

Worst-Case Lattice Sampler with Truncated Gadgets and Applications

Corentin Jeudy¹  and Olivier Sanders¹ 

corentin.jeudy@orange.com, olivier.sanders@orange.com

Orange Labs, Applied Crypto Group, Cesson-Sévigné, France

Abstract. Gadget-based samplers have proven to be a key component of several cryptographic primitives, in particular in the area of privacy-preserving mechanisms. Most constructions today follow the approach introduced by Micciancio and Peikert (MP) yielding preimages whose dimension linearly grows with that of the gadget. To improve performance, some papers have proposed to truncate the gadget but at the cost of an important feature of the MP sampler, namely the ability to invert *arbitrary* syndromes. Technically speaking, they replace the *worst-case* MP sampler by an *average-case* sampler that can only be used in specific contexts. Far from being a mere theoretical restriction, it prevents the main applications of gadget-based samplers from using truncated variants and thus from benefiting from the associated performance gains.

In this paper, we solve this problem by describing a worst-case sampler that still works with truncated gadgets. Its main strength is that it retains the main characteristics of the MP sampler while providing flexibility in the choice of the truncation parameter. As a consequence, it can be used as a plug-in replacement for all applications relying on the MP sampler so far, leading to performance improvements up to 30 % as illustrated by several examples in this paper. Our sampler is supported by a thorough security analysis that addresses the hurdles met by previous works and its practicality is demonstrated by a concrete implementation.

Keywords: Lattice-Based Cryptography · Trapdoors · Preimage Sampling · Advanced Signatures

1 Introduction

A trapdoor function f is a function that is easy to evaluate and hard to invert, except for the entity knowing a specific information, the trapdoor, that enables efficient inversion. It has proven to be particularly useful in the context of digital signatures where one can roughly define the signature as $\sigma = f^{-1}(u)$ for some appropriate u that depends on the message m to be signed. Such a signature (which requires trapdoor knowledge to be computed) can then be publicly verified by testing whether $f(\sigma) = u$. Obviously, some adaptation is necessary to meet the security requirements of digital signatures but this basic idea has underlain many practical constructions since the RSA algorithm [RSA78].

In the lattice setting, a prominent example of trapdoor functions is the one by Ajtai [Ajt96], in particular since its secure adaptation to digital signature by [GPV08]. In the latter, they introduce the notion of trapdoor preimage samplable function (TPSF) which should allow for randomizing the inversion of the function. The TPSF in question is defined by a public matrix $\mathbf{C} \in \mathbb{Z}_q^{d \times m}$ and consists in computing $f(\mathbf{x}) = \mathbf{C}\mathbf{x}$ for short vectors \mathbf{x} . For proper \mathbf{C} , inversion is hard under the ISIS problem, except for the entity owning the corresponding trapdoor. When used in the context of digital signatures, inversion should ideally (1) be possible for any syndrome $\mathbf{u} \in \mathbb{Z}_q^d$, (2) should not leak any information on the trapdoor and (3) should result in preimages \mathbf{x} that are as small as possible. Unfortunately, achieving all these three features simultaneously is very difficult in practice. In particular, as we shall see, improving one of these features often impacts negatively the other ones.

For example, let us set $\mathbf{C} = [\mathbf{A}|\mathbf{T}\mathbf{G} - \mathbf{A}\mathbf{R}]$, as is done in the seminal work by Micciancio and Peikert [MP12], where \mathbf{A} is random, \mathbf{T} is invertible, \mathbf{R} is short and $\mathbf{G} = [1|b|\dots|b^{k-1}] \otimes \mathbf{I}_d \in \mathbb{Z}^{d \times dk}$ is the base- b gadget matrix where $k = \lceil \log_b q \rceil$. With the knowledge of \mathbf{R} , one can easily compute a low-norm preimage for any syndrome $\mathbf{u} \in \mathbb{Z}_q^n$ by generating \mathbf{z} as the base- b decomposition of the vector $\mathbf{T}^{-1}\mathbf{u}$ and by returning $\mathbf{x} = [\mathbf{R}\mathbf{z}^T | \mathbf{z}^T]^T$. Indeed, $\mathbf{C}\mathbf{x} = \mathbf{A}\mathbf{R}\mathbf{z} + \mathbf{T}\mathbf{G}\mathbf{z} - \mathbf{A}\mathbf{R}\mathbf{z} = \mathbf{T}\mathbf{G}\mathbf{z} = \mathbf{u}$ by construction of \mathbf{G} . The matrix \mathbf{R} is then the trapdoor corresponding to \mathbf{C} and the resulting trapdoor function fully satisfies condition (1) and, more or less, condition (3) (we will discuss this point further). However, inversion clearly fails to satisfy condition (2) as the resulting preimage \mathbf{x} leaks information on \mathbf{R} . In other words, using this approach to generate signatures would result in a totally insecure scheme where the signer's secret key would progressively leak with issued signatures. This is what led the authors of [MP12] to devise a much more elaborate preimage sampling algorithm which completely erases the dependency on \mathbf{R} of the preimage \mathbf{x} by imposing a Gaussian distribution on \mathbf{z} and by introducing a well-crafted perturbation \mathbf{p} . The resulting construction thus perfectly satisfies conditions (1) and (2) but, when it comes to (3) we note that \mathbf{x} has a rather large dimension, essentially because of its \mathbf{z} component. Of course, one could try to decrease this dimension by resorting to a larger decomposition base b but it would then increase the norm of \mathbf{z} . The latter option is thus more a tradeoff than a real solution. In the end, it means that using the TPSF of [MP12] in a signature scheme will result in rather large signatures which are not competitive with standardized alternatives (e.g. Falcon [PFH+20], Dilithium [DKL+18]), at least on the size metric.

Fortunately, the other features of the Micciancio-Peikert (MP) sampler have proven to be extremely useful for other applications, in particular in privacy-preserving authentication mechanisms. Indeed, the latter extensively rely on zero-knowledge proofs that are notoriously hard to combine with constructions in the random oracle model such as [PFH+20,DKL+18]. Conversely, the MP sampler can yield standard model signature, in part because it satisfies condition (1). As a consequence, it has served as a core building block of countless constructions such as group signatures [dPLS18,LNPS21,LNP22], anonymous

credentials [JRS23,LLLW23,AGJ+24], blind signatures [JS24], etc. Obviously, the efficiency problem remains but it is somehow compensated by the smooth interaction with zero-knowledge proof that the MP sampler enables.

To improve the performance of the MP sampler, Chen et al. [CGM19] introduced the notion of *approximate* trapdoors. The latter still allow to invert functions $\mathbf{x} \mapsto \mathbf{C}\mathbf{x}$, but up to some error. Concretely, given some $\mathbf{u} \in \mathbb{Z}_q^d$, they allow to generate short vectors \mathbf{x} such that $\mathbf{C}\mathbf{x} = \mathbf{u} + \mathbf{e}$ where \mathbf{e} is also a short vector. Interestingly, this approach does not significantly weaken the underlying computational assumption as the authors show that security can still be reduced to the ISIS assumption.

As demonstrated in [CGM19], the MP sampler lends itself well to approximate trapdoors. More precisely, if we replace \mathbf{G} by a truncated version $\mathbf{G}_H = [b^\ell | \dots | b^{k-1}] \otimes \mathbf{I}_d$ where all the lower entries have been dropped, one is still able to invert the upper part of any $\mathbf{u} \in \mathbb{Z}_q^d$. Concretely, one can now compute, for all $\mathbf{u} \in \mathbb{Z}_q^d$, a short \mathbf{z} such that $\mathbf{G}_H \mathbf{z} = \mathbf{u} + \mathbf{e}$ for some small $\mathbf{e} \in \mathbb{Z}_q^d$. One can thus proceed as previously and compute $\mathbf{x} = [\mathbf{R}\mathbf{z}^T | \mathbf{z}^T]^T$. Assuming that \mathbf{T} does not increase the norm too much (which can easily be enforced in practice), this is a valid approximate preimage with error $\mathbf{T}\mathbf{e}$. From the performance standpoint, the gains are very significant as truncated gadgets lead to a smaller preimage with fewer dimensions. From the security standpoint the situation is however more complex. Indeed, here again, one cannot directly use this naive preimage sampling as it trivially leaks information on \mathbf{R} . The authors then resort to the same perturbation approach as in [MP12] but face a very specific issue related to the error $\mathbf{T}\mathbf{e}$. The latter indeed depends on the perturbation which makes $\mathbf{T}\mathbf{e}$ very difficult to simulate without the trapdoor (a necessary step to prove condition (2)), at least for general syndromes \mathbf{u} . Actually, the authors were only able to prove simulatability of their preimage and error for uniform and reprogrammable syndromes \mathbf{u} , which does not fully satisfy condition (1). The corresponding TPSF is then categorized as “average-case”. While this is not a problem for GPV-like signatures proven in the random oracle model, these conditions are not met in the case of privacy-preserving authentication mechanisms where the syndrome is likely to be adversarially controlled to some extent. This is particularly frustrating because it means that the substantial improvements induced by approximate trapdoors are inaccessible to the main applications of MP trapdoors that must still use the full gadget \mathbf{G} and thus inherit the associated performance limitations.

1.1 Our Contributions

In this paper, we solve this problem by describing a sampler that uses a truncated gadget¹, as in [CGM19], but that can still compute a preimage for *any* syndrome $\mathbf{u} \in \mathbb{Z}_q^d$ without leaking any information on the trapdoor. In other words, our

¹ We talk about truncated gadgets and not approximate trapdoors because our sampler actually produces exact preimages as we shall see

sampler fully satisfies conditions (1) and (2) and drastically improves the performance (condition (3)). This provides an answer to the open question formulated in [CGM19] to use truncated gadgets like in their approximate trapdoor framework while satisfying condition (1). Indeed, all known TPSFs with truncated gadgets [CGM19, YJW23, JRS24] were average-case as their security only held when the inverted syndrome were uniform and reprogrammable. Our sampler thus provides the first “worst-case” TPSF with truncated gadget. Better still, as our sampler retains the general structure of the MP sampler, it acts as a plug-in replacement that readily leads to more efficient schemes for all applications. Concretely, when plugged in the standard model signature derived from [MP12], one gets a signature which is 30% smaller without impacting security. Given the central role of the [MP12] sampler in privacy-preserving authentication mechanisms, satisfying condition (1) while truncating the gadget impacts positively such mechanisms. This is because, beyond the signature component of such systems, they employ zero-knowledge frameworks (such as the one from [LNP22]) which are particularly sensitive to the witness dimension.

To achieve such results, we start by introducing a generic sampler, which encompasses previous variants of the exact MP sampler, and we prove that it satisfies conditions (1) and (2) under carefully identified constraints. We then propose a specific instantiation of this generic sampler that meets such constraints and that also allows to truncate the gadget to improve performance. We then apply our *truncated sampler* to several designs of advanced signatures to showcase its concrete impact. We finally implement our sampler to show that its computational efficiency in the standard model signature use-case is similar to the original MP sampler. It incurs only a mild overhead during signing but the smaller signature dimension results in faster verification and will also positively impact zero-knowledge proofs generation and verification, a major step of the applications mentioned above. We give more details on the concrete performance in Section 5.2. Let us now present a more technical overview of our contributions.

Our Approach. We start with the MP sampler which adapts the naive base- b decomposition sketched above by first sampling a perturbation vector \mathbf{p} following some Gaussian distribution \mathcal{D} and then invert the syndrome $\mathbf{u} - \mathbf{C}\mathbf{p}$ instead of just \mathbf{u} . As a consequence, the resulting $\mathbf{x} = [\mathbf{R}\mathbf{z}^T | \mathbf{z}^T]^T$, where $\mathbf{G}\mathbf{z} = \mathbf{T}^{-1}(\mathbf{u} - \mathbf{C}\mathbf{p})$, is a preimage of $\mathbf{u} - \mathbf{C}\mathbf{p}$ which means that $\mathbf{v} = \mathbf{x} + \mathbf{p}$ is a preimage of \mathbf{u} . By carefully selecting the parameters of \mathcal{D} and the one used to generate \mathbf{z} , one can ensure that \mathbf{v} is distributed independently of the trapdoor \mathbf{R} , thanks to a convolution result by Peikert [Pei10]. As discussed above, [CGM19] follows the same approach but drops some entries of the gadget matrix which leads to an error $\mathbf{T}\mathbf{G}_L\mathbf{z}_L$ where $\mathbf{G}_L = [1|b|\dots|b^{\ell-1}] \otimes \mathbf{I}_d$ and \mathbf{z}_L is constituted of the components of \mathbf{z} matching the columns of \mathbf{G}_L . As this error depends on \mathbf{p} , we cannot directly rely on the same argument as [MP12]. Worse, the thorough analysis in [CGM19] suggests that proving simulatability in the general case is hard, hence the restrictions introduced by the authors on the distribution of \mathbf{u} .

To circumvent this problem, we actually transform the approximate sampler by [CGM19] into an exact one so as to rely on a security argument closer to the one in [MP12]. Note that, from a purely functional standpoint, this is not very difficult. Indeed, if \mathbf{C} is in Hermite Normal form (HNF) $[\mathbf{I}|\mathbf{C}']$ for some $\mathbf{C}' \in \mathbb{Z}_q^{d \times m-d}$, then any approximate preimage \mathbf{x} of \mathbf{u} yields an exact one $\mathbf{x}' = \mathbf{x} - [\mathbf{e}^T | \mathbf{0}]^T$. Indeed, $[\mathbf{I}|\mathbf{C}']\mathbf{x} = \mathbf{u} + \mathbf{e}$ implies that $[\mathbf{I}|\mathbf{C}']\mathbf{x} - \mathbf{e} = \mathbf{u}$ and hence $[\mathbf{I}|\mathbf{C}'](\mathbf{x} - [\mathbf{e}^T | \mathbf{0}]^T) = \mathbf{u}$. By defining appropriate bounds on the error \mathbf{e} , one ensures that \mathbf{x}' remains small and so that it is a valid preimage. Actually, this argument is essentially the one provided in [CGM19] to argue that the approximate ISIS problem is not that different from the regular version of ISIS. Of course, this does not solve the simulatability issue of [CGM19] as the problem regarding the error $\mathbf{e} = \mathbf{T}\mathbf{G}_L\mathbf{z}_L$ is directly transferred to \mathbf{x}' .

This rearrangement is thus not sufficient but it has brought us closer to the spirit of the original MP sampler. For sake of clarity, let us rewrite the associated matrix $[\mathbf{A}|\mathbf{T}\mathbf{G} - \mathbf{A}\mathbf{R}]$ as $[[\mathbf{I}|\mathbf{A}']|\mathbf{T}\mathbf{G} - (\mathbf{R}_1 + \mathbf{A}'\mathbf{R}_2)]$ where $[\mathbf{I}|\mathbf{A}']$ is the HNF form of \mathbf{A} and $\mathbf{R} = [\mathbf{R}_1^T | \mathbf{R}_2^T]^T$. By definition of \mathbf{G} , we can split it into $\mathbf{G} = [\mathbf{G}_L | \mathbf{G}_H]$, and thus $\mathbf{z} = [\mathbf{z}_L^T | \mathbf{z}_H^T]^T$. If we assume that the perturbation $\mathbf{p} = \mathbf{0}$ for the moment, our preimage \mathbf{x}' is then exactly

$$\begin{bmatrix} \mathbf{R}_1\mathbf{z}_H + \mathbf{T}\mathbf{G}_L\mathbf{z}_L \\ \mathbf{R}_2\mathbf{z}_H \\ \mathbf{z}_H \end{bmatrix} = \begin{bmatrix} \mathbf{T}\mathbf{G}_L & \mathbf{R}_1 \\ \mathbf{0} & \mathbf{R}_2 \\ \mathbf{0} & \mathbf{I}_{d(k-\ell)} \end{bmatrix} \begin{bmatrix} \mathbf{z}_L \\ \mathbf{z}_H \end{bmatrix},$$

and we define \mathbf{M} as the matrix before $[\mathbf{z}_L^T | \mathbf{z}_H^T]^T$. From a theoretical standpoint, we are not that far from the MP sampler where the resulting preimage was $[\mathbf{R}^T | \mathbf{I}]^T \mathbf{z}$. Actually, by noticing that in the case of the MP sampler $\mathbf{G}_L = \emptyset$, one can note that our sampler involving \mathbf{M} is a generalization of the MP one. To remove the leakage introduced by \mathbf{M} , we thus aim at resorting to the same solution, namely using a well-crafted perturbation \mathbf{p} . However, this is quite complex in our context for the following two reasons. First, \mathbf{G}_L is rather large, at least compared to \mathbf{R}_1 and \mathbf{R}_2 , which would require very large perturbation parameters if we directly apply the MP approach. Second, deriving good parameters for this perturbation requires at some point to precisely bound the covariance of \mathbf{x}' and derive appropriate smoothing conditions, which is quite difficult in our case as we shall see. To address these problems, we revisit the original MP approach by providing a much more generic result that we can directly use to prove security in our case but that could also serve to analyze other variants of the MP sampler. The strength of our result is that it does not sacrifice efficiency for genericity. Indeed, when used in the original MP setting, it allows to derive even better (albeit very slightly) parameters.

The first step of our result is to write the matrix \mathbf{M} as $\mathbf{K} \cdot \mathbf{L}$ where \mathbf{K} only contains public elements (e.g. the matrix \mathbf{G}_L) and \mathbf{L} contains all the information we want to hide (namely \mathbf{R}_1 , \mathbf{R}_2 and \mathbf{T} ²). We then identify a set of precise requirements related to our perturbation and demonstrate that satisfying them is sufficient to prove security of our sampler. This demonstration is quite technical

² The tag must be hidden to prove security of MP-like signatures as we will explain.

and constitutes our first main contribution. The point is that these requirements depend on \mathbf{L} and not \mathbf{K} . This approach thus allows to completely remove the contribution of \mathbf{K} when defining the perturbation parameters while ensuring that everything will work when we will recombine \mathbf{M} . We thus sample a preimage with \mathbf{L} , which we then project with the public matrix \mathbf{K} . In our case, this concretely allows to define perturbation parameters that only depends on \mathbf{R}_1 , \mathbf{R}_1 and \mathbf{T} , and not \mathbf{G}_L , which thus solves our first problem mentioned above.

However, the second problem remains. Our covariance depends on \mathbf{L} which has a rather complex form in our general case. We nevertheless manage to identify very concrete parameters that make our perturbation tightly satisfies the requirements above. Our corresponding proof is, here again, very technical and constitutes another contribution of this paper. Putting everything together, we are thus able to securely sample preimages of *any* syndrome with truncated gadgets, leading to improved performance.

As mentioned above, this readily leads to a 30% gain on the signature size when plugged in a standard model signature scheme. We also show other immediate impacts of our result in the area of privacy-preserving mechanisms. We indeed consider the prominent cases of group signatures [CvH91], anonymous credentials [Cha85] and blind signatures [Cha82], and apply our result to the most efficient constructions to date that are based on some versions of the MP sampler, yielding significant improvements in all cases. We note that our generic and truncated sampler naturally extend to the ring setting as we detail in Section 4.3, which is used in all the latter constructions. The resulting sizes and improvement ratios are summarized in Table 1.1. A more complete description and comparison is given in Section 5.

	Original	Ours	Improv.
Standard model signature [AGJ ⁺ 24]	6.72 _{KB}	4.82 _{KB}	28.1 %
Group signature [LNPS21,LNP22]	98.02 _{KB}	82.65 _{KB}	15.7 %
Anonymous credentials [AGJ ⁺ 24]	79.58 _{KB}	71.46 _{KB}	10.2 %
Blind signature [JS24]	41.12 _{KB}	36.28 _{KB}	11.8 %

Table 1.1. Comparison of state-of-the-art standard model signatures, group signatures, anonymous credentials, blind signatures using the MP sampler and our new sampler. For the anonymous credentials, the size corresponds to that of the credential presentation proof (here a zero-knowledge proof of a standard model signature) which is the most relevant metric.

The discrepancy between the 30% figure obtained for standard model signature and the ones indicated for privacy-preserving mechanisms is due to the fact that preimage sampling is just one of the building blocks of the latter, which decreases the relative weight of our contributions. In absolute terms, our shorter preimages lead to an improvement up to 15 KB, which is noticeable for practical

applications. Given that our truncated sampler has been designed as a plug-in replacement of the MP sampler, we stress that one can expect similar gains for any constructions that are currently using the latter. The genericity of our security result could also lead to other variants of the MP sampler that would be more appropriate in some contexts, leading to further gains.

2 Preliminaries

We use $\mathbb{N}, \mathbb{Z}, \mathbb{R}$ to respectively denote the set of natural integers, the ring of integers, and the field of reals. For two integers $a \leq b$, we define $[a, b] = \{a, \dots, b\}$ and $[b] = [1, b]$. For a positive integer q , we define $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$. For a finite set S , we call $U(S)$ the uniform distribution over S . For $\delta > 1$ and two distributions $\mathcal{P}_1, \mathcal{P}_2$ of same support S , we write $\mathcal{P}_1 \approx_\delta \mathcal{P}_2$ if for all $x \in S$, $\mathcal{P}_1(x) \in [1/\delta, \delta]\mathcal{P}_2(x)$.

2.1 Linear Algebra

Vectors and matrices are written in bold lowercase and uppercase letters respectively. The transpose of a matrix \mathbf{A} is denoted by \mathbf{A}^T . We define the regular ℓ^p norms of \mathbb{R}^d by $\|\mathbf{a}\|_p = (\sum_{i \in [d]} |a_i|^p)^{1/p}$, and $\|\mathbf{a}\|_\infty = \max_{i \in [d]} |a_i|$. For a set $S \subseteq \mathbb{R}^d$, we write $\text{Span}_{\mathbb{R}}(S)$ to be the subspace of \mathbb{R}^d generated by S . For a matrix $\mathbf{A} = [\mathbf{a}_1 | \dots | \mathbf{a}_m] \in \mathbb{R}^{d \times m}$, we define $\text{Span}_{\mathbb{R}}(\mathbf{A}) = \text{Span}_{\mathbb{R}}(\{\mathbf{a}_i; i \in [m]\}) \subseteq \mathbb{R}^d$. We also define $\ker(\mathbf{A}) = \{\mathbf{x} \in \mathbb{R}^m : \mathbf{A}\mathbf{x} = \mathbf{0}\}$. The set of real unitary matrices of dimension d is denoted by $O_d(\mathbb{R})$.

Singular Value Decomposition. For $\mathbf{A} \in \mathbb{R}^{d \times m}$, the singular value decomposition (SVD) gives the existence of $\mathbf{U} \in O_d(\mathbb{R})$, $\mathbf{V} \in O_m(\mathbb{R})$ and $\mathbf{D} \in \mathbb{R}^{d \times m}$ rectangular diagonal with non-negative diagonal entries in non-increasing order, such that $\mathbf{A} = \mathbf{U}\mathbf{D}\mathbf{V}^T$. When $m > d$, $\mathbf{D} = [\mathbf{D}' | \mathbf{0}_{d \times (m-d)}]$ with $\mathbf{D}' = \text{diag}(s_1(\mathbf{A}), \dots, s_d(\mathbf{A}))$, where $s_1(\mathbf{A}) \geq \dots \geq s_d(\mathbf{A}) \geq 0$ are called the singular values of \mathbf{A} . When $m \leq d$, $\mathbf{D} = \begin{bmatrix} \mathbf{D}' \\ \mathbf{0}_{(d-m) \times m} \end{bmatrix}$ with $\mathbf{D}' = \text{diag}(s_1(\mathbf{A}), \dots, s_m(\mathbf{A}))$.

Moore-Penrose Pseudoinverse. For a matrix $\mathbf{A} \in \mathbb{R}^{d \times m}$, there exists a unique matrix $\mathbf{A}^+ \in \mathbb{R}^{m \times d}$ verifying the four Moore-Penrose conditions: (1) $\mathbf{A}\mathbf{A}^+\mathbf{A} = \mathbf{A}$, (2) $\mathbf{A}^+\mathbf{A}\mathbf{A}^+ = \mathbf{A}^+$, (3) $(\mathbf{A}\mathbf{A}^+)^T = \mathbf{A}\mathbf{A}^+$, (4) $(\mathbf{A}^+\mathbf{A})^T = \mathbf{A}^+\mathbf{A}$. The matrix \mathbf{A}^+ is called the Moore-Penrose pseudoinverse, or just pseudoinverse for short, of \mathbf{A} . The pseudoinverse operator is an involution, i.e., $(\mathbf{A}^+)^+ = \mathbf{A}$, and it commutes with the transpose operator, i.e., $(\mathbf{A}^T)^+ = (\mathbf{A}^+)^T = \mathbf{A}^{+T}$. It holds that $\mathbf{A}\mathbf{A}^+$ is the orthogonal projector onto $\text{Span}_{\mathbb{R}}(\mathbf{A})$, while $\mathbf{A}^+\mathbf{A}$ is the orthogonal projector onto $\text{Span}_{\mathbb{R}}(\mathbf{A}^T)$. We also have that $\text{Span}_{\mathbb{R}}(\mathbf{A}^+) = \text{Span}_{\mathbb{R}}(\mathbf{A}^T)$ and $\ker(\mathbf{A}^+) = \ker(\mathbf{A}^T)$. When \mathbf{A} is invertible, then $\mathbf{A}^+ = \mathbf{A}^{-1}$.

