal{R},\chi'}^{g\cdot m}$ and $\mathbf{e}_i^{\intercal} \cdot \mathbf{T}_i^{-1} = (\mathbf{e}_{i0}^{\intercal}, \mathbf{e}_{i1}^{\intercal})$. Now for $\mathbf{r} \leftarrow \mathcal{D}_{\mathcal{R},\chi}^{m}$ we have

$$\begin{aligned}\mathbf{b}^{\intercal} \cdot \mathbf{r} &= \mathbf{x}^{\intercal} \cdot \mathbf{r} + \mathbf{e}^{\intercal} \cdot \mathbf{r}\\ &= \sum_{i=0}^{g-1} (\mathbf{x}_i^{\intercal} + \mathbf{e}_i^{\intercal}) \cdot \mathbf{T}_i^{-1} \cdot \mathbf{T}_i \cdot \mathbf{r}_i\\ &= \sum_{i=0}^{g-1} \underbrace{\left(\alpha_i - \mathbf{x}_{i1}^{\intercal} \cdot \mathbf{U}_{01} \cdot \mathbf{U}_{00}^{-1} + \mathbf{e}_{i0}^{\intercal}\right) \cdot \mathbf{r}_{i0}}_{\text{RHS}} + \underbrace{(\mathbf{x}_{i1}^{\intercal} + \mathbf{e}_{i1}^{\intercal}) \cdot \mathbf{r}_{i1}}_{\text{LHS}}.\end{aligned}$$

We now adapt the classic leftover hash lemma (LHL) for the LHS with matrix $(\mathbf{x}_{i1}^{\intercal} + \mathbf{e}_{i1}^{\intercal})$ and vector $\mathbf{r}_{i1}$. The RHS is a function of $\mathbf{x}_{i1}$ and parameters that are public or independent of $\mathbf{r}_{i1}$, so the distribution of the sum is still uniform and independent. However, we need to modify the LHL to argue about the distance $\mathsf{SD}\big(\left(\mathbf{x}_{i1}, \mathbf{e}_{i1}, (\mathbf{x}_{i1} + \mathbf{e}_{i1})^{\intercal} \cdot \mathbf{r}_{i1} \bmod \mathfrak{q}_i\right), \left(\mathbf{x}_{i1}, \mathbf{e}_{i1}, \mathcal{U}(\mathbb{F}_{(i)})\right)\big)$ for $\mathbf{x}_{i1} \leftarrow \mathbb{F}_{(i)}^{m-\kappa}$ and $\mathbf{e}_{i1} \leftarrow \mathcal{D}_{\mathcal{R},\chi'}^{m-\kappa}$ where $\mathbf{x}_{i1}, \mathbf{e}_{i1}$ are revealed separately. Leaking the two values separately does not increase the statistical distance as they can be simulated from a single uniform vector. For $\tilde{\mathbf{x}}_{i1} \leftarrow \mathbb{F}_{(i)}^{m-\kappa}$

$$\mathsf{SD}\left(\left(\tilde{\mathbf{x}}_{i1},\ \tilde{\mathbf{x}}_{i1}^{\intercal} \cdot \mathbf{r}_{i1} \bmod \mathfrak{q}_i\right),\quad \left(\tilde{\mathbf{x}}_{i1},\ \mathcal{U}(\mathbb{F}_{(i)})\right)\right) \le 2^{-\Omega(\varphi m)}$$

holds by leftover hash lemma from Lemma 3. Then we sample $\tilde{\mathbf{e}}_{i1} \leftarrow \mathcal{D}_{\mathcal{R},\chi'}^{m-\kappa}$ and set $\mathbf{e}_{i1} = \tilde{\mathbf{e}}_{i1}$ and $\mathbf{x}_{i1} = \tilde{\mathbf{x}}_{i1} - \tilde{\mathbf{e}}_{i1}$ and the statement we need follows. Then $\forall i = 0, \ldots, g-1$

$$\mathsf{SD}\left(\left(\{\mathbf{x}_{j1}, \mathbf{e}_{j1}\}_j,\ (\mathbf{x}_{i1} + \mathbf{e}_{i1})^{\intercal} \cdot \mathbf{r}_{i1} \bmod \mathfrak{q}_i\right),\ \left(\{\mathbf{x}_{j1}, \mathbf{e}_{j1}\}_j,\ \mathcal{U}(\mathbb{F}_{(i)})\right)\right) \le \mathsf{negl}(\lambda).$$

Since $\{\mathbf{r}_{i1}\}_i$ are sampled independently, the random variables $(\mathbf{x}_{i1} + \mathbf{e}_{i1})^{\intercal} \cdot \mathbf{r}_{i1}$ are independent too. Then the sum $\sum_{i=1}^{g-1} (\mathbf{x}_{i1} + \mathbf{e}_{i1})^{\intercal} \cdot \mathbf{r}_{i1}$ is distributed uniformly in the ring $\mathcal{R}_q$. □

The following is a direct corollary of Lemmas 21 and 26.

Corollary 1. *Let $g, \varphi, \kappa, m, n, \log q, \chi, \chi' \in \mathsf{poly}(\lambda)$, $a \ge 1$ be a real number, and $\{\mathsf{diag}(\mathbf{\Sigma}_i, g)\}_{i\in[\kappa]}$ be positive semi-definite matrices in $\mathbb{R}^{\varphi g m \times \varphi g m}$ with blocks equal to $\mathbf{\Sigma}_i$ on the diagonal satisfying the following constraints:*

- *$\mathcal{R}_q$ is a ring of degree φ splitting into g fields,*
- *$n \le m$, $\kappa < m$,*
- *$\varphi \cdot (m - \kappa) \ge \Omega(\lambda)$,*
- *there exists $\varepsilon \le \mathsf{negl}(\lambda)$ so that for all $i \in [\kappa]$ it holds*

$$\max\left\{\eta_A, 2\sqrt{\varphi} \cdot (a^{\kappa} q^n)^{1/(m-\kappa)}\right\} \le s_{\min}\left(\sqrt{\mathbf{\Sigma}_i}\right),$$
$$\min\left\{q^{1/g}/\sqrt{m}, a \cdot s_{\min}\left(\sqrt{\mathbf{\Sigma}_i}\right)\right\} \ge s_{\max}\left(\sqrt{\mathbf{\Sigma}_i}\right),$$

  where $\eta_A \ge 8\,\varphi\sqrt{m} \cdot q^{n/m + 2/(\varphi\cdot m)}$,
- *$ng/q^{\varphi\cdot(m-n+1)/g} \le \mathsf{negl}(\lambda)$,*
- *$\chi > 8\varphi\sqrt{m-\kappa} \cdot q^{1/(m-\kappa)+2/\varphi\cdot(m-\kappa)}$*

It holds that

$$\left\{ \left(\mathbf{A}, \mathbf{U}, \mathbf{b}, v\right) \;\middle|\; \begin{array}{l} \mathbf{A} \leftarrow \mathcal{R}_q^{n\times gm}\\ \mathbf{u}_i \leftarrow \mathcal{D}_{\Lambda_q^{\perp}(\mathbf{A}), \sqrt{\mathsf{diag}(\mathbf{\Sigma}_i, g)}}\ \forall i \in [\kappa]\\ \mathbf{U} \coloneqq (\mathbf{u}_i)_{i\in[\kappa]}\\ \mathbf{b} \leftarrow \mathcal{U}(\mathbf{U}^{\perp}) + \mathcal{D}_{\mathcal{R},\chi'}^{gm}\\ \mathbf{r} \leftarrow \mathcal{D}_{\mathcal{R},\chi}^{gm};\ v \coloneqq \mathbf{b}^{\intercal} \cdot \mathbf{r} \bmod q \end{array} \right\} \approx_s \left\{ \left(\mathbf{A}, \mathbf{U}, \mathbf{b}, v\right) \;\middle|\; \begin{array}{l} \mathbf{A} \leftarrow \mathcal{R}_q^{n\times gm}\\ \mathbf{u}_i \leftarrow \mathcal{D}_{\Lambda_q^{\perp}(\mathbf{A}), \sqrt{\mathsf{diag}(\mathbf{\Sigma}_i, g)}}\ \forall i \in [\kappa]\\ \mathbf{U} \coloneqq (\mathbf{u}_i)_{i\in[\kappa]}\\ \mathbf{b} \leftarrow \mathcal{U}(\mathbf{U}^{\perp}) + \mathcal{D}_{\mathcal{R},\chi'}^{gm}\\ v \leftarrow \mathcal{R}_q \end{array} \right\}.$$

Proof. Write $\mathbf{A} = [\mathbf{A}_0, \ldots, \mathbf{A}_{g-1}]$ where $\mathbf{A}_j \leftarrow \mathcal{R}_q^{n \times m}$ for $j \in [g]$. With overwhelming probability in λ, all $\mathbf{A}_j$ are primitive by Lemma 8, and $s_{\min}(\sqrt{\mathsf{diag}\,(\mathbf{\Sigma}_i, g)}) = s_{\min}(\sqrt{\mathbf{\Sigma}_i}) > \eta_\varepsilon(\Lambda_q^\perp(\mathbf{A}_j))$ for all $i \in [\kappa]$ by choice of parameters. Split $\mathbf{U} = [\mathbf{U}_0 || \ldots || \mathbf{U}_{g-1}]$ where $\mathbf{U}_j = (\mathbf{u}_{ij})_{i \in [\kappa]}$. Then $\sum_{j=0}^{g-1} \mathbf{A}_j \cdot \mathbf{U}_j = \sum_{j=0}^{g-1} \alpha_j = 0 \bmod q$ for some $(\alpha_j)_{j \in [g]}$. Then conditioned on the ith coordinate α_{ij} of α_j we have $\mathbf{u}_{ij} \leftarrow \mathcal{D}_{\Lambda_q^{\alpha_{ij}}(\mathbf{A}_j), \sqrt{\mathbf{\Sigma}_i}}$. The covariance matrix of the full vectors $\mathbf{u}_i$ is block diagonal so that we split the preimages on g vectors as above. Conditioned on these, every $\mathbf{U}_j$ is primitive with overwhelming probability in λ by Lemma 26. Then the claim follows from Lemma 21. □

We are ready to prove that our partial trapdoor construction in Fig. 2 has indistinguishability. The high level proof structure is similar to that of one-wayness (Theorem 3), with some challenges unique to κ-MLWE, which we highlight below.

As in the proof of one-wayness (Theorem 3), we observe that most components in the view of the adversary $\mathcal{A}$ in the indistinguishability experiments follow the distribution D_0 in Fig. 3, except that the adversary $\mathcal{A}$ additionally obtains the challenge partial preimages and images $\left(\mathbf{x}_j^*\right)_{j \in T^* \setminus \{i^*\}}, \mathbf{y}^*$ and the challenge LWE sample $(\mathbf{c}_0, c_1)$. The challenge partial preimages and images can be simulated as in D_0, while the challenge LWE sample can be generated by post-processing the rest of the components. Using Lemma 16, we can therefore hop to a pair of hybrid experiments $\mathsf{Hyb}_{b,4}$ where most components of the adversary's view is generated as in distribution D_4 in Fig. 4.

The main part of the proof is to show that this pair of hybrid experiments $\mathsf{Hyb}_{0,4}$ and $\mathsf{Hyb}_{1,4}$ are computationally indistinguishable under the κ-MLWE assumption. Note that given a κ-MLWE instance $\left(\tilde{\mathbf{A}}_0, \mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}}, \mathbf{b}\right)$, the components

$$\left(\mathbf{A}, \left(\mathbf{U}_j, \mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}\right)_{j \in \mathcal{C}}, (\mathbf{x}_{\ell,j}, \mathbf{y}_\ell)_{\ell \in [L], j \in T_\ell}, \left(\mathbf{x}_j^*\right)_{j \in T^* \setminus \{i^*\}}\right)$$

which are common in the adversary's view of both hybrid experiments can be simulated by our reduction as in distribution D_4 in Fig. 4. The remaining component $(\mathbf{y}^*, \mathbf{c}_0, c_1)$, where $\mathbf{y}^*$ is common in both hybrid experiments and $(\mathbf{c}_0, c_1)$ is dependent on the hidden bit b, require some care.

Note that in the hybrid experiments the distribution of $\mathbf{y}^*$ is induced by $\mathbf{x}_{i^*}^*$ sampled from a Gaussian distribution over $\Lambda(\mathbf{U}_{i^*})$. We show that $\mathbf{x}_{i^*}^*$ can instead be replaced by a Gaussian $\mathbf{r}$ sampled from the full module $\mathcal{R}^{2gmt}$ without noticeably affecting the distribution of $\mathbf{y}^*$. This sets us up to argue about the distribution of the challenge LWE sample $(\mathbf{c}_0, c_1)$ via Lemma 21. In more detail, we show that if the hidden bit in the given κ-MLWE instance is $b = 0$, i.e. $\mathbf{b}$ is an LWE sample, then our reduction perfectly simulates $\mathsf{Hyb}_{0,4}$. On the other hand, if the hidden bit is $b = 1$, then our reduction almost simulates $\mathsf{Hyb}_{1,4}$ in that c_1 is uniformly random due to Lemma 21, but $\mathbf{c}_0$ is of the form $\mathbf{c}_0 \leftarrow \mathcal{U}(\mathbf{U}^\perp) + \mathcal{D}_{\mathcal{R},\chi_0}^{2gmt} + \mathbf{s}_{\neq 0}^\mathsf{T} \cdot \tilde{\mathbf{A}}_{\neq 0} \bmod q$ instead of $\mathbf{s}_0^\mathsf{T} \tilde{\mathbf{A}}_0 + \mathbf{s}_{\neq 0}^\mathsf{T} \cdot \tilde{\mathbf{A}}_{\neq 0} + \mathbf{e}^\mathsf{T} \bmod q$. We are therefore left with two cases: Either the above gap is computationally unnoticeable, in which case our reduction can solve the given κ-MLWE instance $\left(\tilde{\mathbf{A}}_0, \mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}}, \mathbf{b}\right)$ with non-negligible probability, or the gap is computationally noticeable, in which case we can still write down another reduction from κ-MLWE.

Theorem 4. *Let parameters be as in Table 2. The partial lattice trapdoor scheme in Fig. 2 has $(nt, 2gmt, \chi_0, \chi_1, \sigma)$-indistinguishability, if for $\kappa = 2m \cdot (t-1)$ the κ-$\mathsf{MLWE}_{\mathcal{R}_q, n, 2gmt, \{\mathsf{diag}(\tilde{\mathbf{S}}_i, g)\}_i, \chi_0}$ assumption holds.*

Proof. Let $\mathcal{A}$ be a PPT adversary winning the indistinguishability experiment defined in Fig. 1 with a non-negligible probability. We build a PPT algorithm $\mathcal{B}^\mathcal{A}$ solving the κ-$\mathsf{MLWE}_{\mathcal{R}_q, n, 2gmt, \{\mathsf{diag}(\tilde{\mathbf{S}}_i, g)\}_i, \chi_0}$ problem with a non-negligible probability, where $\kappa = 2m \cdot (t-1)$. As before, we suppose that $\mathcal{A}$ has declared $\mathcal{C} \subset_{t-1} [k]$ to be the set of $t-1$ corrupt parties in the security experiment and w.l.o.g. assume that $0 \notin \mathcal{C}$. Consider the following hybrid experiments parametrised by a bit $b \in \{0, 1\}$:[23]

- $\mathsf{Hyb}_{b,0}$: The real indistinguishability experiments as defined in Fig. 1. Note that for any possible sequence of queries $(T_1, \ldots, T_L, (T^*, i^*))$ the view

$$\left(\mathbf{A}, \left(\mathsf{ptd}_j = (\mathbf{U}_j, \mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j})\right)_{j \in \mathcal{C}}, (\mathbf{x}_{\ell,j}, \mathbf{y}_\ell)_{\ell \in [L], j \in T_\ell}, \left(\mathbf{x}_j^*\right)_{j \in T^* \setminus \{i^*\}}, \mathbf{y}^*, \mathbf{c}_0, c_1\right) \tag{9}$$

[23] We only define $\mathsf{Hyb}_{b,0}$ and $\mathsf{Hyb}_{b,4}$ since they naturally correspond to distributions D_0 and D_4 in Figs. 3 and 4 respectively, which are statistically close in λ by Lemma 16.

of the adversary $\mathcal{A}$ almost follows the distribution D_0 defined in Fig. 3, except for two differences:

1. The adversary $\mathcal{A}$ additionally obtains the challenge partial preimages and images $\left(\mathbf{x}_j^*\right)_{j\in T^*\setminus\{i^*\}}, \mathbf{y}^*$.
2. The adversary $\mathcal{A}$ additionally obtains the challenge LWE sample $(\mathbf{c}_0, c_1)$.