If \mathbf{A} has linearly independent columns, then $\mathbf{A}^+ = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T$ and thus $\mathbf{A}^+\mathbf{A} = \mathbf{I}_m$. If \mathbf{A} has linearly independent rows, $\mathbf{A}^+ = \mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1}$, and thus $\mathbf{A}\mathbf{A}^+ = \mathbf{I}_d$. More generally, if $\mathbf{A} = \mathbf{U}\mathbf{D}\mathbf{V}^T$ is an SVD of \mathbf{A} , then $\mathbf{A}^+ = \mathbf{V}\mathbf{D}^+\mathbf{U}^T$. As opposed to the inverse, in general $(\mathbf{A}\mathbf{B})^+ \neq \mathbf{B}^+\mathbf{A}^+$. In certain specific cases, the equality holds: for example, if $\mathbf{B} = \mathbf{A}^T$, or if \mathbf{A} has orthonormal columns,

or if \mathbf{B} has orthonormal rows. From that we can deduce that for any pairs of unitary matrices, we have $(\mathbf{UAV}^T)^+ = \mathbf{VA}^+\mathbf{U}^T$. Finally, for all $\alpha \neq 0$, it holds that $(\alpha\mathbf{A})^+ = \alpha^{-1}\mathbf{A}^+$.

Square Roots. A symmetric matrix $\mathbf{S} \in \mathbb{R}^{d \times d}$ is said positive semi-definite if for all $\mathbf{x} \in \mathbb{R}^d$, $\mathbf{x}^T \mathbf{S} \mathbf{x} \geq 0$. In that case, we write $\mathbf{S} \geq 0$. It is said positive definite if for all $\mathbf{x} \in \mathbb{R}^d \setminus \{\mathbf{0}\}$, $\mathbf{x}^T \mathbf{S} \mathbf{x} > 0$, in which case we write $\mathbf{S} > 0$. We write $\mathcal{S}_d^+(\mathbb{R})$ (resp. $\mathcal{S}_d^{++}(\mathbb{R})$) the set of symmetric positive semi-definite (resp. definite) matrices of $\mathbb{R}^{d \times d}$. The Loewner order gives a partial ordering on $\mathcal{S}_d^+(\mathbb{R})$, and $\mathcal{S}_d^{++}(\mathbb{R})$. Concretely, we write $\mathbf{S} \geq \mathbf{R}$ (resp. $\mathbf{S} > \mathbf{R}$) if $\mathbf{S} - \mathbf{R}$ is positive semi-definite (resp. positive definite). We note that the Loewner order is compatible with the inverse on $\mathcal{S}_d^{++}(\mathbb{R})$, that is for $\mathbf{S}, \mathbf{R} \in \mathcal{S}_d^{++}(\mathbb{R})$, $\mathbf{S} \geq \mathbf{R} \Rightarrow \mathbf{R}^{-1} \geq \mathbf{S}^{-1}$. The same holds for strict inequalities. On the other hand, this property is not true in general over $\mathcal{S}_d^+(\mathbb{R})$ using pseudoinverses. The implication is true if and only if \mathbf{S} and \mathbf{R} have the same kernel. Finally, for any $\mathbf{S} \in \mathcal{S}_d^+(\mathbb{R})$, we write $\sqrt{\mathbf{S}}$ any full-rank matrix such that $\mathbf{S} = \sqrt{\mathbf{S}}\sqrt{\mathbf{S}}^T$. Note that $\sqrt{\mathbf{S}}$ need not be square.

2.2 Lattices

A real d -dimensional lattice \mathcal{L} is a finitely generated free \mathbb{Z} -module, accompanied with a Euclidean norm on $\text{Span}_{\mathbb{R}}(\mathcal{L})$. There exists a finite family $(\mathbf{b}_1, \dots, \mathbf{b}_r) \in \mathcal{L}^r$ of linearly independent vectors of \mathbb{R}^d such that $\mathcal{L} = \bigoplus_{i \in [r]} \mathbb{Z} \mathbf{b}_i$. We write $\mathbf{B} = [\mathbf{b}_1 | \dots | \mathbf{b}_r]$ and \mathbf{B} is called a basis of \mathcal{L} . To specify a basis, we usually write $\mathcal{L} = \mathcal{L}(\mathbf{B})$. The integer r is called the rank of the lattice \mathcal{L} and is independent on the choice of basis. For a lattice \mathcal{L} , we define $\lambda_1(\mathcal{L}) = \min_{\mathbf{x} \in \mathcal{L} \setminus \{\mathbf{0}\}} \|\mathbf{x}\|_2$.

The dual lattice of \mathcal{L} is defined by $\mathcal{L}^* = \{\mathbf{x} \in \text{Span}_{\mathbb{R}}(\mathcal{L}) : \forall \mathbf{y} \in \mathcal{L}, \mathbf{x}^T \mathbf{y} \in \mathbb{Z}\}$. If \mathbf{B} is a basis of \mathcal{L} , it holds that $\mathbf{B}^{+T} = (\mathbf{B}^T)^+$ is a basis of \mathcal{L}^* . Because \mathbf{B} has linearly independent columns, it holds that $\mathbf{B}^{+T} = \mathbf{B}(\mathbf{B}^T \mathbf{B})^{-1}$. More generally, we have the following lemma.

Lemma 2.1. *Let $d, m \in \mathbb{N}^\times$ and $\mathbf{A} \in \mathbb{R}^{d \times m}$ be a matrix with linearly independent rows. Then, let $\mathcal{L} \subset \text{Span}_{\mathbb{R}}(\mathbf{A}^T)$ be an m -dimensional lattice. It holds that $(\mathbf{A}\mathcal{L})^* = \mathbf{A}^{+T} \mathcal{L}^*$.*

Proof. Because \mathbf{A} has linearly independent rows, we have $\mathbf{A}^+ = \mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1}$. Let $\mathbf{x} \in (\mathbf{A}\mathcal{L})^*$. We define $\mathbf{y} = \mathbf{A}^T \mathbf{x}$. Let $\mathbf{y}' \in \mathcal{L}$. We have $\mathbf{y}^T \mathbf{y}' = \mathbf{x}^T (\mathbf{A} \mathbf{y}') \in \mathbb{Z}$ by definition of $(\mathbf{A}\mathcal{L})^*$. So $\mathbf{y} \in \mathcal{L}^*$. Yet, we have $\mathbf{x} = \mathbf{A}^{+T} \mathbf{A}^T \mathbf{x} = \mathbf{A}^{+T} \mathbf{y}$ which shows that $\mathbf{x} \in \mathbf{A}^{+T} \mathcal{L}^*$.

Reciprocally, let $\mathbf{y} = \mathbf{A}^{+T} \mathbf{x}$ with $\mathbf{x} \in \mathcal{L}^*$. Let $\mathbf{y}' = \mathbf{A} \mathbf{x}'$ be in $\mathbf{A}\mathcal{L}$ with $\mathbf{x}' \in \mathcal{L}$. We have $\mathbf{y}^T \mathbf{y}' = \mathbf{x}^T \mathbf{A}^+ \mathbf{A} \mathbf{x}'$. Yet $\mathbf{A}^+ \mathbf{A}$ is the orthogonal projector onto $\text{Span}_{\mathbb{R}}(\mathbf{A}^T)$. Because $\mathbf{x}' \in \mathcal{L} \subset \text{Span}_{\mathbb{R}}(\mathbf{A}^T)$, it holds that $\mathbf{A}^+ \mathbf{A} \mathbf{x}' = \mathbf{x}'$ and thus that $\mathbf{y}^T \mathbf{y}' = \mathbf{x}^T \mathbf{x}' \in \mathbb{Z}$ by definition of \mathcal{L}^* . So $\mathbf{x} \in (\mathbf{A}\mathcal{L})^*$. \square

For a matrix $\mathbf{A} \in \mathbb{Z}_q^{d \times m}$ and $\mathbf{u} \in \mathbb{Z}_q^d$, we define the q -ary lattice $\mathcal{L}_q^\perp(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A} \mathbf{x} = \mathbf{0} \pmod{q\mathbb{Z}}\}$, and the lattice coset $\mathcal{L}_q^{\mathbf{u}}(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A} \mathbf{x} = \mathbf{u} \pmod{q\mathbb{Z}}\}$.

2.3 Gaussian Measures

For $\mathbf{S} \in \mathcal{S}_d^+(\mathbb{R})$, and $\mathbf{c} \in \mathbb{R}^d$, we define the Gaussian function $\rho_{\sqrt{\mathbf{S}}, \mathbf{c}}$ by

$$\forall \mathbf{x} \in \mathbb{R}^d, \rho_{\sqrt{\mathbf{S}}, \mathbf{c}}(\mathbf{x}) = \begin{cases} \exp(-\pi(\mathbf{x} - \mathbf{c})^T \mathbf{S}^+ (\mathbf{x} - \mathbf{c})) & \text{if } \mathbf{x} - \mathbf{c} \in \text{Span}_{\mathbb{R}}(\mathbf{S}) \\ 0 & \text{otherwise} \end{cases}$$

Note that the expression only depends on \mathbf{S} and not a specific choice of square root, which is why we index the function by $\sqrt{\mathbf{S}}$. For any d -dimensional lattice \mathcal{L} such that $(\mathcal{L} - \mathbf{c}) \cap \text{Span}_{\mathbb{R}}(\mathbf{S}) \neq \emptyset$, we define the discrete Gaussian distribution by its probability mass function $\mathcal{D}_{\mathcal{L}, \sqrt{\mathbf{S}}, \mathbf{c}} : \mathbf{x} \in \mathcal{L} \mapsto \rho_{\sqrt{\mathbf{S}}, \mathbf{c}}(\mathbf{x}) / \rho_{\sqrt{\mathbf{S}}, \mathbf{c}}(\mathcal{L})$. When $\mathbf{c} = \mathbf{0}$, we omit the subscript, and when $\mathbf{S} = s^2 \mathbf{I}_d$ for $s > 0$, we replace the subscript $\sqrt{\mathbf{S}}$ by s .

As coined by Micciancio and Regev [MR07], we define the *smoothing parameter* of a lattice \mathcal{L} , parameterized by some $\varepsilon > 0$, by $\eta_\varepsilon(\mathcal{L}) = \min\{s > 0 : \rho_{1/s}(\mathcal{L}^*) = 1 + \varepsilon\}$. A recent work by Espitau, Wallet and Yu [EWY23] gives an exact expression of $\eta_\varepsilon(\mathcal{L})$ with tight approximations for remarkable lattices.

Lemma 2.2 ([EWY23, Lem. 5]). *Let \mathcal{L} be a lattice and $\varepsilon > 0$. It holds that*

$$\eta_\varepsilon(\mathcal{L}) = \frac{1}{\lambda_1(\mathcal{L}^*)} \sqrt{\frac{1}{\pi} \ln \left(\frac{\kappa(\mathcal{L}^*)}{\varepsilon} (1 + o_\varepsilon(1)) \right)},$$

where $\kappa(\mathcal{L}^*) = |\{\mathbf{x} \in \mathcal{L}^* : \|\mathbf{x}\|_2 = \lambda_1(\mathcal{L}^*)\}|$ is the kissing number of \mathcal{L}^* . In particular, it holds that for any $d \in \mathbb{N}^\times$, $\eta_\varepsilon(\mathbb{Z}^d) \approx \sqrt{\ln(2d/\varepsilon)}/\pi$.

For non-spherical Gaussian distributions, if $\mathbf{S} \in \mathcal{S}_d^+(\mathbb{R})$ and $\mathcal{L} \subset \text{Span}_{\mathbb{R}}(\mathbf{S})$ is a lattice within the span of \mathbf{S} , we say that $\sqrt{\mathbf{S}} \geq \eta_\varepsilon(\mathcal{L})$ if $1 \geq \eta_\varepsilon(\sqrt{\mathbf{S}}^+ \mathcal{L})$. We have another characterization of $\sqrt{\mathbf{S}}$ exceeding the smoothing parameter, which we give here.

Lemma 2.3. *Let $d \in \mathbb{N}^\times$, $\varepsilon > 0$, $\mathbf{S} \in \mathcal{S}_d^+(\mathbb{R})$ and $\mathcal{L} \subset \text{Span}_{\mathbb{R}}(\mathbf{S})$ a lattice. It holds that $\sqrt{\mathbf{S}} \geq \eta_\varepsilon(\mathcal{L})$ if and only if $\rho_{\sqrt{\mathbf{S}}^+}(\mathcal{L}^*) \leq 1 + \varepsilon$.*

Proof. To show the equivalence, it suffices to have $\rho_1((\sqrt{\mathbf{S}}^+ \mathcal{L})^*) = \rho_{\sqrt{\mathbf{S}}^+}(\mathcal{L}^*)$. Let $\mathbf{R} \in \mathbb{R}^{d \times r}$ be a full rank square root of \mathbf{S} . We have $\text{rank}(\mathbf{R}^+) = \text{rank}(\mathbf{R}^T) = \text{rank}(\mathbf{R}) = r$. So \mathbf{R}^+ has linearly independent rows. Additionally, $\text{Span}_{\mathbb{R}}(\mathbf{R}^{+T}) = \text{Span}_{\mathbb{R}}(\mathbf{R}) = \text{Span}_{\mathbb{R}}(\mathbf{S})$. So $\mathcal{L} \subset \text{Span}_{\mathbb{R}}(\mathbf{R}^{+T})$. By Lemma 2.1 for $\mathbf{A} = \mathbf{R}^+$, it

holds that $(\mathbf{R}^+\mathcal{L})^* = \mathbf{R}^T\mathcal{L}^*$. Hence, we have

$$\begin{aligned}
\rho_1((\mathbf{R}^+\mathcal{L})^*) &= \sum_{\mathbf{x} \in \mathcal{L}^*} \rho_1(\mathbf{R}^T\mathbf{x}) \\
&= \sum_{\mathbf{x} \in \mathcal{L}^*} \exp(-\pi\mathbf{x}^T\mathbf{R}\mathbf{R}^T\mathbf{x}) && \text{(because } \mathbf{R}^T\mathbf{x} \in \text{Span}_{\mathbb{R}}(\mathbf{I}_r) = \mathbb{R}^r\text{)} \\
&= \sum_{\mathbf{x} \in \mathcal{L}^*} \exp(-\pi\mathbf{x}^T(\mathbf{S}^+)^+\mathbf{x}) \\
&= \sum_{\mathbf{x} \in \mathcal{L}^*} \rho_{\sqrt{\mathbf{S}^+}}(\mathbf{x}) && \text{(because } \mathcal{L}^* \subset \text{Span}_{\mathbb{R}}(\mathbf{S}) = \text{Span}_{\mathbb{R}}(\mathbf{S}^+)\text{)} \\
&= \rho_{\sqrt{\mathbf{S}^+}}(\mathcal{L}^*).
\end{aligned}$$

It thus shows the equivalence, which argues that the definition $\sqrt{\mathbf{S}} \geq \eta_\varepsilon(\mathcal{L})$ does not depend on the specific choice of full rank square root. \square

Additionally, if $\mathbf{S} \in \mathcal{S}_d^{++}(\mathbb{R})$ and \mathcal{L} is full rank, $\sqrt{\mathbf{S}} \geq \eta_\varepsilon(\mathcal{L})$ is equivalent to $\mathbf{S} \geq \eta_\varepsilon(\mathcal{L})^2\mathbf{I}_d$ for the Loewner order. This does not hold true if \mathbf{S} is singular.

Lemma 2.4 ([MR07]). *Let $d \in \mathbb{N}^\times$, \mathcal{L} a lattice, $\varepsilon > 0$, $\mathbf{S} \in \mathcal{S}_d^+(\mathbb{R})$ such that $\sqrt{\mathbf{S}} \geq \eta_\varepsilon(\mathcal{L})$. Then for any $\mathbf{c} \in \text{Span}_{\mathbb{R}}(\mathcal{L})$, it holds that $\rho_{\sqrt{\mathbf{S}}}(\mathcal{L} + \mathbf{c}) \in [(1 - \varepsilon)/(1 + \varepsilon), 1]\rho_{\sqrt{\mathbf{S}}}(\mathcal{L})$.*

We will also need the following fact on the smoothing parameter of strict sublattices.

Lemma 2.5. *Let \mathcal{L}_1 be a lattice and $\mathcal{L}_2 \subset \mathcal{L}_1$ be a strict sublattice of \mathcal{L}_1 such that $\text{rank}(\mathcal{L}_1) = \text{rank}(\mathcal{L}_2)$. Then for any $\varepsilon > 0$, it holds that $\eta_\varepsilon(\mathcal{L}_1) < \eta_\varepsilon(\mathcal{L}_2)$.*

Proof. Because \mathcal{L}_2 is a strict sublattice of same rank, it holds that \mathcal{L}_1^* is a strict sublattice of \mathcal{L}_2^* . Indeed, we have $\text{Span}_{\mathbb{R}}(\mathcal{L}_1) = \text{Span}_{\mathbb{R}}(\mathcal{L}_2)$ because $\mathcal{L}_2 \subset \mathcal{L}_1$ and both have the same rank. Let $\mathbf{x} \in \mathcal{L}_1^*$. Then, $\mathbf{x} \in \text{Span}_{\mathbb{R}}(\mathcal{L}_1) = \text{Span}_{\mathbb{R}}(\mathcal{L}_2)$. Now let $\mathbf{y} \in \mathcal{L}_2$. It then holds that $\mathbf{y} \in \mathcal{L}_1$ and therefore $\mathbf{x}^T\mathbf{y} \in \mathbb{Z}$, thus proving that $\mathbf{x} \in \mathcal{L}_2^*$. If the inclusion is not strict, i.e., $\mathcal{L}_1^* = \mathcal{L}_2^*$, then $\mathcal{L}_1 = \mathcal{L}_2$ which yields a contradiction. Because we now have the strict inclusion $\mathcal{L}_1^* \subset \mathcal{L}_2^*$, it naturally holds that for all $s > 0$, $\rho_{1/s}(\mathcal{L}_1^*) < \rho_{1/s}(\mathcal{L}_2^*)$. If we choose $s = \eta_\varepsilon(\mathcal{L}_1)$, we then get $1 + \varepsilon = \rho_{1/s}(\mathcal{L}_1^*) < \rho_{1/s}(\mathcal{L}_2^*)$. By definition of the smoothing parameter, noticing that this is a minimum by continuity and not just an infimum, we then get that $s < \eta_\varepsilon(\mathcal{L}_2)$ as desired. \square

2.4 Algebraic Number Theory

We now give the necessary notions in algebraic number theory. A number field $K = \mathbb{Q}(\zeta)$ is a field extension of \mathbb{Q} of finite degree n adjoining an algebraic number ζ . The set of algebraic integers in K is a ring R called the ring of integers of K . We also define $K_{\mathbb{R}} = K \otimes_{\mathbb{Q}} \mathbb{R}$. For any $q \geq 2$, we define $R_q = R/qR$. A

popular choice of number field is the class of power-of-two cyclotomic fields, for which the degree n is a power-of-two and which are isomorphic to $\mathbb{Q}[x]/\langle x^n + 1 \rangle$. The ring of integers in this case is identified with $\mathbb{Z}[x]/\langle x^n + 1 \rangle$.

The following is stated for $K_{\mathbb{R}}$ but holds also for K and R . Elements of $K_{\mathbb{R}}$ can naturally be embedded into the Euclidean space \mathbb{R}^n by their coefficient vector when seen as a polynomial in ζ or x . We use τ to denote this coefficient embedding, i.e., for $r = \sum_{i=0}^{n-1} r_i \zeta^i$, $\tau(r) = [r_0 | \dots | r_{n-1}]^T$. We also define the conjugate r^* of the element r by being $r^* = r(\zeta^{-1})$. In the power-of-two cyclotomic field of degree n , it holds that $\tau(r^*) = [r_0 | -r_{n-1} | \dots | -r_1]^T$. We denote by $K_{\mathbb{R}}^+$ the set of elements $r \in K_{\mathbb{R}}$ such that $r^* = r$, and by $K_{\mathbb{R}}^{++}$ the set of elements $r \in K_{\mathbb{R}}^+$ that can be written as $r = ss^*$. We then define the multiplication matrix map M_τ defined by the relation $\tau(rs) = M_\tau(r)\tau(s)$ for all pairs of elements r, s . In power-of-two cyclotomic fields, $M_\tau(r)$ corresponds to the nega-circulant matrix with first column $\tau(r)$. We define the usual ℓ^p norms over $K_{\mathbb{R}}$ with respect to the embedding τ , i.e., $\|r\|_p := \|\tau(r)\|_p$. We define $T_1 = \tau^{-1}(\{0, 1\}^n)$.

These notations extend to vectors and/or matrices in the natural way by concatenation, except that the conjugate of a matrix actually corresponds to the conjugate transpose. For a matrix $\mathbf{A} \in K_{\mathbb{R}}^{d \times m}$, we define its spectral norm through its embedding to $\mathbb{R}^{nd \times nm}$ via M_τ as $\|\mathbf{A}\|_2 = \|M_\tau(\mathbf{A})\|_2$.

We extend the notations $\mathcal{L}_q^\perp(\mathbf{A})$ and $\mathcal{L}_q^u(\mathbf{A})$ for matrices over R_q as $\mathcal{L}_q^u(\mathbf{A}) = \{\mathbf{x} \in R^m : \mathbf{A}\mathbf{x} = \mathbf{u} \bmod qR\}$ and $\mathcal{L}_q^\perp(\mathbf{A}) = \mathcal{L}_q^0(\mathbf{A})$. Gaussian distributions over the ring are defined through their coefficient embedding, i.e., $\mathcal{D}_{\mathcal{M}, \sqrt{\mathbf{S}}} = \tau^{-1}(\mathcal{D}_{\tau(\mathcal{M}), \sqrt{\mathbf{S}}})$ where $\tau(\mathcal{M})$ is an module lattice corresponding to an R -module \mathcal{M} . We also define the centered binomial distribution \mathcal{B}_1 over R obtained by sampling each coefficient according to ψ_1 , with $\psi_1(-1) = \psi_1(1) = 1/4$ and $\psi_1(0) = 1/2$.

2.5 Hardness Assumptions

The security of our standard model construction relies on the *Module Short Integer Solution* (M-SIS) and *Module Learning With Errors* (M-LWE) problems formalized in [LS15].

Definition 2.1 (M-SIS). Let $R = \mathbb{Z}[x]/\langle x^n + 1 \rangle$ with n a power-of-two. Let d, m, q be positive integers and $\beta > 0$ with $m > d$. The *Module Short Integer Solution* problem $\text{M-SIS}_{n,d,m,q,\beta}$ asks to find $\mathbf{x} \in \mathcal{L}_q^\perp([\mathbf{I}_d | \mathbf{A}']) \setminus \{\mathbf{0}\}$ such that $\|\mathbf{x}\|_2 \leq \beta$, given $\mathbf{A}' \leftarrow U(R_q^{d \times m-d})$. The advantage of an adversary \mathcal{A} is $\text{Adv}_{\text{M-SIS}}[\mathcal{A}] = \mathbb{P}_{\mathbf{x} \leftarrow \mathcal{A}(\mathbf{A}')} [\mathbf{x} \in \mathcal{L}_q^\perp([\mathbf{I}_d | \mathbf{A}']) \wedge 0 < \|\mathbf{x}\|_2 \leq \beta]$.

Definition 2.2 (M-LWE). Let $R = \mathbb{Z}[x]/\langle x^n + 1 \rangle$ with n a power-of-two. Let d, m, k, q be positive integers and \mathcal{D}_r a distribution on R . The *Module Learning With Errors* problem $\text{M-LWE}_{n,d,m,q,\mathcal{D}_r}^k$ asks to distinguish between the following distributions: (1) $(\mathbf{A}', [\mathbf{I}_m | \mathbf{A}']\mathbf{R} \bmod qR)$, where $\mathbf{A}' \sim U(R_q^{m \times d})$ and $\mathbf{R} \sim \mathcal{D}_r^{d+m \times k}$, and (2) $(\mathbf{A}', \mathbf{B})$, where $\mathbf{A}' \sim U(R_q^{m \times d})$ and $\mathbf{B} \sim U(R_q^{m \times k})$. We define $\text{Adv}_{\text{M-LWE}}[\mathcal{A}] = |\mathbb{P}[\mathcal{A}(\mathbf{A}', [\mathbf{I}_m | \mathbf{A}']\mathbf{R}) = 1] - \mathbb{P}[\mathcal{A}(\mathbf{A}', \mathbf{B}) = 1]|$ as the advantage of an adversary \mathcal{A} .

When the parameters are clear from the context, we define the hardness bound of the problem $P \in \{\text{M-LWE}, \text{M-SIS}\}$ as $\varepsilon_P = \sup_{\mathcal{A}} \text{Adv}_P[\mathcal{A}]$. We recall that a standard hybrid argument shows that $\text{M-LWE}_{n,d,m,q,\mathcal{D}_r}^k$ is at least as hard as $\text{M-LWE}_{n,d,m,q,\mathcal{D}_r}^1$ at the expense of a loss factor k in the reduction.

3 Generic Sampler: Towards Worst-Case TPSF with Truncated Gadgets

We now describe our sampler which generalizes the one in [MP12] as we will see in Section 4.1. A concrete instantiation achieving the performance improvement mentioned in Section 1 is presented in Section 4.2.

We recall that our ultimate goal is to generate short vectors \mathbf{x} such that $\mathbf{A}_T \mathbf{x} = \mathbf{u}$ for any syndrome $\mathbf{u} \in \mathbb{Z}_q^d$, where $\mathbf{A}_T = [\mathbf{I}_d | \mathbf{A} | \mathbf{T} \mathbf{G}_H - (\mathbf{R}_1 + \mathbf{A} \mathbf{R}_2)]$, \mathbf{G}_H is a truncated gadget matrix and \mathbf{T} , \mathbf{R}_1 and \mathbf{R}_2 are parameters. To this end, we introduce the `GenericSampler` that generates short \mathbf{v}' such that $\mathbf{A}_T \mathbf{K} \mathbf{v}' = \mathbf{u}$, where \mathbf{K} is defined below. In other words, $\mathbf{K} \mathbf{v}'$ is exactly the vector \mathbf{x} we seek, assuming some mild conditions on \mathbf{K} .