As discussed in the proof of one-wayness (Theorem 3), the partial preimages and image $\left(\mathbf{x}_j^*\right)_{j\in T^*\setminus\{i^*\}}, \mathbf{y}^*$ can be generated in the same way that $(\mathbf{x}_{\ell,j}, \mathbf{y}_\ell)_{\ell\in[L], j\in T_\ell}$ are generated, except that $\mathbf{x}_{i^*}$ is withheld from the adversary $\mathcal{A}$. For the challenge LWE samples, they can be generated as described in Fig. 3, i.e. by sampling $\mathbf{s} \leftarrow \mathcal{R}_q^n,\ \mathbf{e} \leftarrow \chi_0^m$, setting $\mathbf{c}_0 \coloneqq \mathbf{s}^\mathsf{T}\cdot\mathbf{A} + \mathbf{e}^\mathsf{T} \bmod q$, and setting $c_1 \coloneqq \mathbf{s}^\mathsf{T}\cdot\mathbf{y}^* + e \bmod q$ with $e \leftarrow \chi_1$ when $b = 0$ and $c_1 \leftarrow \mathcal{R}_q$ when $b = 1$.

– $\mathsf{Hyb}_{b,4}$: Modified indistinguishability experiments where the view Eq. (9) of $\mathcal{A}$, except for the $(\mathbf{c}_0, c_1)$ component, is instead generated as in distribution D_4 defined in Fig. 4, with the same treatment to $\left(\mathbf{x}_j^*, \mathbf{y}^*\right)_{j\in T^*\setminus\{i^*\}}$ mentioned above in Hyb_0. After that, the challenge LWE sample $(\mathbf{c}_0, c_1)$ is generated as described in the real experiments Fig. 1 and recalled above.

By Lemma 16, we can conclude that $\mathsf{Hyb}_{b,0}$ is statistically close in λ to $\mathsf{Hyb}_{b,4}$.

Reduction. We next show that if the adversary $\mathcal{A}$ succeeds in distinguishing $\mathsf{Hyb}_{0,4}$ and $\mathsf{Hyb}_{1,4}$ with non-negligible probability, we can build a PPT algorithm $\mathcal{B}^{\mathcal{A}}$ which distinguishes a κ-MLWE instance $\left(\tilde{\mathbf{A}}_0, \mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}}, \mathbf{b}\right)$ with non-negligible probability. To recall, $\tilde{\mathbf{A}}_0 \leftarrow \mathcal{R}_q^{n\times 2gmt}$ is a uniformly random matrix, $(\mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}})$ are Gaussian with varying Gaussian widths subject to $\tilde{\mathbf{A}}_0\cdot(\mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}}) \equiv \mathbf{0}$, and either $\mathbf{b}^\mathsf{T} \equiv \mathbf{s}_0^\mathsf{T}\cdot\tilde{\mathbf{A}}_0 + \mathbf{e}_0^\mathsf{T} \bmod q$ (Case $b = 0$) or $\mathbf{b} \leftarrow \mathcal{R}_q^{2gmt}$ (Case $b = 1$). The task of $\mathcal{B}^{\mathcal{A}}$ is to recover the hidden bit b.

First, we note that $\mathcal{B}^{\mathcal{A}}$ given κ-MLWE instance $\left(\tilde{\mathbf{A}}_0, \mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}}, \mathbf{b}\right)$ can simulate most of $\mathsf{Hyb}_{b,4}$, namely the values

$$\left(\mathbf{A}, \left(\mathbf{U}_j, \mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}\right)_{j\in\mathcal{C}}, (\mathbf{x}_{\ell,j}, \mathbf{y}_\ell)_{\ell\in[L], j\in T_\ell}\right)$$

by running the subroutine $(\mathbf{A}, \left(\mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}\right)_{j\in\mathcal{C}}) \leftarrow \mathsf{ExtendA}(\mathbf{A}, \tilde{\mathbf{A}}_0, \mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}})$ defined in Fig. 4 and then following the procedures in distribution D_4. Then, $\mathcal{B}^{\mathcal{A}}$ could simulate the challenge partial preimages and images $\left(\mathbf{x}_j^*\right)_{j\in T^*\setminus\{i^*\}}, \mathbf{y}^*$ and the challenge LWE sample $(\mathbf{c}_0, c_1)$ based on the internal randomness used to generate the above. To be precise, we detail the procedures below:

Public matrix A. For $j \in \mathcal{C}$, sample $(\mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}) \leftarrow \mathsf{TrapGen}(1^\lambda, 1^n, 1^m)$. Then sample random $\tilde{\mathbf{A}}_{\neq 0} \in \mathcal{R}_q^{n(t-1)\times 2gmt}$ subject to $\tilde{\mathbf{A}}_{\neq 0}\cdot\mathbf{U}_{\mathtt{Cor}} \equiv \left(\mathsf{diag}\left(\{\mathbf{C}_j\}_{j\in\mathcal{C}}\right), \mathbf{G}_{n(t-1)}\right)$ and let

$$\mathbf{A} \coloneqq \left[\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n\right]\begin{bmatrix}\tilde{\mathbf{A}}_0\\ \tilde{\mathbf{A}}_{\neq 0}\end{bmatrix} \bmod q.$$

Note that $\tilde{\mathbf{A}}_{\neq 0}$ is efficiently sampleable by Section 3.5.

Partial trapdoors of corrupt parties. For each $j \in \mathcal{C}$, the partial trapdoor is $\mathsf{ptd}_j = (\mathbf{U}_j, \mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j})$, where $\mathbf{U}_j \coloneqq \mathbf{U}_{\mathtt{Cor}}\cdot(\hat{\mathbf{e}}_{\mathcal{C},j} \otimes \mathbf{I}_n)$, i.e. the j-th column chunk of $\mathbf{U}_{\mathtt{Cor}}$.

Partial trapdoors of non-corrupt parties. Let $\mathbf{T} \coloneqq \begin{bmatrix}\mathbf{0}_{2gmt\times\tilde{m}} & \mathbf{U}_{\mathtt{Hon}}\\ -\mathbf{I}_{\tilde{m}} & \mathbf{0}_{\tilde{m}\times m(t-1)}\end{bmatrix}$. For each $j \in [k] \notin \mathcal{C}$, let $\mathbf{B}_j \coloneqq \left[\mathbf{H}_j\mathbf{A}\middle|\mathbf{1}_t \otimes -\mathbf{G}_n\right]$ so that $\mathbf{B}_j\,\mathbf{T} \equiv \hat{\mathbf{H}}_j\mathbf{G}_{nt} \bmod q$, where $\mathbf{H}_j$ and $\hat{\mathbf{H}}_j$ are defined in the beginning of Section 5.4, i.e. $\mathbf{T}$ is a gadget trapdoor of $\mathbf{B}_j$ with tag matrix $\hat{\mathbf{H}}_j$. Run

$$\begin{bmatrix}\mathbf{U}_j\\ \mathbf{W}_j\end{bmatrix} \leftarrow (\mathsf{SampPre}(\mathbf{B}_j, \mathbf{T}, \hat{\mathbf{H}}_j, \mathbf{0}, \mathbf{S}_{\mathtt{Cor}}))^m$$

to obtain $\mathbf{U}_j \in \mathcal{R}_q^{2gmt\times m}$ and discard $\mathbf{W}_j$.

Preimage oracle answers. To answer each query $\mathsf{PSampPreO}(T)$: For $j \in T$, sample $\mathbf{x}_j \leftarrow \mathcal{D}_{\Lambda(\mathbf{U}_j),\sigma}$. Let $\mathbf{y} \coloneqq \mathbf{A}\sum_{j\in T}\mathbf{x}_j \bmod q$. Return $\left((\mathbf{x}_j)_{j\in T}, \mathbf{y}\right)$.

Hints $\left(\mathbf{x}_j^*\right)_{j\in T^*\setminus\{i^*\}}$ and challenge $\mathbf{y}^*$. Upon receiving (T^*, i^*) from $\mathcal{A}$, simulate $\left(\left(\mathbf{x}_j^*\right)_{j\in T^*\setminus\{i^*\}}, \mathbf{y}^*\right)$ as follows:

– For $j \in T^* \setminus \{i^*\}$, sample $\mathbf{x}_j^* \leftarrow \mathcal{D}_{\Lambda(\mathbf{U}_j),\sigma}$. Compute $\mathbf{z}_j$ satisfying $\mathbf{A} \cdot \mathbf{x}_j^* \equiv \mathbf{v}_j \otimes \mathbf{z}_j$.
– Sample a Gaussian vector $\mathbf{r} \leftarrow \mathcal{D}_{\mathcal{R},\chi_2}^{2gmt}$, let $\mathbf{y}_0^* \coloneqq \tilde{\mathbf{A}}_0 \cdot \mathbf{r} \bmod q$.
– Solve, by linear algebra, for $(\mathbf{y}_{\neq 0}^*, \mathbf{z}_{i^*}) \in \mathcal{R}_q^{n(t-1)} \times \mathcal{R}_q^n$ satisfying

$$(\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n) \begin{bmatrix} \mathbf{y}_0^* \\ \mathbf{y}_{\neq 0}^* \end{bmatrix} \equiv \sum_{j \in T^* \setminus \{i^*\}} \mathbf{v}_j \otimes \mathbf{z}_j + \mathbf{v}_{i^*} \otimes \mathbf{z}_{i^*} \bmod q. \tag{10}$$

– Let $\mathbf{y}^* \coloneqq (\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n) \begin{bmatrix} \mathbf{y}_0^* \\ \mathbf{y}_{\neq 0}^* \end{bmatrix} \bmod q$.

Using the guarantee that $i^* \notin \mathcal{C}$ (as defined in the indistinguishability experiment Fig. 1), we show that the system of linear equations (with variables $(\mathbf{y}_{\neq 0}^*, \mathbf{z}_{i^*}) \in \mathcal{R}_q^{n(t-1)} \times \mathcal{R}_q^n$) in Eq. (10) has a unique solution. First, note that $\mathbf{V}_{\{0\}\cup\mathcal{C}} \in \mathcal{R}_q^{t\times t}$ is invertible since $0 \notin \mathcal{C}$. Multiplying the inverse of $(\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n)$ on both sides, we obtain

$$\begin{bmatrix} \mathbf{y}_0^* \\ \mathbf{y}_{\neq 0}^* \end{bmatrix} \equiv \underbrace{\sum_{j \in T^* \setminus \{i^*\}} \mathbf{V}_{\{0\}\cup\mathcal{C}}^{-1} \mathbf{v}_j \otimes \mathbf{z}_j}_{\begin{bmatrix} \mathbf{t}_0^* \\ \mathbf{t}_{\neq 0}^* \end{bmatrix} \coloneqq} + \underbrace{\mathbf{V}_{\{0\}\cup\mathcal{C}}^{-1} \mathbf{v}_{i^*}}_{\begin{bmatrix} a_0^* \\ \mathbf{a}_{\neq 0}^* \end{bmatrix} \coloneqq} \otimes \mathbf{z}_{i^*} \bmod q.$$

Note that if $a_0^* \in \mathcal{R}_q^\times$, then $(\mathbf{y}_{\neq 0}^*, \mathbf{z}_{i^*})$ is uniquely determined by

$$\mathbf{z}_{i^*} \equiv (\mathbf{y}_0^* - \mathbf{t}_0^*) \cdot (a_0^*)^{-1} \bmod q \qquad \text{and} \qquad \mathbf{y}_{\neq 0}^* \equiv \mathbf{t}_{\neq 0}^* + \mathbf{a}_{\neq 0}^* \otimes \mathbf{z}_{i^*} \bmod q.$$

Suppose towards contradiction that $a_0^* \notin \mathcal{R}_q^\times$. Then there must exist an ideal $I | q\mathcal{R}$ such that $a_0^* \equiv 0 \bmod I$. Hence

$$\mathbf{V}_{\{0\}\cup\mathcal{C}} \begin{bmatrix} 0 \\ \mathbf{a}_{\neq 0}^* \end{bmatrix} = \mathbf{V}_{\mathcal{C}} \cdot \mathbf{a}_{\neq 0}^* \equiv \mathbf{v}_{i^*} \bmod I.$$

This means that $\mathbf{v}_{i^*}$ is linearly dependent on the columns of $\mathbf{V}_{\mathcal{C}}$ modulo I, and hence $\mathbf{V}_{\mathcal{C}\cup\{i^*\}}$ is singular modulo I. However, since $v_j, v_{j'} \in \mathcal{R}_q$ for distinct $j, j' \in [k]$ are picked so that $v_j - v_{j'} \in \mathcal{R}_q^\times$, we have $v_j - v_{j'} \in (\mathcal{R}/I)^\times$ and hence $\mathbf{V}_{\mathcal{C}\cup\{i^*\}}$ is invertible modulo I, a contradiction.

LWE challenge $(\mathbf{c}_0, c_1)$. To simulate $(\mathbf{c}_0, c_1)$, sample $\mathbf{s}_{\neq 0} \leftarrow \mathcal{R}_q^{n(t-1)}$ and $e \leftarrow \chi_1$ and let

$$\mathbf{c}_0^\mathsf{T} = \mathbf{b}^\mathsf{T} + \mathbf{s}_{\neq 0}^\mathsf{T} \cdot \tilde{\mathbf{A}}_{\neq 0} \bmod q \qquad \text{and} \qquad c_1 \coloneqq \mathbf{b}^\mathsf{T} \cdot \mathbf{r} + \mathbf{s}_{\neq 0}^\mathsf{T} \cdot \mathbf{y}_{\neq 0} + e \bmod q.$$

The reduction passes all of the above to $\mathcal{A}$, then return whatever $\mathcal{A}$ returns.

Analysis. The distribution of

$$\left(\mathbf{A}, \left(\mathbf{U}_j, \mathbf{C}_j, \mathsf{td}_{\mathbf{C}_j}\right)_{j\in\mathcal{C}}, (\mathbf{x}_{\ell,j}, \mathbf{y}_\ell)_{\ell\in[L], j\in T_\ell}\right)$$

simulated by $\mathcal{B}^{\mathcal{A}}$ is clearly identical to that defined in $\mathsf{Hyb}_{b,4}$ for both $b \in \{0,1\}$. Below we analyse the distributions of $\left(\left(\mathbf{x}_j^*\right)_{j\in T^*\setminus\{i^*\}}, \mathbf{y}^*\right)$ and $(\mathbf{c}_0, c_1)$.

The partial preimages $\left(\mathbf{x}_j^*\right)_{j\in T^*\setminus\{i^*\}}$ are simulated in the same way as the answers to $\mathsf{PSampPreO}$ queries, hence they are well-distributed for the same reasons. We argue that $\mathbf{y}^*$ is also well-distributed. First, $\mathbf{y}_0^* = \tilde{\mathbf{A}}_0 \cdot \mathbf{r} \bmod q$ and since $\chi_2 > 8\varphi\sqrt{2gmt} \cdot q^{n/2gm+2/(\varphi\cdot 2gmt)}$, $\mathbf{y}_0^*$ is statistically close to uniform over $\mathcal{R}_q^n$ by Lemma 12. Recall that $\mathbf{y}^* \equiv \sum_{j\in T^*} \mathbf{v}_j \otimes \mathbf{z}_j + \mathbf{v}_{i^*} \otimes \mathbf{z}_{i^*} \bmod q$, where $(\mathbf{z}_j)_{j\in T^*\setminus\{i^*\}}$ are fixed by $(\mathbf{x}_j^*)_{j\in T^*\setminus\{i^*\}}$, which are available to $\mathcal{A}$, and $\mathbf{z}_{i^*}$ is uniquely determined by the system of linear equations in Eq. (10), as analysed above. Furthermore, we notice from the above analysis that the relation between $\mathbf{y}_0^*$ and $\mathbf{z}_{i^*}$ is bijective, and thus a uniformly random $\mathbf{y}_0^*$ induces a uniformly random $\mathbf{z}_{i^*}$. Following the proof of statistical closeness between D_1 and D_2 (Lemma 18), we can conclude that $\mathbf{v}_{i^*} \otimes \mathbf{z}_{i^*} \bmod q$ with a uniformly random $\mathbf{z}_{i^*}$ is statistically close to $\mathbf{A} \cdot \mathbf{x}_{i^*}^* \bmod q$ where $\mathbf{x}_{i^*}^* \leftarrow \Lambda(\mathbf{U}_{i^*})$, where the latter is as in $\mathsf{Hyb}_{b,4}$.

It remains to argue that $(\mathbf{c}_0, c_1)$ is distributed as in $\mathsf{Hyb}_{b,4}$ when the hidden bit in the κ-$\mathsf{MLWE}_{\mathcal{R}_q,n,2gmt,\{\mathsf{diag}(\tilde{\mathbf{S}}_i,g)\}_i,\chi_0}$ instance is b. First, we consider the $b=0$ case where $\mathbf{b}^\intercal = \mathbf{s}_0^\intercal \tilde{\mathbf{A}}_0 + \mathbf{e}^\intercal \bmod q$. For $\mathbf{c}_0$, we have

$$\begin{aligned}\mathbf{c}_0^\intercal &= (\mathbf{s}_0^\intercal, \mathbf{s}_{\neq 0}^\intercal)\begin{bmatrix}\tilde{\mathbf{A}}_0\\ \tilde{\mathbf{A}}_{\neq 0}\end{bmatrix} + \mathbf{e}^\intercal \bmod q\\ &= \underbrace{(\mathbf{s}_0^\intercal, \mathbf{s}_{\neq 0}^\intercal)(\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n)^{-1}}_{\mathbf{s}^\intercal}(\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n)\begin{bmatrix}\tilde{\mathbf{A}}_0\\ \tilde{\mathbf{A}}_{\neq 0}\end{bmatrix} + \mathbf{e}^\intercal = \mathbf{s}^\intercal \cdot \mathbf{A} + \mathbf{e} \bmod q,\end{aligned}$$

where $\mathbf{s}$ is uniformly distributed since both $\mathbf{s}_0, \mathbf{s}_{\neq 0}$ are, and multiplication by $(\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n)^{-1}$ is a bijective map. For c_1, we have

$$\mathbf{b}^\intercal \mathbf{r} + e = (\mathbf{s}_0^\intercal \cdot \tilde{\mathbf{A}}_0 + \mathbf{e}^\intercal) \cdot \mathbf{r} + e \approx_s \mathbf{s}_0^\intercal \cdot \tilde{\mathbf{A}}_0 \cdot \mathbf{r} + e = \mathbf{s}_0^\intercal \cdot \mathbf{y}_0^* + e \bmod q$$

where the statistical closeness follows from noise drowning (Lemma 9), since $\lambda^{\omega(1)}\chi_0\chi_2 \cdot 2gmt \leq \chi_1$. Hence

$$\begin{aligned}c_1 &\approx_s \mathbf{s}_0^\intercal \cdot \mathbf{y}_0^* + \mathbf{s}_{\neq 0}^\intercal \cdot \mathbf{y}_{\neq 0}^* + e \bmod q\\ &= (\mathbf{s}_0^\intercal, \mathbf{s}_{\neq 0}^\intercal)(\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n)^{-1}(\mathbf{V}_{\{0\}\cup\mathcal{C}} \otimes \mathbf{I}_n)\begin{bmatrix}\tilde{\mathbf{y}}_0\\ \mathbf{y}_{\neq 0}\end{bmatrix} + e = \mathbf{s}^\intercal \cdot \mathbf{y}^* + e \bmod q.\end{aligned}$$

Jointly, $(\mathbf{c}_0, c_1)$ is distributed statistically close to that in $\mathsf{Hyb}_{0,4}$.

Next, we consider the $b=1$ case where $\mathbf{b} \leftarrow \mathcal{U}(\mathbf{U}^\perp) + \mathcal{D}_{\mathcal{R},\chi_0}^{2gmt}$ for $\mathbf{U} = (\mathbf{U}_{\mathtt{Cor}}, \mathbf{U}_{\mathtt{Hon}})$. Recall that

$$\mathbf{c}_0^\intercal = \mathbf{b}^\intercal + \mathbf{s}_{\neq 0}^\intercal \cdot \tilde{\mathbf{A}}_{\neq 0} \bmod q \qquad \text{and} \qquad c_1 = \mathbf{b}^\intercal \cdot \mathbf{r} + e \bmod q.$$

By Corollary 1, the above distribution is statistically close to

$$\mathbf{c}_0^\intercal = \mathbf{b}^\intercal + \mathbf{s}_{\neq 0}^\intercal \cdot \tilde{\mathbf{A}}_{\neq 0} \bmod q \qquad \text{and} \qquad c_1 \leftarrow \mathcal{R}_q.$$

Note that this distribution of $(\mathbf{c}_0, c_1)$ is almost identical to that in $\mathsf{Hyb}_{1,4}$, except that in the former we have $\mathbf{b} \leftarrow \mathcal{U}(\mathbf{U}^\perp) + \mathcal{D}_{\mathcal{R},\chi_0}^{2gmt}$ and in the latter $\mathbf{b}^\intercal = \mathbf{s}_0^\intercal \tilde{\mathbf{A}}_0 + \mathbf{e}^\intercal \bmod q$ – precisely as in the κ-$\mathsf{MLWE}_{\mathcal{R}_q,n,2gmt,\{\mathsf{diag}(\tilde{\mathbf{S}}_i,g)\}_i,\chi_0}$ problem.

Therefore, we have that either $\mathcal{B}^{\mathcal{A}}$ solves the given κ-$\mathsf{MLWE}_{\mathcal{R}_q,n,2gmt,\{\mathsf{diag}(\tilde{\mathbf{S}}_i,g)\}_i,\chi_0}$ problem with non-negligible probability, and we are done, or there exists a PPT algorithm which distinguishes $(\mathbf{c}_0, c_1)$ as simulated above from those as sampled in $\mathsf{Hyb}_{1,4}$ with non-negligible probability, which implies another PPT algorithm for κ-$\mathsf{MLWE}_{\mathcal{R}_q,n,2gmt,\{\mathsf{diag}(\tilde{\mathbf{S}}_i,g)\}_i,\chi_0}$. □

6 Example Applications

We showcase simple constructions of threshold signatures and threshold identity-based encryption (IBE) as direct applications of partial trapdoors.

6.1 Definitions

We formally define the syntax of threshold signatures and threshold identity-based encryption (IBE), and the security notions that we consider for these primitives.

Threshold Signatures. A t-out-of-k threshold signature scheme, or (t,k)-signature scheme, for the number of signers k, threshold $t \leq k$ and message space $\mathcal{M}$ consists of PPT algorithms $(\mathsf{MKGen}, \mathsf{KGen}, \mathsf{Sign}, \mathsf{Rec}, \mathsf{Vf})$:

- $(\mathsf{mpk}, \mathsf{msk}) \leftarrow \mathsf{MKGen}(1^\lambda)$: The master key generation algorithm generates a master public key mpk and a master secret key msk.
- $\mathsf{sk}_j \leftarrow \mathsf{KGen}(\mathsf{mpk}, \mathsf{msk}, j)$: The key generation algorithm, on input mpk and msk and an index $j \in [k]$, generates a signing key for signer j.

$\mathsf{Exp}_{\Pi,\mathcal{A}}(1^\lambda)$

$\left(\mathcal{C}, (T_\ell, \mu_\ell)_{\ell\in[L]}, (T^*, i^*, \mu^*)\right) \leftarrow \mathcal{A}$

// set $\mathcal{C}$ of corrupt signers, signature queries $(T_\ell, \mu_\ell)_{\ell\in[L]}$, challenge (T^*, i^*, μ^*)

assert $\mathcal{C} \subset_{<t} [k]$; **assert** $i^* \in T^* \setminus \mathcal{C}$

if $\exists\, \ell, \ell' \in [L] : \mu_\ell = \mu_{\ell'}$: **return** $\perp$

$(\mathsf{mpk}, \mathsf{msk}) \leftarrow \mathsf{MKGen}(1^\lambda)$

for $j \in [k] : \mathsf{sk}_j \leftarrow \mathsf{KGen}(\mathsf{mpk}, \mathsf{msk}, j)$

for $\ell \in [L] : \mathsf{sig}_{\ell,j} \leftarrow \mathsf{Sign}\left(\mathsf{mpk}, \mathsf{sk}_j, T_\ell, \mu_\ell\right)\ \forall j \in T_\ell$

$\mathsf{sig}^*_j \leftarrow \mathsf{Sign}\left(\mathsf{mpk}, \mathsf{sk}_j, T^*, \mu^*\right)\ \forall j \in T^* \setminus \{i^*\}$

$\mathsf{sig}^* \leftarrow \mathcal{A}\left(\mathsf{mpk}, (\mathsf{sk}_j)_{j\in\mathcal{C}}, \left(\mathsf{sig}_{\ell,j}\right)_{\ell\in[L], j\in T_\ell}, \left(\mathsf{sig}^*_j\right)_{j\in T^*\setminus\{i^*\}}\right)$

return $(1 \leftarrow \mathsf{Vf}(\mathsf{mpk}, \mathsf{sig}^*, \mu^*))$

Fig. 5. Highly-selective security for t-out-of-k threshold signature scheme Π.

$\mathsf{sig}_j \leftarrow \mathsf{Sign}(\mathsf{mpk}, \mathsf{sk}_j, T, \mu)$ The signing algorithm, on input mpk, a signing key sk_j from signer j, a set $T \subseteq_t [k]$, and a message $\mu \in \mathcal{M}$, generates a partial signature sig_j for μ.

$\mathsf{sig} \leftarrow \mathsf{Rec}\left(\left(\mathsf{sig}_j\right)_{j\in T}\right)$: The reconstruction algorithm inputs a tuple of partial signatures sig_j from a set T of signers, and reconstructs a full signature sig.

$b \leftarrow \mathsf{Vf}(\mathsf{mpk}, \mathsf{sig}, \mu)$: The verification algorithm inputs mpk, a full signature sig and a message μ, and outputs a bit $b \in \{0, 1\}$.

Definition 9 (Correctness). *A (t, k)-signature scheme is said to be correct if, for any $(\mathsf{mpk}, \mathsf{msk}) \in \mathsf{MKGen}(1^\lambda)$, $j \in [k]$, $\mathsf{sk}_j \in \mathsf{KGen}(\mathsf{mpk}, \mathsf{msk}, j)$, $\mu \in \mathcal{M}$, and any set $T \subseteq_t [k]$, it holds that*

$$\Pr\left[1 \leftarrow \mathsf{Vf}(\mathsf{mpk}, \mathsf{sig}, \mu) \;\middle|\; \begin{array}{l} \mathsf{sig}_j \leftarrow \mathsf{Sign}\left(\mathsf{mpk}, \mathsf{sk}_j, T, \mu\right)\ \forall j \in T \\ \mathsf{sig} \leftarrow \mathsf{Rec}\left(\left(\mathsf{sig}_j\right)_{j\in T}\right) \end{array}\right] \geq 1 - \mathsf{negl}(\lambda).$$

Definition 10 (Security). *A (t, k)-signature scheme $\Pi = (\mathsf{MKGen}, \mathsf{KGen}, \mathsf{Sign}, \mathsf{Rec}, \mathsf{Vf})$ is said to be highly-selectively secure, if for any PPT $\mathcal{A}$ it holds that*

$$\Pr\left[\mathsf{Exp}_{\Pi,\mathcal{A}}(1^\lambda) = 1\right] \leq \mathsf{negl}(\lambda),$$

where $\mathsf{Exp}_{\Pi,\mathcal{A}}$ is defined in Fig. 5.

Threshold Identity-based Encryption. A t-out-of-k threshold identity-based encryption, or (t, k)-IBE, for the number of authorities k, threshold $t \leq k$, identity space $\mathcal{ID}$ and message space $\mathcal{M}$ consists of PPT algorithms $(\mathsf{MKGen}, \mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$:

$\left(\mathsf{mpk}, \left(\mathsf{msk}_j\right)_{j\in[k]}\right) \leftarrow \mathsf{MKGen}(1^\lambda)$: The master key generation algorithm generates a master public key mpk and a tuple of master secret keys mpk_j for each authority $j \in [k]$.

$\mathsf{sk}_j \leftarrow \mathsf{KGen}(\mathsf{mpk}, \mathsf{msk}_j, T, \mathsf{id})$: The key generation algorithm, on input mpk, msk_j from an authority j, a set $T \subseteq_t [k]$, and an identity id, generates a partial secret key sk_j for id.

$\mathsf{ctxt} \leftarrow \mathsf{Enc}(\mathsf{mpk}, \mathsf{id}, \mu)$: The encryption algorithm inputs mpk, $\mathsf{id} \in \mathcal{ID}$, and a message $\mu \in \mathcal{M}$, and outputs a ciphertext ctxt.

$\mu' \leftarrow \mathsf{Dec}\left(\mathsf{ctxt}, \left(\mathsf{sk}_j\right)_{j\in T}\right)$: The decryption algorithm, on input a ciphertext ctxt and a tuple of partial secret keys sk_j from a set T of authorities, outputs μ'.

$\mathsf{Exp}^{b}_{\Pi,\mathcal{A}}(1^\lambda)$

$\left(\mathcal{C}, (T_\ell, \mathsf{id}_\ell)_{\ell\in[L]}, (T^*, i^*, \mathsf{id}^*)\right) \leftarrow \mathcal{A}$

// set $\mathcal{C}$ of corrupt authorities, key queries $(T_\ell, \mathsf{id}_\ell)_{\ell\in[L]}$, challenge $(T^*, i^*, \mathsf{id}^*)$

assert $\mathcal{C} \subset_{<t} [k]$; **assert** $i^* \in T^* \setminus \mathcal{C}$

if $\exists \ell, \ell' \in [L] : \mathsf{id}_\ell = \mathsf{id}_{\ell'}$: **return** $\perp$

$\left(\mathsf{mpk}, (\mathsf{msk}_j)_{j\in[k]}\right) \leftarrow \mathsf{MKGen}(1^\lambda)$

for $\ell \in [L] : \mathsf{sk}_{\ell,j} \leftarrow \mathsf{KGen}\left(\mathsf{mpk}, \mathsf{msk}_j, T_\ell, \mathsf{id}_\ell\right)\ \forall j \in T_\ell$

$\mathsf{sk}^*_j \leftarrow \mathsf{KGen}\left(\mathsf{mpk}, \mathsf{msk}_j, T^*, \mathsf{id}^*\right)\ \forall j \in T^* \setminus \{i^*\}$

$(\mu_0, \mu_1) \leftarrow \mathcal{A}\left(\mathsf{mpk}, (\mathsf{msk}_j)_{j\in\mathcal{C}}, \left(\mathsf{sk}_{\ell,j}\right)_{\ell\in[L], j\in T_\ell}, \left(\mathsf{sk}^*_j\right)_{j\in T^*\setminus\{i^*\}}\right)$

$\mathsf{ctxt} \leftarrow \mathsf{Enc}(\mathsf{mpk}, \mathsf{id}^*, \mu_b)\ ;\ b' \leftarrow \mathcal{A}(\mathsf{ctxt})$

return b'

Fig. 6. Highly-selective security for t-out-of-k threshold IBE scheme Π.

$\mathsf{MKGen}(1^\lambda)$

$(\mathbf{A}, \mathsf{td}) \leftarrow \mathsf{TrapGen}(1^\lambda, 1^{nt}, 1^{2mt})$

return $(\mathsf{mpk}, \mathsf{msk}) \coloneqq (\mathbf{A}, \mathsf{td})$

$\mathsf{KGen}(\mathsf{mpk}, \mathsf{msk}, j)$

$\mathsf{ptd}_j \leftarrow \mathsf{PT.PTrapGen}(\mathbf{A}, \mathsf{td}, j)$

$\mathsf{sk}_j \coloneqq \mathsf{ptd}_j$; **return** sk_j

$\mathsf{Rec}(\left(\mathsf{sig}_j\right)_{j\in T})$

$\mathbf{x} \leftarrow \mathsf{PT.Rec}((\mathbf{x}_j)_{j\in T})$; $\mathsf{sig} \coloneqq \mathbf{x}$; **return** sig

$\mathsf{Sign}(\mathsf{mpk}, \mathsf{sk}_j, T, \mu)$

$\mathbf{y} \coloneqq \mathsf{H}(\mu) \in \mathcal{R}_q^{nt}$

$\mathbf{x}_j \leftarrow \mathsf{PT.PSampPre}(\mathbf{A}, \mathsf{ptd}_j, T, \mathbf{y}, \sigma)$

$\mathsf{sig}_j \coloneqq \mathbf{x}_j$; **return** sig_j

$\mathsf{Vf}(\mathsf{mpk}, \mathsf{sig}, \mu)$

$\mathbf{y} \coloneqq \mathsf{H}(\mu)$

$b_0 \coloneqq (\mathbf{A} \cdot \mathbf{x} = \mathbf{y} \bmod q)$

$b_1 \coloneqq (\|\mathbf{x}\| \le \beta)$

return $b_0 \wedge b_1$

Fig. 7. t-out-of-k threshold signature construction.