Algorithm 3.1: `GenericSampler`($\mathbf{R}_1, \mathbf{R}_2, \mathbf{A}, \mathbf{u}, \mathbf{T}, \mathbf{K}, \mathbf{L}, \mathbf{S}, \mathbf{S}_G, \mathbf{G}_H$)

Input: trapdoor $\mathbf{R}_1, \mathbf{R}_2 \in \mathbb{Z}^{d \times d(k-\ell)}$, matrix $\mathbf{A} \in \mathbb{Z}_q^{d \times d}$, syndrome $\mathbf{u} \in \mathbb{Z}_q^d$, covariance matrices \mathbf{S}, \mathbf{S}_G , tag $\mathbf{T} \in GL_d(\mathbb{Z}_q)$, matrices $\mathbf{K} \in \mathbb{Z}^{d(2+k-\ell) \times r}$ and $\mathbf{L} \in \mathbb{Z}^{r \times dk}$ such that $\mathbf{K} \mathbf{L} = \begin{bmatrix} \mathbf{T} \mathbf{G}_L & \mathbf{R}_1 \\ \mathbf{0} & \mathbf{R}_2 \\ \mathbf{0} & \mathbf{I}_{d(k-\ell)} \end{bmatrix} \pmod{q\mathbb{Z}}$, matrix $\mathbf{G}_H \in \mathbb{Z}^{d \times d(k-\ell)}$ such that $\mathbf{G} = [\mathbf{G}_L | \mathbf{G}_H]$ is primitive.

1. $\mathbf{p} \leftarrow \mathcal{D}_{\mathbb{Z}^r, \sqrt{\mathbf{S}_p}}$. $\triangleright \mathbf{S}_p = \mathbf{S} - \mathbf{L} \mathbf{S}_G \mathbf{L}^T$
2. $\mathbf{w} \leftarrow \mathbf{T}^{-1}(\mathbf{u} - [\mathbf{I}_d | \mathbf{A} | \mathbf{T} \mathbf{G}_H - (\mathbf{R}_1 + \mathbf{A} \mathbf{R}_2)] \mathbf{K} \mathbf{p}) \pmod{q\mathbb{Z}}$.
3. $\mathbf{z} \leftarrow \mathcal{D}_{\mathcal{L}_q^w(\mathbf{G}), \sqrt{\mathbf{S}_G}}$.
4. $\mathbf{v}' \leftarrow \mathbf{p} + \mathbf{L} \mathbf{z}$.

Output: \mathbf{v}' .

One can note that `GenericSampler` provides some flexibility in the choice of \mathbf{K} and \mathbf{L} as we only enforce a condition on their product $\mathbf{K} \mathbf{L}$. Looking ahead, this choice should however take into account the following two facts:

- $\mathbf{K} \mathbf{v}'$ will be made public in concrete applications so \mathbf{K} should not depend on secret information;
- A matrix \mathbf{L} with large norm will require a large perturbation \mathbf{p} , leading `GenericSampler` to produce larger \mathbf{v}' .

Roughly speaking, one should then try to move as much elements as possible in \mathbf{K} and only use \mathbf{L} to hide secret information. The latter will not be leaked by the distribution \mathcal{P} produced by the sampler, as formally stated by the theorem below under certain conditions on $\mathbf{K}, \mathbf{L}, \mathbf{S}, \mathbf{S}_G$ which we clearly identify. In particular, we show that \mathcal{P} is close to $\mathcal{D}_{\mathcal{L}_q^u(\mathbf{A}_T \mathbf{K}), \sqrt{\mathbf{S}}}$, hence such that $\mathbf{K} \text{Supp}(\mathcal{P}) = \mathcal{L}_q^u(\mathbf{A}_T)$.

Theorem 3.1. *Let d, q, k, ℓ, r be positive integers with $\ell < k$, and $\varepsilon \in (0, 1)$. Let $\mathbf{R}_1, \mathbf{R}_2$ be in $\mathbb{Z}^{d \times d(k-\ell)}$, $\mathbf{A} \in \mathbb{Z}_q^{d \times d}$, $\mathbf{u} \in \mathbb{Z}_q^d$ and $\mathbf{T} \in GL_d(\mathbb{Z}_q)$. Let $\mathbf{S} \in \mathcal{S}_r^+(\mathbb{R})$ and $\mathbf{S}_G \in \mathcal{S}_{dk}^{++}(\mathbb{R})$. We let $\mathbf{K} \in \mathbb{Z}^{d(2+k-\ell) \times r}$ and $\mathbf{L} \in \mathbb{Z}^{r \times dk}$. We denote by \mathcal{P} the output distribution of `GenericSampler`($\mathbf{R}_1, \mathbf{R}_2, \mathbf{A}, \mathbf{u}, \mathbf{T}, \mathbf{K}, \mathbf{L}, \mathbf{S}, \mathbf{S}_G, \mathbf{G}_H$). We assume the following 5 conditions.*

- ① $\mathbf{KL} = \begin{bmatrix} \mathbf{TG}_L & \mathbf{R}_1 \\ \mathbf{0} & \mathbf{R}_2 \\ \mathbf{0} & \mathbf{I}_{d(k-\ell)} \end{bmatrix} \bmod q\mathbb{Z}$
- ② $\text{rank}(\mathbf{L}) = dk$
- ③ $\mathbf{S}_G \geq \eta_\varepsilon (\mathcal{L}_q^\perp(\mathbf{G}))^2 \mathbf{I}_{dk}$
- ④ $\mathbf{S} > \mathbf{L} \mathbf{S}_G \mathbf{L}^T$
- ⑤ $\mathbf{S} \geq \mathbf{L} \left(\mathbf{S}_G + \left(\frac{1}{\eta_\varepsilon (\mathbb{Z}^{dk})^2} \mathbf{I}_{dk} - \mathbf{S}_G^{-1} \right)^{-1} \right) \mathbf{L}^T$

Under these conditions, it holds that $\text{Supp}(\mathcal{P}) = \mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_T \mathbf{K})$, and

$$\forall \mathbf{x} \in \mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_T \mathbf{K}), \mathcal{P}(\mathbf{x}) \in \left[\left(\frac{1-\varepsilon}{1+\varepsilon} \right)^2, \left(\frac{1+\varepsilon}{1-\varepsilon} \right)^2 \right] \mathcal{D}_{\mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_T \mathbf{K}), \sqrt{\mathbf{S}}}(\mathbf{x}),$$

where $\mathbf{A}_T = [\mathbf{I}_d | \mathbf{A} | \mathbf{TG}_H - (\mathbf{R}_1 + \mathbf{AR}_2)] \bmod q\mathbb{Z}$. It then holds that $\mathbf{K}\mathcal{P}$ is close to $\mathbf{KD}_{\mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_T \mathbf{K}), \sqrt{\mathbf{S}}}$ which is supported on $\mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_T)$.

One of the main pitfalls of our proof is to handle pseudoinverses which can sometimes defeat the intuition we may have with inverses. In particular, we insist that $(\mathbf{AB})^+$ is not equal to $\mathbf{B}^+ \mathbf{A}^+$ in general, and that $\mathbf{A} \geq \mathbf{B}$ does not necessarily imply $\mathbf{B}^+ \geq \mathbf{A}^+$ either. This prevents us from using some of the key arguments of the proof of [MP12, Thm. 5.5] where bounds on $\Sigma_{\mathbf{p}}^+$ and $\Sigma_{\mathbf{y}}^+$ are derived from those on $\Sigma_{\mathbf{p}}$ and $\Sigma_{\mathbf{y}}$.

Proof. We start by checking that the support of \mathcal{P} is indeed $\mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_T \mathbf{K})$.

Let $\mathbf{x} = \mathbf{p} + \mathbf{Lz}$ be in $\text{Supp}(\mathcal{P})$ outputted by the sampler. By construction, $\mathbf{z} \in \mathcal{L}_q^{\mathbf{w}}(\mathbf{G})$ with $\mathbf{w} = \mathbf{T}^{-1}(\mathbf{u} - \mathbf{A}_T \mathbf{Kp}) \bmod q\mathbb{Z}$. As result, we have $\mathbf{A}_T \mathbf{Kx} = \mathbf{A}_T \mathbf{Kp} + \mathbf{A}_T \mathbf{KLz} \bmod q\mathbb{Z}$. By condition ①, it holds that

$$\mathbf{A}_T \mathbf{KL} = [\mathbf{TG}_L | \mathbf{R}_1 + \mathbf{AR}_2 + (\mathbf{TG}_H - (\mathbf{R}_1 + \mathbf{AR}_2))] = \mathbf{TG} \bmod q\mathbb{Z}$$

Hence, we have $\mathbf{A}_T \mathbf{Kx} = \mathbf{A}_T \mathbf{Kp} + \mathbf{TGz} \bmod q\mathbb{Z} = \mathbf{u} \bmod q\mathbb{Z}$. So $\mathbf{x} \in \mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_T \mathbf{K})$. Reciprocally, let $\mathbf{x} \in \mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_T \mathbf{K})$. Let $\mathbf{z} \in \mathbb{Z}^{dk}$ and define $\mathbf{p} = \mathbf{x} - \mathbf{Lz}$. Then $\mathbf{p} \in \mathbb{Z}^r$ because \mathbf{L} has integer coefficients and $\mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_T \mathbf{K}) \subseteq \mathbb{Z}^r$, and we can check that $\mathbf{TGz} = \mathbf{u} - \mathbf{A}_T \mathbf{Kp} \bmod q\mathbb{Z}$. Indeed

$$\mathbf{TGz} = \mathbf{A}_T \mathbf{KLz} = \mathbf{A}_T \mathbf{K}(\mathbf{x} - \mathbf{p}) = \mathbf{A}_T \mathbf{Kx} - \mathbf{A}_T \mathbf{Kp} = \mathbf{u} - \mathbf{A}_T \mathbf{Kp} \bmod q\mathbb{Z}$$

as desired. So because \mathbf{p} is in the support of $\mathcal{D}_{\mathbb{Z}^r, \sqrt{\mathbf{S}_p}}$ and \mathbf{z} in that of $\mathcal{D}_{\mathcal{L}_q^{\mathbf{w}}(\mathbf{G}), \sqrt{\mathbf{G}}}$, we indeed have that $\mathbf{p} + \mathbf{Lz} = \mathbf{x}$ is in $\text{Supp}(\mathcal{P})$. Hence $\text{Supp}(\mathcal{P}) = \mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_T \mathbf{K})$. We also note that it is direct to see that $\mathbf{Kx} \in \mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_T)$ if and only if $\mathbf{x} \in \mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_T \mathbf{K})$.

We now look at the distribution. First, let us define $\mathbf{S}_y = \mathbf{L}\mathbf{S}_G\mathbf{L}^T$, and $\mathbf{S}_p = \mathbf{S} - \mathbf{S}_y$. Also, let $\mathbf{P} = \mathbf{L}\mathbf{L}^+$, and define $\mathbf{S}_3 = (\mathbf{P}(\mathbf{S}_p^+ + \mathbf{S}_y^+)\mathbf{P})^+$. We finally define $V = \text{Span}_{\mathbb{R}}(\mathbf{L}) \subset \mathbb{R}^r$. By properties of the pseudo-inverse, $\mathbf{P} = \mathbf{L}\mathbf{L}^+$ is the orthogonal projector onto V . Condition ② yields that \mathbf{L} has linearly independent columns, and we can therefore express $\mathbf{L}^+ = (\mathbf{L}^T\mathbf{L})^{-1}\mathbf{L}^T$, and most importantly get $\mathbf{L}^+\mathbf{L} = \mathbf{I}_{dk}$. We then define $\mathcal{L} = \mathbb{Z}^r \cap V$. It holds that $\mathcal{L}(\mathbf{L}) \subseteq \mathcal{L}$ and $dk = \text{rank}(\mathcal{L}(\mathbf{L})) = \text{rank}(\mathcal{L})$.

Note that condition ② yields that \mathbf{L} can indeed be used to define a lattice $\mathcal{L}(\mathbf{L})$ because it has linearly independent columns. As such $\mathcal{L}(\mathbf{L}) = \mathbf{L}\mathbb{Z}^{dk} \subset \mathbf{L}\mathbb{R}^{dk} = V$. Additionally, $\mathbf{L} \in \mathbb{Z}^{r \times dk}$ is integral and therefore $\mathbf{L}\mathbb{Z}^{dk} \subseteq \mathbb{Z}^r$. Hence $\mathcal{L}(\mathbf{L}) \subseteq \mathcal{L}$.

Because \mathbf{L} has integer coefficients, it holds that its columns are in \mathcal{L} and thus that $V = \text{Span}_{\mathbb{R}}(\mathbf{L}) \subseteq \text{Span}_{\mathbb{R}}(\mathcal{L})$. As a result, we directly get $dk \leq \text{rank}(\mathcal{L})$. Also, due to how \mathcal{L} is defined, $\text{rank}(\mathcal{L}) \leq dk$. Indeed, assume towards contradiction that $\text{rank}(\mathcal{L}) > dk$. Then, there exists $dk + 1$ vectors $\mathbf{v}_1, \dots, \mathbf{v}_{dk+1}$ of \mathcal{L} that are linearly independent. Yet $\mathcal{L} \subset V$ so $(\mathbf{v}_1, \dots, \mathbf{v}_{dk+1})$ is a linearly independent family in a vector space of dimension dk which is a contradiction. We conclude that $dk = \text{rank}(\mathcal{L}(\mathbf{L})) = \text{rank}(\mathcal{L})$.

As a result, we get that $\text{Span}_{\mathbb{R}}(\mathcal{L}) = V$. It also holds that $\text{Span}_{\mathbb{R}}(\mathbf{S}_y) = V$ where we recall that $\mathbf{S}_y = \mathbf{L}\mathbf{S}_G\mathbf{L}^T$. Additionally, we have³ $\mathbf{S}_3 = \mathbf{L}(\mathbf{L}^T\mathbf{S}_p^{-1}\mathbf{L} + \mathbf{S}_G^{-1})^{-1}\mathbf{L}^T$, and therefore $\text{Span}_{\mathbb{R}}(\mathbf{S}_3) = V$.

First, $\mathbf{S}_y = \mathbf{L}(\mathbf{S}_G\mathbf{L}^T)$ so it directly follows that $\text{Span}_{\mathbb{R}}(\mathbf{S}_y) \subseteq V$. Reciprocally, let $\mathbf{x} = \mathbf{L}\mathbf{y} \in V$ for some $\mathbf{y} \in \mathbb{R}^{dk}$. Because of condition ③, it holds that \mathbf{S}_G is positive definite and therefore invertible. So $\mathbf{x} = \mathbf{L}\mathbf{y} = \mathbf{L}\mathbf{S}_G(\mathbf{S}_G^{-1}\mathbf{y})$. Then, using the fact that $\mathbf{L}^+\mathbf{L} = \mathbf{I}_{dk}$, we have $\mathbf{S}_G^{-1}\mathbf{y} = \mathbf{L}^T\mathbf{L}^+\mathbf{S}_G^{-1}\mathbf{y}$. Defining $\mathbf{z} = \mathbf{L}^+\mathbf{S}_G^{-1}\mathbf{y}$, it holds that $\mathbf{x} = \mathbf{L}\mathbf{S}_G\mathbf{L}^T\mathbf{z}$ which belongs to $\text{Span}_{\mathbb{R}}(\mathbf{S}_y)$, proving that $V \subseteq \text{Span}_{\mathbb{R}}(\mathbf{S}_y)$.

To prove the expression of \mathbf{S}_3 , we start by showing that $\mathbf{P}\mathbf{S}_y^+\mathbf{P} = \mathbf{S}_y^+$. For that we first prove that $\mathbf{X} = \mathbf{L}^+\mathbf{S}_G^{-1}\mathbf{L}^+$ is the Moore-Penrose pseudoinverse of \mathbf{S}_y . We now verify the four Moore-Penrose conditions.

1. $\mathbf{S}_y\mathbf{X}\mathbf{S}_y = \mathbf{L}\mathbf{S}_G\mathbf{L}^T \cdot \mathbf{L}^+\mathbf{S}_G^{-1}\mathbf{L}^+ \cdot \mathbf{L}\mathbf{S}_G\mathbf{L}^T = \mathbf{L}\mathbf{S}_G\mathbf{L}^T$, using the fact that $\mathbf{L}^+\mathbf{L} = \mathbf{I}_{dk}$.
2. $\mathbf{X}\mathbf{S}_y\mathbf{X} = \mathbf{L}^+\mathbf{S}_G^{-1}\mathbf{L}^+ \cdot \mathbf{L}\mathbf{S}_G\mathbf{L}^T \cdot \mathbf{L}^+\mathbf{S}_G^{-1}\mathbf{L}^+ = \mathbf{L}^+\mathbf{S}_G^{-1}\mathbf{L}^+$, again due to $\mathbf{L}^+\mathbf{L} = \mathbf{I}_{dk}$.
- 3&4. $\mathbf{S}_y\mathbf{X} = \mathbf{L}\mathbf{S}_G\mathbf{L}^T \cdot \mathbf{L}^+\mathbf{S}_G^{-1}\mathbf{L}^+ = \mathbf{L}\mathbf{L}^+ = \mathbf{P}$ and $\mathbf{X}\mathbf{S}_y = \mathbf{L}^+\mathbf{S}_G^{-1}\mathbf{L}^+ \cdot \mathbf{L}\mathbf{S}_G\mathbf{L}^T = (\mathbf{L}\mathbf{L}^+)^T = \mathbf{P}^T = \mathbf{P}$. So $\mathbf{S}_y\mathbf{X}$ and $\mathbf{X}\mathbf{S}_y$ are indeed symmetric.

It shows that $\mathbf{X} = \mathbf{S}_y^+$. Hence $\mathbf{P}\mathbf{S}_y^+\mathbf{P} = \mathbf{P}^T\mathbf{S}_y^+\mathbf{P} = \mathbf{L}^+\mathbf{L}^T\mathbf{L}^+\mathbf{S}_G^{-1}\mathbf{L}^+\mathbf{L}\mathbf{L}^+ = \mathbf{L}^+\mathbf{S}_G^{-1}\mathbf{L}^+$ by property of the pseudoinverse. So $\mathbf{P}\mathbf{S}_y^+\mathbf{P} = \mathbf{S}_y^+$. Let us now look at the expression of \mathbf{S}_3 . We note that in most cases, we do not have

³ We insist that in general $(\mathbf{A}\mathbf{B})^+ \neq \mathbf{B}^+\mathbf{A}^+$ which requires care in computing \mathbf{S}_3 .

$\mathbf{S}_3 = \mathbf{P}^+(\mathbf{S}_p^{-1} + \mathbf{S}_y^+)^+\mathbf{P}^+$. Instead, we rely on the fact that $\mathbf{P} = \mathbf{L}\mathbf{L}^+$ and $\mathbf{P} = \mathbf{P}^T = \mathbf{L}^{+T}\mathbf{L}^T$ to “factor out” \mathbf{L} and \mathbf{L}^T instead of \mathbf{P} . It holds that

$$\begin{aligned}\mathbf{S}_3 &= (\mathbf{P}(\mathbf{S}_p^{-1} + \mathbf{S}_y^+)\mathbf{P})^+ \\ &= (\mathbf{L}^{+T}\mathbf{L}^T\mathbf{S}_p^{-1}\mathbf{L}\mathbf{L}^+ + \mathbf{L}^{+T}\mathbf{S}_G^{-1}\mathbf{L}^+)^+ \\ &= (\mathbf{L}^{+T} \cdot (\mathbf{L}^T\mathbf{S}_p^{-1}\mathbf{L} + \mathbf{S}_G^{-1}) \cdot \mathbf{L}^+)^+\end{aligned}$$

We define $\mathbf{S}' = \mathbf{L}^T\mathbf{S}_p^{-1}\mathbf{L} + \mathbf{S}_G^{-1}$. It clearly holds that \mathbf{S}' is symmetric, and positive definite. Indeed, for $\mathbf{x} \neq \mathbf{0}$, $\mathbf{x}^T\mathbf{S}'\mathbf{x} \geq \mathbf{x}^T\mathbf{S}_G^{-1}\mathbf{x} > 0$. It is therefore invertible. Using the same method we used to derive \mathbf{S}_y^+ , we can show that $(\mathbf{L}^{+T}\mathbf{S}'\mathbf{L}^+)^+ = \mathbf{L}\mathbf{S}'^{-1}\mathbf{L}^T$ using the fact that $\mathbf{L}^+\mathbf{L} = \mathbf{I}_{dk}$ and \mathbf{S}' invertible. So $\mathbf{S}_3 = \mathbf{L}(\mathbf{L}^T\mathbf{S}_p^{-1}\mathbf{L} + \mathbf{S}_G^{-1})^{-1}\mathbf{L}^T$. Finally, using the same calculation we did for proving that $\text{Span}_{\mathbb{R}}(\mathbf{S}_y) = V$, we obtain that $\text{Span}_{\mathbb{R}}(\mathbf{S}_3) = V$ (by virtually substituting \mathbf{S}_G with \mathbf{S}'^{-1} in the above calculation).

Now let $\mathbf{x} \in \mathcal{L}_q^u(\mathbf{A}_T\mathbf{K}) = \text{Supp}(\mathcal{P})$ be a possible output of the sampler, i.e., there exists $\mathbf{p} \in \mathbb{Z}^r$ and $\mathbf{z} \in \mathcal{L}_q^w(\mathbf{G}) \subseteq \mathbb{Z}^{dk}$ such that $\mathbf{x} = \mathbf{p} + \mathbf{L}\mathbf{z}$. Then, we have $\mathbf{p} = \mathbf{x} - \mathbf{L}\mathbf{z} \in \mathbf{x} + \mathcal{L}(\mathbf{L}) \subseteq \mathbf{x} + \mathcal{L}$. Also, once \mathbf{p} and \mathbf{x} are fixed, there exists a unique \mathbf{z}' such that $\mathbf{x} = \mathbf{p} + \mathbf{L}\mathbf{z}'$, i.e., $\mathbf{z}' = \mathbf{z}$. This is because \mathbf{L} has linearly independent columns. This link between, $\mathbf{x}, \mathbf{p}, \mathbf{z}$ also entails that

$$\rho_{\sqrt{\mathbf{S}_G}}(\mathbf{z}) = \rho_{\sqrt{\mathbf{S}_y}}(\mathbf{x} - \mathbf{p}). \quad (1)$$

By the expression derived for \mathbf{S}_y^+ , we have $\mathbf{L}^T\mathbf{S}_y^+\mathbf{L} = \mathbf{L}^T\mathbf{L}^{+T}\mathbf{S}_G^{-1}\mathbf{L}^+\mathbf{L} = \mathbf{S}_G^{-1}$. It yields

$$\begin{aligned}\rho_{\sqrt{\mathbf{S}_y}}(\mathbf{x} - \mathbf{p}) &= \begin{cases} \exp(-\pi(\mathbf{x} - \mathbf{p})^T\mathbf{S}_y^+(\mathbf{x} - \mathbf{p})) & \text{if } \mathbf{x} - \mathbf{p} \in \text{Span}_{\mathbb{R}}(\mathbf{S}_y) \\ 0 & \text{otherwise} \end{cases} \\ &= \begin{cases} \exp(-\pi\mathbf{z}^T\mathbf{L}\mathbf{S}_y^+\mathbf{L}\mathbf{z}) & \text{if } \mathbf{L}\mathbf{z} \in \text{Span}_{\mathbb{R}}(\mathbf{S}_y) \\ 0 & \text{otherwise} \end{cases} \\ &= \begin{cases} \exp(-\pi\mathbf{z}^T\mathbf{S}_G^{-1}\mathbf{z}) & \text{if } \mathbf{L}\mathbf{z} \in \text{Span}_{\mathbb{R}}(\mathbf{S}_y) \\ 0 & \text{otherwise} \end{cases}\end{aligned}$$

Because $\text{Span}_{\mathbb{R}}(\mathbf{S}_y) = V$, $\mathbf{L}\mathbf{z}$ always belongs to $\text{Span}_{\mathbb{R}}(\mathbf{S}_y)$, meaning that the 0 case never occurs. The last quantity then corresponds exactly to $\rho_{\sqrt{\mathbf{S}_G}}(\mathbf{z})$, and thence $\rho_{\sqrt{\mathbf{S}_y}}(\mathbf{x} - \mathbf{p}) = \rho_{\sqrt{\mathbf{S}_G}}(\mathbf{z})$.