Definition 11 (Correctness). *A (t,k)-IBE is said to be correct if, for any $\mathsf{id} \in \mathcal{ID}$, $\mu \in \mathcal{M}$, set $T \subseteq_t [k]$, and $\left(\mathsf{mpk}, (\mathsf{msk}_j)_{j\in[k]}\right) \in \mathsf{MKGen}(1^\lambda)$, it holds that*

$$\Pr\left[\mu = \mu' \;\middle|\; \begin{array}{l} \mathsf{sk}_j \leftarrow \mathsf{KGen}(\mathsf{mpk}, \mathsf{msk}_j, T, \mathsf{id})\ \forall j \in T \\ \mathsf{ctxt} \leftarrow \mathsf{Enc}(\mathsf{mpk}, \mathsf{id}, \mu) \\ \mu' \leftarrow \mathsf{Dec}\left(\mathsf{ctxt}, \left(\mathsf{sk}_j\right)_{j\in T}\right) \end{array}\right] \ge 1 - \mathsf{negl}(\lambda).$$

Definition 12 (Security). *A (t,k)-IBE scheme $\Pi = (\mathsf{MKGen}, \mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ is said to be highly-selectively secure, if for any PPT $\mathcal{A}$ it holds that*

$$\left|\Pr\left[\mathsf{Exp}^0_{\Pi,\mathcal{A}}(1^\lambda) = 1\right] - \Pr\left[\mathsf{Exp}^1_{\Pi,\mathcal{A}}(1^\lambda) = 1\right]\right| \le \mathsf{negl}(\lambda),$$

where $\mathsf{Exp}^b_{\Pi,\mathcal{A}}$ is defined in Fig. 6.

6.2 Threshold Signatures Construction

Let $n, m, t, k, \log q, \tilde{s}, \beta, \sigma \in \mathsf{poly}(\lambda)$ where $m = n \cdot (\lceil \log q \rceil + 1)$ and $t \le k$. Let $\mathsf{H} : \mathcal{M} \to \mathcal{R}_q^{nt}$ be a hash function modelled by a random oracle, $(\mathsf{TrapGen}, \mathsf{SampPre})$ a lattice trapdoor scheme, PT the (t,k)-partial lattice trapdoor scheme in Fig. 2. In Fig. 7 we construct a (t,k)-signature scheme for the message space $\mathcal{M}$.

Theorem 5. *The (t,k)-signature scheme in Fig. 7 is correct, if the partial lattice trapdoor scheme* PT *is $(nt, 2mt, \tilde{s}, \beta, \sigma)$-correct.*

Theorem 5 is immediate so we omit the proof.

Theorem 6. *The (t,k)-signature scheme in Fig. 7 is highly-selectively secure, if the partial lattice trapdoor scheme* PT *is $(nt, 2mt, \beta', \sigma)$-one-way and $\beta \leq \beta'$.*

Proof. Suppose there exists a PPT $\mathcal{A}$ against the highly-selective security of the threshold signature scheme. We construct a PPT $\mathcal{B}^{\mathcal{A}}$ against the β'-one-wayness of PT. Let $\mathcal{B}^{\mathcal{A}}$ proceed as follows:

- Initialise list H. Receive from $\mathcal{A}$ set $\mathcal{C} \subset [k]$ of corrupt users, set $(T_\ell, \mu_\ell)_{\ell \in [L]}$ of all signature queries, and challenge (T^*, i^*, μ^*). W.l.o.g. assume $|\mathcal{C}| = t-1$.
- Pass $\mathcal{C}$ to the PT security challenger, and receive $\mathbf{A}$ and partial trapdoors $(\mathbf{U}_j)_{j \in \mathcal{C}}$. Let $\mathsf{mpk} \coloneqq \mathbf{A}$ and $\mathsf{sk}_j \coloneqq \mathbf{U}_j$ for all $j \in \mathcal{C}$.
- For each $\ell \in [L]$, query $\mathsf{PSampPreO}(T_\ell)$ to obtain $\left((\mathbf{x}_{\ell,j})_{j \in T}, \mathbf{y}_\ell\right)$, where $\mathbf{y}_\ell$ is uniform. Let $\mathsf{sig}_{\ell,j} \coloneqq \mathbf{x}_{\ell,j}$. Programme $\mathsf{H}[\mu_\ell] \coloneqq \mathbf{y}_\ell$.
- Pass (T^*, i^*) to PT challenger, and receive $\left(\left(\mathbf{x}_j^*\right)_{j \in T^* \setminus \{i^*\}}, \mathbf{y}^*\right)$, where $\mathbf{y}^*$ is uniform. Let $\mathsf{sig}_j^* \coloneqq \mathbf{x}_j^*$. Programme $\mathsf{H}[\mu^*] \coloneqq \mathbf{y}^*$.
- Upon any $\mathsf{H}[\mu]$ query from $\mathcal{A}$, if $\mathsf{H}[\mu] = \bot$, sample a random vector $\mathbf{y} \leftarrow \mathcal{R}_q^{nt}$, let $H[\mu] \coloneqq \mathbf{y}$. Answer with $H[\mu]$.
- Pass $\mathsf{mpk}, (\mathsf{sk}_j)_{j \in \mathcal{C}}, \left(\mathsf{sig}_{\ell,j}\right)_{\ell \in [L], j \in T_\ell}, \left(\mathsf{sig}_j^*\right)_{j \in T^* \setminus \{i^*\}}$ to $\mathcal{A}$, then return whatever $\mathcal{A}$ returns.

It is straightforward that mpk, the corrupt signing keys $\left(\mathsf{sk}_j\right)_{j \in \mathcal{C}}$, the answers to signature queries, and the challenge $\left(\mathsf{sig}_j^*\right)_{j \in T^* \setminus i^*}$ are all distributed identically to the real signature scheme by design. If $\mathcal{A}$ returns a valid forgery sig^*, this means $\mathbf{A} \cdot \mathsf{sig}^* = \mathbf{y}^* \bmod q$ and $\|\mathsf{sig}^*\| \leq \beta \leq \beta'$, hence $\mathcal{B}^{\mathcal{A}}$ breaks the β'-one-wayness of PT with sig^*. The success probability of $\mathcal{B}^{\mathcal{A}}$ is thus the same as $\mathcal{A}$. □

Remark 4. We note that a malicious party in the signing set T can always abort the protocol by going offline. Yet, we note that our scheme has identifiable aborts. The contribution of every party in the protocol is a short preimage $\mathbf{x}_j$ of $\mathbf{v}_j \otimes \mathbf{z}_j$ where the sets $\{\mathbf{v}_j\}$ and $\mathbf{z}_j$ can be computed by anyone who knows the signing set T. Therefore, to verify the contribution $\mathbf{x}_j$ of a participant, it is enough to check the norm bound and the image of $\mathbf{x}_j$. If the conditions are met then $\sum \mathbf{x}_j$ results in a valid signature.

6.3 Threshold Identity-Based Encryption Construction

Let $n, m, t, k, \log q, \tilde{s}, \beta, \sigma \in \mathsf{poly}(\lambda)$ where $m = n(\lceil \log q \rceil + 1)$ and $t \leq k$. Let χ_0, χ_1 be distributions parametrised by λ. Let $\mathsf{H} : \mathcal{ID} \to \mathcal{R}_q^{nt}$ be a hash function modelled by a random oracle, $(\mathsf{TrapGen}, \mathsf{SampPre})$ a lattice trapdoor scheme, PT the (t,k) partial lattice trapdoor scheme in Fig. 2. In Fig. 8 we construct a (t,k)-IBE scheme for the identity space $\mathcal{ID}$ and message space $\mathcal{M} = \{0,1\}$.

Theorem 7. *If the partial lattice trapdoor scheme* PT *is $(nt, 2mt, \tilde{s}, \beta, \sigma)$-correct and $q > 2(\chi_0 + \chi_1 \beta \sqrt{2mt\varphi})$, then the (t,k)-IBE scheme in Fig. 8 is correct.*

Proof. During decryption, we obtain $\mathbf{x} \leftarrow \mathsf{PT.Rec}((\mathbf{x}_j)_{j \in T})$ which, by the β-correctness of PT, satisfies $\mathbf{Ax} = \mathbf{y} = \mathsf{H}(\mathsf{id})$ and $\|\mathbf{x}\| \leq \beta$. Then

$$\begin{aligned} c = c_1 - \mathbf{c}_0^{\mathsf{T}}\mathbf{x} &= \mathbf{s}^{\mathsf{T}} \cdot \mathbf{y} + e + \lfloor q/2 \rfloor \mu - (\mathbf{s}^{\mathsf{T}} \cdot \mathbf{A} + \mathbf{e}^{\mathsf{T}}) \cdot \mathbf{x} \bmod q \\ &= \mathbf{s}^{\mathsf{T}} \cdot \mathbf{y} + e + \lfloor q/2 \rfloor \mu - \mathbf{s}^{\mathsf{T}} \cdot \mathbf{y} - \mathbf{e}^{\mathsf{T}} \cdot \mathbf{x} \bmod q \\ &= e - \mathbf{e}^{\mathsf{T}} \cdot \mathbf{x} + \lfloor q/2 \rfloor \mu \bmod q. \end{aligned}$$

With overwhelming probability $\|e - \mathbf{e}^{\mathsf{T}} \cdot \mathbf{x}\| \leq \chi_0 + \chi_1 \beta \sqrt{2mt\varphi}$, hence decryption correctness is guaranteed as long as $q > 2(\chi_0 + \chi_1 \beta \sqrt{2mt\varphi})$. □

$\underline{\mathsf{MKGen}(1^\lambda)}$

$(\mathbf{A}, \mathsf{td}) \leftarrow \mathsf{TrapGen}(1^\lambda, 1^{nt}, 1^{2mt})$
$\mathsf{mpk} \coloneqq \mathbf{A}$
for $j \in [k]$:
  $\mathsf{ptd}_j \leftarrow \mathsf{PT.PTrapGen}(\mathbf{A}, \mathsf{td}, j)$
  $\mathsf{msk}_j \coloneqq \mathsf{ptd}_j$
return $\left(\mathsf{mpk}, (\mathsf{msk}_j)_{j\in[k]}\right)$

$\underline{\mathsf{KGen}(\mathsf{mpk}, \mathsf{msk}_j, T, \mathsf{id})}$

$\mathbf{y} \coloneqq \mathsf{H}(\mathsf{id}) \in \mathcal{R}_q^{nt}$
$\mathbf{x}_j \leftarrow \mathsf{PT.PSampPre}(\mathbf{A}, \mathsf{msk}_j, T, \mathbf{y}, \sigma)$
$\mathsf{sk}_j \coloneqq \mathbf{x}_j$; **return** sk_j

$\underline{\mathsf{Enc}(\mathsf{mpk}, \mathsf{id}, \mu)}$

$\mathbf{y} \coloneqq \mathsf{H}(\mathsf{id})$
$\mathbf{s} \leftarrow \mathcal{R}_q^{nt}$; $\mathbf{e}_0 \leftarrow \chi^{2mt}$; $e_1 \leftarrow \chi$
$\mathbf{c}_0^\intercal \coloneqq \mathbf{s}^\intercal \mathbf{A} + \mathbf{e}_0^\intercal \bmod q$
$c_1 \coloneqq \mathbf{s}^\intercal \mathbf{y} + e_1 + \lfloor q/2 \rfloor \mu \bmod q$
$\mathsf{ctxt} \coloneqq (\mathbf{c}_0, c_1)$; **return** ctxt

$\underline{\mathsf{Dec}\left(\mathsf{ctxt}, (\mathsf{sk}_j)_{j\in T}\right)}$

$\mathbf{x} \leftarrow \mathsf{PT.Rec}((\mathbf{x}_j)_{j\in T})$
$c \coloneqq c_1 - \mathbf{c}_0^\intercal \mathbf{x} \bmod q$
if $c > q/4$: **return** 1
return 0

Fig. 8. t-out-of-k threshold IBE construction.

Theorem 8. *The (t, k)-IBE scheme in Fig. 8 is highly-selectively secure, if the partial lattice trapdoor scheme* PT *has $(nt, 2mt, \chi_0, \chi_1, \sigma)$-indistinguishability.*

Proof. thm:IBE-security Define the following hybrids:

$\mathsf{Hyb}_{b,0}$: The real highly-selective experiment encrypting μ_b.
$\mathsf{Hyb}_{b,1}$: The ciphertext component c_1 is swapped to uniform over $\mathcal{R}_q$.

We have $\mathsf{Hyb}_{b,0} \equiv \mathsf{Hyb}_{b,1}$ since c_1 is uniform in both hybrids. Below we show that $\mathsf{Hyb}_{b,0} \approx_c \mathsf{Hyb}_{b,1}$ if PT has (χ_0, χ_1)-indistinguishability.

Suppose there exists a PPT $\mathcal{A}$ distinguishing $\mathsf{Hyb}_{b,0}$ and $\mathsf{Hyb}_{b,1}$ with non-negligible probability. We construct a PPT $\mathcal{B}^{\mathcal{A}}$ against the (χ_0, χ_1)-indistinguishability of PT. Let $\mathcal{B}^{\mathcal{A}}$ proceed as follows:

- Initialise list H. Receive from $\mathcal{A}$ set $\mathcal{C} \subset [k]$ of corrupt users, set $(T_\ell, \mathsf{id}_\ell)_{\ell\in[L]}$ of all secret key queries, and challenge $(T^*, i^*, \mathsf{id}^*)$. W.l.o.g. assume $|\mathcal{C}| = t - 1$.
- Pass $\mathcal{C}$ to the PT indistinguishability challenger, and receive $\mathbf{A}$ and partial trapdoors $(\mathbf{U}_j)_{j\in\mathcal{C}}$. Let $\mathsf{mpk} \coloneqq \mathbf{A}$ and $\mathsf{msk}_j \coloneqq \mathbf{U}_j$ for all $j \in \mathcal{C}$.
- For each $\ell \in [L]$, query $\mathsf{PSampPreO}(T_\ell)$ to obtain $\left((\mathbf{x}_{\ell,j})_{j\in T}, \mathbf{y}_\ell\right)$, where $\mathbf{y}_\ell$ is uniform. Let $\mathsf{sk}_{\ell,j} \coloneqq \mathbf{x}_{\ell,j}$. Programme $\mathsf{H}[\mathsf{id}_\ell] \coloneqq \mathbf{y}_\ell$.
- Pass (T^*, i^*) to PT challenger, and receive $\left(\left(\mathbf{x}_j^*\right)_{j\in T^*\setminus\{i^*\}}, \mathbf{y}^*, \mathbf{c}_0, \mathbf{c}_1\right)$, where $\mathbf{y}^*$ is uniform. Let $\mathsf{sk}_j^* \coloneqq \mathbf{x}_j^*$. Programme $\mathsf{H}[\mathsf{id}^*] \coloneqq \mathbf{y}^*$.
- Upon any $\mathsf{H}[\mathsf{id}]$ query from $\mathcal{A}$, if $\mathsf{H}[\mathsf{id}] = \perp$, sample a random vector $\mathbf{y} \leftarrow \mathcal{R}_q^{nt}$, let $H[\mathsf{id}] \coloneqq \mathbf{y}$. Answer with $H[\mathsf{id}]$.
- Pass $\mathsf{mpk}, (\mathsf{msk}_j)_{j\in\mathcal{C}}, \left(\mathsf{sk}_{\ell,j}\right)_{\ell\in[L], j\in T_\ell}, \left(\mathsf{sk}_j^*\right)_{j\in T^*\setminus\{i^*\}}$ to $\mathcal{A}$, and receive (μ_0, μ_1).
- Let $c_1 \coloneqq c_1 + \lfloor q/2 \rfloor \mu_b \bmod q$. Pass $\mathsf{ctxt} \coloneqq (\mathbf{c}_0, c_1)$ to $\mathcal{A}$, and return whatever $\mathcal{A}$ returns.

It is straightforward to verify that mpk, the corrupt master secret keys $(\mathsf{msk}_j)_{j\in\mathcal{C}}$, the answers to secret key queries, and the challenge $\left(\mathsf{sk}_j^*\right)_{j\in T^*\setminus i^*}$ are distributed identically to corresponding values in $\mathsf{Hyb}_{b,0}$ and $\mathsf{Hyb}_{b,1}$. If c_1 from the PT challenger equals $\mathbf{s}^\intercal \mathbf{y}^* + e \bmod q$, then $\mathcal{B}^{\mathcal{A}}$ perfectly simulates c_1 in $\mathsf{Hyb}_{b,0}$. Else, c_1 is uniformly random, hence $\mathcal{B}^{\mathcal{A}}$ perfectly simulates c_1 in $\mathsf{Hyb}_{b,1}$. We conclude that the success probability of $\mathcal{B}^{\mathcal{A}}$ is same as $\mathcal{A}$. □

6.4 Open Problems

We discuss application specific limitations and future directions.

Highly-selective Security. For both of our threshold signature and IBE, we only prove a notion that we dub "highly-selective security" where the adversary has to provide all of its queries upfront. This is because we are unable to programme the random oracle value $\mathsf{H}(\mu)$ (resp. $\mathsf{H}(\mathsf{id})$) before the partial signature (resp. identity key) query for a message μ (resp. identity id) is known, for otherwise we cannot sample appropriate partial preimages for the set T.

Moreover, similar to GPV we do not prove security against the same message (resp. identity) being queried more than once. As in GPV we may address this limitation by de-randomising the signing algorithm when the same message is queried again for the same signing set T. If the same message is queried but for different signing sets T_1, T_2 then we are unable to argue security (and our security experiment forbids it); we note that we are not aware of an attack either.

In the setting of signatures, the obvious fix is to call $\mathsf{H}(\mu, T)$ and to derandomise the signature generation algorithm. We did not adopt this approach above, because it squashes any hope for achieving anonymity. Still, this change would recover the usual notion of adaptive security with static corruptions for threshold signatures.

A similar fix does not apply in IBE since the encrypter does not know which set of parties T will be decrypting the message.

Robustness. Full robustness of the signature scheme is an interesting direction for future work. We could plausibly use techniques of error correction for Shamir secret sharing schemes and subtractive sets [AL21].

Acknowledgements

Martin Albrecht's work is supported by UKRI grant EP/Y02432X/1. Russell W. F. Lai and Ivy K. Y. Woo are supported by the Research Council of Finland projects No. 358951 and 358950 respectively. Oleksandra Lapiha was supported by the EPSRC and the UK Government as part of the Centre for Doctoral Training in Cyber Security for the Everyday at Royal Holloway, University of London (EP/S021817/1).

References

ABB10a. Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (H)IBE in the standard model. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 553–572. Springer, Berlin, Heidelberg, May / June 2010. 1

ABB10b. Shweta Agrawal, Dan Boneh, and Xavier Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 98–115. Springer, Berlin, Heidelberg, August 2010. 4, A, 25

ACL+22. Martin R. Albrecht, Valerio Cini, Russell W. F. Lai, Giulio Malavolta, and Sri Aravinda Krishnan Thyagarajan. Lattice-based SNARKs: Publicly verifiable, preprocessing, and recursively composable - (extended abstract). In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 102–132. Springer, Cham, August 2022. 2.7, 4

ACX21. Thomas Attema, Ronald Cramer, and Chaoping Xing. A note on short invertible ring elements and applications to cyclotomic and trinomials number fields. *Mathematical Cryptology*, 1(1):45–70, Jun. 2021. 9, 5.2

AGHS13. Shweta Agrawal, Craig Gentry, Shai Halevi, and Amit Sahai. Discrete Gaussian leftover hash lemma over infinite domains. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part I*, volume 8269 of *LNCS*, pages 97–116. Springer, Berlin, Heidelberg, December 2013. 27

Ajt96. Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *28th ACM STOC*, pages 99–108. ACM Press, May 1996. 2.1, 3.1

AL21. Martin R. Albrecht and Russell W. F. Lai. Subtractive sets over cyclotomic rings - limits of Schnorr-like arguments over lattices. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part II*, volume 12826 of *LNCS*, pages 519–548, Virtual Event, August 2021. Springer, Cham. 1.2, 9, 2.7, 6.4

AR13. Divesh Aggarwal and Oded Regev. A note on discrete gaussian combinations of lattice vectors. Draft. Available at http://arxiv.org/pdf/1308.2405v1.pdf, 2013. B

ASY22. Shweta Agrawal, Damien Stehlé, and Anshu Yadav. Round-optimal lattice-based threshold signatures, revisited. In Mikolaj Bojanczyk, Emanuela Merelli, and David P. Woodruff, editors, *ICALP 2022*, volume 229 of *LIPIcs*, pages 8:1–8:20. Schloss Dagstuhl, July 2022. 1.2

Ban93. Wojciech Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296:625–635, 1993. 22

BF11. Dan Boneh and David Mandell Freeman. Linearly homomorphic signatures over binary fields and new tools for lattice-based signatures. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, *PKC 2011*, volume 6571 of *LNCS*, pages 1–16. Springer, Berlin, Heidelberg, March 2011. 1, 5, 2.4, 14, 2, 4, 4, 2, 26, 9

BGG+14. Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 533–556. Springer, Berlin, Heidelberg, May 2014. 1, 2.6

BGG+18. Dan Boneh, Rosario Gennaro, Steven Goldfeder, Aayush Jain, Sam Kim, Peter M. R. Rasmussen, and Amit Sahai. Threshold cryptosystems from threshold fully homomorphic encryption. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 565–596. Springer, Cham, August 2018. 1.2, 1.2, 1.2, 2.7

BGI15. Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 337–367. Springer, Berlin, Heidelberg, April 2015. 1.2

BGI16a. Elette Boyle, Niv Gilboa, and Yuval Ishai. Breaking the circuit size barrier for secure computation under DDH. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 509–539. Springer, Berlin, Heidelberg, August 2016. 1.2

BGI16b. Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing: Improvements and extensions. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 1292–1303. ACM Press, October 2016. 1.2

BJRW23. Katharina Boudgoust, Corentin Jeudy, Adeline Roux-Langlois, and Weiqiang Wen. On the hardness of module learning with errors with short distributions. *Journal of Cryptology*, 36(1):1, January 2023. 8

BKL+25. Cecilia Boschini, Darya Kaviani, Russell W. F. Lai, Giulio Malavolta, Akira Takahashi, and Mehdi Tibouchi. Ringtail: Practical two-round threshold signatures from learning with errors. In *2025 IEEE Symposium on Security and Privacy, SP 2025*. IEEE Computer Society, 2025. 1.2, 1.2, 3

BKP13. Rikke Bendlin, Sara Krehbiel, and Chris Peikert. How to share a lattice trapdoor: Threshold protocols for signatures and (H)IBE. In Michael J. Jacobson, Jr., Michael E. Locasto, Payman Mohassel, and Reihaneh Safavi-Naini, editors, *ACNS 13International Conference on Applied Cryptography and Network Security*, volume 7954 of *LNCS*, pages 218–236. Springer, Berlin, Heidelberg, June 2013. 1.2, 1.2

CLM23. Valerio Cini, Russell W. F. Lai, and Giulio Malavolta. Lattice-based succinct arguments from vanishing polynomials - (extended abstract). In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part II*, volume 14082 of *LNCS*, pages 72–105. Springer, Cham, August 2023. 2.7

COS+22. Ilaria Chillotti, Emmanuela Orsini, Peter Scholl, Nigel P. Smart, and Barry van Leeuwen. Scooby: Improved multi-party homomorphic secret sharing based on FHE. In Clemente Galdi and Stanislaw Jarecki, editors, *SCN 22*, volume 13409 of *LNCS*, pages 540–563. Springer, Cham, September 2022. 1.2

DKL+23. Nico Döttling, Dimitris Kolonelos, Russell W. F. Lai, Chuanwei Lin, Giulio Malavolta, and Ahmadreza Rahimi. Efficient laconic cryptography from learning with errors. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part III*, volume 14006 of *LNCS*, pages 417–446. Springer, Cham, April 2023. 1, 2.7

DKM+24. Rafaël Del Pino, Shuichi Katsumata, Mary Maller, Fabrice Mouhartem, Thomas Prest, and Markku-Juhani O. Saarinen. Threshold raccoon: Practical threshold signatures from standard lattice assumptions. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part II*, volume 14652 of *LNCS*, pages 219–248. Springer, Cham, May 2024. 1.2, 1.2, 2.7, 3

DKW21. Pratish Datta, Ilan Komargodski, and Brent Waters. Decentralized multi-authority ABE for DNFs from LWE. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part I*, volume 12696 of *LNCS*, pages 177–209. Springer, Cham, October 2021. 2.7

DLN+21. Julien Devevey, Benoît Libert, Khoa Nguyen, Thomas Peters, and Moti Yung. Non-interactive CCA2-secure threshold cryptosystems: Achieving adaptive security in the standard model without pairings. In Juan Garay, editor, *PKC 2021, Part I*, volume 12710 of *LNCS*, pages 659–690. Springer, Cham, May 2021. 1.2

EKT24. Thomas Espitau, Shuichi Katsumata, and Kaoru Takemure. Two-round threshold signature from algebraic one-more learning with errors. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VII*, volume 14926 of *LNCS*, pages 387–424. Springer, Cham, August 2024. 1.2, 1.2

ENP24. Thomas Espitau, Guilhem Niot, and Thomas Prest. Flood and submerse: Distributed key generation and robust threshold signature from lattices. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VII*, volume 14926 of *LNCS*, pages 425–458. Springer, Cham, August 2024. 2.7

GKPV10. Shafi Goldwasser, Yael Tauman Kalai, Chris Peikert, and Vinod Vaikuntanathan. Robustness of the learning with errors assumption. In Andrew Chi-Chih Yao, editor, *ICS 2010*, pages 230–240. Tsinghua University Press, January 2010. 9

GKS24. Kamil Doruk Gür, Jonathan Katz, and Tjerand Silde. Two-round threshold lattice-based signatures from threshold homomorphic encryption. In Markku-Juhani Saarinen and Daniel Smith-Tone, editors, *Post-Quantum Cryptography - 15th International Workshop, PQCrypto 2024, Part II*, pages 266–300. Springer, Cham, June 2024. 1.2

GMPW20. Nicholas Genise, Daniele Micciancio, Chris Peikert, and Michael Walter. Improved discrete gaussian and subgaussian analysis for lattice cryptography. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part I*, volume 12110 of *LNCS*, pages 623–651. Springer, Cham, May 2020. 1

GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008. 1, 1, 1.1, 2.1, 2.3, 2.6, 5, 5.3, 5.4

GVW15. Sergey Gorbunov, Vinod Vaikuntanathan, and Daniel Wichs. Leveled fully homomorphic signatures from standard lattices. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *47th ACM STOC*, pages 469–477. ACM Press, June 2015. 1

HLL23. Yao-Ching Hsieh, Huijia Lin, and Ji Luo. Attribute-based encryption for circuits of unbounded depth from lattices. In *64th FOCS*, pages 415–434. IEEE Computer Society Press, November 2023. 2.7

HMP06. Shlomo Hoory, Avner Magen, and Toniann Pitassi. Monotone circuits for the majority function. In Josep Díaz, Klaus Jansen, José D. P. Rolim, and Uri Zwick, editors, *Approximation, Randomization, and Combinatorial Optimization*, volume 4110 of *Lecture Notes in Computer Science*, pages 410–425. Springer, 2006. 1.2

KNSW20. Elena Kirshanova, Huyen Nguyen, Damien Stehlé, and Alexandre Wallet. On the smoothing parameter and last minimum of random orthogonal lattices. *DCC*, 88(5):931–950, 2020. 10

KRT24. Shuichi Katsumata, Michael Reichle, and Kaoru Takemure. Adaptively secure 5 round threshold signatures from MLWE/MSIS and DL with rewinding. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VII*, volume 14926 of *LNCS*, pages 459–491. Springer, Cham, August 2024. 2.7, 3

LDK+22. Vadim Lyubashevsky, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Peter Schwabe, Gregor Seiler, Damien Stehlé, and Shi Bai. CRYSTALS-DILITHIUM. Technical report, National Institute of Standards and Technology, 2022. available at https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022. 1.2

LNP22. Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Maxime Plançon. Lattice-based zero-knowledge proofs and applications: Shorter, simpler, and more general. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 71–101. Springer, Cham, August 2022. 1

LPR13. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for ring-LWE cryptography. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 35–54. Springer, Berlin, Heidelberg, May 2013. 3, 17, 12

LPSS14. San Ling, Duong Hieu Phan, Damien Stehlé, and Ron Steinfeld. Hardness of k-LWE and applications in traitor tracing. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 315–334. Springer, Berlin, Heidelberg, August 2014. 1.1, 5, 2.5, 4, 4, 1, 4, 19, 28, 29

MP12. Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Berlin, Heidelberg, April 2012. 1, 1.2, 2.1, 7, 2.3, 2.4, 6, 3.5, 3.5, 13, 23

MR04. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. In *45th FOCS*, pages 372–381. IEEE Computer Society Press, October 2004. 4

PR06. Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In Shai Halevi and Tal Rabin, editors, *TCC 2006*, volume 3876 of *LNCS*, pages 145–166. Springer, Berlin, Heidelberg, March 2006. 7

Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005. 3.1, 24

Sto00. Arne Storjohann. *Algorithms for Matrix Canonical Forms*. PhD thesis, ETH Zurich, 2000. https://cs.uwaterloo.ca/~astorjoh/diss2up.pdf. 3.5, 18

Val84. Leslie G. Valiant. Short monotone formulae for the majority function. *J. Algorithms*, 5(3):363–366, 1984. 1.2

Wee21. Hoeteck Wee. ABE for DFA from LWE against bounded collusions, revisited. In Kobbi Nissim and Brent Waters, editors, *TCC 2021, Part II*, volume 13043 of *LNCS*, pages 288–309. Springer, Cham, November 2021. 1.2, 1.2

Wee22. Hoeteck Wee. Optimal broadcast encryption and CP-ABE from evasive lattice assumptions. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 217–241. Springer, Cham, May / June 2022. 1

WW23. Hoeteck Wee and David J. Wu. Succinct vector, polynomial, and functional commitments from lattices. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part III*, volume 14006 of *LNCS*, pages 385–416. Springer, Cham, April 2023. 1.1, 2.4, 2.4, 2.7, 5.4

A Generalisations of Existing Lemmas to Ring Setting

The goal of this section is to prove Lemma 26 and Theorem 9 which generalise their counterparts over $\mathbb{Z}$ in the literature to the ring setting and to allow Gaussian samples with different parameters. Readers wanting to focus on the $\mathbb{Z}$ setting can safely ignore this section.

More specifically, Lemma 26 states that a Gaussian matrix $\mathbf{U}$, whose columns are sampled independently from discrete Gaussian distributions with potentially different parameters, is likely to have linearly independent columns over each factor field of $\mathcal{R}_q$. Theorem 9 states that sampling a uniform $\mathbf{A}$ then sampling a Gaussian $\mathbf{U}$ subject to $\mathbf{A} \cdot \mathbf{U} = \mathbf{V} \bmod q$ is statistically close to sampling in reverse order.

We first recall some classic lemmas from the literature.

Lemma 22 ([Ban93, Lemma 1.5]). *For any n-dimensional lattice Λ, $\mathbf{c} \in \mathbb{R}^n$, $r > 1/\sqrt{2\,\pi}$,*

$$\rho((\Lambda + \mathbf{c}) \setminus r\sqrt{n}\,\mathcal{B}) < 2\,C_r^n \cdot \rho(\Lambda)$$

where $r\sqrt{n}\,\mathcal{B} \subseteq \mathbb{R}^n$ denotes the set of all vectors in $\mathbb{R}^n$ with norm at most $r\sqrt{n}$ and $C_r = r\sqrt{2\pi e} \cdot e^{-\pi \cdot r^2}$.