We now have

$$\begin{aligned}
\mathcal{P}(\mathbf{x}) &= \sum_{\mathbf{p} \in \mathbf{x} + \mathcal{L}} \mathcal{D}_{\mathbb{Z}^n, \sqrt{\mathbf{S}_p}}(\mathbf{p}) \mathcal{D}_{\mathcal{L}_q^w(\mathbf{G}), \sqrt{\mathbf{S}_G}}(\mathbf{z}) \\
&= \sum_{\mathbf{p} \in \mathbf{x} + \mathcal{L}} \frac{\rho_{\sqrt{\mathbf{S}_p}}(\mathbf{p}) \rho_{\sqrt{\mathbf{S}_G}}(\mathbf{z})}{\rho_{\sqrt{\mathbf{S}_p}}(\mathbb{Z}^n) \rho_{\sqrt{\mathbf{S}_G}}(\mathcal{L}_q^w(\mathbf{G}))} \\
&= \sum_{\mathbf{p} \in \mathbf{x} + \mathcal{L}} \frac{\rho_{\sqrt{\mathbf{S}_p}}(\mathbf{p}) \rho_{\sqrt{\mathbf{S}_y}}(\mathbf{p} - \mathbf{x})}{\rho_{\sqrt{\mathbf{S}_p}}(\mathbb{Z}^n) \rho_{\sqrt{\mathbf{S}_G}}(\mathcal{L}_q^w(\mathbf{G}))} \tag{Eq. (1)}
\end{aligned}$$

We observe that condition [4](#) implies that $\mathbf{S}_p = \mathbf{S} - \mathbf{L}\mathbf{S}_G\mathbf{L}^T > 0$. As a result, \mathbf{S}_p is invertible and therefore $\text{Span}_{\mathbb{R}}(\mathbf{S}_p) = \mathbb{R}^n$. Then by [\[MP12, Fact 5.6\]](#) applied to $\Sigma_1 = \mathbf{S}_p$, $\Sigma_2 = \mathbf{S}_y$, $V_1 = \mathbb{R}^n$, $V_2 = V$ and $V_3 = V$, we have

$$\rho_{\sqrt{\mathbf{S}_p}}(\mathbf{p}) \rho_{\sqrt{\mathbf{S}_y}}(\mathbf{p} - \mathbf{x}) = \rho_{\sqrt{\mathbf{S}}}(\mathbf{x}) \rho_{\sqrt{\mathbf{S}_3}}(\mathbf{p} - \mathbf{c}),$$

where $\mathbf{c} \in \mathbf{v} + V$ such that $\mathbf{S}_3^+(\mathbf{c} - \mathbf{v}) = \mathbf{S}_p^+(\mathbf{0} - \mathbf{v}) + \mathbf{S}_y^+(\mathbf{x} - \mathbf{v})$ and with \mathbf{v} the unique element of $(\mathbf{x} + V) \cap V^\perp$. More precisely, we have $\mathbf{v} = \mathbf{x} - \mathbf{P}\mathbf{x}$, and

$$\begin{aligned}
\mathbf{c} &= \mathbf{c} - \mathbf{P}\mathbf{c} + \mathbf{P}\mathbf{c} \\
&= \mathbf{v} + \mathbf{S}_3\mathbf{S}_3^+\mathbf{c} \\
&= \mathbf{v} + \mathbf{S}_3(\mathbf{S}_y^+(\mathbf{x} - \mathbf{v}) - \mathbf{S}_p^+\mathbf{v}) \\
&= \mathbf{v} + \mathbf{S}_3(\mathbf{S}_y^+\mathbf{x} - \mathbf{S}_p^+\mathbf{v}) \\
&= \mathbf{S}_3\mathbf{S}_y^+\mathbf{x} + (\mathbf{I}_n - \mathbf{S}_3\mathbf{S}_p^+)(\mathbf{x} - \mathbf{P}\mathbf{x})
\end{aligned}$$

The second equality holds due to the fact that $\mathbf{c} - \mathbf{P}\mathbf{c} \in (\mathbf{v} + V) \cap V^\perp = (\mathbf{x} + V) \cap V^\perp$. So $\mathbf{c} - \mathbf{P}\mathbf{c} = \mathbf{v} = \mathbf{x} - \mathbf{P}\mathbf{x}$. It also uses the fact that $\mathbf{S}_3\mathbf{S}_3^+$ is the orthogonal projector onto $\text{Span}_{\mathbb{R}}(\mathbf{S}_3) = V$, which by uniqueness implies $\mathbf{S}_3\mathbf{S}_3^+ = \mathbf{P}$. The third equality holds by the equality verified by \mathbf{c} and the fact that $\mathbf{v} \in V^\perp = \ker(\mathbf{S}_3^+) = \ker(\mathbf{S}_y^+)$. Going back to our computation of $\mathcal{P}(\mathbf{x})$, we have

$$\begin{aligned}
\mathcal{P}(\mathbf{x}) &= \sum_{\mathbf{p} \in \mathbf{x} + \mathcal{L}} \frac{\rho_{\sqrt{\mathbf{S}}}(\mathbf{x}) \rho_{\sqrt{\mathbf{S}_3}}(\mathbf{p} - \mathbf{c})}{\rho_{\sqrt{\mathbf{S}_p}}(\mathbb{Z}^n) \rho_{\sqrt{\mathbf{S}_G}}(\mathcal{L}_q^w(\mathbf{G}))} \\
&= \mathcal{D}_{\mathcal{L}_q^u(\mathbf{A}_T\mathbf{K}), \sqrt{\mathbf{S}}}(\mathbf{x}) \frac{\rho_{\sqrt{\mathbf{S}}}(\mathcal{L}_q^u(\mathbf{A}_T\mathbf{K}))}{\rho_{\sqrt{\mathbf{S}_p}}(\mathbb{Z}^n)} \sum_{\mathbf{p} \in \mathbf{x} + \mathcal{L}} \frac{\rho_{\sqrt{\mathbf{S}_3}}(\mathbf{p} - \mathbf{c})}{\rho_{\sqrt{\mathbf{S}_G}}(\mathcal{L}_q^w(\mathbf{G}))} \\
&\in \left[1, \frac{1 + \varepsilon}{1 - \varepsilon}\right] \mathcal{D}_{\mathcal{L}_q^u(\mathbf{A}_T\mathbf{K}), \sqrt{\mathbf{S}}}(\mathbf{x}) \frac{\rho_{\sqrt{\mathbf{S}}}(\mathcal{L}_q^u(\mathbf{A}_T\mathbf{K})) \rho_{\sqrt{\mathbf{S}_3}}(\mathcal{L} + \mathbf{x} - \mathbf{c})}{\rho_{\sqrt{\mathbf{S}_p}}(\mathbb{Z}^n) \rho_{\sqrt{\mathbf{S}_G}}(\mathcal{L}_q^\perp(\mathbf{G}))} \\
&\in \left[\frac{1 - \varepsilon}{1 + \varepsilon}, \frac{1 + \varepsilon}{1 - \varepsilon}\right] \mathcal{D}_{\mathcal{L}_q^u(\mathbf{A}_T\mathbf{K}), \sqrt{\mathbf{S}}}(\mathbf{x}) \frac{\rho_{\sqrt{\mathbf{S}}}(\mathcal{L}_q^u(\mathbf{A}_T\mathbf{K})) \rho_{\sqrt{\mathbf{S}_3}}(\mathcal{L})}{\rho_{\sqrt{\mathbf{S}_p}}(\mathbb{Z}^n) \rho_{\sqrt{\mathbf{S}_G}}(\mathcal{L}_q^\perp(\mathbf{G}))}
\end{aligned}$$

where the second to last equation uses condition (3) to argue that $\sqrt{\mathbf{S}_{\mathbf{G}}} \geq \eta_\varepsilon(\mathcal{L}_q^\perp(\mathbf{G}))$ combined with Lemma 2.4. The last equation is argued by condition (5), the expression of \mathbf{S}_3 , and Lemma 2.4 as well. We detail how below.

First, because \mathbf{G} is primitive, i.e., $\mathbf{G}\mathbb{Z}^{dk} = \mathbb{Z}^d$, it holds that $\mathcal{L}_q^\perp(\mathbf{G})$ is a strict sublattice of \mathbb{Z}^{dk} , of same rank dk . Inclusion is strict because \mathbf{G} being primitive implies $\text{Vol}(\mathcal{L}_q^\perp(\mathbf{G})) = q^d \neq 1 = \text{Vol}(\mathbb{Z}^{dk})$. By Lemma 2.5 and condition (3), it holds that $\mathbf{S}_{\mathbf{G}} \geq \eta_\varepsilon(\mathcal{L}_q^\perp(\mathbf{G}))^2 \mathbf{I}_{dk} > \eta_\varepsilon(\mathbb{Z}^{dk})^2 \mathbf{I}_{dk}$. As such, we have $\mathbf{M} := \frac{1}{\eta_\varepsilon(\mathbb{Z}^{dk})^2} \mathbf{I}_{dk} - \mathbf{S}_{\mathbf{G}}^{-1} > 0$. So \mathbf{M}^{-1} appearing in condition (5) indeed exists.

Condition (5) can then be written as $\mathbf{S}_{\mathbf{p}} \geq \mathbf{L}\mathbf{M}^{-1}\mathbf{L}^T$. We now let $\delta > 0$ be a positive real and $\mathbf{P}_\perp = \mathbf{I}_r - \mathbf{P}$ be the orthogonal projector onto $V^\perp = \ker(\mathbf{L}^T)$. Because \mathbf{P}_\perp is the matrix of an orthogonal projector, it holds that $\mathbf{P}_\perp \in \mathcal{S}_r^+(\mathbb{R})$. We then have

$$\mathbf{S}_{\mathbf{p},\delta} := \mathbf{S}_{\mathbf{p}} + \delta\mathbf{P}_\perp \geq \mathbf{L}\mathbf{M}^{-1}\mathbf{L}^T + \delta\mathbf{P}_\perp. \quad (2)$$

Because $\delta\mathbf{P}_\perp \in \mathcal{S}_r^+(\mathbb{R})$, we have $\mathbf{S}_{\mathbf{p},\delta} \geq \mathbf{S}_{\mathbf{p}} > 0$, where the last inequality comes from condition (4). We now show that $\mathbf{L}\mathbf{M}^{-1}\mathbf{L}^T + \delta\mathbf{P}_\perp$ is invertible. Let $\mathbf{x} \in \ker(\mathbf{L}\mathbf{M}^{-1}\mathbf{L}^T + \delta\mathbf{P}_\perp)$ and decompose it uniquely onto $V \oplus V^\perp$ as $\mathbf{x} = \mathbf{x}_V + \mathbf{x}_{V^\perp}$ with $\mathbf{x}_V \in V$ and $\mathbf{x}_{V^\perp} \in V^\perp$. We then have

$$\begin{aligned} \mathbf{0} &= (\mathbf{L}\mathbf{M}^{-1}\mathbf{L}^T + \delta\mathbf{P}_\perp)\mathbf{x} \\ &= \mathbf{L}\mathbf{M}^{-1}\mathbf{L}^T(\mathbf{x}_V + \mathbf{x}_{V^\perp}) + \delta\mathbf{P}_\perp(\mathbf{x}_V + \mathbf{x}_{V^\perp}) \\ &= \mathbf{L}\mathbf{M}^{-1}\mathbf{L}^T(\mathbf{x}_V + \mathbf{x}_{V^\perp}) + \delta\mathbf{x}_{V^\perp} \\ &= \mathbf{L}\mathbf{M}^{-1}\mathbf{L}^T\mathbf{x}_V + \delta\mathbf{x}_{V^\perp}, \end{aligned}$$

where the second to last equality holds by definition of \mathbf{P}_\perp being the orthogonal projector onto V^\perp , and the last equality follows from the fact that $V^\perp = \ker(\mathbf{L}^T)$ and so $\mathbf{L}^T\mathbf{x}_{V^\perp} = \mathbf{0}$. Then, by definition of V , we have that $\mathbf{L}\mathbf{M}^{-1}\mathbf{L}^T\mathbf{x}_V \in V$, and also $\delta\mathbf{x}_{V^\perp} \in V^\perp$. Because V and V^\perp are complementary subspaces, the above equality implies that $\mathbf{L}\mathbf{M}^{-1}\mathbf{L}^T\mathbf{x}_V = \mathbf{0}$ and $\delta\mathbf{x}_{V^\perp} = \mathbf{0}$. It directly yields $\mathbf{x}_{V^\perp} = \mathbf{0}$. Using the fact that $\mathbf{L}^+\mathbf{L} = \mathbf{I}_{dk}$, we have $\mathbf{L}^T\mathbf{x}_V = \mathbf{M}\mathbf{L}^+\mathbf{0} = \mathbf{0}$. So $\mathbf{x}_V \in \ker(\mathbf{L}^T) = V^\perp$, which proves that $\mathbf{x}_V \in V \cap V^\perp = \{\mathbf{0}\}$. Hence, $\mathbf{x}_V = \mathbf{0}$ and therefore $\mathbf{x} = \mathbf{0}$. We have then proven that the kernel of $\mathbf{L}\mathbf{M}^{-1}\mathbf{L}^T + \delta\mathbf{P}_\perp$ is trivial, which proves that it is invertible.

Both matrices involved in Equation (2) being invertible, we get

$$\mathbf{S}_{\mathbf{p},\delta}^{-1} \leq (\mathbf{L}\mathbf{M}^{-1}\mathbf{L}^T + \delta\mathbf{P}_\perp)^{-1}.$$

It then implies

$$\mathbf{L}^T\mathbf{S}_{\mathbf{p},\delta}^{-1}\mathbf{L} \leq \mathbf{L}^T(\mathbf{L}\mathbf{M}^{-1}\mathbf{L}^T + \delta\mathbf{P}_\perp)^{-1}\mathbf{L}. \quad (3)$$

We now prove that the right-hand side is exactly \mathbf{M} . For that we show that for all $\mathbf{x} \in V$, $(\mathbf{L}\mathbf{M}^{-1}\mathbf{L}^T + \delta\mathbf{P}_\perp)^{-1}\mathbf{x} = \mathbf{L}^+\mathbf{M}\mathbf{L}^+\mathbf{x}$, or equivalently that for

all $\mathbf{x} \in V$, $\mathbf{x} = (\mathbf{L}\mathbf{M}^{-1}\mathbf{L}^T + \delta\mathbf{P}_\perp)\mathbf{L}^{+T}\mathbf{M}\mathbf{L}^+\mathbf{x}$. Let $\mathbf{x} \in V$. We first have

$$\mathbf{L}\mathbf{M}^{-1}\mathbf{L}^T\mathbf{L}^{+T}\mathbf{M}\mathbf{L}^+\mathbf{x} = \mathbf{L}\mathbf{M}^{-1}(\mathbf{L}^+\mathbf{L})^T\mathbf{M}\mathbf{L}^+\mathbf{x} = \mathbf{L}\mathbf{L}^+\mathbf{x} = \mathbf{P}\mathbf{x} = \mathbf{x},$$

where we use the fact that $\mathbf{L}^+\mathbf{L} = \mathbf{I}_{dk}$ and that $\mathbf{L}\mathbf{L}^+ = \mathbf{P}$ is the orthogonal projector onto V . For the second term, we notice that $\mathbf{L}^{+T}\mathbf{M}\mathbf{L}^+\mathbf{x}$ belongs to $\text{Span}_{\mathbb{R}}(\mathbf{L}^{+T}) = \text{Span}_{\mathbb{R}}(\mathbf{L}) = V$. Hence $\mathbf{P}_\perp\mathbf{L}^{+T}\mathbf{M}\mathbf{L}^+\mathbf{x} = \mathbf{0}$. As a result, we indeed obtain

$$\forall \mathbf{x} \in V, (\mathbf{L}\mathbf{M}^{-1}\mathbf{L}^T + \delta\mathbf{P}_\perp)^{-1}\mathbf{x} = \mathbf{L}^{+T}\mathbf{M}\mathbf{L}^+\mathbf{x}.$$

A direct consequence is that

$$\forall \mathbf{x} \in \mathbb{R}^{dk}, (\mathbf{L}\mathbf{M}^{-1}\mathbf{L}^T + \delta\mathbf{P}_\perp)^{-1}\mathbf{L}\mathbf{x} = \mathbf{L}^{+T}\mathbf{M}\mathbf{L}^+\mathbf{L}\mathbf{x} = \mathbf{L}^{+T}\mathbf{M}\mathbf{x},$$

again using $\mathbf{L}^+\mathbf{L} = \mathbf{I}_{dk}$. Left multiplying by \mathbf{L}^T yields $\mathbf{L}^T(\mathbf{L}\mathbf{M}^{-1}\mathbf{L}^T + \delta\mathbf{P}_\perp)^{-1}\mathbf{L} = \mathbf{M}$. Equation (3) then becomes $\mathbf{L}^T\mathbf{S}_{\mathbf{p},\delta}^{-1}\mathbf{L} \leq \mathbf{M}$, or equivalently $\mathbf{L}^T\mathbf{S}_{\mathbf{p},\delta}^{-1}\mathbf{L} + \mathbf{S}_{\mathbf{G}}^{-1} \leq \frac{1}{\eta_\varepsilon(\mathbb{Z}^{dk})^2}\mathbf{I}_{dk}$. Using the definition of the Loewner order, the fact that $\lim_{\delta \rightarrow 0}\mathbf{S}_{\mathbf{p},\delta} = \mathbf{S}_{\mathbf{p}}$ and the continuity of $\mathbf{A} \mapsto \mathbf{A}^{-1}$ over $GL_r(\mathbb{R})$, we get

$$\mathbf{L}^T\mathbf{S}_{\mathbf{p}}^{-1}\mathbf{L} + \mathbf{S}_{\mathbf{G}}^{-1} \leq \frac{1}{\eta_\varepsilon(\mathbb{Z}^{dk})^2}\mathbf{I}_{dk}. \quad (4)$$

We indeed have that for all $\mathbf{x} \in \mathbb{R}^{dk}$, $\mathbf{x}^T(\mathbf{L}^T\mathbf{S}_{\mathbf{p},\delta}^{-1}\mathbf{L} + \mathbf{S}_{\mathbf{G}}^{-1})\mathbf{x} \leq \frac{1}{\eta_\varepsilon(\mathbb{Z}^{dk})^2}\|\mathbf{x}\|_2^2$. We naturally have $\lim_{\delta \rightarrow 0}\mathbf{S}_{\mathbf{p},\delta} = \mathbf{S}_{\mathbf{p}}$ by construction, and by continuity $\lim_{\delta \rightarrow 0}\mathbf{S}_{\mathbf{p},\delta}^{-1} = \mathbf{S}_{\mathbf{p}}^{-1}$ because the limit $\mathbf{S}_{\mathbf{p}}$ is also invertible due to condition (4). Then, by compatibility of inequalities and limits, we get $\mathbf{x}^T(\mathbf{L}^T\mathbf{S}_{\mathbf{p}}^{-1}\mathbf{L} + \mathbf{S}_{\mathbf{G}}^{-1})\mathbf{x} \leq \frac{1}{\eta_\varepsilon(\mathbb{Z}^{dk})^2}\|\mathbf{x}\|_2^2$ as desired.

Earlier, we have shown that $\mathbf{S}_3 = \mathbf{L}\mathbf{S}'^{-1}\mathbf{L}^T$ for $\mathbf{S}' = \mathbf{L}^T\mathbf{S}_{\mathbf{p}}^{-1}\mathbf{L} + \mathbf{S}_{\mathbf{G}}^{-1}$. We now show that Equation (4) is equivalent to $\mathbf{y}^T\mathbf{S}_3\mathbf{y} \geq \eta_\varepsilon(\mathbb{Z}^{dk})^2\mathbf{y}^T\mathbf{L}\mathbf{L}^T\mathbf{y}$ for all \mathbf{y} . First, because \mathbf{S}' and $\eta_\varepsilon(\mathbb{Z}^{dk})^{-2}\mathbf{I}_{dk}$ are both positive definite and thus invertible, it holds that Equation (4) is equivalent to $\mathbf{S}'^{-1} \geq \eta_\varepsilon(\mathbb{Z}^{dk})^2\mathbf{I}_{dk}$. Because of condition (2), it holds that $\text{Span}_{\mathbb{R}}(\mathbf{L}^T) = \mathbb{R}^{dk}$. Hence, we have the following equivalences.

$$\begin{aligned} \frac{1}{\eta_\varepsilon(\mathbb{Z}^{dk})^2}\mathbf{I}_{dk} \geq \mathbf{S}' &\Leftrightarrow \mathbf{S}'^{-1} \geq \eta_\varepsilon(\mathbb{Z}^{dk})^2\mathbf{I}_{dk} \\ &\Leftrightarrow \forall \mathbf{x} \in \mathbb{R}^{dk}, \mathbf{x}^T(\mathbf{S}'^{-1} - \eta_\varepsilon(\mathbb{Z}^{dk})^2\mathbf{I}_{dk})\mathbf{x} \geq 0 \\ &\Leftrightarrow \forall \mathbf{y} \in \mathbb{R}^r, (\mathbf{L}^T\mathbf{y})^T(\mathbf{S}'^{-1} - \eta_\varepsilon(\mathbb{Z}^{dk})^2\mathbf{I}_{dk})(\mathbf{L}^T\mathbf{y}) \geq 0 \\ &\Leftrightarrow \forall \mathbf{y} \in \mathbb{R}^r, \mathbf{y}^T(\mathbf{L}\mathbf{S}'^{-1}\mathbf{L}^T - \eta_\varepsilon(\mathbb{Z}^{dk})^2\mathbf{L}\mathbf{L}^T)\mathbf{y} \geq 0 \\ &\Leftrightarrow \forall \mathbf{y} \in \mathbb{R}^r, \mathbf{y}^T\mathbf{S}_3\mathbf{y} \geq \eta_\varepsilon(\mathbb{Z}^{dk})^2\mathbf{y}^T\mathbf{L}\mathbf{L}^T\mathbf{y} \end{aligned}$$

As a result, we have the following

$$\begin{aligned}
1 + \varepsilon &= \sum_{\mathbf{y} \in \mathbb{Z}^{dk}} \exp(-\pi \eta_\varepsilon (\mathbb{Z}^{dk})^2 \mathbf{y}^T \mathbf{y}) \\
&= \sum_{\mathbf{y} \in \mathbb{Z}^{dk}} \exp(-\pi \eta_\varepsilon (\mathbb{Z}^{dk})^2 \mathbf{y}^T \mathbf{L}^+ \mathbf{L} \mathbf{L}^T \mathbf{L}^+ \mathbf{y}) \quad (\text{as } \mathbf{L}^+ \mathbf{L} = \mathbf{I}_{dk}) \\
&= \sum_{\mathbf{y} \in \mathbb{Z}^{dk}} \exp(-\pi \eta_\varepsilon (\mathbb{Z}^{dk})^2 (\mathbf{L}^+ \mathbf{y})^T \mathbf{L} \mathbf{L}^T (\mathbf{L}^+ \mathbf{y})) \\
&= \sum_{\mathbf{y} \in \mathcal{L}(\mathbf{L})^*} \exp(-\pi \eta_\varepsilon (\mathbb{Z}^{dk})^2 \mathbf{y}^T \mathbf{L} \mathbf{L}^T \mathbf{y}) \\
&\geq \sum_{\mathbf{y} \in \mathcal{L}^*} \exp(-\pi \eta_\varepsilon (\mathbb{Z}^{dk})^2 \mathbf{y}^T \mathbf{L} \mathbf{L}^T \mathbf{y}) \quad (\text{as } \mathcal{L}^* \subseteq \mathcal{L}(\mathbf{L})^*) \\
&\geq \sum_{\mathbf{y} \in \mathcal{L}^*} \exp(-\pi \mathbf{y}^T \mathbf{S}_3 \mathbf{y})
\end{aligned}$$

The second to last inequality is due to the fact that $\mathcal{L}^* \subseteq \mathcal{L}(\mathbf{L})^*$ as a consequence of $\mathcal{L}(\mathbf{L})$ being a sublattice of \mathcal{L} of same rank. Then, because $\mathcal{L}^* \subset V = \text{Span}_{\mathbb{R}}(\mathbf{S}_3) = \text{Span}_{\mathbb{R}}(\mathbf{S}_3^+)$, the last quantity is equal to $\rho_{\sqrt{\mathbf{S}_3^+}}(\mathcal{L}^*)$. Using Lemma 2.3, we thus have $\sqrt{\mathbf{S}_3} \geq \eta_\varepsilon(\mathcal{L})$. Then, we observe that $\mathbf{c} \in \mathbf{x} + V$ and therefore that $\mathbf{x} - \mathbf{c} \in V = \text{Span}_{\mathbb{R}}(\mathcal{L})$. Lemma 2.4 thus concludes the final equation.

We therefore have $\mathcal{P}(\mathbf{x}) \in [a, b] \cdot c \cdot \mathcal{D}_{\mathcal{L}_q^u(\mathbf{A}_T \mathbf{K}), \sqrt{\mathbf{S}}}(\mathbf{x})$. Summing over all $\mathbf{x} \in \mathcal{L}_q^u(\mathbf{A}_T \mathbf{K}) = \text{Supp}(\mathcal{P})$, it holds that $ac \leq 1 \leq bc$ and therefore that $c \in [1/b, 1/a]$. Hence, $\mathcal{P}(\mathbf{x}) \in [a/b, b/a] \mathcal{D}_{\mathcal{L}_q^u(\mathbf{A}_T \mathbf{K}), \sqrt{\mathbf{S}}}(\mathbf{x})$. Plugging $a = (1 - \varepsilon)/(1 + \varepsilon)$ and $b = (1 + \varepsilon)/(1 - \varepsilon)$ gives the result. \square

We note that because the inequality in condition ④ must be strict, it is not necessarily implied by condition ⑤ due to the fact that $\mathbf{L} \mathbf{M}^{-1} \mathbf{L}^T$ is not invertible. We later use $\mathbf{S}_G = s_G^2 \mathbf{I}_{dk}$. In that case condition ④ comes down to $\mathbf{S} > s_G^2 \mathbf{L} \mathbf{L}^T$ and condition ⑤ becomes $\mathbf{S} \geq \frac{s_G^4}{s_G^2 - \eta_\varepsilon(\mathbb{Z}^{dk})^2} \mathbf{L} \mathbf{L}^T$.

4 Instantiating the Generic Sampler

We now go over several instantiations of our generic sampler to showcase its full potential. As a warm-up, we start by showing that the Micciancio-Peikert sampler [MP12] is actually a specific case of our sampler. We even obtain a tighter analysis of [MP12] allowing for a small gain in parameters, which shows that our sampler did not trade performance for genericity.

In a second step, we show how to leverage the specific features of our sampler, and in particular the \mathbf{K} -projection, to truncate the gadget matrix \mathbf{G} and thereby improve performance. This provides in the process the first worst-case analysis of gadget-based samplers with truncated gadgets.