Lemma 23 ([MP12, Lemma 2.5]). *Let $\Lambda \subseteq \mathbb{R}^n$ be a lattice, $\mathbf{\Sigma} \geq 0$, and $\mathbf{c} \in \mathbb{R}^n$. It holds that $\rho_{\sqrt{\mathbf{\Sigma}}}(\Lambda + \mathbf{c}) \leq \rho_{\sqrt{\mathbf{\Sigma}}}(\Lambda)$. Furthermore, if $\sqrt{\mathbf{\Sigma}} \geq \eta_\varepsilon(\Lambda)$ for some $\varepsilon > 0$ and $\mathbf{c} \in \mathsf{span}(\Lambda)$, then $\rho_{\sqrt{\mathbf{\Sigma}}}(\Lambda + \mathbf{c}) \geq \frac{1-\varepsilon}{1+\varepsilon} \cdot \rho_{\sqrt{\mathbf{\Sigma}}}(\Lambda)$.*

The following is a direct corollary of Lemma 23.

Corollary 2 (Generalisation of [BF11, Lemma 4.4]). *Let $\mathbf{A} \in \mathcal{R}_q^{n \times m}$ be primitive and $\sqrt{\mathbf{\Sigma}} \geq \eta_\varepsilon(\Lambda_q^\perp(\mathbf{A}))$ for some $\varepsilon > 0$. For any $\mathbf{v} \in \mathcal{R}_q^n$, it holds that*

$$\frac{1-\varepsilon}{1+\varepsilon} \cdot \rho_{\sqrt{\mathbf{\Sigma}}}(\Lambda_q^\perp(\mathbf{A})) \leq \rho_{\sqrt{\mathbf{\Sigma}}}(\Lambda_q^{\mathbf{v}}(\mathbf{A})) \leq \rho_{\sqrt{\mathbf{\Sigma}}}(\Lambda_q^\perp(\mathbf{A})).$$

Furthermore, if $\varepsilon = \mathsf{negl}(\lambda)$, then

$$\rho_{\sqrt{\mathbf{\Sigma}}}(\mathcal{R}^m) = \left|\mathcal{R}_q^n\right| \cdot \rho_{\sqrt{\mathbf{\Sigma}}}(\Lambda_q^\perp(\mathbf{A})) \cdot (1 - \mathsf{negl}(\lambda)).$$

The following is immediate from Corollary 2 and its proof omitted.

Corollary 3. *Let $\mathbf{A} \in \mathcal{R}_q^{n \times m}$ be primitive and $\sqrt{\mathbf{\Sigma}} \geq \eta_\varepsilon(\Lambda_q^\perp(\mathbf{A}))$ for some $\varepsilon > 0$. Let $\sigma^* = s_{\min}(\sqrt{\mathbf{\Sigma}})$. If $\varepsilon = \mathsf{negl}(\lambda)$, then*

$$\rho_{\sigma^*}(\mathcal{R}^m) \leq \left|\mathcal{R}_q^n\right| \cdot \rho_{\sqrt{\mathbf{\Sigma}}}(\Lambda_q^\perp(\mathbf{A})) \cdot (1 - \mathsf{negl}(\lambda)).$$

Lemma 24 ([Reg05, Claim 3.8]). *For any n-dimensional lattice Λ, $\varepsilon > 0$, $\sigma \geq \eta_\varepsilon(\Lambda)$, and $\mathbf{c} \in \mathbb{R}^n$, it holds that $\rho_\sigma(\Lambda + \mathbf{c}) \in [1 \pm \varepsilon] \cdot \sigma^n \cdot \det(\Lambda)^{-1}$. Furthermore, when $\mathbf{c} = \mathbf{0}$, $\rho_\sigma(\Lambda) \geq \sigma^n \cdot \det(\Lambda)^{-1}$.*

The following is a direct corollary of Lemma 24.

Corollary 4 (Generalisation of [ABB10b, Full version, Lemma 21]). *For any $i \geq 0$, $\varepsilon > 0$, $\sigma \geq \eta_\varepsilon(\mathcal{R}^i)$, it holds that*

$$\sigma^{\varphi i} \cdot \det(\mathcal{R}^i)^{-1} \leq \rho_\sigma(\mathcal{R}^i) \leq (1+\varepsilon) \cdot \sigma^{\varphi i} \cdot \det(\mathcal{R}^i)^{-1}.$$

We also generalise a lemma from [ABB10b] which upper-bounds the Gaussian weights of q-ary lattices. We generalise it to cover $\mathfrak{q}$-ary lattices for ideals $\mathfrak{q} \subset \mathcal{R}$.

Lemma 25 (Generalisation of [ABB10b, Full version, Lemma 31]). *Let m, N be positive integers, k be a non-negative integer, $\mathbf{\Sigma} \in \mathbb{R}^{\varphi m \times \varphi m}$ be positive semi-definite, $\sigma = s_{\max}(\sqrt{\mathbf{\Sigma}})$, and $\mathfrak{q} \subset \mathcal{R}$ be an ideal satisfying the following:*

$$\mathcal{N}(\mathfrak{q}) = N, \qquad k \leq m, \qquad \varphi m \geq \Omega(\lambda), \qquad \sigma \cdot \sqrt{m} \leq N^{1/\varphi}.$$

For any $\mathbf{U} \in \mathcal{R}^{m \times k}$, it holds that

$$\rho_{\sqrt{\mathbf{\Sigma}}}(\Lambda_{\mathfrak{q}}(\mathbf{U})) \leq \rho_\sigma(\mathcal{R}^k)/(1 - \mathsf{negl}(\lambda))$$

where if $k = 0$, i.e. $\mathbf{U}$ is the empty matrix then we take $\Lambda_{\mathfrak{q}}(\mathbf{U}) = \mathfrak{q}^m$.

Proof. Below, for $\ell \geq 0$, write $\mathcal{B}_\ell \subseteq \mathbb{H}^\ell$ for the set of vectors in $\mathbb{H}^\ell$ with norm at most 1.

Without loss of generality, we can assume that there exists a k-by-k submatrix $\mathbf{W}$ of $\mathbf{U}$ so that $\mathbf{W}$ is either the empty matrix (i.e. $k = 0$) or invertible over $\mathcal{R}_\mathfrak{q} = \mathcal{R}/\mathfrak{q}$, for otherwise there exists $\mathbf{U}' \in \mathcal{R}^{m\times k}$ with such a property and $\rho_\sigma(\Lambda_\mathfrak{q}(\mathbf{U})) \leq \rho_\sigma(\Lambda_\mathfrak{q}(\mathbf{U}'))$. The case $k = 0$ is clear. For $k > 0$, consider prime factorisation $\mathfrak{q} = \prod_j \mathfrak{q}_j$ and the Chinese remainder representation $\mathbf{U}_j \coloneqq \mathbf{U} \bmod \mathfrak{q}_j \in \mathcal{R}_{\mathfrak{q}_j}^{m\times k}$ of $\mathbf{U}$. For each j, let $\mathbf{U}'_j \in \mathcal{R}_{\mathfrak{q}_j}^{m\times k}$ be a matrix obtained from $\mathbf{U}_j$ by picking the maximal set of $\mathcal{R}_{\mathfrak{q}_j}$-linearly independent columns and filling the rest with arbitrary columns which are $\mathcal{R}_{\mathfrak{q}_j}$-linearly independent of the previous ones. Let $\mathbf{U}' \in \mathcal{R}_\mathfrak{q}$ be the unique matrix (modulo $\mathfrak{q}$) satisfying $\mathbf{U}' = \mathbf{U}'_j \bmod \mathfrak{q}$. It is clear that $\Lambda_\mathfrak{q}(\mathbf{U}) \subseteq \Lambda_\mathfrak{q}(\mathbf{U}')$ and that every k-by-k submatrix of $\mathbf{U}'$ is invertible modulo $\mathfrak{q}$.

Let $\mathbf{P} \in \{0,1\}^{k\times m}$ be a partial permutation matrix, i.e. each row and column contains at most a single 1, such that $\mathbf{P}\cdot\mathbf{U} = \mathbf{W}$ is a submatrix of $\mathbf{U}$ which is either the empty matrix or is invertible over $\mathcal{R}_\mathfrak{q}$. We observe the following two properties:

1. For all $\mathbf{x} \in \Lambda_\mathfrak{q}(\mathbf{U})$, since $\|\mathbf{P}\cdot\mathbf{x}\| \leq \|\mathbf{x}\|$, we have $\rho_{\sqrt{\mathbf{\Sigma}}}(\mathbf{x}) \leq \rho_\sigma(\mathbf{P}\cdot\mathbf{x})$.[24]
2. The map $\mathbf{x} \mapsto \mathbf{P}\cdot\mathbf{x} \bmod \mathfrak{q}$ is injective over $\mathcal{C} \coloneqq \Lambda_\mathfrak{q}(\mathbf{U}) \cap \mathcal{R}_\mathfrak{q}^m$ where we identify elements in $\mathcal{R}_\mathfrak{q}^m$ by their shortest representatives. Indeed, let $\mathbf{x} = \mathbf{U}\cdot\mathbf{d} \bmod \mathfrak{q}$ and $\mathbf{x}' = \mathbf{U}\cdot\mathbf{d}' \bmod \mathfrak{q}$ for some $\mathbf{d}, \mathbf{d}' \in \mathcal{R}_\mathfrak{q}^k$. If $\mathbf{P}\cdot\mathbf{x} = \mathbf{P}\cdot\mathbf{x}' \bmod \mathfrak{q}$, then $\mathbf{W}\cdot(\mathbf{d}-\mathbf{d}') = \mathbf{0} \bmod \mathfrak{q}$, which implies $\mathbf{d} = \mathbf{d}' \bmod \mathfrak{q}$ and hence $\mathbf{x} = \mathbf{x}' \bmod \mathfrak{q}$.

Using the two properties above, we obtain

$$\rho_{\sqrt{\mathbf{\Sigma}}}(\mathcal{C}) = \sum_{\mathbf{x}\in\mathcal{C}} \rho_{\sqrt{\mathbf{\Sigma}}}(\mathbf{x}) \leq \sum_{\mathbf{x}\in\mathcal{C}} \rho_\sigma(\mathbf{P}\cdot\mathbf{x}) \leq \sum_{\mathbf{y}\in\mathcal{R}^k} \rho_\sigma(\mathbf{y}) = \rho_\sigma(\mathcal{R}^k) \tag{11}$$

where the first and second inequalities follow from the first and second properties respectively.

Next, let $r \coloneqq N^{1/\varphi}/(2\sigma\sqrt{m})$. We show that $\Lambda_\mathfrak{q}(\mathbf{U}) \cap r\sqrt{\varphi m}\cdot\sigma\mathcal{B}_{\varphi m} \subseteq \mathcal{C}$ and hence

$$\rho_{\sqrt{\mathbf{\Sigma}}}(\Lambda_\mathfrak{q}(\mathbf{U}) \cap r\sqrt{\varphi m}\cdot\sigma\mathcal{B}_{\varphi m}) \leq \rho_{\sqrt{\mathbf{\Sigma}}}(\mathcal{C}). \tag{12}$$

Suppose towards a contradiction that $\Lambda_\mathfrak{q}(\mathbf{U}) \cap r\sqrt{\varphi m}\cdot\sigma\mathcal{B}_{\varphi m} \not\subseteq \mathcal{C}$. Then there exists distinct $\mathbf{a}, \mathbf{b} \in r\sqrt{\varphi m}\cdot\sigma\mathcal{B}_{\varphi m}$ such that $\mathbf{a} = \mathbf{b} \bmod \mathfrak{q}$ or equivalently $\mathbf{0} \neq \mathbf{a}-\mathbf{b} \in \mathfrak{q}^m$. Since $\mathbf{a}$ and $\mathbf{b}$ are contained in the ball of radius $r\sigma\sqrt{\varphi m}$, so does $\mathbf{a}-\mathbf{b}$, and thus $\|\mathbf{a}-\mathbf{b}\| \leq 2r\sigma\sqrt{\varphi m} = \sqrt{\varphi}\cdot N^{1/\varphi}$. By the inequality of arithmetic and geometric mean, we have

$$\lambda_1(\mathfrak{q}) \geq \sqrt{\varphi}\cdot\mathcal{N}(\mathfrak{q})^{1/\varphi} = \sqrt{\varphi}\cdot N^{1/\varphi}.$$

Therefore, $\|\mathbf{a}-\mathbf{b}\| \leq \sqrt{\varphi}\cdot N^{1/\varphi} \leq \lambda_1(\mathfrak{q}) = \lambda_1(\mathfrak{q}^m)$, contradicting $\mathbf{a}-\mathbf{b}$ being a non-zero vector in $\mathfrak{q}^m$.

Finally, using $r = N^{1/\varphi}/(2\sigma\sqrt{m}) > 1/\sqrt{2\pi}$ from above, let

$$\begin{aligned} C_r &= r\cdot\sqrt{2\pi e}\cdot\exp(-\pi\cdot r^2) \\ &= \sqrt{2\pi e}\cdot\frac{N^{1/\varphi}}{2\sigma\sqrt{m}}\cdot\exp\left(-\pi\cdot\left(\frac{N^{1/\varphi}}{2\sigma\sqrt{m}}\right)^2\right) \\ &\leq \frac{\sqrt{2\pi e}}{2}\cdot\exp(-\pi/4) \approx 0.942, \end{aligned}$$

where the last inequality holds since $\frac{N^{1/\varphi}}{\sigma\sqrt{m}} > 1$ by the assumed parameters, and the function $\sqrt{2\pi e}\cdot x\cdot\exp\left(-\pi\cdot x^2\right)$ is decreasing for $x \geq \sqrt{1/(2\pi)}$. Applying Lemma 22, we have that

$$\frac{\rho_{\sqrt{\mathbf{\Sigma}}}(\Lambda_\mathfrak{q}(\mathbf{U}) \cap r\sqrt{\varphi m}\cdot\sigma\mathcal{B}_m)}{\rho_{\sqrt{\mathbf{\Sigma}}}(\Lambda_\mathfrak{q}(\mathbf{U}))} \geq \frac{\rho_{\sqrt{\mathbf{\Sigma}}}(\Lambda_\mathfrak{q}(\mathbf{U}) \cap r\sqrt{\varphi m}\cdot\sqrt{\mathbf{\Sigma}}\mathcal{B}_m)}{\rho_{\sqrt{\mathbf{\Sigma}}}(\Lambda_\mathfrak{q}(\mathbf{U}))}$$

[24] Write $\mathbf{\Sigma} = \mathbf{U}^\mathsf{T}\mathbf{D}\mathbf{U}$ as its singular decomposition, then $\rho_{\sqrt{\mathbf{\Sigma}}}(\mathbf{x}) = \exp(-\pi\mathbf{x}^\mathsf{T}\mathbf{\Sigma}^{-1}\mathbf{x}) \leq \exp(-\pi\mathbf{x}^\mathsf{T}\mathbf{U}^{-1}(\sigma^{-2}\mathbf{I})\mathbf{U}^{-\mathsf{T}}\mathbf{x}) = \exp(-\pi\|\mathbf{x}\|^2/\sigma^2) \leq \exp(-\pi\|\mathbf{P}\mathbf{x}\|^2/\sigma^2) = \rho_\sigma(\mathbf{P}\mathbf{x})$, where we recall $\sigma = s_{\max}(\sqrt{\mathbf{\Sigma}})$.

$$
\begin{aligned}
&= \frac{\rho((\sqrt{\mathbf{\Sigma}})^{-1}\Lambda_{\mathfrak{q}}(\mathbf{U}) \cap r\sqrt{\varphi m}\cdot\mathcal{B}_m)}{\rho((\sqrt{\mathbf{\Sigma}})^{-1}\Lambda_{\mathfrak{q}}(\mathbf{U}))} \\
&\geq 1 - 2C_r^{\varphi m} \quad \geq 1 - \mathsf{negl}(\lambda).
\end{aligned}
$$

Chaining together with Eqs. (11) and (12), we arrive at

$$\rho_{\sqrt{\mathbf{\Sigma}}}(\Lambda_{\mathfrak{q}}(\mathbf{U})) \leq \rho_\sigma(\mathcal{R}^k)/(1-\mathsf{negl}(\lambda)). \qquad \square$$

Given the above, we are ready to prove Lemma 26 and Theorem 9.