4.1 Tighter Analysis of the Micciancio-Peikert Sampler

To motivate the genericity of our sampler, we briefly show how to instantiate it to recover the sampler from [MP12]. We can indeed set $\ell = 0$ so that $\mathbf{G} = \mathbf{G}_H$ and $\mathbf{G}_L = \emptyset$. We then let $r = d(2+k)$ and $\mathbf{L} = [\mathbf{R}_1^T | \mathbf{R}_2^T | \mathbf{I}_{dk}]^T$ with $\mathbf{K} = \mathbf{I}_{d(2+k)}$. By setting $\mathbf{S} = s^2 \mathbf{I}_{d(2+k)}$ and $\mathbf{S}_G = s_G^2 \mathbf{I}_{dk}$, we recover the exact sampler from [MP12]. We now give the following corollary which yields tighter parameter conditions than [MP12] so as to meet the five conditions of Theorem 3.1.

Corollary 4.1. *Let d, q, b be positive integers such that $b \leq \sqrt{q}$, and let $k = \lceil \log_b q \rceil$. Let $\mathbf{R}_1, \mathbf{R}_2$ be in $\mathbb{Z}^{d \times dk}$, $\mathbf{A} \in \mathbb{Z}_q^{d \times d}$, $\mathbf{u} \in \mathbb{Z}_q^d$ and $\mathbf{T} \in GL_d(\mathbb{Z}_q)$. Let s, s_G be positive reals. We define $\mathbf{R} = [\mathbf{R}_1^T | \mathbf{R}_2^T]^T$ and $\mathbf{L} = [\mathbf{R}^T | \mathbf{I}_{dk}]^T$, and $\mathbf{G} = [\mathbf{I}_d | b \mathbf{I}_d | \dots | b^{k-1} \mathbf{I}_d]$. Let $\varepsilon \in (0, 1)$ and define \mathcal{P} the output distribution of $\text{GenericSampler}(\mathbf{R}_1, \mathbf{R}_2, \mathbf{A}, \mathbf{u}, \mathbf{T}, \mathbf{I}, \mathbf{L}, s^2 \mathbf{I}, s_G^2 \mathbf{I}, \mathbf{G})$. Assuming*

$$s_G \geq \eta_\varepsilon(\mathbb{Z}^{dk}) \sqrt{b^2 + 1}, \quad \text{and} \quad s \geq \frac{s_G^2}{\sqrt{s_G^2 - \eta_\varepsilon(\mathbb{Z}^{dk})^2}} \sqrt{1 + \|\mathbf{R}\|_2^2}$$

it holds that

$$\forall \mathbf{x} \in \mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_T), \mathcal{P}(\mathbf{x}) \in \left[\left(\frac{1-\varepsilon}{1+\varepsilon} \right)^2, \left(\frac{1+\varepsilon}{1-\varepsilon} \right)^2 \right] \mathcal{D}_{\mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_T), \sqrt{\mathbf{S}}(\mathbf{x})},$$

where $\mathbf{A}_T = [\mathbf{I}_d | \mathbf{A} | \mathbf{T} \mathbf{G} - (\mathbf{R}_1 + \mathbf{A} \mathbf{R}_2)] \bmod q \mathbb{Z}$.

Proof. First, due to how we set \mathbf{L} and \mathbf{K} , conditions ① and ② trivially hold. Then, by [MP12], $\mathcal{L}_q^\perp(\mathbf{G})$ has a basis \mathbf{B} whose Gram-Schmidt norms are at most $\sqrt{b^2 + 1}$. Using the smoothing bound of [GPV08, Thm. 3.1], it then holds that $\eta_\varepsilon(\mathcal{L}_q^\perp(\mathbf{G})) \leq \eta_\varepsilon(\mathbb{Z}^{dk}) \sqrt{b^2 + 1} \leq s_G$. Hence, by setting $\mathbf{S}_G = s_G^2 \mathbf{I}_{dk}$, it holds that condition ③ of Theorem 3.1 holds.

We let $n = d(2+k)$. There exists $\mathbf{U} \in O_n(\mathbb{R})$, $\mathbf{V} \in O_{dk}(\mathbb{R})$ and $\mathbf{D}' \in \mathbb{R}^{dk \times dk}$ diagonal such that

$$\mathbf{L} = \mathbf{U} \begin{bmatrix} \mathbf{D}' \\ \mathbf{0} \end{bmatrix} \mathbf{V}^T$$

Hence, $\mathbf{L} \mathbf{L}^T = \mathbf{U} \mathbf{D} \mathbf{U}^T$ with \mathbf{D} diagonal. By re-ordering without loss of generality, it holds that $\mathbf{D} = \text{diag}(1 + s_1(\mathbf{R})^2, \dots, 1 + s_{2d}(\mathbf{R})^2, 1, \dots, 1, 0, \dots, 0)$ where the 1 is repeated $d(k-2)$ times and the 0 is repeated $2d$ times, and that $\mathbf{D}' = \text{diag}(\sqrt{1 + s_1(\mathbf{R})^2}, \dots, \sqrt{1 + s_{2d}(\mathbf{R})^2}, 1, \dots, 1)$.

We can indeed compute the characteristic polynomial of $\mathbf{L} \mathbf{L}^T$ by using Schur's formula, see e.g. [Ber11, Fact. 2.14.13], where the computation is performed over the field $\mathbb{R}(X)$ of rational fractions to deal with inverses of matrices of

polynomials.

$$\begin{aligned}
\chi_{\mathbf{L}\mathbf{L}^T} &= \det \left(\begin{bmatrix} X\mathbf{I}_{2d} - \mathbf{R}\mathbf{R}^T & -\mathbf{R} \\ -\mathbf{R}^T & X\mathbf{I}_{dk} - \mathbf{I}_{dk} \end{bmatrix} \right) \\
&= \det((X-1)\mathbf{I}_{dk}) \det((X\mathbf{I}_{2d} - \mathbf{R}\mathbf{R}^T) - (-\mathbf{R})(X-1)^{-1}\mathbf{I}_{dk}(-\mathbf{R}^T)) \\
&= (X-1)^{dk} \det(X\mathbf{I}_{2d} - X(X-1)^{-1}\mathbf{R}\mathbf{R}^T) \\
&= X^{2d}(X-1)^{d(k-2)} \det((X-1)\mathbf{I}_{2d} - \mathbf{R}\mathbf{R}^T) \\
&= X^{2d}(X-1)^{d(k-2)} \chi_{\mathbf{I}_{2d} + \mathbf{R}\mathbf{R}^T} \\
&= X^{2d}(X-1)^{d(k-2)} \prod_{i \in [2d]} (X - (1 + s_i(\mathbf{R})^2))
\end{aligned}$$

Because $\mathbf{S} = \mathbf{U}\mathbf{S}\mathbf{U}^T$, it holds that condition ④ holds if and only if $s^2 > s_{\mathbf{G}}^2 \|\mathbf{D}\|_2 = s_{\mathbf{G}}^2 \sqrt{1 + \|\mathbf{R}\|_2^2}$, while condition ⑤ holds if and only if $s^2 \geq s_{\mathbf{G}}^4 / (s_{\mathbf{G}}^2 - \eta_\varepsilon(\mathbb{Z}^{dk})^2) \sqrt{1 + \|\mathbf{R}\|_2^2}$. In this case, because $s_{\mathbf{G}}^4 / (s_{\mathbf{G}}^2 - \eta_\varepsilon(\mathbb{Z}^{dk})^2) > s_{\mathbf{G}}^2$, both conditions are verified by our choice of s . \square

Setting $s_{\mathbf{G}}$ and s at their identified lower bounds gives $s = \eta_\varepsilon(\mathbb{Z}^{dk})(b + 1/b) \sqrt{1 + \|\mathbf{R}\|_2^2}$. The condition given in [MP12] is $s \geq r\sqrt{b^2 + 3} \sqrt{1 + \|\mathbf{R}\|_2^2}$ with $r \geq \eta_\varepsilon(\mathbb{Z}^{dk})$ for their last smoothing argument to hold. As a result, our condition is slightly tighter.

We can also study the specific case where $\mathbf{S} = \text{diag}(s_1^2 \mathbf{I}, s_2^2 \mathbf{I})$ which was used for example in [AGJ+24]. We determine the exact conditions on s_1, s_2 which are slightly different than those announced in [AGJ+24]. Nevertheless, our updated parameter conditions are very close and thus do not change the performance or security claims.

Corollary 4.2. *Let d, q, b be positive integers such that $b \leq \sqrt{q}$, $\varepsilon \in (0, 1)$, and let $k = \lceil \log_b q \rceil$. Let \mathbf{R} be in $\mathbb{Z}^{2d \times dk}$, $\mathbf{A} \in \mathbb{Z}_q^{d \times 2d}$, $\mathbf{u} \in \mathbb{Z}_q^d$ and $\mathbf{T} \in GL_d(\mathbb{Z}_q)$. We define $\mathbf{R} = [\mathbf{R}_1^T | \mathbf{R}_2^T]^T$, $\mathbf{L} = [\mathbf{R}^T | \mathbf{I}_{dk}]^T$, and $\mathbf{G} = [\mathbf{I}_d | b\mathbf{I}_d | \dots | b^{k-1}\mathbf{I}_d]$. We then set $s_{\mathbf{G}} = \eta_\varepsilon(\mathbb{Z}^{dk}) \sqrt{b^2 + 1}$ and*

$$s_1 = \sqrt{1 + \frac{1}{c^2 - 1}} \cdot \frac{s_{\mathbf{G}}^2}{\sqrt{s_{\mathbf{G}}^2 - \eta_\varepsilon(\mathbb{Z}^{dk})^2}} \|\mathbf{R}\|_2, \quad \text{and} \quad s_2 = c \frac{s_{\mathbf{G}}^2}{\sqrt{s_{\mathbf{G}}^2 - \eta_\varepsilon(\mathbb{Z}^{dk})^2}},$$

for some $c > 1$. We finally let $\mathbf{S} = \text{diag}(s_1^2 \mathbf{I}_{2d}, s_2^2 \mathbf{I}_{dk})$, and define \mathcal{P} the output distribution of $\text{GenericSampler}(\mathbf{R}_1, \mathbf{R}_2, \mathbf{A}, \mathbf{u}, \mathbf{T}, \mathbf{I}, \mathbf{L}, \mathbf{S}, s_{\mathbf{G}}^2 \mathbf{I}, \mathbf{G})$. It then holds that

$$\forall \mathbf{x} \in \mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_{\mathbf{T}}), \mathcal{P}(\mathbf{x}) \in \left[\left(\frac{1 - \varepsilon}{1 + \varepsilon} \right)^2, \left(\frac{1 + \varepsilon}{1 - \varepsilon} \right)^2 \right] \mathcal{D}_{\mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_{\mathbf{T}}), \sqrt{\mathbf{S}}}(\mathbf{x}),$$

where $\mathbf{A}_{\mathbf{T}} = [\mathbf{I}_d | \mathbf{A} | \mathbf{T}\mathbf{G} - (\mathbf{R}_1 + \mathbf{A}\mathbf{R}_2)] \bmod q\mathbb{Z}$.

Setting $c = \sqrt{2}$, we get $s_2 = \sqrt{2}\eta_\varepsilon(\mathbb{Z}^{dk})(b + 1/b)$ and $s_1 = \sqrt{2}\eta_\varepsilon(\mathbb{Z}^{dk})(b + 1/b)\|\mathbf{R}\|_2 = s_2\|\mathbf{R}\|_2$.

Proof. As for Corollary 4.1, conditions ①, ② and ③ are verified because the form of \mathbf{S} does not affect those. Let us now study conditions ④ and ⑤. There exists $\mathbf{U} \in O_{2d}(\mathbb{R})$, $\mathbf{V} \in O_{dk}(\mathbb{R})$ and $\mathbf{D}' = [\mathbf{D}|\mathbf{0}]$ where $\mathbf{D} = \text{diag}(r_1, \dots, r_{2d})$ be a singular value decomposition of \mathbf{R} , where $r_i = s_i(\mathbf{R})$ for clarity. We then straightforwardly have $\mathbf{L} = \text{diag}(\mathbf{U}, \mathbf{V})\mathbf{L}'\mathbf{V}^T$ with $\mathbf{L}' = [\mathbf{D}'^T|\mathbf{I}_{dk}]^T$. We call $\mathbf{W} = \text{diag}(\mathbf{U}, \mathbf{V})$. Because \mathbf{W} and \mathbf{V}^T are unitary, and because $\mathbf{S} = \mathbf{W}\mathbf{S}\mathbf{W}^T$, it directly holds that we can work directly on \mathbf{L}' instead of \mathbf{L} . More formally, condition ④ is equivalent to $\mathbf{S} > s_{\mathbf{G}}^2\mathbf{L}'\mathbf{L}'^T$ and condition ⑤ is equivalent to $\mathbf{S} \geq s_{\mathbf{G}}^4/(s_{\mathbf{G}}^2 - \eta_\varepsilon(\mathbb{Z}^{dk})^2)\mathbf{L}'\mathbf{L}'^T$. We thus prove both conditions at once by using a free parameter a . We define $\mathbf{S}'_{\mathbf{p}} = \mathbf{S} - a\mathbf{L}'\mathbf{L}'^T$. We can write it blockwise as follows.

$$\mathbf{S}'_{\mathbf{p}} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ s_1^2\mathbf{I}_{2d} - a\mathbf{D}^2 & -a\mathbf{D}' \\ -a\mathbf{D}'^T & (s_2^2 - a)\mathbf{I}_{dk} \\ \mathbf{B}^T & \mathbf{C} \end{bmatrix}$$

We note here that if $s_2^2 \leq a$, then $\mathbf{S}'_{\mathbf{p}}$ is not positive semi-definite. Hence, for both conditions we need to assume $s_2^2 > a$ in which case $\mathbf{C} > 0$. We can then use the Schur complement characterization, that is $\mathbf{S}'_{\mathbf{p}} > 0$ (resp. ≥ 0) if and only if $\mathbf{C} > 0$ and $\mathbf{S}'_{\mathbf{p}}/\mathbf{C} = \mathbf{A} - \mathbf{B}\mathbf{C}^{-1}\mathbf{B}^T > 0$ (resp. ≥ 0). We thus compute $\mathbf{S}'_{\mathbf{p}}/\mathbf{C} = s_1^2\mathbf{I}_{2d} - a\mathbf{D}^2 - a^2/(s_2^2 - a)\mathbf{D}'\mathbf{D}'^T = s_1^2\mathbf{I}_{2d} - s_2^2a/(s_2^2 - a)\mathbf{D}^2$. It then holds that $\mathbf{S}'_{\mathbf{p}}/\mathbf{C} > 0$ (resp. ≥ 0) if and only if $s_1^2 > \|\mathbf{R}\|_2^2 s_2^2 a / (s_2^2 - a)$ (resp. $s_1^2 \geq \|\mathbf{R}\|_2^2 s_2^2 a / (s_2^2 - a)$). Overall, condition ④ is equivalent to

$$s_1 > \frac{1}{\sqrt{1 - \frac{s_{\mathbf{G}}^2}{s_2^2}}} s_{\mathbf{G}} \|\mathbf{R}\|_2, \quad \text{and} \quad s_2 > s_{\mathbf{G}}$$

while condition ⑤ is equivalent to

$$s_1 \geq \frac{1}{\sqrt{1 - \frac{s_{\mathbf{G}}^4}{s_2^2(s_{\mathbf{G}}^2 - \eta^2)}}} \frac{s_{\mathbf{G}}^2}{\sqrt{s_{\mathbf{G}}^2 - \eta^2}} \|\mathbf{R}\|_2, \quad \text{and} \quad s_2 > \frac{s_{\mathbf{G}}^2}{\sqrt{s_{\mathbf{G}}^2 - \eta^2}} \quad (5)$$

with $\eta = \eta_\varepsilon(\mathbb{Z}^{dk})$. The latter set of conditions imply the former because $s_{\mathbf{G}}^2/(s_{\mathbf{G}}^2 - \eta^2) > 1$. It then holds that conditions ④ and ⑤ are mutually verified if and only if Equation (5) holds. Then, setting $s_2 = cs_{\mathbf{G}}^2/\sqrt{s_{\mathbf{G}}^2 - \eta^2}$, the lower bound on s_1 in Equation (5) is exactly $\sqrt{1 + 1/(c^2 - 1)} \cdot \|\mathbf{R}\|_2 s_{\mathbf{G}}^2/\sqrt{s_{\mathbf{G}}^2 - \eta^2}$. The parameters in the corollary statement thus imply conditions ④ and ⑤ of Theorem 3.1. The particular case of $c = \sqrt{2}$ comes up when setting c so that $c = \sqrt{1 + 1/(c^2 - 1)}$. We conclude on \mathcal{P} by applying Theorem 3.1. \square

4.2 Leveraging the Projection for Truncated Gadget

One of the novelties of our generic sampler is the matrix \mathbf{K} which projects the output samples onto $\mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_{\mathbf{T}})$. In the case of the regular [MP12] sampler, we have seen that it comes down to choosing $\mathbf{K} = \mathbf{I}$ which therefore discards this feature. We now give an instantiation of our generic sampler which leverages the freedom given by the matrix \mathbf{K} in order to truncate the gadget matrix.

The idea is to use the relation between \mathbf{K} and \mathbf{L} to define a full column rank matrix \mathbf{L} that would allow for finer sampling. From a security perspective, if \mathbf{K} does not depend on the secret trapdoor $(\mathbf{R}_1, \mathbf{R}_2)$, then analyzing the distribution \mathcal{P} of our sampler is sufficient and we do not have to study the distribution of $\mathbf{K}\mathcal{P}$ per se. In our case, we use \mathbf{K} to factor out the low-order gadget \mathbf{G}_L so that the security analysis and resulting Gaussian parameters do not depend on b^ℓ . As our goal is also to provide an analysis for advanced signatures with truncated gadgets, which are generally reliant on the use of *hidden* tags, we must consider $\mathbf{T} \neq \mathbf{I}_d$ as opposed to previous samplers with truncated gadgets, e.g., [CGM19, YJW23, JRS24]. As the tag is multiplied to \mathbf{G}_L , the form of \mathbf{T} as well as its size naturally impacts the proof. We show however that only $\|\mathbf{T}\|_2$ impacts the parameters, and we can still extract a meaningful result where this tag contribution is only additive. In typical use cases where $\mathbf{T} = \text{diag}(t_1, \dots, t_d)$, this will only feature $\|\mathbf{T}\|_2 = \max_{i \in [d]} |t_i|$. If the tag does not have to be hidden or if $\mathbf{T} = \mathbf{I}_d$, our analysis and parameter constraints become slightly simpler by putting the tag within the matrix \mathbf{K} . We thus consider \mathbf{K} and \mathbf{L} of this form.

$$\mathbf{K} = \begin{bmatrix} \mathbf{G}_L & & \\ & \mathbf{I}_d & \\ & & \mathbf{I}_{d(k-\ell)} \end{bmatrix}, \quad \text{and} \quad \mathbf{L} = \begin{bmatrix} \mathbf{T} & \mathbf{0} & \mathbf{R}_1 \\ \mathbf{0} & \mathbf{I}_{\ell-1} \otimes \mathbf{T} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{R}_2 \\ \mathbf{0} & \mathbf{0} & \mathbf{I}_{d(k-\ell)} \end{bmatrix} \in \mathbb{Z}^{d+dk \times dk}, \quad (6)$$

with $\mathbf{T} \in GL_d(\mathbb{Z}_q)$ and $\mathbf{G}_L = [\mathbf{I}_d | b\mathbf{I}_d | \dots | b^{\ell-1}\mathbf{I}_d]$. Using the fact that $\mathbf{G}_L(\mathbf{I}_\ell \otimes \mathbf{T}) = ([1|b|\dots|b^{\ell-1}] \otimes \mathbf{I}_d)(\mathbf{I}_\ell \otimes \mathbf{T}) = [\mathbf{T}|b\mathbf{T}|\dots|b^{\ell-1}\mathbf{T}] = \mathbf{T}\mathbf{G}_L$, a block matrix calculation shows that condition ① of Theorem 3.1 is verified. Then, ② also holds whenever $\mathbf{T} \in GL_d(\mathbb{Z}_q)$. Finally, choosing $\mathbf{S}_{\mathbf{G}} = s_{\mathbf{G}}^2 \mathbf{I}_{dk}$ with $s_{\mathbf{G}} = \eta_\varepsilon(\mathbb{Z}^{dk})\sqrt{b^2+1}$ yield ③. In that case, we show that choosing the covariance $\mathbf{S} = \text{diag}(s_1^2 \mathbf{I}_d, s_2^2 \mathbf{I}_{d(\ell-1)}, s_3^2 \mathbf{I}_d, s_4^2 \mathbf{I}_{d(k-\ell)})$ with

$$[s_1, s_2, s_3, s_4] = \frac{s_{\mathbf{G}}^2}{\sqrt{s_{\mathbf{G}}^2 - \eta_\varepsilon(\mathbb{Z}^{dk})^2}} \left[\sqrt{\|\mathbf{T}\|_2^2 + c^2 \|\mathbf{R}_1\|_2^2}, \|\mathbf{T}\|_2, c \|\mathbf{R}_2\|_2, c \right]$$

with $c = \sqrt{3}$ is sufficient to obtain the last two conditions of Theorem 3.1. We start by giving the parameter constraints on the s_i which are both sufficient and necessary for conditions ④ and ⑤, before showing the above formulas meet these constraints. Of course, $\|\mathbf{T}\|_2$ and $\|\mathbf{R}_i\|_2$ should be replaced by worst-case bounds so that the Gaussian parameters are fixed for any chosen $\mathbf{T}, \mathbf{R}_1, \mathbf{R}_2$. We generalize our analysis to tags and trapdoors over the reals as we only look at the spectral properties. We also mention that it would hold for complex matrices

simply by replacing the transpose operator with the Hermitian operator, i.e., conjugate transpose.

Lemma 4.1. *Let d, k, ℓ be positive integers with $\ell < k$, and $a > 0$. Then, let $\mathbf{T} \in GL_d(\mathbb{R})$, and $\mathbf{R}_1, \mathbf{R}_2$ be in $\mathbb{R}^{d \times d(k-\ell)}$. We define \mathbf{L} as follows.*

$$\mathbf{L} = \begin{bmatrix} \mathbf{T} & \mathbf{0} & \mathbf{R}_1 \\ \mathbf{0} & \mathbf{I}_{\ell-1} \otimes \mathbf{T} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{R}_2 \\ \mathbf{0} & \mathbf{0} & \mathbf{I}_{d(k-\ell)} \end{bmatrix} \in \mathbb{R}^{d+dk \times dk},$$

and $\mathbf{S} = \text{diag}(s_1^2 \mathbf{I}_d, s_2^2 \mathbf{I}_{d(\ell-1)}, s_3^2 \mathbf{I}_d, s_4^2 \mathbf{I}_{d(k-\ell)})$. If $\mathbf{R}_2 = \mathbf{U}_2 \mathbf{D}'_2 \mathbf{V}_2^T$ is a singular value decomposition of \mathbf{R}_2 with \mathbf{U}_2 and \mathbf{V}_2 unitary, it holds that $\mathbf{S} > a \mathbf{L} \mathbf{L}^T$ if and only if

$$\begin{aligned} s_1^2 &> a \|\mathbf{T} \mathbf{T}^T + \mathbf{R}_1 \mathbf{V}_2 \mathbf{D}'_2 \mathbf{V}_2^T \mathbf{R}_1^T\|_2 & s_2^2 &> a \|\mathbf{T}\|_2^2 \\ s_3 &> \frac{s_4^2 a}{s_4^2 - a} \|\mathbf{R}_2\|_2^2 & s_4^2 &> a, \end{aligned}$$

where $\mathbf{D} = \text{diag} \left(\text{diag} \left(\frac{s_3^2 s_4^2}{s_4^2 (s_3^2 - a s_i (\mathbf{R}_2)^2) - s_3^2 a} \right), \frac{s_4^2}{s_4^2 - a} \mathbf{I}_{d(k-\ell-1)} \right)$. The first condition on s_1 is implied by

$$s_1^2 > a \left(\|\mathbf{T}\|_2^2 + \|\mathbf{R}_1\|_2^2 \cdot \frac{s_3^2 s_4^2}{s_4^2 (s_3^2 - a \|\mathbf{R}_2\|_2^2) - s_3^2 a} \right)$$

We also have the same conditions for $\mathbf{S} \geq a \mathbf{L} \mathbf{L}^T$ where the inequalities on s_1 and s_2 do not need to be strict.

Proof. We start by computing $\mathbf{S}' = \mathbf{S} - a \mathbf{L} \mathbf{L}^T$ and then use the characterization of $\mathcal{S}^{++}(\mathbb{R})$ (resp. $\mathcal{S}^+(\mathbb{R})$) using Schur complements. For clarity, for $\mathbf{r} = [r_1, \dots, r_n]$, and $s \in \mathbb{R}$, we abbreviate $\text{diag}(\text{diag}(\mathbf{r}), s \mathbf{I})$ by $\text{diag}(r_i, s)$ where the dimensions of \mathbf{r} and \mathbf{I} are implicit. We also define $\mathbf{Q} = \mathbf{T} \mathbf{T}^T$ and $\mathbf{Q}_{\ell-1} = \mathbf{I}_{\ell-1} \otimes \mathbf{Q}$. By computing $\mathbf{L} \mathbf{L}^T$, we get that \mathbf{S}' equals

$$\begin{bmatrix} \begin{array}{cc|cc} \mathbf{A} & & & \mathbf{B} \\ s_1^2 \mathbf{I}_d - a \mathbf{Q} - a \mathbf{R}_1 \mathbf{R}_1^T & \mathbf{0} & -a \mathbf{R}_1 \mathbf{R}_2^T & -a \mathbf{R}_1 \\ \mathbf{0} & s_2^2 \mathbf{I}_{d(\ell-1)} - a \mathbf{Q}_{\ell-1} & \mathbf{0} & \mathbf{0} \\ \hline -a \mathbf{R}_2 \mathbf{R}_1^T & \mathbf{0} & s_3^2 \mathbf{I}_d - a \mathbf{R}_2 \mathbf{R}_2^T & -a \mathbf{R}_2 \\ -a \mathbf{R}_1^T & \mathbf{0} & -a \mathbf{R}_2^T & (s_4^2 - a) \mathbf{I}_{d(k-\ell)} \end{array} \\ \mathbf{B}^T & & & \mathbf{C} \end{bmatrix}$$

We first note that if \mathbf{C} is not positive definite, then the specific form of \mathbf{S}' yields that it is not positive semi-definite. As we are interested in equivalent conditions to \mathbf{S}' being positive definite (resp. positive semi-definite), then the Schur complement characterization gives us that $\mathbf{S}' > 0$ (resp. ≥ 0) if and only

if $\mathbf{C} > 0$ and $\mathbf{S}'/\mathbf{C} = \mathbf{A} - \mathbf{BC}^{-1}\mathbf{B}^T > 0$ (resp. ≥ 0). Let us first look at the condition $\mathbf{C} > 0$. As we will need to invert \mathbf{C} , it is not sufficient to again use the Schur complement characterization as in the proof of Corollary 4.2. Let $\mathbf{U}_2 \in O_d(\mathbb{R})$, $\mathbf{V}_2 \in O_{d(k-\ell)}(\mathbb{R})$ and $\mathbf{D}'_2 = [\mathbf{D}_2 | \mathbf{0}]$ with $\mathbf{D}_2 = \text{diag}(s_1(\mathbf{R}_2), \dots, s_d(\mathbf{R}_2))$ be a singular value decomposition of \mathbf{R}_2 , that is such that $\mathbf{R}_2 = \mathbf{U}_2 \mathbf{D}'_2 \mathbf{V}_2^T$ and $s_i(\mathbf{R}_2)$ the singular values of \mathbf{R}_2 in non-increasing order. A simple factorization then shows that

$$\mathbf{C} = \mathbf{W}_2 \mathbf{C}_D \mathbf{W}_2^T = \mathbf{W}_2 \begin{bmatrix} s_3^2 \mathbf{I}_d - a \mathbf{D}_2^2 & -a \mathbf{D}'_2 \\ -a \mathbf{D}'_2{}^T & (s_4^2 - a) \mathbf{I}_{d(k-\ell)} \end{bmatrix} \mathbf{W}_2^T,$$

with $\mathbf{W}_2 = \text{diag}(\mathbf{U}_2, \mathbf{V}_2)$. We can then compute the characteristic polynomial of \mathbf{C}_D . We have

$$\begin{aligned} \chi_{\mathbf{C}_D} &= \det(X \mathbf{I}_{d+d(k-\ell)} - \mathbf{C}_D) \\ &= \det((X - (s_4^2 - a)) \mathbf{I}_{d(k-\ell)}) \cdot \det((X - s_3^2) \mathbf{I}_d + a \mathbf{D}_2^2 - \frac{a^2}{X - (s_4^2 - a)} \mathbf{D}'_2{}^T \mathbf{D}'_2) \\ &= (X - (s_4^2 - a))^{d(k-\ell)} \cdot \det((X - s_3^2) \mathbf{I}_d + a(1 - \frac{a}{X - (s_4^2 - a)}) \mathbf{D}_2^2) \\ &= (X - (s_4^2 - a))^{d(k-\ell-1)} \prod_{i=1}^d P_i, \end{aligned}$$

where $P_i = (X - s_3^2)(X - (s_4^2 - a)) + a(X - s_4^2)s_i(\mathbf{R}_2)^2 = X^2 - (s_3^2 + s_4^2 - a - as_i(\mathbf{R}_2)^2)X + s_3^2(s_4^2 - a) - as_4^2s_i(\mathbf{R}_2)^2$. Each P_i has non negative discriminant, and as such its roots are positive if and only if $s_3^2(s_4^2 - a) - as_4^2s_i(\mathbf{R}_2)^2 > 0$. We then conclude that \mathbf{C} is positive definite if and only if

$$s_4^2 > a, \quad \text{and} \quad \forall i \in [d], s_3^2(s_4^2 - a) - as_4^2s_i(\mathbf{R}_2)^2 > 0.$$

This is then equivalent to

$$s_4^2 > a, \quad \text{and} \quad s_3^2 > \frac{s_4^2 a}{s_4^2 - a} \|\mathbf{R}_2\|_2^2.$$

We now need to compute \mathbf{S}_p/\mathbf{C} which involves the inverse of \mathbf{C} . As we have $\mathbf{C}^{-1} = \mathbf{W}_2 \mathbf{C}_D^{-1} \mathbf{W}_2^T$, we reduce it to computing \mathbf{C}_D^{-1} . The matrix \mathbf{C}_D is a 2×2 block matrix and we can therefore use the inverse formula of [LS02]. More precisely, we have $\mathbf{C}_D = \begin{bmatrix} \mathbf{E} & \mathbf{F} \\ \mathbf{F}^T & \mathbf{H} \end{bmatrix}$, and its inverse is therefore given by

$$\mathbf{C}_D^{-1} = \begin{bmatrix} \mathbf{E}^{-1} + \mathbf{E}^{-1} \mathbf{F} (\mathbf{H} - \mathbf{F}^T \mathbf{E}^{-1} \mathbf{F})^{-1} \mathbf{F}^T \mathbf{E}^{-1} & -\mathbf{E}^{-1} \mathbf{F} (\mathbf{H} - \mathbf{F}^T \mathbf{E}^{-1} \mathbf{F})^{-1} \\ -(\mathbf{H} - \mathbf{F}^T \mathbf{E}^{-1} \mathbf{F})^{-1} \mathbf{F}^T \mathbf{E}^{-1} & (\mathbf{H} - \mathbf{F}^T \mathbf{E}^{-1} \mathbf{F})^{-1} \end{bmatrix},$$

assuming the involved inverses exist. We will see when computing these matrices that the prior conditions yield the invertibility of the intermediate matrices,

thus allowing this computation. For simplicity, we call $\mathbf{X} := (\mathbf{H} - \mathbf{F}^T \mathbf{E}^{-1} \mathbf{F})^{-1}$, and $\mathbf{D}_{C1}, \mathbf{D}_{C2}, \mathbf{D}_{C3}$ such that $\mathbf{C}_D^{-1} = \begin{bmatrix} \mathbf{D}_{C1} & \mathbf{D}_{C2} \\ \mathbf{D}_{C2}^T & \mathbf{D}_{C3} \end{bmatrix}$. First of all, we have that $\mathbf{E} = \text{diag}(s_3^2 - as_i(\mathbf{R}_2)^2)$. By the condition identified on s_3 and the fact that $s_4/\sqrt{s_4^2 - a} > 1$, \mathbf{E} is indeed invertible. So $\mathbf{E}^{-1} = \text{diag}(1/(s_3^2 - as_i(\mathbf{R}_2)^2))$. Then, we have

$$\begin{aligned} \mathbf{H} - \mathbf{F}^T \mathbf{E}^{-1} \mathbf{F} &= (s_4^2 - a) \mathbf{I}_{d(k-\ell)} - a^2 \mathbf{D}_2'^T \mathbf{E}^{-1} \mathbf{D}_2' \\ &= \text{diag} \left(s_4^2 - a - \frac{a^2 s_i(\mathbf{R}_2)^2}{s_3^2 - as_i(\mathbf{R}_2)^2}, s_4^2 - a \right) \\ &= \text{diag} \left(\frac{s_4^2(s_3^2 - as_i(\mathbf{R}_2)^2) - s_3^2 a}{s_3^2 - as_i(\mathbf{R}_2)^2}, s_4^2 - a \right). \end{aligned}$$

The conditions on s_3 and s_4 are equivalent to this matrix being invertible. We can then proceed in the computation with

$$\mathbf{D}_{C3} = \mathbf{X} = \text{diag} \left(\frac{s_3^2 - as_i(\mathbf{R}_2)^2}{s_4^2(s_3^2 - as_i(\mathbf{R}_2)^2) - s_3^2 a}, \frac{1}{s_4^2 - a} \right).$$

It then gives

$$\begin{aligned} \mathbf{D}_{C2} &= -\mathbf{E}^{-1} \mathbf{F} \mathbf{X} = \left[\text{diag} \left(\frac{as_i(\mathbf{R}_2)}{s_4^2(s_3^2 - as_i(\mathbf{R}_2)^2) - s_3^2 a} \right) \middle| \mathbf{0} \right] \\ \mathbf{D}_{C2}^T &= -\mathbf{X} \mathbf{F}^T \mathbf{E}^{-1} = \left[\text{diag} \left(\frac{as_i(\mathbf{R}_2)}{s_4^2(s_3^2 - as_i(\mathbf{R}_2)^2) - s_3^2 a} \right) \right. \\ &\quad \left. \mathbf{0} \right]. \end{aligned}$$

Finally, we have $\mathbf{E}^{-1} \mathbf{F} \mathbf{X} \mathbf{F}^T \mathbf{E}^{-1} = \text{diag} \left(\frac{a^2 s_i(\mathbf{R}_2)^2}{(s_4^2(s_3^2 - as_i(\mathbf{R}_2)^2) - s_3^2 a)(s_3^2 - as_i(\mathbf{R}_2)^2)} \right)$, and thus the top left term simplifies to

$$\mathbf{D}_{C1} = \mathbf{E}^{-1} + \mathbf{E}^{-1} \mathbf{F} \mathbf{X} \mathbf{F}^T \mathbf{E}^{-1} = \text{diag} \left(\frac{s_4^2 - a}{s_4^2(s_3^2 - as_i(\mathbf{R}_2)^2) - s_3^2 a} \right)$$

We can then compute the block-wise product and get that $\mathbf{B} \mathbf{C}^{-1} \mathbf{B}^T = \text{diag}(\mathbf{Y}, \mathbf{0})$ where

$$\mathbf{Y} = a^2 \mathbf{R}_1 \mathbf{V}_2 (\mathbf{D}_2'^T \mathbf{D}_{C1} \mathbf{D}_2' + \mathbf{D}_{C2}^T \mathbf{D}_2' + \mathbf{D}_2'^T \mathbf{D}_{C2} + \mathbf{D}_{C3}) \mathbf{V}_2^T \mathbf{R}_1^T$$

As a result, $\mathbf{S}'/\mathbf{C} = \text{diag}(s_1^2 \mathbf{I}_d - a \mathbf{Q} - a \mathbf{R}_1 \mathbf{R}_1^T - \mathbf{Y}, s_2^2 \mathbf{I}_{d(\ell-1)} - a \mathbf{Q}_{\ell-1})$. By simplifying the expression of \mathbf{Y} , the first block of \mathbf{S}'/\mathbf{C} becomes

$$s_1^2 \mathbf{I}_d - a \left(\mathbf{T} \mathbf{T}^T + \mathbf{R}_1 \mathbf{V}_2 \text{diag} \left(\frac{s_3^2 s_4^2}{s_4^2(s_3^2 - as_i(\mathbf{R}_2)^2) - s_3^2 a}, \frac{s_4^2}{s_4^2 - a} \right) \mathbf{V}_2^T \mathbf{R}_1^T \right)$$

Then because \mathbf{S}'/\mathbf{C} is block diagonal, it is positive definite (resp. semi-definite) if and only if each block is positive definite (resp. semi-definite). We thus get

that the necessary and sufficient conditions are

$$\begin{aligned} s_1^2 &> a \left\| \mathbf{T}\mathbf{T}^T + \mathbf{R}_1\mathbf{V}_2 \text{diag} \left(\frac{s_3^2 s_4^2}{s_4^2 (s_3^2 - a s_i (\mathbf{R}_2)^2) - s_3^2 a}, \frac{s_4^2}{s_4^2 - a} \right) \mathbf{V}_2^T \mathbf{R}_1^T \right\|_2 \\ s_2^2 &> a \|\mathbf{I}_{\ell-1} \otimes \mathbf{T}\mathbf{T}^T\|_2 = a \|\mathbf{T}\|_2^2, \end{aligned}$$

as claimed, and where the inequalities do not need to be strict for the positive semi-definite case. For the sufficient condition on s_1 , we simply use the triangle inequality and submultiplicativity of the spectral norm, the fact that $\|\mathbf{V}_2 \mathbf{D} \mathbf{V}_2^T\|_2 = \|\mathbf{D}\|_2$, and the fact that the maximal entry of \mathbf{D} is achieved at its first diagonal entry, i.e., for $i = 1$ where $s_1(\mathbf{R}_2) = \|\mathbf{R}_2\|_2$. \square

Notice that the relaxed condition on s_1 is tight and perfectly matches the exact condition when $\mathbf{T} = \mathbf{t}\mathbf{I}_d$ and $\mathbf{R}_1^T \mathbf{R}_1$ and $\mathbf{R}_2^T \mathbf{R}_2$ have the same eigenvectors (i.e., $\mathbf{V}_1 = \mathbf{V}_2$) for example. We can then apply Lemma 4.1 to $a = s_{\mathbf{G}}^2$ for condition (4) and $a = s_{\mathbf{G}}^4 / (s_{\mathbf{G}}^2 - \eta_\varepsilon(\mathbb{Z}^{dk})^2)$ for condition (5). This gives the following corollary.

Corollary 4.3. *Let d, k, ℓ be positive integers with $\ell < k$, $s_{\mathbf{G}} > 0$, and $\varepsilon \in (0, 1)$. Then, let $\mathbf{T} \in GL_d(\mathbb{R})$, and $\mathbf{R}_1, \mathbf{R}_2$ be in $\mathbb{R}^{d \times d(k-\ell)}$. Let*

$$\begin{aligned} s_1 &= \frac{s_{\mathbf{G}}^2}{\sqrt{s_{\mathbf{G}}^2 - \eta_\varepsilon(\mathbb{Z}^{dk})^2}} \sqrt{\|\mathbf{T}\|_2^2 + 3\|\mathbf{R}_1\|_2^2} & s_2 &= \frac{s_{\mathbf{G}}^2}{\sqrt{s_{\mathbf{G}}^2 - \eta_\varepsilon(\mathbb{Z}^{dk})^2}} \|\mathbf{T}\|_2 \\ s_3 &= \sqrt{3} \frac{s_{\mathbf{G}}^2}{\sqrt{s_{\mathbf{G}}^2 - \eta_\varepsilon(\mathbb{Z}^{dk})^2}} \|\mathbf{R}_2\|_2 & s_4 &= \sqrt{3} \frac{s_{\mathbf{G}}^2}{\sqrt{s_{\mathbf{G}}^2 - \eta_\varepsilon(\mathbb{Z}^{dk})^2}}, \end{aligned}$$

Then $\mathbf{S} > s_{\mathbf{G}}^2 \mathbf{L}\mathbf{L}^T$ and $\mathbf{S} \geq \frac{s_{\mathbf{G}}^4}{s_{\mathbf{G}}^2 - \eta_\varepsilon(\mathbb{Z}^{dk})^2} \mathbf{L}\mathbf{L}^T$, where \mathbf{L} is defined in Equation (6), and $\mathbf{S} = \text{diag}(s_1^2 \mathbf{I}_d, s_2^2 \mathbf{I}_{d(\ell-1)}, s_3^2 \mathbf{I}_d, s_4^2 \mathbf{I}_{d(k-\ell)})$.

Proof. We first apply Lemma 4.1 with $a_1 = s_{\mathbf{G}}^2$ and strict inequalities for s_1, s_2 with the relaxed condition on s_1 , and then with $a_2 = s_{\mathbf{G}}^4 / (s_{\mathbf{G}}^2 - \eta_\varepsilon(\mathbb{Z}^{dk})^2)$ with large inequalities (and also the relaxed condition on s_1). It gives us two sets of conditions. Because $a_2 > a_1$, we can easily check that the more restrictive conditions are the latter ones. We thus let $c_3, c_4 > 1$ be free variables and select

$$\begin{aligned} s_1^2 &= a_2 \left(\|\mathbf{T}\|_2^2 + \frac{s_3^2 s_4^2}{s_4^2 (s_3^2 - a_2 \|\mathbf{R}_2\|_2^2) - s_3^2 a_2} \|\mathbf{R}_1\|_2^2 \right) & s_2^2 &= a_2 \|\mathbf{T}\|_2^2 \\ s_3^2 &= c_3^2 \frac{s_4^2 a_2}{s_4^2 - a_2} \|\mathbf{R}_2\|_2^2 & s_4^2 &= c_4^2 a_2, \end{aligned}$$

As such they verify the conditions of Lemma 4.1, thus showing that $\mathbf{S} > a_1 \mathbf{L}\mathbf{L}^T$ and $\mathbf{S} \geq a_2 \mathbf{L}\mathbf{L}^T$. By setting $c_3 = \sqrt{c_4^2 - 1}$, assuming $c_4 > \sqrt{2}$, then $c_3 > 1$ and $c_3 s_4 / \sqrt{s_4^2 - a_2} = c_3 c_4 / \sqrt{c_4^2 - 1} = c_4$. Then, we can compute the factor of $\|\mathbf{R}_1\|_2^2$ in the expression of s_1 . We have

$$\frac{s_3^2 s_4^2}{s_4^2 (s_3^2 - a_2 \|\mathbf{R}_2\|_2^2) - s_3^2 a_2} = \frac{1}{1 - a_2 \|\mathbf{R}_2\|_2^2 / s_3^2 - a_2 / s_4^2} = \frac{1}{1 - 2/c_4^2}$$

If we now set c_4 so that $c_4^2 = 1/(1 - 2/c_4^2)$, we get $c_4 = \sqrt{3}$ which is indeed larger than $\sqrt{2}$. Hence $c_3 = \sqrt{2} > 1$. So the parameter given in the corollary statement indeed yield condition ④ and ⑤ of Theorem 3.1. \square

The conclusion of Corollary 4.3 is that setting the s_i as above is sufficient so that the matrices \mathbf{K} and \mathbf{L} from Equation (6), as well as \mathbf{S} and $\mathbf{S}_{\mathbf{G}} = s_{\mathbf{G}}^2 \mathbf{I}_{dk}$, satisfy all five conditions of Theorem 3.1. By replacing $\|\mathbf{R}_i\|_2$ by a worst-case bound enforced during the key generation, and $\|\mathbf{T}\|_2$ by a worst-case bound as well, it holds that we can leverage the \mathbf{K} -projection to truncate the gadget to \mathbf{G}_H , while hiding the effective tag \mathbf{T} . We summarize it in Algorithm 4.1 and Theorem 4.1. It then provides the first worst-case analysis of trapdoor sampling with truncated gadget, i.e., for an arbitrary (possibly adversarial) syndrome \mathbf{u} .

Algorithm 4.1: TruncatedSampler($\mathbf{R}_1, \mathbf{R}_2, \mathbf{A}, \mathbf{u}, \mathbf{T}, s_{\mathbf{G}}, s_1, s_2, s_3, s_4$)

Input: Trapdoor $\mathbf{R}_1, \mathbf{R}_2 \in \mathbb{Z}^{d \times d(k-\ell)}$, Matrix $\mathbf{A} \in \mathbb{Z}_q^{d \times d}$, Syndrome $\mathbf{u} \in \mathbb{Z}_q^d$, tag matrix $\mathbf{T} \in GL_d(\mathbb{Z}_q)$, Gaussian parameters $s_{\mathbf{G}}, s_1, s_2, s_3, s_4$.

1. $\mathbf{p} \leftarrow \mathcal{D}_{\mathbb{Z}^{d(k+1)}, \sqrt{s_{\mathbf{p}}}}$. $\triangleright \mathbf{S}_{\mathbf{p}} = \mathbf{S} - s_{\mathbf{G}}^2 \mathbf{L}\mathbf{L}^T$
2. Parse $\mathbf{p} = [\mathbf{p}_L^T | \mathbf{p}_{1,2}^T | \mathbf{p}_2^T]^T$ with $\mathbf{p}_L \in \mathbb{Z}^{d\ell}$, $\mathbf{p}_{1,2} \in \mathbb{Z}^d$ and $\mathbf{p}_2 \in \mathbb{Z}^{d(k-\ell)}$.
3. $\mathbf{w} \leftarrow \mathbf{T}^{-1}(\mathbf{u} - \mathbf{G}_L \mathbf{p}_L - \mathbf{A} \mathbf{p}_{1,2} - (\mathbf{T} \mathbf{G}_H - (\mathbf{R}_1 + \mathbf{A} \mathbf{R}_2)) \mathbf{p}_2) \bmod q\mathbb{Z}$.
4. $\mathbf{z} \leftarrow \mathcal{D}_{\mathcal{L}_q^{\mathbf{w}}(\mathbf{G}), s_{\mathbf{G}}}$.
5. $\mathbf{v}' \leftarrow \mathbf{p} + \mathbf{L}\mathbf{z}$.
6. Parse $\mathbf{v}' = [\mathbf{v}_L^T | \mathbf{v}_{1,2}^T | \mathbf{v}_2^T]^T$ with $\mathbf{v}_L \in \mathbb{Z}^{d\ell}$, $\mathbf{v}_{1,2} \in \mathbb{Z}^d$ and $\mathbf{v}_2 \in \mathbb{Z}^{d(k-\ell)}$.
7. $\mathbf{v}_{1,1} \leftarrow \mathbf{G}_L \mathbf{v}_L$.

Output: $\mathbf{v} = [\mathbf{v}_{1,1}^T | \mathbf{v}_{1,2}^T | \mathbf{v}_2^T]^T \in \mathbb{Z}^{d(2+k-\ell)}$.

Theorem 4.1. Let d, q, b, ℓ be positive integers and $\varepsilon \in (0, 1)$, such that $b \leq \sqrt{q}$, and $\ell < k = \lceil \log_b q \rceil$. Let $\mathbf{R}_1, \mathbf{R}_2$ be in $\mathbb{Z}^{d \times d(k-\ell)}$, $\mathbf{A} \in \mathbb{Z}_q^{d \times d}$, $\mathbf{u} \in \mathbb{Z}_q^d$ and $\mathbf{T} \in GL_d(\mathbb{Z}_q)$. Define $\mathbf{G}_L = [1|b|\dots|b^{\ell-1}] \otimes \mathbf{I}_d$ and $\mathbf{G}_H = [b^\ell|\dots|b^{k-1}] \otimes \mathbf{I}_d$. Finally, let $w \geq \|\mathbf{T}\|_2$, $s_{\mathbf{G}} = \eta_\varepsilon(\mathbb{Z}^{dk}) \sqrt{b^2 + 1}$ and $\mathbf{S} = \text{diag}(s_1^2 \mathbf{I}_d, s_2^2 \mathbf{I}_{d(\ell-1)}, s_3^2 \mathbf{I}_d, s_4^2 \mathbf{I}_{d(k-\ell)})$ where

$$\begin{aligned} s_1 &= \eta_\varepsilon(\mathbb{Z}^{dk}) \left(b + \frac{1}{b}\right) \sqrt{w^2 + 3\|\mathbf{R}_1\|_2^2} & s_2 &= \eta_\varepsilon(\mathbb{Z}^{dk}) \left(b + \frac{1}{b}\right) w \\ s_3 &= \sqrt{3}\eta_\varepsilon(\mathbb{Z}^{dk}) \left(b + \frac{1}{b}\right) \|\mathbf{R}_2\|_2 & s_4 &= \sqrt{3}\eta_\varepsilon(\mathbb{Z}^{dk}) \left(b + \frac{1}{b}\right), \end{aligned}$$

The distribution \mathcal{P} of TruncatedSampler($\mathbf{R}_1, \mathbf{R}_2, \mathbf{A}, \mathbf{u}, \mathbf{T}, s_{\mathbf{G}}, s_1, s_2, s_3, s_4$) is such that $\text{Supp}(\mathcal{P}) = \mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_{\mathbf{T}})$ and

$$\forall \mathbf{x} \in \mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_{\mathbf{T}}), \mathcal{P}(\mathbf{x}) \in \left[\left(\frac{1-\varepsilon}{1+\varepsilon}\right)^2, \left(\frac{1+\varepsilon}{1-\varepsilon}\right)^2 \right] \cdot \mathbf{KD}_{\mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_{\mathbf{T}}\mathbf{K}), \sqrt{\mathbf{S}}}(\mathbf{x}),$$

where $\mathbf{A}_{\mathbf{T}} = [\mathbf{I}_d | \mathbf{A} | \mathbf{T} \mathbf{G}_H - (\mathbf{R}_1 + \mathbf{A} \mathbf{R}_2)] \bmod q\mathbb{Z}$, and $\mathbf{K} = \text{diag}(\mathbf{G}_L, \mathbf{I}_d, \mathbf{I}_{d(k-\ell)})$.

Our method of sampling with truncated gadget actually perturbs the temporary preimage before computing the ‘‘preimage error’’ with \mathbf{G}_L . Nevertheless,

our fine-grained analysis allows us to aim for a Gaussian distribution with four parameters so that most of the terms in the preimage error (those corresponding to parameter s_2) are very small and do not depend on the size of the trapdoor. This limits the size of the preimage error after the \mathbf{K} -projection. In particular, a subsequent step would be to argue that $\mathbf{K}\mathcal{D}_{\mathcal{L}_q^u(\mathbf{A}_T\mathbf{K}),\sqrt{\mathbf{S}}} \approx \mathcal{D}_{\mathcal{L}_q^u(\mathbf{A}_T),\sqrt{\mathbf{S}'}}$, using e.g., [GMPW20, Thm 4.6], with $\mathbf{S}' = \text{diag}(s_{1,1}^2\mathbf{I}_d, s_3^2\mathbf{I}_d, s_4^2\mathbf{I}_{d(k-\ell)})$ and where

$$\begin{aligned} s_{1,1} &= \sqrt{s_1^2 + \sum_{i=1}^{\ell-1} b^{2i} s_2^2} = \sqrt{s_1^2 + s_2^2 b^2 \frac{b^{2(\ell-1)} - 1}{b^2 - 1}} \\ &= \eta_\varepsilon(\mathbb{Z}^{dk}) \left(b + \frac{1}{b}\right) \sqrt{\frac{b^{2\ell} - 1}{b^2 - 1} \|\mathbf{T}\|_2^2 + 3\|\mathbf{R}_1\|_2^2} \end{aligned}$$

The case $\ell = 1$ thus comes for free, which we could already observe as the $s_2^2\mathbf{I}_{d(\ell-1)}$ block is empty. For $\ell > 1$, the increase stays mild as long as $\|\mathbf{T}\|_2 \cdot b^{\ell-1}$ stays below or close to $\sqrt{3}\|\mathbf{R}_1\|_2$. For example, when $\mathbf{T} = \text{diag}(t_1, \dots, t_d)$ with $t_i \in \{-1, 1\}$, the tag space is of size 2^d and $\|\mathbf{T}\|_2 = 1$. Most importantly, the term corresponding to the preimage error in $\approx b^{\ell-1}$ is additive with respect to $\|\mathbf{R}_1\|_2$ instead of multiplicative in [CGM19] (as they essentially set $s_2 = s_1$).