Lemma 26 (Generalisation of [BF11, Lemma 4.5]). *Let g, k, m, n, q be positive integers, $a \geq 1$ be a real number, $\mathbf{A} \in \mathcal{R}_q^{n\times m}$ be primitive, and $\{\mathbf{\Sigma}_i\}_{i\in[k]}$ be positive semi-definite matrices in $\mathbb{R}^{\varphi m\times\varphi m}$ satisfying the following constraints:*

- *$k, n \leq m$,*
- *$\varphi(m-k) \geq \Omega(\lambda)$, and*
- *q is an unramified prime that factors as $\langle q\rangle = \prod_{j=1}^g \mathfrak{q}_j$ with norm $\mathcal{N}(\mathfrak{q}_j) = q^{\varphi/g}$ in $\mathcal{R}$, and*
- *there exists $\varepsilon \leq \mathsf{negl}(\lambda)$ so that for all $i \in [k]$ it holds*

$$\max\left\{\eta_\varepsilon(\Lambda_q^\perp(\mathbf{A})), 2\sqrt{\varphi}\cdot(a^k q^n)^{1/(m-k)}\right\} \leq s_{\min}\left(\sqrt{\mathbf{\Sigma}_i}\right), \quad s_{\max}\left(\sqrt{\mathbf{\Sigma}_i}\right) \leq \min\left\{q^{1/g}/\sqrt{m}, a\cdot s_{\min}\left(\sqrt{\mathbf{\Sigma}_i}\right)\right\}.$$

Furthermore, let $\{\mathbf{v}_i\}_{i\in[k]} \in \mathcal{R}_q^n$ be arbitrary images. For $i \in [k]$, let $\mathbf{u}_i \leftarrow \mathcal{D}_{\Lambda_q^{\mathbf{v}_i}(\mathbf{A}),\sqrt{\mathbf{\Sigma}_i}}$. With overwhelming probability in λ, the vectors $\{\mathbf{u}_i\}_{i\in[k]}$ are $\mathcal{R}_{\mathfrak{q}_j}$-linearly independent for all $j \in [g]$.

Proof. For $i \in \{1,\ldots,k\}$, let $\mathbf{U}_i \in \mathcal{R}^{m\times i}$ be the matrix with columns given by $(\mathbf{u}_j)_{j\in[i]}$, and define $\mathbf{U}_0$ to be the empty matrix. For each $j \in [g]$, define $\Lambda_{\mathfrak{q}_j}(\mathbf{U}_0) \coloneqq \mathfrak{q}_j^m$. By a union bound and Corollary 2, the probability that $\{\mathbf{u}_i\}_{i\in[k]}$ is not $\mathcal{R}_{\mathfrak{q}_j}$-linearly independent for all $j \in [g]$ is at most

$$
\begin{aligned}
\sum_{i\in[k]}\sum_{j\in[g]} \Pr\left[\mathbf{u}_i \in \Lambda_{\mathfrak{q}_j}(\mathbf{U}_i)\right] &\leq \sum_{i\in[k]}\sum_{j\in[g]} \frac{\rho_{\sqrt{\mathbf{\Sigma}_i}}(\Lambda_{\mathfrak{q}_j}(\mathbf{U}_i))}{\rho_{\sqrt{\mathbf{\Sigma}_i}}(\Lambda_q^{\mathbf{v}_i}(\mathbf{A}))} \\
&\leq (1+\mathsf{negl}(\lambda))\cdot\sum_{i\in[k]}\sum_{j\in[g]} \frac{\rho_{\sqrt{\mathbf{\Sigma}_i}}(\Lambda_{\mathfrak{q}_j}(\mathbf{U}_i))}{\rho_{\sqrt{\mathbf{\Sigma}_i}}(\Lambda_q^\perp(\mathbf{A}))}.
\end{aligned} \tag{13}
$$

Below, denote by $\sigma_{i,\max} \coloneqq s_{\max}(\sqrt{\mathbf{\Sigma}_i})$ and $\sigma_{i,\min} \coloneqq s_{\min}(\sqrt{\mathbf{\Sigma}_i})$ the maximum and minimum singular values of $\sqrt{\mathbf{\Sigma}_i}$ respectively. We recall the following properties proven in separate lemmas:

1. For all $j \in [g]$, the norm of $\mathfrak{q}_j$ is given by $\mathcal{N}(\mathfrak{q}_j) = q^{\varphi/g}$. Since $\sigma_{i,\max} \leq q^{1/g}/\sqrt{m} = \mathcal{N}(\mathfrak{q}_j)^{1/\varphi}/\sqrt{m}$, by Lemma 25, for $i \in \{0,\ldots,k\}$ and $j \in [g]$,

$$\rho_{\sqrt{\mathbf{\Sigma}_i}}(\Lambda_{\mathfrak{q}_j}(\mathbf{U}_i)) \leq \rho_{\sigma_{i,\max}}(\mathcal{R}^i)/(1-\mathsf{negl}(\lambda)).$$

2. For all $i \in [k]$, since $\sqrt{\mathbf{\Sigma}_i} \geq \eta_\varepsilon(\Lambda_q^\perp(\mathbf{A}))$, by Corollary 3, we have

$$\frac{1}{\rho_{\sqrt{\mathbf{\Sigma}_i}}(\Lambda_q^\perp(\mathbf{A}))} \leq \left|\mathcal{R}_q^n\right|\cdot\frac{1-\mathsf{negl}(\lambda)}{\rho_{\sigma_{i,\min}}(\mathcal{R}^m)}.$$

3. For $i \in [k]$, since $\sigma_{i,\min} \geq \eta_\varepsilon(\Lambda_q^\perp(\mathbf{A}))$, we have $\sigma_{i,\min} \geq \eta_\varepsilon(\mathcal{R}^{i'})$ for all $0 \leq i' \leq m$. By Corollary 4,

$$\sigma_{i,\min}^{\varphi i'}\cdot\det(\mathcal{R}^{i'})^{-1} \leq \rho_{\sigma_{i,\min}}(\mathcal{R}^{i'}) \leq (1+\mathsf{negl}(\lambda))\cdot\sigma_{i,\min}^{\varphi i'}\cdot\det(\mathcal{R}^{i'})^{-1},$$

and analogously for $\sigma_{i,\max}$.

Combining the first and second properties and $g \leq \varphi$, the quantity in Eq. (13) is upper bounded by

$$\left|\mathcal{R}_q^n\right| \cdot \varphi \cdot \sum_{i\in[k]} \frac{\rho_{\sigma_{i,\max}}(\mathcal{R}^i)}{\rho_{\sigma_{i,\min}}(\mathcal{R}^m)} \cdot (1 + \mathsf{negl}(\lambda)). \tag{14}$$

Define the shorthands $\sigma \coloneqq \min\{\sigma_{i,\min}\}_{i\in[k]}$ and $r = \frac{\sigma^\varphi}{\det(\mathcal{R})}$. By the third property, the quantity in Eq. (14) is at most

$$\begin{aligned}
&\left|\mathcal{R}_q^n\right| \cdot \varphi \cdot \sum_{i\in[k]} \frac{\sigma_{i,\max}^{\varphi i} \cdot \det(\mathcal{R})^{-i}}{\sigma_{i,\min}^{\varphi m} \cdot \det(\mathcal{R})^{-m}} \cdot (1 + \mathsf{negl}(\lambda)) \\
\leq\ & q^{\varphi n} \cdot \varphi \cdot \sum_{i\in[k]} a^{\varphi i} \left(\frac{\det(\mathcal{R})}{\sigma_{i,\min}^{\varphi}}\right)^{m-i} \cdot (1 + \mathsf{negl}(\lambda)) \\
\leq\ & q^{\varphi n} \cdot \varphi \cdot \sum_{i\in[k]} a^{\varphi i} \cdot r^{i-m} \cdot (1 + \mathsf{negl}(\lambda)) \\
=\ & q^{\varphi n} \cdot \varphi \cdot \sum_{i\in[k]} r^{-m} \cdot (a^\varphi r)^i \cdot (1 + \mathsf{negl}(\lambda)) \\
=\ & q^{\varphi n} \cdot \varphi \cdot r^{-m} \cdot \frac{(a^\varphi r)^k - 1}{(a^\varphi r) - 1} \cdot (1 + \mathsf{negl}(\lambda))
\end{aligned}$$

where the first inequality is due to $\sigma_{i,\max} \leq a \cdot \sigma_{i,\min}$, and the second due to $\sigma \leq \sigma_{i,\min}$, for all $i \in [k]$.

Recall that the discriminant $\Delta_\mathcal{K}$ of $\mathcal{K}$ satisfies $\Delta_\mathcal{K} = \det(\mathcal{R})^2 \leq \varphi^\varphi$. Since $\sigma \geq 2\sqrt{\varphi}(a^k q^n)^{1/(m-k)}$, we have

$$a^\varphi r = a^\varphi \cdot \frac{\sigma^\varphi}{\det(\mathcal{R})} \geq a^\varphi \cdot \frac{\sigma^\varphi}{\varphi^{\varphi/2}} \geq a^\varphi \cdot \frac{(2\sqrt{\varphi}(a^k q^n)^{1/(m-k)})^\varphi}{\varphi^{\varphi/2}} = 2^\varphi a^\varphi (a^k q^n)^{\varphi/(m-k)} \geq 2.$$

We can therefore further upper-bound the above quantity by

$$\begin{aligned}
& q^{\varphi n} \cdot \varphi \cdot a^{\varphi k} \cdot r^{k-m} \cdot (1 + \mathsf{negl}(\lambda)) \\
=\ & q^{\varphi n} \cdot \varphi \cdot a^{\varphi k} \cdot \left(\frac{\det(\mathcal{R})}{\sigma^\varphi}\right)^{m-k} \cdot (1 + \mathsf{negl}(\lambda)) \\
\leq\ & q^{\varphi n} \cdot \varphi \cdot a^{\varphi k} \cdot \left(\frac{\varphi^{\varphi/2}}{\sigma^\varphi}\right)^{m-k} \cdot (1 + \mathsf{negl}(\lambda)) \\
=\ & \varphi \cdot \left(\sqrt{\varphi}(a^k q^n)^{1/(m-k)}/\sigma\right)^{\varphi(m-k)} \cdot (1 + \mathsf{negl}(\lambda)) \quad \leq\ \mathsf{negl}(\lambda)
\end{aligned}$$

where the first inequality is due to $\det(\mathcal{R}) \leq \varphi^{\varphi/2}$, and the second is due to $\sigma \geq 2\sqrt{\varphi} \cdot (a^k q^n)^{1/(m-k)}$ and $2^{-\varphi(m-k)} \leq \mathsf{negl}(\lambda)$. □

Theorem 9 (Generalisation of [BF11, Theorem 4.3]). *Let g, k, m, n, q be positive integers, $a \geq 1$ be a real number, and $\{\mathbf{\Sigma}_i\}_{i\in[k]}$ be positive semi-definite matrices in $\mathbb{R}^{\varphi m \times \varphi m}$ satisfying the following constraints:*

- *$k, n \leq m$,*
- *$\varphi(m-k) \geq \Omega(\lambda)$,*
- *q be an unramified prime that factors as $\langle q\rangle = \prod_{j=1}^g \mathfrak{q}_j$ with norm $\mathcal{N}(\mathfrak{q}_j) = q^{\varphi/g}$ in $\mathcal{R}$,*
- *$ng/q^{\varphi(m-n+1)/g} \leq \mathsf{negl}(\lambda)$, and*
- *there exists $\varepsilon \leq \mathsf{negl}(\lambda)$ so that for all $i \in [k]$ it holds*

$$\max\left\{\eta_A, 2\sqrt{\varphi} \cdot (a^k q^n)^{1/(m-k)}\right\} \leq s_{\min}\left(\sqrt{\mathbf{\Sigma}_i}\right), \quad s_{\max}\left(\sqrt{\mathbf{\Sigma}_i}\right) \leq \min\left\{q^{1/g}/\sqrt{m}, a \cdot s_{\min}\left(\sqrt{\mathbf{\Sigma}_i}\right)\right\},$$

where $\eta_A \geq 8\,\varphi\sqrt{m} \cdot q^{n/m + 2/(\varphi \cdot m)}$.

For any $\mathbf{V} \in \mathcal{R}_q^{n\times k}$, the following distributions are statistically close in λ:

$$\mathcal{D}_0 \coloneqq \left\{ (\mathbf{A}, \mathbf{U}) \;\middle|\; \begin{array}{l} \mathbf{A} \leftarrow \mathcal{R}_q^{n\times m} \\ \mathbf{u}_i \leftarrow \mathcal{D}_{\Lambda_q^{\mathbf{v}_i}(\mathbf{A}), \sqrt{\mathbf{\Sigma}_i}} \quad \forall i \in [k] \end{array} \right\}, \; \mathcal{D}_1 \coloneqq \left\{ (\mathbf{A}, \mathbf{U}) \;\middle|\; \begin{array}{l} \mathbf{u}_i \leftarrow \mathcal{D}_{\mathcal{R}^m, \sqrt{\mathbf{\Sigma}_i}} \quad \forall i \in [k] \\ \mathbf{A} \leftarrow \mathcal{R}_q^{n\times m} : \mathbf{A}\cdot\mathbf{U} \equiv \mathbf{V} \end{array} \right\}.$$

Proof. We calculate the statistical distance directly. First, consider $\mathcal{D}_0$. Clearly, for any $\mathbf{A} \in \mathcal{R}_q^{n\times m}$ and any $\mathbf{U} \in \mathcal{R}^{m\times k}$ satisfying $\mathbf{A}\cdot\mathbf{U} \equiv \mathbf{V} \bmod q$,

$$\begin{aligned} \Pr[\mathcal{D}_0 = (\mathbf{A}, \mathbf{U})] &= \frac{1}{q^{\varphi nm}} \prod_{i\in[k]} \frac{\rho_{\sqrt{\mathbf{\Sigma}_i}}(\mathbf{u}_i)}{\rho_{\sqrt{\mathbf{\Sigma}_i}}(\Lambda_q^{\mathbf{v}_i}(\mathbf{A}))} \\ &\in [1 \pm \mathsf{negl}(\lambda)] \cdot \frac{1}{q^{\varphi nm}} \prod_{i\in[k]} \frac{\rho_{\sqrt{\mathbf{\Sigma}_i}}(\mathbf{u}_i)}{\rho_{\sqrt{\mathbf{\Sigma}_i}}(\Lambda_q^{\perp}(\mathbf{A}))} \\ &= [1 \pm \mathsf{negl}(\lambda)] \cdot \frac{q^{\varphi nk}}{q^{\varphi nm}} \prod_{i\in[k]} \frac{\rho_{\sqrt{\mathbf{\Sigma}_i}}(\mathbf{u}_i)}{\rho_{\sqrt{\mathbf{\Sigma}_i}}(\mathcal{R}^m)} \end{aligned}$$

where the inclusion is due to Corollary 2.

Next, consider $\mathcal{D}_1$. For any $\mathbf{U} \in \mathcal{R}^{m\times k}$ having $\mathcal{R}_{\mathfrak{q}_j}$-linearly independent columns for all $j \in [g]$, there are exactly $q^{\varphi n(m-k)}$ possible choices of $\mathbf{A}$ which satisfies $\mathbf{A}\cdot\mathbf{U} \equiv \mathbf{V} \bmod q$. Therefore, for each of these choices of $(\mathbf{A}, \mathbf{U})$,

$$\Pr[\mathcal{D}_1 = (\mathbf{A}, \mathbf{U})] = \frac{q^{\varphi nk}}{q^{\varphi nm}} \prod_{i\in[k]} \frac{\rho_{\sqrt{\mathbf{\Sigma}_i}}(\mathbf{u}_i)}{\rho_{\sqrt{\mathbf{\Sigma}_i}}(\mathcal{R}^m)}.$$

Let $\mathcal{P} \subseteq \mathcal{R}_q^{n\times m}$ denote the set of all primitive matrices, and let $\mathcal{T}_{\mathbf{A},\mathbf{V}} \subseteq \mathcal{R}^{m\times k}$ denote those matrices $\mathbf{U}$ with k $\mathcal{R}_q$-linearly independent columns and satisfying $\mathbf{A}\cdot\mathbf{U} = \mathbf{V} \bmod q$. The statistical distance between $\mathcal{D}_0$ and $\mathcal{D}_1$, when restricted to $\mathbf{A} \in \mathcal{P}$ and $\mathbf{U} \in \mathcal{T}_{\mathbf{A},\mathbf{V}}$, denoted by Δ, satisfies the following:

$$\begin{aligned} \Delta &\coloneqq \frac{1}{2} \sum_{\mathbf{A}\in\mathcal{P}} \sum_{\mathbf{U}\in\mathcal{T}_{\mathbf{A},\mathbf{V}}} |\Pr[\mathcal{D}_0 = (\mathbf{A}, \mathbf{U})] - \Pr[\mathcal{D}_1 = (\mathbf{A}, \mathbf{U})]| \\ &\le \frac{1}{2} \sum_{\mathbf{A}\in\mathcal{P}} \sum_{\mathbf{U}\in\mathcal{T}_{\mathbf{A},\mathbf{V}}} \frac{q^{\varphi nk}}{q^{\varphi nm}} \prod_{i\in[k]} \frac{\rho_{\sqrt{\mathbf{\Sigma}_i}}(\mathbf{u}_i)}{\rho_{\sqrt{\mathbf{\Sigma}_i}}(\mathcal{R}^m)} \cdot \mathsf{negl}(\lambda) \\ &\le \frac{1}{2} \sum_{\mathbf{A}\in\mathcal{P}} \frac{q^{\varphi nk}}{q^{\varphi nm}} \prod_{i\in[k]} \frac{\rho_{\sqrt{\mathbf{\Sigma}_i}}(\Lambda_q^{\mathbf{v}_i}(\mathbf{A}))}{\rho_{\sqrt{\mathbf{\Sigma}_i}}(\mathcal{R}^m)} \cdot \mathsf{negl}(\lambda) \\ &\le \frac{1}{2} \sum_{\mathbf{A}\in\mathcal{P}} \frac{q^{\varphi nk}}{q^{\varphi nm}} \prod_{i\in[k]} \frac{\rho_{\sqrt{\mathbf{\Sigma}_i}}(\Lambda_q^{\perp}(\mathbf{A}))}{\rho_{\sqrt{\mathbf{\Sigma}_i}}(\mathcal{R}^m)} \cdot \mathsf{negl}(\lambda) \\ &\le \frac{1}{2} \sum_{\mathbf{A}\in\mathcal{P}} \frac{q^{\varphi nk}}{q^{\varphi nm}} \prod_{i\in[k]} (|\mathcal{R}_q^n|(1 - \mathsf{negl}(\lambda)))^{-1} \cdot \mathsf{negl}(\lambda) \\ &= \frac{1}{2} \sum_{\mathbf{A}\in\mathcal{P}} \frac{1}{q^{\varphi nm}} \cdot (1 - \mathsf{negl}(\lambda))^{-k} \mathsf{negl}(\lambda) \quad \le \mathsf{negl}(\lambda). \end{aligned}$$

In the above, the first inequality is clear, the second is due to $\prod_{i\in[k]} \rho_{\sqrt{\mathbf{\Sigma}_i}}(\Lambda_q^{\mathbf{v}_i}(\mathbf{A}))$ containing all terms in $\sum_{\mathbf{U}\in\mathcal{T}_{\mathbf{A},\mathbf{V}}} \prod_{i\in[k]} \rho_{\sqrt{\mathbf{\Sigma}_i}}(\mathbf{u}_i)$ plus some more, the third and fourth are due to Corollary 2, and the final one is due to $|\mathcal{P}| \le q^{\varphi nm}$.

By Lemma 8, $\mathbf{A} \leftarrow \mathcal{R}_q^{n\times m}$ is primitive with probability at least

$$1 - \frac{ng}{q^{\varphi\cdot(m-n+1)/g}} \ge 1 - \mathsf{negl}(\lambda),$$

since by assumption $ng/q^{\varphi(m-n+1)/g} \leq \mathsf{negl}(\lambda)$. By the parameter choices and Lemma 3, $s_{\min}(\sqrt{\mathbf{\Sigma}_i}) \geq \eta_\varepsilon(\Lambda_q^\perp(\mathbf{A}))$ with overwhelming probability in λ for all $i \in [k]$. Then by Lemma 26, for a primitive $\mathbf{A}$ and the parameter choices, the vectors $\{\mathbf{u}_i\}_{i\in[k]}$ where $\mathbf{u}_i \leftarrow \mathcal{D}_{\Lambda_q^{\mathbf{v}_i}(\mathbf{A}),\sqrt{\mathbf{\Sigma}_i}}$ have $\mathcal{R}_{\mathsf{q}_j}$-linearly independent columns for all $j \in [g]$ with overwhelming probability in λ. It follows that the statistical distance between $\mathcal{D}_0$ and $\mathcal{D}_1$ (unrestricted) is also negligible in λ. □

B κ-MSIS and κ-MLWE Problems and Reductions

In this section we give definitions for the ring version of κ-MSIS and κ-MLWE assumptions. We also give proofs for the adapted lemmas required for the proof of Theorem 1.

Definition 13 (κ-MSIS Assumption). *Let $\varphi, n, m, \log q, \kappa, s_{\max}(\mathbf{\Sigma}), \beta \in \mathsf{poly}(\lambda)$ with $n, \kappa \leq m$. An instance of κ-$\mathsf{MSIS}_{\mathcal{R}_q,n,m,\mathbf{\Sigma},\beta}$ problem is a uniform matrix $\mathbf{A} \leftarrow \mathcal{R}_q^{n\times m}$ and κ vectors $\mathbf{e}_i \leftarrow \mathcal{D}_{\Lambda_q^\perp(\mathbf{A}),\mathbf{\Sigma}}$ for $i \in [\kappa]$. A solution to the problem is a vector $\mathbf{v} \in \mathcal{R}_q^m$ such that:*

$$(1)\quad \mathbf{A}\cdot\mathbf{v} = \mathbf{0} \bmod q, \qquad (2)\quad \|\mathbf{v}\| \leq \beta, \qquad (3)\quad \mathbf{v} \notin \mathcal{K}\text{-}\mathsf{span}(\mathbf{e}_0,\ldots,\mathbf{e}_{\kappa-1}).$$

The κ-$\mathsf{MSIS}_{\mathcal{R}_q,n,m,\mathbf{\Sigma},\beta}$ assumption states that any PPT $\mathcal{A}$ finds a solution with probability $\leq \mathsf{negl}(\lambda)$.

Definition 14 (κ-MLWE Assumption). *Let $n, m, \log q, \kappa, \chi, s_{\max}(\mathbf{\Sigma}) \in \mathsf{poly}(\lambda)$ with $n, \kappa \leq m$. The κ-$\mathsf{MLWE}_{\mathcal{R}_q,n,m,\mathbf{\Sigma},\chi}$ assumption states that for any PPT $\mathcal{A}$,*

$$\left|\Pr\left[b=1 \;\middle|\; \begin{array}{l}\mathbf{A}\leftarrow\mathcal{R}_q^{n\times m}\\ \mathbf{U}\leftarrow(\mathcal{D}_{\Lambda_q^\perp(\mathbf{A}),\mathbf{\Sigma}})^\kappa\\ \mathbf{s}\leftarrow\mathcal{R}_q^n,\ \mathbf{e}\leftarrow\mathcal{D}_{\mathcal{R},\chi}^m\\ \mathbf{b}^\mathsf{T}\coloneqq\mathbf{s}^\mathsf{T}\cdot\mathbf{A}+\mathbf{e}^\mathsf{T} \bmod q\\ b\leftarrow\mathcal{A}(\mathbf{A},\mathbf{U},\mathbf{b})\end{array}\right] - \Pr\left[b=1 \;\middle|\; \begin{array}{l}\mathbf{A}\leftarrow\mathcal{R}_q^{n\times m}\\ \mathbf{U}\leftarrow(\mathcal{D}_{\Lambda_q^\perp(\mathbf{A}),\mathbf{\Sigma}})^\kappa\\ \mathbf{b}\leftarrow\mathcal{U}(\mathbf{U}^\perp)+\mathcal{D}_{\mathcal{R},\chi}^m\\ b\leftarrow\mathcal{A}(\mathbf{A},\mathbf{U},\mathbf{b})\end{array}\right]\right| \leq \mathsf{negl}(\lambda).$$

The Lemmas 27 and 28 are used in the proof of Lemma 29. We give indications on how to adapt the proofs below the statement. The Lemma 29 is the central adaptation we need for Theorem 1, there we give a proof in full.

Lemma 27 (Generalised from [AGHS13, Lem. 8 (eprint)]). *Let $m > 2n > 2$, and $\min_i(\sigma_i) \geq C\cdot\sqrt{n}$ for an absolute constant C. Let $\mathbf{X} \leftarrow \prod_{i=0}^{n-1}\mathcal{D}_{\mathbb{Z}^m,\sigma_i}$. Then*

$$\Omega(\min(\sigma_i)\sqrt{m}) \leq s_{\min}(\mathbf{X}^\mathsf{T}) \leq s_{\max}(\mathbf{X}^\mathsf{T}) \leq O(\max(\sigma_i)\sqrt{m})$$

with probability $1-2^{-\Omega(m)}$.

The generalisation uses standard inequalities of the form $s_{\max}(\mathbf{A}\cdot\mathbf{B}) \leq s_{\max}(\mathbf{A})\cdot s_{\max}(\mathbf{B})$ and $s_{\min}(\mathbf{A}\cdot\mathbf{B}) \geq s_{\min}(\mathbf{A})\cdot s_{\min}(\mathbf{B})$ for rectangular matrices $\mathbf{A}$ and $\mathbf{B}$.

Lemma 28 (Generalised from [LPSS14, Lem. 7 (eprint)]). *Let $\sigma_i > 0$ for $i = 0,\ldots,n-1$. Let $n \geq 100$, $\varepsilon \in (0, 1/1000)$, $\min(\sigma_i) \geq 9\sqrt{\ln(2n(1+1/\varepsilon))/\pi}$, $m \geq 30n\log(n\max(\sigma_i))$, and $\mathbf{S} \in \mathbb{R}^{m\times m}$ a full-rank matrix s.t. $s_{\min}(\mathbf{S}) \geq 10n\max(\sigma_i)\log^{3/2}(nm\max(\sigma_i)/\varepsilon)$. Let $\mathbf{X} \leftarrow \prod_{i=0}^{n-1}\mathcal{D}_{\mathbb{Z}^m,\sigma_i}$ then*

- $\mathbb{Z}^n = \mathbf{X}^\mathsf{T}\cdot\mathbb{Z}^m$ *with probability* $1-2^{-\Omega(n)}$,
- $\mathsf{SD}(\mathbf{X}^\mathsf{T}\cdot\mathcal{D}_{\mathbb{Z}^m,\mathbf{S},\mathbf{c}}, \mathcal{D}_{\mathbb{Z}^n,\mathbf{S}',\mathbf{X}^\mathsf{T}\mathbf{c}}) \leq 2^{-\Omega(n)}$ *with* $\mathbf{S}' = \sqrt{\mathbf{X}^\mathsf{T}\cdot\mathbf{S}\cdot\mathbf{S}^\mathsf{T}\cdot\mathbf{X}}$.

The generalisation requires updating the norm bounds for the elements in $\mathbf{X}$ throughout the proof in [AR13, Thm. 5.1 (arxiv)].

Lemma 29 (Generalised from [LPSS14, Lem. 16 (eprint)]). *Let $\sigma_i > 0, \sigma'_i > 0$ for all $i = 0, \ldots, \kappa - 1$. Let $n \geq 100$, $m \geq \Omega(n \log(n \max_i(\sigma_i)))$, $\min_i(\sigma_i) > \Omega(\sqrt{nm \log m})$ and*

$$\min_i(\sigma'_i) \geq \Omega(n\sqrt{m} \max(\sigma_i)^2 \log^{3/2}(nm \max(\sigma_i))).$$

Let $\mathbf{S} = \mathsf{diag}\left(\sigma'_0, \ldots, \sigma'_{\kappa-1}\right)$. Let $\mathbf{X} \leftarrow \prod_{i=0}^{n-1} \mathcal{D}_{\mathbb{Z}^m, \sigma_i}$. For any $\mathbf{c} \in \mathbb{Z}^n$ we can efficiently sample $\mathbf{r} \in \mathbb{Z}^m$ such that for $\mathbf{x}' = \mathbf{c} + \mathbf{X}^{\mathsf{T}}\mathbf{r}$

- $\|\mathbf{r}\| = O(\max_i(\sigma'_i)/\min_i(\sigma_i))$ *with probability* $1 - 2^{-\Omega(n)}$,
- $\mathsf{SD}((\mathbf{X}, \mathbf{x}'), (\prod_{i=1}^{\kappa} \mathcal{D}_{\mathbb{Z}^m, \sigma_i}, \mathcal{D}_{\mathbb{Z}^n, \mathbf{S}, \mathbf{c}})) \leq 2^{-\Omega(n)}$.

Proof. Set the distribution of $\mathbf{r} \in \mathbb{Z}^m$ to $\mathcal{D}_{\mathbb{Z}^m, \mathbf{\Sigma}}$ where $\mathbf{\Sigma} = [(\mathbf{X}^{\mathsf{T}})^+ \cdot \mathbf{S} \mid \mathbf{M}_{ker}]$ where $\mathbf{M}_{ker} \in \mathbb{Q}^{(m-n)\times(m-n)}$ is a basis of $\ker(\mathbf{X}^{\mathsf{T}})$. We find $\mathbf{M}_{ker}$ by solving $\mathbf{X}^{\mathsf{T}} \cdot \mathbf{v} = 0$ equation using Gaussian elimination over $\mathbb{Q}$, then we scale the matrix down to obtain $\mathbf{M}_{ker}$ of norm 1. We obtain

$$\mathbf{X}^{\mathsf{T}} \cdot \mathbf{\Sigma} \cdot \mathbf{\Sigma}^{\mathsf{T}} \cdot \mathbf{X} = \begin{bmatrix}\mathbf{S}\ \mathbf{0}\end{bmatrix} \cdot \begin{bmatrix}\mathbf{S}^T \\ \mathbf{0}\end{bmatrix} = \mathbf{S} \cdot \mathbf{S}^T.$$

Therefore, $\mathbf{S}$ is also a square root of the matrix above. We check the conditions of Lemma 28. The matrix $\mathbf{\Sigma}$ is a full rank by construction. For the lowest singular value of $\mathbf{\Sigma}$ we know $s_{\min}(\mathbf{\Sigma}) \geq s_{\min}((\mathbf{X}^{\mathsf{T}})^+ \cdot \mathbf{S}) - s_{\max}(\mathbf{M}_{ker}) \geq \frac{\min_i(\sigma'_i)}{\max_i(\sigma_i)\sqrt{m}} - 1$. We then verify that $\frac{\min_i(\sigma'_i)}{\max_i(\sigma_i)\sqrt{m}} - 1 \geq 10n \max(\sigma_i) \log^{3/2}(nm \max(\sigma_i)/\varepsilon)$ is compatible with our parameters. Therefore by Lemma 28

$$\mathbf{X}^{\mathsf{T}} \cdot \mathcal{D}_{\mathbb{Z}^m, \mathbf{\Sigma}} = \mathcal{D}_{\mathbf{X}^{\mathsf{T}}\mathbb{Z}^m, \mathbf{X}^{\mathsf{T}}\mathbf{\Sigma}} = \mathcal{D}_{\mathbb{Z}^n, \mathbf{S}}.$$

It remains to upper bound the norm of the vector $\mathbf{r}$.

$$\begin{aligned} s_{\max}(\mathbf{\Sigma}) &\leq \sqrt{\|(\mathbf{X}^{\mathsf{T}})^+\|^2 \cdot \|\mathbf{S}\|^2 + 1} \\ &= \sqrt{\frac{\max_i(\sigma'_i)^2}{s_{\min}(\mathbf{X}^{\mathsf{T}})^2} + 1} \\ &\leq \sqrt{\frac{\max_i(\sigma'_i)^2}{\Omega(\min_i(\sigma_i)^2 m)} + 1} = O(\max_i(\sigma'_i)/\min_i(\sigma_i) \cdot \sqrt{m}) \end{aligned}$$

in the first inequality we use $s_{\max}([\mathbf{A} \mid \mathbf{B}] \leq s_{\max}(\mathbf{A}) + s_{\max}(\mathbf{B}))$ and that the matrix norm is sub-multiplicative. The third equality hold by construction of the Moore-Penrose pseudoinverse using the SVD of matrix $\mathbf{X}^{\mathsf{T}}$. The last inequality follows from Lemma 27 with overwhelming probability. By the Gaussian concentration bound in Lemma 6 $\|\mathbf{r}\| \leq O(\max_i(\sigma'_i)/\min_i(\sigma_i) \cdot \sqrt{m}) \cdot \sqrt{m}$ with overwhelming probability. □