4.3 The Ring Setting

The above sampler with truncated gadget of Algorithm 4.1 naturally transfers to the ring setting to enable faster computations and reduced storage. The analysis naturally extends as the only relevant quantities are the norms of $\mathbf{T}, \mathbf{R}_1, \mathbf{R}_2$ which are defined from their embedding $M_\tau(\mathbf{T}), M_\tau(\mathbf{R}_1), M_\tau(\mathbf{R}_2)$ in the ring case. Because the covariance matrices are also defined with respect to M_τ , we can essentially replace these matrices by their embedding in the results of Sections 3, 4.1 and 4.2 and directly obtain the same guarantees where the dimension d is replaced by nd . The actual structure of $M_\tau(\cdot)$ is then only relevant in estimating $\|M_\tau(\cdot)\|_2$. We explicit the result over rings for completeness.

Algorithm 4.2: RingTruncatedSampler($\mathbf{R}_1, \mathbf{R}_2, \mathbf{A}, \mathbf{u}, \mathbf{T}, s_G, s_1, s_2, s_3, s_4$)

Input: Trapdoor $\mathbf{R}_1, \mathbf{R}_2 \in R^{d \times d(k-\ell)}$, Matrix $\mathbf{A} \in R_q^{d \times d}$, Syndrome $\mathbf{u} \in R_q^d$, tag matrix $\mathbf{T} \in GL_d(R_q)$, Gaussian parameters s_G, s_1, s_2, s_3, s_4 .

1. $\mathbf{p} \leftarrow \mathcal{D}_{R^{d(k+1)}, \sqrt{M_\tau(\mathbf{S}_p)}}$. $\triangleright \mathbf{S}_p = \mathbf{S} - s_G^2 \mathbf{L}\mathbf{L}^*$
2. Parse $\mathbf{p} = [\mathbf{p}_L^T | \mathbf{p}_{1,2}^T | \mathbf{p}_2^T]^T$ with $\mathbf{p}_L \in R^{d\ell}$, $\mathbf{p}_{1,2} \in R^d$ and $\mathbf{p}_2 \in R^{d(k-\ell)}$.
3. $\mathbf{w} \leftarrow \mathbf{T}^{-1}(\mathbf{u} - \mathbf{G}_L \mathbf{p}_L - \mathbf{A} \mathbf{p}_{1,2} - (\mathbf{T} \mathbf{G}_H - (\mathbf{R}_1 + \mathbf{A} \mathbf{R}_2)) \mathbf{p}_2) \bmod qR$.
4. $\mathbf{z} \leftarrow \mathcal{D}_{\mathcal{L}_q^w(\mathbf{G}), s_G}$.
5. $\mathbf{v}' \leftarrow \mathbf{p} + \mathbf{L}\mathbf{z}$.
6. Parse $\mathbf{v}' = [\mathbf{v}_L^T | \mathbf{v}_{1,2}^T | \mathbf{v}_2^T]^T$ with $\mathbf{v}_L \in \mathbb{Z}^{d\ell}$, $\mathbf{v}_{1,2} \in R^d$ and $\mathbf{v}_2 \in R^{d(k-\ell)}$.
7. $\mathbf{v}_{1,1} \leftarrow \mathbf{G}_L \mathbf{v}_L$.

Output: $\mathbf{v} = [\mathbf{v}_{1,1}^T | \mathbf{v}_{1,2}^T | \mathbf{v}_2^T]^T \in R^{d(2+k-\ell)}$.

Theorem 4.2. Let R be the ring of integer of a number field of degree n . Let d, q, b, ℓ be positive integers and $\varepsilon \in (0, 1)$, such that $b \leq \sqrt{q}$, and $\ell < k = \lceil \log_b q \rceil$. Let $\mathbf{R}_1, \mathbf{R}_2$ be in $R^{d \times d(k-\ell)}$, $\mathbf{A} \in R_q^{d \times d}$, $\mathbf{u} \in R_q^d$ and $\mathbf{T} \in GL_d(R_q)$. Define $\mathbf{G}_L = [1|b|\dots|b^{\ell-1}] \otimes \mathbf{I}_d$ and $\mathbf{G}_H = [b^\ell|\dots|b^{k-1}] \otimes \mathbf{I}_d$. Finally, let $w \geq \|\mathbf{T}\|_2$, $s_{\mathbf{G}} = \eta_\varepsilon(\mathbb{Z}^{ndk})\sqrt{b^2 + 1}$ and $\mathbf{S} = \text{diag}(s_1^2 \mathbf{I}_d, s_2^2 \mathbf{I}_{d(\ell-1)}, s_3^2 \mathbf{I}_d, s_4^2 \mathbf{I}_{d(k-\ell)})$ where

$$\begin{aligned} s_1 &= \eta_\varepsilon(\mathbb{Z}^{ndk}) \left(b + \frac{1}{b} \right) \sqrt{w^2 + 3\|\mathbf{R}_1\|_2^2} & s_2 &= \eta_\varepsilon(\mathbb{Z}^{ndk}) \left(b + \frac{1}{b} \right) w \\ s_3 &= \sqrt{3}\eta_\varepsilon(\mathbb{Z}^{ndk}) \left(b + \frac{1}{b} \right) \|\mathbf{R}_2\|_2 & s_4 &= \sqrt{3}\eta_\varepsilon(\mathbb{Z}^{ndk}) \left(b + \frac{1}{b} \right), \end{aligned}$$

The distribution \mathcal{P} of RingTruncatedSampler($\mathbf{R}_1, \mathbf{R}_2, \mathbf{A}, \mathbf{u}, \mathbf{T}, s_{\mathbf{G}}, s_1, s_2, s_3, s_4$) is such that $\text{Supp}(\mathcal{P}) = \mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_{\mathbf{T}})$ and

$$\forall \mathbf{x} \in \mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_{\mathbf{T}}), \mathcal{P}(\mathbf{x}) \in \left[\left(\frac{1-\varepsilon}{1+\varepsilon} \right)^2, \left(\frac{1+\varepsilon}{1-\varepsilon} \right)^2 \right] \cdot \mathbf{KD}_{\mathcal{L}_q^{\mathbf{u}}(\mathbf{A}_{\mathbf{T}}\mathbf{K}), \sqrt{M_\tau(\mathbf{S})}}(\mathbf{x}),$$

with $\mathbf{A}_{\mathbf{T}} = [\mathbf{I}_d | \mathbf{A} | \mathbf{T}\mathbf{G}_H - (\mathbf{R}_1 + \mathbf{A}\mathbf{R}_2)] \text{ mod } qR$, and $\mathbf{K} = \text{diag}(\mathbf{G}_L, \mathbf{I}_d, \mathbf{I}_{d(k-\ell)})$.

In particular, for the case of $\mathbf{T} = \text{diag}(\mathbf{t}_1, \dots, \mathbf{t}_d)$ in a power-of-two cyclotomic ring, we have $\|M_\tau(\mathbf{T})\|_2 = \max_{i \in [d]} \|M_\tau(\mathbf{t}_i)\|_2 \leq \max_{i \in [d]} \|\mathbf{t}_i\|_1$.

4.3.1 Sampling the Perturbation. The perturbation sampling step can be made fairly efficient in power-of-two cyclotomic rings by using the ring sampler of [GM18], which exploits the tower structure, and the recursive Schur complement convolution sampling [GM18, Lem. 4.3]. We also observe that based on the specific shape of our matrix $\mathbf{S}_{\mathbf{p}}$, a lot of elements can be sampled independently or precomputed offline. As will be the case in all the constructions covered in Section 5, we are interested in simple tag matrices $\mathbf{T} = \mathbf{t}\mathbf{I}_d$ with \mathbf{t} a short ring element such that $\mathbf{t} \in R_q^\times$. The perturbation covariance matrix $\mathbf{S}_{\mathbf{p}}$ (over the ring) is then

$$\mathbf{S}_{\mathbf{p}} = \begin{bmatrix} s_1^2 \mathbf{I}_d - a(\mathbf{t}\mathbf{t}^* \mathbf{I}_d + \mathbf{R}_1 \mathbf{R}_1^*) & \mathbf{0} & -a \mathbf{R}_1 \mathbf{R}_2^* & -a \mathbf{R}_1 \\ \mathbf{0} & (s_2^2 - a\mathbf{t}\mathbf{t}^*) \mathbf{I}_{d(\ell-1)} & \mathbf{0} & \mathbf{0} \\ -a \mathbf{R}_2 \mathbf{R}_1^* & \mathbf{0} & s_3^2 \mathbf{I}_d - a \mathbf{R}_2 \mathbf{R}_2^* & -a \mathbf{R}_2 \\ -a \mathbf{R}_1^* & \mathbf{0} & -a \mathbf{R}_2^* & (s_4^2 - a) \mathbf{I}_{d(k-\ell)} \end{bmatrix}$$

where $a = s_{\mathbf{G}}^2$. If we parse the vector \mathbf{p} into $[\mathbf{p}_{L,1}^T | \mathbf{p}_{L,2}^T | \mathbf{p}_{1,2}^T | \mathbf{p}_2^T]^T$ with $\mathbf{p}_{L,1}, \mathbf{p}_{1,2} \in R^d$, $\mathbf{p}_{L,2} \in R^{d(\ell-1)}$ and $\mathbf{p}_2 \in R^{d(k-\ell)}$, we can see that $\mathbf{p}_{L,2}$ is independent of the others. As a result, we can sample $\mathbf{p}_{L,2} \leftarrow \mathcal{D}_{R^{d(\ell-1)}, \sqrt{M_\tau(s_2^2 - a\mathbf{t}\mathbf{t}^*)}}$ independently, and then $(\mathbf{p}_{L,1}, \mathbf{p}_{1,2}, \mathbf{p}_2)$ from $\mathcal{D}_{R^{2d+d(k-\ell)}, \sqrt{M_\tau(\mathbf{S}'_{\mathbf{p}})}}$ with

$$\mathbf{S}'_{\mathbf{p}} = \begin{bmatrix} s_1^2 \mathbf{I}_d - a(\mathbf{t}\mathbf{t}^* \mathbf{I}_d + \mathbf{R}_1 \mathbf{R}_1^*) & -a \mathbf{R}_1 \mathbf{R}_2^* & -a \mathbf{R}_1 \\ -a \mathbf{R}_2 \mathbf{R}_1^* & s_3^2 \mathbf{I}_d - a \mathbf{R}_2 \mathbf{R}_2^* & -a \mathbf{R}_2 \\ -a \mathbf{R}_1^* & -a \mathbf{R}_2^* & (s_4^2 - a) \mathbf{I}_{d(k-\ell)} \end{bmatrix}$$

The sampling of $\mathbf{p}_{L,2}$ consists in sampling $d(\ell - 1)$ independent elements from $\mathcal{D}_{R, \sqrt{M_r(s_2^2 - att^*)}}$ using the efficient ring sampler of [GM18] as $\text{SampleFz}(s_2^2 - att^*, 0)$. Notice that $s_2^2 - att^*$ is clearly auto-adjoint, and because of the condition on s_2 of Theorem 4.2, we have $s_2^2 - att^* \in K_{\mathbb{R}}^{++}$ as needed for the sampling. Now let us look at the sampling with covariance $\mathbf{S}'_{\mathbf{p}}$. Here, we use the recursive Schur complement convolution sampling based on [GM18, Lem. 4.3] and used in previous implemented works [BEP⁺21, AGJ⁺24]. The idea is to compute elements $f_i \in K_{\mathbb{R}}^{++}$ from the successive Schur complements in order to sample the whole vector. We can first sample $\mathbf{p}_2 \leftarrow \mathcal{D}_{R^{d(k-\ell)}, \sqrt{s_4^2 - a}}$ (that is sampling over $\mathbb{Z}^{nd(k-\ell)}$ directly) and then update the sampling of $(\mathbf{p}_{L,1}, \mathbf{p}_{1,2})$ to account for the dependency. For that we must sample $(\mathbf{p}_{L,1}, \mathbf{p}_{1,2})$ with covariance $\mathbf{S}''_{\mathbf{p}} = \mathbf{S}_{\mathbf{p}} / (s_4^2 - a) \mathbf{I}_{d(k-\ell)}$ and center \mathbf{c} where

$$\mathbf{S}''_{\mathbf{p}} = \begin{bmatrix} (s_1^2 - att^*) \mathbf{I}_d - \frac{as_4^2}{s_4^2 - a} \mathbf{R}_1 \mathbf{R}_1^* & -\frac{as_4^2}{s_4^2 - a} \mathbf{R}_1 \mathbf{R}_2^* \\ -\frac{as_4^2}{s_4^2 - a} \mathbf{R}_2 \mathbf{R}_1^* & s_3^2 \mathbf{I}_d - \frac{as_4^2}{s_4^2 - a} \mathbf{R}_2 \mathbf{R}_2^* \end{bmatrix} \text{ and } \mathbf{c} = -\frac{a}{s_4^2 - a} \begin{bmatrix} \mathbf{R}_1 \\ \mathbf{R}_2 \end{bmatrix} \mathbf{p}_2$$

For that, we use the approach of [GM18] by computing successive Schur complements and updated centers, in its iterative formulation as in [AGJ⁺24, Alg. 3.1] for example.

1. $(\mathbf{S}_{2d}, \mathbf{c}_{2d}) \leftarrow (\mathbf{S}''_{\mathbf{p}}, \mathbf{c})$
2. **for** $i = 2d, \dots, 1$ **do**
3. Write $\mathbf{S}_i, \mathbf{c}_i$ as $\mathbf{S}_i = \begin{bmatrix} \mathbf{S}'_i & \mathbf{s}_i \\ \mathbf{s}_i^* & f_i \end{bmatrix}$ and $\mathbf{c}_i = \begin{bmatrix} \mathbf{c}'_i \\ d_i \end{bmatrix}$.
4. $p_i \leftarrow \mathcal{D}_{R, \sqrt{M_r(f_i)}, d_i}$ ▷ with $\text{SampleFz}(f_i, d_i)$
5. $\mathbf{c}'_{i-1} \leftarrow \mathbf{c}'_i + f_i^{-1} (p_i - d_i) \mathbf{s}_i$.
6. $\mathbf{S}_{i-1} \leftarrow \mathbf{S}'_i - f_i^{-1} \mathbf{s}_i \mathbf{s}_i^*$.

and define $[\mathbf{p}_{L,1}^T | \mathbf{p}_{1,2}^T]^T = [p_1, \dots, p_{2d}]^T$. We now note that in the first d iterations of the loop (from $i = 2d$ to $i = d + 1$ included), the f_i and $f_i^{-1} \mathbf{s}_i$ only depend on $\mathbf{R}_1, \mathbf{R}_2$ and can thus be precomputed at key generation. In particular, the key generation would precompute all the f_{d+1}, \dots, f_{2d} and the $f_{2d}^{-1} \mathbf{s}_{2d}, \dots, f_{d+1}^{-1} \mathbf{s}_{d+1}$, as well as the part of \mathbf{S}_d which does not depend on \mathbf{t} . Indeed, the update rule of \mathbf{S}_i and the form of $\mathbf{S}''_{\mathbf{p}}$ clearly yields $\mathbf{S}_d = F(\mathbf{R}_1, \mathbf{R}_2) - att^* \mathbf{I}_d$, where $F(\mathbf{R}_1, \mathbf{R}_2)$ is a function of $s_1, s_4, a, \mathbf{R}_1, \mathbf{R}_2$ which can be precomputed at key generation. For the remaining d iterations, the Schur complements depend on \mathbf{tt}^* which may differ at each signing. The corresponding f_i would then be computed online. All things considered, we end up with the following perturbation sampler, which limits as much as possible the online covariance computations.

Algorithm 4.3: SamplePerturb($\mathbf{R}_1, \mathbf{R}_2, \mathbf{t}, s_{\mathbf{G}}, s_1, s_2, s_3, s_4$)

- Input:** Trapdoor $\mathbf{R}_1, \mathbf{R}_2 \in R^{d \times d(k-\ell)}$, tag $\mathbf{t} \in R_q^\times$, Gaussian parameters $s_{\mathbf{G}}, s_1, s_2, s_3, s_4$.
- Precomputed:** $(f_i)_{i \in [d+1, 2d]} \in K_{\mathbb{R}}^{++}$, $(f_i^{-1} \mathbf{s}_i)_{i \in [d+1, 2d]}$, and $F(\mathbf{R}_1, \mathbf{R}_2) \in K_{\mathbb{R}}^{d \times d}$.
1. **for** $i = 1, \dots, d(\ell - 1)$ **do**

2. $[\mathbf{p}_{L,2}]_i \leftarrow \text{SampleFz}(s_2^2 - s_{\mathbf{G}}^2 \mathbf{t} \mathbf{t}^*, 0)$ $\triangleright \mathcal{D}_{R, \sqrt{M_\tau(s_2^2 - s_{\mathbf{G}}^2 \mathbf{t} \mathbf{t}^*)}}$
3. $\mathbf{p}_2 \leftarrow \mathcal{D}_{R^{d(k-\ell)}, \sqrt{s_4^2 - s_{\mathbf{G}}^2}}$
4. $\mathbf{c}_{2d} \leftarrow -\frac{s_{\mathbf{G}}^2}{s_4^2 - s_{\mathbf{G}}^2} \begin{bmatrix} \mathbf{R}_1 \\ \mathbf{R}_2 \end{bmatrix} \mathbf{p}_2$
5. **for** $i = 2d, \dots, d+1$ **do**
6. Parse $\mathbf{c}_i = [\mathbf{c}'_i | d_i]^T$.
7. $[\mathbf{p}_{1,2}]_{i-d} \leftarrow \text{SampleFz}(f_i, d_i)$.
8. $\mathbf{c}_{i-1} \leftarrow \mathbf{c}'_i + ([\mathbf{p}_{1,2}]_{i-d} - d_i) f_i^{-1} \mathbf{s}_i$.
9. $\mathbf{S}_d \leftarrow F(\mathbf{R}_1, \mathbf{R}_2) - s_{\mathbf{G}}^2 \mathbf{t} \mathbf{t}^* \mathbf{I}_d$.
10. **for** $i = d, \dots, 1$ **do**
11. Parse $\mathbf{S}_i = \begin{bmatrix} \mathbf{S}'_i & \mathbf{s}_i \\ \mathbf{s}_i^* & f_i \end{bmatrix}$ and $\mathbf{c}_i = \begin{bmatrix} \mathbf{c}'_i \\ d_i \end{bmatrix}$.
12. $[\mathbf{p}_{L,1}]_i \leftarrow \text{SampleFz}(f_i, d_i)$.
13. $\mathbf{c}_{i-1} \leftarrow \mathbf{c}'_i + ([\mathbf{p}_{L,1}]_i - d_i) f_i^{-1} \mathbf{s}_i$.
14. $\mathbf{S}_{i-1} \leftarrow \mathbf{S}'_i - f_i^{-1} \mathbf{s}_i \mathbf{s}_i^*$.

Output: $\mathbf{p} = [\mathbf{p}_{L,1}^T | \mathbf{p}_{L,2}^T | \mathbf{p}_{1,2}^T | \mathbf{p}_2^T]^T \in R^{d(k+1)}$.

The precomputed elements can be stored within a single matrix $\mathbf{F} \in K_{\mathbb{R}}^{2d \times 2d}$ as

$$\mathbf{F} = \begin{bmatrix} \begin{array}{|c|c|} \hline F(\mathbf{R}_1, \mathbf{R}_2) & \mathbf{f}_{d+1} \\ \hline \end{array} & \begin{array}{|c|} \hline \mathbf{f}_{2d} \\ \hline \end{array} \\ \hline \begin{array}{|c|c|} \hline \mathbf{s}_{d+1}^* & f_{d+1} \\ \hline \end{array} & \begin{array}{|c|} \hline \mathbf{f}_{2d} \\ \hline \end{array} \\ \hline \dots & \dots \\ \hline \begin{array}{|c|c|} \hline \mathbf{s}_{2d}^* & f_{2d} \\ \hline \end{array} & \begin{array}{|c|} \hline \mathbf{f}_{2d} \\ \hline \end{array} \end{bmatrix}, \text{ with } \mathbf{f}_i = f_i^{-1} \mathbf{s}_i$$

The vectors \mathbf{s}_i^* in gray are necessary to compute $F(\mathbf{R}_1, \mathbf{R}_2)$ and are therefore stored by default, but they are no longer needed in the sampling procedure, and can thus be discarded after key generation. From the description of Algorithm 4.3, we can see it does not differ much compared to the elliptic perturbation sampler described in [AGJ⁺24, Alg. 3.1] except for the additional sampling of $\mathbf{p}_{L,2}$ and the fact that the covariance for the $\mathbf{p}_{L,1}$ depends on \mathbf{t} . We can therefore easily adapt the security and precision analysis. Based on the analysis of [GM18] and the parameter chosen in Theorem 4.2, each sample from SampleFz is within $[1/e^{n-1}, e^{n-1}]$ of the ideal distribution where $\epsilon = (1 + \epsilon)/(1 - \epsilon)$, if we assume a perfect integer sampler for $\mathcal{D}_{\mathbb{Z}, s, c}$. Then, the sampler involves $2d$ Schur complement convolutions [GM18, Lem. 4.3], noticing that the sampling of $\mathbf{p}_{L,2}$ is perfectly independent from the rest. Combined with the $d(\ell-1) + 2d = d(\ell+1)$ calls to SampleFz , it holds that the distribution of \mathbf{p} is within $[1/e^{d(\ell+1)(n-1)+2d}, \epsilon^{d(\ell+1)(n-1)+2d}]$ of $\mathcal{D}_{R^{d(k+1)}, \sqrt{M_\tau(\mathbf{S}_p)}}$. If one performs

the gadget sampling in step 4 of Algorithm 4.2 using the Klein sampler [Kle00] as is done in [AGJ+24], the distribution outputted by Algorithm 4.2 is within $[\delta^{-1}, \delta]$ of the ideal distribution with

$$\delta = \left(\frac{1 + \varepsilon/ndk}{1 - \varepsilon/ndk} \right)^{ndk} \left(\frac{1 + \varepsilon}{1 - \varepsilon} \right)^{d(\ell+1)(n-1)+2d+2}$$

The floating-point precision analysis carried in [AGJ+24] can be adapted simply to our sampler, and it shows that the standard 53 bits precision is sufficient to incur no noticeable security loss. A concrete performance assessment of our new perturbation sampler, along with a comparison with the one for full gadget, is provided in Section 5.2.

5 Applications

The MP sampler is the cornerstone of many privacy-preserving authentication mechanisms because it yields a standard model signature that smoothly interacts with zero-knowledge proofs. This signature was described in the original paper [MP12] and was progressively improved, leading to several variants culminating, to our knowledge, to the construction implicitly used in [AGJ+24]. As we discuss in Section 1, all these variants use a full gadget because they cannot fulfil the requirements of the approximate sampler by Chen et al. [CGM19]. Our sampler solves this problem and thus allows to leverage truncated gadgets to build a more efficient standard model signature scheme and in turn more efficient privacy-preserving authentication mechanisms.

In the first part of this section we indeed show that our approach yields a 30% improvement of the size of the standard model signature based on the MP sampler, leading to signatures of around 4.82 KB instead of 6.72 KB [AGJ+24]. We then show that this 2 KB decrease is considerably amplified when we plug the resulting signature in privacy-preserving authentication mechanisms, which is straightforward given their modular structure. We illustrate it in the cases of anonymous credentials [AGJ+24], blind signatures [JS24] and group signatures [LNPS21, LNP22], where the decrease can reach up to 15 KB, but any other primitive requiring to prove knowledge of a signature should greatly benefit from the reduced dimension enabled by our truncated gadget. Obviously, the improvement ratios are smaller than 30% in these applications as signature is just one of the components of such advanced primitives that blend encryption, commitment and zero-knowledge proofs but they are still significant, ranging from 10% to 17%.

5.1 Standard Model Signature

As a first step to benchmark the improvement procured by our sampler, we consider lattice-based signatures in the standard model based on the MP sampler. To our knowledge, the most efficient one based on well-studied lattice assumptions such as M-SIS and M-LWE is described in [AGJ+24]. In a nutshell, it

defines a public key as $\mathbf{B} = \mathbf{R}_1 + \mathbf{A}\mathbf{R}_2 \bmod q$, where \mathbf{R}_1 and \mathbf{R}_2 constitute the secret trapdoor for the sampler. To generate a signature on a message $m \in T_1$, the signer selects a tag \mathbf{t} , samples some vector \mathbf{v}_3 and then uses a tag-friendly gadget-based sampler (such as the MP one [MP12]) to get $(\mathbf{v}_1, \mathbf{v}_2)$ such that

$$[\mathbf{I}_d | \mathbf{A}] \mathbf{v}_1 + (\mathbf{t}\mathbf{G} - (\mathbf{R}_1 + \mathbf{A}\mathbf{R}_2)) \mathbf{v}_2 = \mathbf{u} + m\mathbf{d} - \mathbf{A}_3 \mathbf{v}_3 \bmod q.$$

The third block \mathbf{A}_3 and the vector \mathbf{u} are necessary for the security reduction. However, they, as well as \mathbf{A}' and \mathbf{d} , can be part of the public parameters common to several signers. Thanks to the trapdoor switching technique of [AGJ+24], the third block \mathbf{A}_3 has only k columns. Finally, the message appears in committed form in $m\mathbf{d}$. Whether it is generated as the output of some hash function H or not depends on the use-case, the point is that, in all cases, security only requires collision-resistance from H and not any form of reprogrammability, which avoids the random oracle model.

Plugging our truncated sampler is almost straightforward here and only requires very minor adjustments. The resulting scheme is described below for completeness.

Algorithm 5.1: Setup

Input: Security parameter λ .

1. Choose $d \in \mathbb{N}^\times$.
2. Choose $\kappa \leq n$ to be a power of two.
3. Choose $q \in \mathbb{N}^\times$ prime s.t. $q = 2\kappa + 1 \bmod 4\kappa$ and $q \geq \sqrt{\kappa}^\kappa$.
4. Choose $w \in \mathbb{N}^\times$ st $\binom{n}{w} \geq Q$.
5. Choose $b \in \mathbb{N}^\times \cap [2, \sqrt{q}]$.
6. $\mathcal{T}_w \leftarrow \{\mathbf{t} \in T_1 : \|\mathbf{t}\|_1 = w\}$.
7. $k \leftarrow \lceil \log_b q \rceil$.
8. Choose $\ell < k$.
9. $B_R \leftarrow \frac{3}{4}(\sqrt{nd} + \sqrt{nd(k-\ell)} + 6)$
10. $\mathbf{G}_L = [1 \dots b^{\ell-1}] \otimes \mathbf{I}_d \in R_q^{d \times d\ell}$.
11. $\mathbf{G}_H = [b^\ell \dots b^{k-1}] \otimes \mathbf{I}_d \in R_q^{d \times d(k-\ell)}$.
12. $\mathbf{G} = [\mathbf{G}_L | \mathbf{G}_H]$
13. $\eta \leftarrow \sqrt{\ln(2ndk/\varepsilon)}/\pi$. $\triangleright \eta \approx \eta_\varepsilon(R^{dk})$ by Lemma 2.2
14. $s_G \leftarrow \eta\sqrt{b^2 + 1}$.
15. $s_1 \leftarrow \eta(b + 1/b)\sqrt{4w^2 + 3B_R^2}$.
16. $s_2 \leftarrow \eta(b + 1/b) \cdot 2w$.
17. $s_{1,1} \leftarrow \eta(b + 1/b)\sqrt{4w^2(b^{2\ell} - 1)/(b^2 - 1) + 3B_R^2}$.
18. $s_3 \leftarrow \sqrt{3}\eta(b + 1/b)B_R$.
19. $s_4 \leftarrow \sqrt{3}\eta(b + 1/b)$.
20. $(\alpha_{1,1}, \alpha_{1,2}) \leftarrow (s_{1,1}/(n\sqrt{d/2}), s_3/(n\sqrt{d/2}))$.
21. $M_{1,i} \leftarrow \exp(\pi/\alpha_{1,i}^2)$.
22. $\mathbf{d} \leftarrow U(R_q^d)$.
23. $\mathbf{A} \leftarrow U(R_q^{d \times d})$.
24. $\mathbf{A}_3 \leftarrow U(R_q^{d \times (k-\ell)})$.
25. $\mathbf{u} \leftarrow U(R_q^d)$.

Output: $\text{pp} = (\lambda, n, d, q, w, b, k, \ell, \eta, s_G, s_1, s_2, s_3, s_4, \mathbf{d}, \mathbf{A}, \mathbf{A}_3, \mathbf{u})$.

Algorithm 5.2: KeyGen

Input: Public parameters pp as in Algorithm 5.1.

1. $\mathbf{R}_1, \mathbf{R}_2 \leftarrow \mathcal{B}_1^{d \times d(k-\ell)}$ conditioned on $\|\mathbf{R}_i\|_2 \leq B_R$.
2. $\mathbf{B} \leftarrow \mathbf{R}_1 + \mathbf{A}\mathbf{R}_2 \bmod qR \in R_q^{d \times d(k-\ell)}$.

Output: $\text{pk} = \mathbf{B}$, and $\text{sk} = \mathbf{R}$.

Algorithm 5.3: Sign

Input: Signing key sk , Message $\mathbf{m} \in \{0, 1\}^*$, Public key pk , Public Parameters pp , State st

1. $m \leftarrow \mathcal{H}(\mathbf{m}) \in T_1$.
2. $\mathbf{c} \leftarrow m\mathbf{d} \bmod qR$.
3. $\mathbf{t} \leftarrow F(\text{st}) \in \mathcal{T}_w$.
4. $\mathbf{v}_3 \leftarrow \mathcal{D}_{R^{k-\ell}, s_4}$.
5. $(\mathbf{v}_{1,1}, \mathbf{v}_{1,2}, \mathbf{v}_2) \leftarrow \text{RingTruncatedSampler}(\mathbf{R}_1, \mathbf{R}_2, \mathbf{A}, \mathbf{u} + m\mathbf{d} - \mathbf{A}_3\mathbf{v}_3, \mathbf{t}\mathbf{d}, s_{\mathbf{G}}, s_1, s_2, s_3, s_4)$
6. **if** $\|\mathbf{v}_{1,1}\|_2 > B_{1,1} \vee \|\mathbf{v}_{1,2}\|_2 > B_{1,2} \vee \|\mathbf{v}_2\|_2 > B_2 \vee \|\mathbf{v}_3\|_2 > B_3$ **goto** 4).
7. $\text{st} \leftarrow \text{st} + 1$.

Output: $\text{sig} = (\mathbf{t}, \mathbf{v}_{1,2}, \mathbf{v}_2, \mathbf{v}_3)$.

Algorithm 5.4: Verify

Input: Public key pk , Message $\mathbf{m} \in \{0, 1\}^*$, Signature sig , Public Parameters pp .

1. $\mathbf{v}_{1,1} \leftarrow \mathbf{u} + \mathcal{H}(\mathbf{m})\mathbf{d} - \mathbf{A}\mathbf{v}_{1,2} - (\mathbf{t}\mathbf{G}_H - \mathbf{B})\mathbf{v}_2 - \mathbf{A}_3\mathbf{v}_3 \bmod qR \in R^d$.
2. $b_1 \leftarrow \|\mathbf{v}_{1,1}\|_2 \leq B_{1,1}$. $\triangleright B_{1,1} = c_{nd}s_{1,1}\sqrt{nd}$
3. $b_2 \leftarrow \|\mathbf{v}_{1,2}\|_2 \leq B_{1,2}$. $\triangleright B_{1,2} = c_{nd}s_3\sqrt{nd}$
4. $b_3 \leftarrow \|\mathbf{v}_2\|_2 \leq B_2$. $\triangleright B_2 = c_{nd(k-\ell)}s_4\sqrt{nd(k-\ell)}$
5. $b_4 \leftarrow \|\mathbf{v}_3\|_2 \leq B_3$. $\triangleright B_3 = c_{n(k-\ell)}s_4\sqrt{n(k-\ell)}$
6. $b_5 \leftarrow \mathbf{t} \in \mathcal{T}_w$.

Output: $b_1 \wedge b_2 \wedge b_3 \wedge b_4 \wedge b_5$.

\triangleright 1 if valid, 0 otherwise

The tailcuts c_N in the verification bounds are set so that the Gaussian tailcut is verified with probability at least $1 - p$, for say $p = 2^{-20}$ to avoid too many repetitions, using [Ban93, Lem. 1.5]. We also take a bound B_R on $\|\mathbf{R}_i\|_2$ that is slightly larger than the one from [AGJ⁺24, Lem. 2.2] to limit the number of rejections and thus spectral norm estimations during key generation.

Theorem 5.1. *The signature scheme of Algorithms 5.1, 5.2, 5.3 and 5.4 is unforgeable in the standard model based on the hardness of M-LWE, M-SIS, and the collision resistance of \mathcal{H} . More precisely, the advantage of PPT adversary in breaking the unforgeability of the signature is upper-bounded by*

$$\begin{aligned} \text{Adv}[\mathcal{A}] \lesssim & \varepsilon_{cr}(\mathcal{H}) + 2 \max \left(h^{\text{od}} \left(C(|\mathcal{T}_w| - Q) \varepsilon_{\text{M-SIS}}^{\bullet} \right), \right. \\ & \left. \varepsilon_{\text{M-LWE}} + \frac{1 + \varepsilon}{1 - \varepsilon} \cdot \frac{4C^2 M_{1,1} M_{1,2}}{(1 - p)^4} h^{\text{od}} \left(CQ \varepsilon_{\text{M-SIS}}^{\bullet} \right) \right) \end{aligned}$$

where $\varepsilon_{\text{cr}}(\mathcal{H})$ is the probability of \mathcal{A} finding a collision for \mathcal{H} , $\varepsilon_{\text{M-LWE}}$ the hardness bound of M-LWE $_{n,d,q,B_1}$, and $\varepsilon_{\text{M-SIS}}^{\bullet}$ and $\varepsilon_{\text{M-SIS}}^{\circ}$ are the respective hardness bounds of M-SIS $_{n,d,2d+k-\ell+2,q,\beta_{\bullet}}$ and M-SIS $_{n,d,2d+k-\ell,q,\beta_{\circ}}$. The constant $C \approx 2$ is the one from [AGJ⁺24, Lem. 2.3], and the probability p corresponds to the tailcut probability used to set the Gaussian verification bounds. The function $h^{\circ d}$ corresponds to the d -th composition power of the function h defined by

$$h(x) = (k - \ell)\varepsilon_{\text{M-LWE}} + \delta \left(2(k - \ell)\varepsilon_{\text{M-LWE}} + \delta \left((k - \ell)\varepsilon_{\text{M-LWE}} + x \right)^{\frac{2\lambda-1}{2\lambda}} \right)^{\frac{2\lambda-1}{2\lambda}},$$

with

$$\delta = 1 + Q(\lambda - 1/2) \cdot \left(\left(\frac{1 + \varepsilon}{1 - \varepsilon} \right)^{2d(\ell+1)(n-1)+4d+4} \left(\frac{1 + \varepsilon/ndk}{1 - \varepsilon/ndk} \right)^{2ndk} - 1 \right)^2$$

and with

$$\beta_{\bullet} = \sqrt{\left(\sqrt{B_{1,1}^2 + B_{1,2}^2 + \sqrt{nd}B_2} \right)^2 + B_3^2 + n + 1}$$

$$\beta_{\circ} = \sqrt{\left(2\sqrt{B_{1,1}^2 + B_{1,2}^2 + \sqrt{nd}\sqrt{4B_2^2 + n}} \right)^2 + 4B_3^2}$$

Proof. We follow the blueprint of the standard model signature of [AGJ⁺24], with only a few modification due to our change of sampler. We call G_0 the regular unforgeability game where the challenger generates the public parameters and keys legitimately using Algorithms 5.1 and 5.2, and where it answers signing queries using Algorithm 5.3. We then progressively change G_0 so that the challenger can use the forgery of \mathcal{A} to solve an M-SIS instance. For the subsequent games G_i , we denote by $\text{Adv}_{G_i}[\mathcal{A}]$ the advantage of the adversary \mathcal{A} in producing a valid forgery in the modified game G_i .

Game G_1 . We first change the tag generation. We instead generate all the tags $\{\mathbf{t}^{(i)}; i \in [Q]\}$ at the outset of the game instead of at each signature issuance. The view of \mathcal{A} remains unchanged, yielding that G_1 is identically distributed as G_0 .

Game G_2 . In this game, the challenger proceeds as in G_1 and thus eventually receives $(\mathbf{m}^*, \text{sig}^*)$ where \mathbf{m}^* differs from all the queried messages. It then computes the corresponding digest $m^* = \mathcal{H}(\mathbf{m}^*)$ and aborts if a collision occurs with the hashes of the queried message. In the latter case, \mathcal{A} can readily be used against the collision resistance of \mathcal{H} and we thus get

$$\text{Adv}_{G_1}[\mathcal{A}] \leq \text{Adv}_{G_2}[\mathcal{A}] + \varepsilon_{\text{cr}}(\mathcal{H}).$$

Game G_3 . We now introduce a branching of our security reduction, similarly to what is done in [JS24]. At the outset, the challenger samples $\rho \leftarrow U(\{1, 2\})$. If $\rho = 1$, the challenger expects what we call a type \bullet forgery which corresponds to a forgery where the tag \mathbf{t}^* is not in $\{\mathbf{t}^{(i)}; i \in [Q]\}$. On the other hand, if $\rho = 2$,

the challenger expects a type \bullet forgery which corresponds to a forgery where the tag is among the emitted ones, i.e., there exists $i \in [Q]$ such that $\mathbf{t}^* = \mathbf{t}^{(i)}$. The challenger aborts if the type guess turns out to be wrong. We then get

$$\text{Adv}_{G_2}[\mathcal{A}] = 2\text{Adv}_{G_3}[\mathcal{A}].$$

Depending on the value of ρ , the reduction will now behave differently. We define $\text{Adv}_G^\bullet[\mathcal{A}]$ (resp. $\text{Adv}_G^\circ[\mathcal{A}]$) to be the advantage of \mathcal{A} in producing a type \bullet (resp. \circ) forgery in game G . With the prior modification, if the reduction does not abort, the type guess is correct and thus $\text{Adv}_{G_3}[\mathcal{A}] \leq \max(\text{Adv}_{G_3}^\bullet[\mathcal{A}], \text{Adv}_{G_3}^\circ[\mathcal{A}])$. In the next games, we specify which branch is impacted by the change. The first change we introduce is that after sampling ρ and the set of tags, the reduction samples $\mathbf{t}^+ \leftarrow U(\mathcal{T}_w \setminus \{\mathbf{t}^{(i)}; i \in [Q]\})$ if $\rho = 1$, and otherwise it samples $i^+ \leftarrow U([Q])$ and sets $\mathbf{t}^+ = \mathbf{t}^{(i^+)}$ if $\rho = 2$. Because \mathbf{t}^+ is not used anywhere yet, it does not change the view of \mathcal{A} . We thus have

$$\text{Adv}_{G_2}[\mathcal{A}] \leq 2 \max\left(\text{Adv}_{G_3}^\bullet[\mathcal{A}], \text{Adv}_{G_3}^\circ[\mathcal{A}]\right)$$

Game G_4 ($\rho = 2$). If $\rho = 2$, the challenger hides an M-LWE secret in \mathbf{d} . More precisely, it samples $\mathbf{s}_1, \mathbf{s}_2$ from \mathcal{B}_1^d and defines $\mathbf{d} = \mathbf{s}_1 + \mathbf{A}\mathbf{s}_2 \bmod qR$. Distinguishing between the view in G_3 and G_4 is exactly solving M-LWE $_{n,d,d,q,\mathcal{B}_1}$. As it only impacts the branch $\rho = 2$, we have

$$\text{Adv}_{G_3}^\bullet[\mathcal{A}] = \text{Adv}_{G_4}^\bullet[\mathcal{A}], \text{ and } \text{Adv}_{G_3}^\circ[\mathcal{A}] \leq \text{Adv}_{G_4}^\circ[\mathcal{A}] + \varepsilon_{\text{M-LWE}},$$

where $\varepsilon_{\text{M-LWE}}$ is the hardness bound of M-LWE $_{n,d,d,q,\mathcal{B}_1}$.

Game G_5 ($\rho = 2$). We now hide a short relation in \mathbf{u} when $\rho = 2$. The challenger samples $\mathbf{v}_{1,1} \leftarrow \mathcal{D}_{R^d, s_{1,1}}$, $\mathbf{v}_{1,2} \leftarrow \mathcal{D}_{R^d, s_3}$ and $[\mathbf{v}_2^T | \mathbf{v}_3^T]^T \leftarrow \mathcal{D}_{R^{(k-\ell)(d+1)}, s_4}$, and computes $\mathbf{u} = \mathbf{v}_{1,1} + \mathbf{A}\mathbf{v}_{1,2} + (\mathbf{t}^+ \mathbf{G}_H - \mathbf{B})\mathbf{v}_2 + \mathbf{A}_3\mathbf{v}_3 \bmod qR$. By the regularity lemma of [GPV08, Cor. 5.2], and with our parameter choices satisfying the smoothing condition, we then get that

$$\text{Adv}_{G_4}^\bullet[\mathcal{A}] = \text{Adv}_{G_5}^\bullet[\mathcal{A}], \text{ and } \text{Adv}_{G_4}^\circ[\mathcal{A}] \in \left[\frac{1}{1+\varepsilon}, \frac{1+\varepsilon}{1-\varepsilon} \right] \text{Adv}_{G_5}^\circ[\mathcal{A}],$$

Game G_6 ($\rho = 2$). If $\rho = 2$, we enforce a bound on $m^+ \mathbf{s}_1, m^+ \mathbf{s}_2$ to prepare the rejection sampling argument of the next games. We also enforce the Gaussian tail bound on \mathbf{v}_2 and \mathbf{v}_3 for the last step of the proof. For the first, we use the heuristic bound from [AGJ⁺24, Lem. 2.3]. Concretely, for the i^+ -th query in branch $\rho = 2$, after computing $m^+ = \mathcal{H}(\mathbf{m}^{(i^+)})$ we abort the reduction altogether if $\|\mathbf{s}_i m^+\|_2 > \sqrt{nd} \|m^+\|_2$. The probability of the reduction not aborting is then $1/C^2$ for a small constant $C \approx 2$. We also abort if $\|\mathbf{v}_2\|_2 > B_2$ or $\|\mathbf{v}_3\|_2 > B_3$. Because of our choice of tailcut, it holds that the bound on \mathbf{v}_2 or \mathbf{v}_3 are both verified with probability $(1-p)^2$. It then holds that

$$\text{Adv}_{G_5}^\bullet[\mathcal{A}] = \text{Adv}_{G_6}^\bullet[\mathcal{A}], \text{ and } \text{Adv}_{G_5}^\circ[\mathcal{A}] \leq \frac{C^2}{(1-p)^2} \text{Adv}_{G_6}^\circ[\mathcal{A}],$$

Game G_7 ($\rho = 2$). We keep preparing the rejection sampling step for the i^+ -th query in the branch $\rho = 2$. More precisely, in the i^+ -th query, the reduction samples $(\mathbf{v}_{1,1}^{(+)}, \mathbf{v}_{1,2}^{(+)}, \mathbf{v}_2^{(+)}, \mathbf{v}_3^{(+)})$ legitimately using the preimage sampler, and then rejects based on the value of $(\mathbf{v}_{1,1}, \mathbf{v}_{1,2}, \mathbf{v}_2, \mathbf{v}_3)$ which are so far independent of $(\mathbf{v}_{1,1}^{(+)}, \mathbf{v}_{1,2}^{(+)}, \mathbf{v}_2^{(+)}, \mathbf{v}_3^{(+)})$. More precisely, it samples $u_1, u_2 \leftarrow U([0, 1])$. The reduction continues only if $u_1 \leq 1/M_{1,1}$, $u_2 \leq 1/M_{1,2}$, and $\langle \mathbf{v}_{1,1}, \mathbf{s}_1 m^+ \rangle \geq 0$ and $\langle \mathbf{v}_{1,2}, \mathbf{s}_2 m^+ \rangle \geq 0$, otherwise it aborts. We insist that if the challenger does not abort, it answers the query with $(\mathbf{v}_{1,1}^{(+)}, \mathbf{v}_{1,2}^{(+)}, \mathbf{v}_2^{(+)}, \mathbf{v}_3^{(+)})$ and not the hidden relation of \mathbf{u} . It then holds that

$$\text{Adv}_{G_6}^{\bullet}[\mathcal{A}] = \text{Adv}_{G_7}^{\bullet}[\mathcal{A}], \text{ and } \text{Adv}_{G_6}^{\circ}[\mathcal{A}] \leq 4M_{1,1}M_{1,2}\text{Adv}_{G_7}^{\circ}[\mathcal{A}],$$

Game G_8 ($\rho = 2$). We now finally use the relation hidden in \mathbf{u} to answer the i^+ -th query in branch $\rho = 2$. We need however to perform rejection sampling so that the output satisfies the correct equation involving $m^+ = \mathcal{H}(\mathbf{m}^{(i+)})$, while producing the correct signature distribution. The reduction thus samples $u_1, u_2 \leftarrow U([0, 1])$, computes $\delta_1 = \langle \mathbf{v}_{1,1} + \mathbf{s}_1 m^+, \mathbf{s}_1 m^+ \rangle$ and $\delta_2 = \langle \mathbf{v}_{1,2} + \mathbf{s}_2 m^+, \mathbf{s}_2 m^+ \rangle$. It then aborts the reduction if

$$\begin{aligned} & \delta_1 < 0, \text{ or } u_1 > \frac{1}{M_{1,1}} \exp\left(\frac{\pi}{s_{1,1}^2} \left(\|\mathbf{s}_1 m^+\|_2^2 - 2\delta_1\right)\right), \\ & \text{or } \delta_2 < 0, \text{ or } u_2 > \frac{1}{M_{1,2}} \exp\left(\frac{\pi}{s_3^2} \left(\|\mathbf{s}_2 m^+\|_2^2 - 2\delta_2\right)\right). \end{aligned}$$

If the challenger did not abort, it constructs

$$\mathbf{v}_{1,1}^{(i+)} = \mathbf{v}_{1,1} + \mathbf{s}_1 m^+, \mathbf{v}_{1,2}^{(i+)} = \mathbf{v}_{1,2} + \mathbf{s}_2 m^+ \text{ and } \mathbf{v}_2^{(i+)} = \mathbf{v}_2, \mathbf{v}_3^{(i+)} = \mathbf{v}_3,$$

It also aborts if $\|\mathbf{v}_{1,1}^{(i+)}\|_2 > B_{1,1}$ or $\|\mathbf{v}_{1,2}^{(i+)}\|_2 > B_{1,2}$. If it did not abort, it outputs $(\mathbf{t}^+, \mathbf{v}_{1,2}^{(i+)}, \mathbf{v}_2^{(i+)}, \mathbf{v}_3^{(i+)})$ as the valid signature. Based on the bounds enforced on $\mathbf{s}_i h^+$, and the way the parameters $M_{1,i}$ are set, the rejection sampling argument of [LNS21, Lem. 3.2] shows that the distribution is the same. The subsequent tailcut on $\mathbf{v}_{1,1}^{(i+)}, \mathbf{v}_{1,2}^{(i+)}$ then yield the reduction continuing with probability at least $(1-p)^2$. We then get

$$\text{Adv}_{G_7}^{\bullet}[\mathcal{A}] = \text{Adv}_{G_8}^{\bullet}[\mathcal{A}], \text{ and } \text{Adv}_{G_7}^{\circ}[\mathcal{A}] \leq \frac{1}{(1-p)^2} \text{Adv}_{G_8}^{\circ}[\mathcal{A}].$$

Hybrid Games $G_{j,i}$. For both branches, we now use the hybrid argument used in [AGJ⁺24] to hide the tag \mathbf{t}^+ in the public key \mathbf{B} . The authors rely on a specific partial trapdoor switching method and define hybrid games $G_{j,i}$ for $j \in [d]$ and $i \in [0, 9]$. More precisely, $G_{j,0}$ is essentially G_8 but where $\mathbf{B} = \mathbf{R}_1 + \mathbf{A}\mathbf{R}_2 + \mathbf{t}^+ [b^\ell, \dots, b^{k-1}] \otimes \text{diag}(\mathbf{t}^+, \dots, \mathbf{t}^+, 0, \otimes, 0)$ where \mathbf{t}^+ appears j times. In particular, we note that $G_{1,0} = G_8$. We define $\mathbf{G}_{H,j} = [b^\ell, \dots, b^{k-1}] \otimes \mathbf{e}_j$ where

\mathbf{e}_j j . We then observe that \mathbf{G}_H is an interleaving of the columns of the different $\mathbf{G}_{H,j}$. For a matrix with $d(k-\ell)$ columns (like $\mathbf{R}_1, \mathbf{R}_2$, etc.), we use the index j to denote this submatrix with $k-\ell$ columns. Concretely, if $\mathbf{C} = [\mathbf{c}_1 | \dots | \mathbf{c}_{d(k-\ell)}]$, $\mathbf{C}_j = [\mathbf{c}_j | \mathbf{c}_{d+j} | \dots | \mathbf{c}_{d(k-\ell)+j}]$ for any $j \in [d]$.

Then, $G_{j,1}$ hides a partial gadget in \mathbf{A}_3 as $\mathbf{A}_3 = \mathbf{G}_{H,j} - \mathbf{A}'_3$ for \mathbf{A}'_3 drawn uniformly. In $G_{j,2}$, the challenger hides a relation in \mathbf{A}_3 as $\mathbf{A}_3 = \mathbf{G}_{H,j} - (\mathbf{R}'_{1,j} + \mathbf{A}\mathbf{R}'_{2,j})$ under $\text{M-LWE}_{n,d,d,q,\mathcal{B}_1}^{k-\ell}$. In $G_{j,3}$, the challenger uses the partial trapdoor $\mathbf{R}'_{1,j}, \mathbf{R}'_{2,j}$ instead of $\mathbf{R}_{1,j}, \mathbf{R}_{2,j}$ to produce signatures (except for the i^+ -th query, if $\rho = 2$, which remains unchanged), which is argued by the trapdoor switching lemma of [AGJ⁺24, Lem. 4.1] updated to our new sampler. Then, $G_{j,4}$ simulates the partial public key $\mathbf{B}_j = \mathbf{R}_{1,j} + \mathbf{A}\mathbf{R}_{2,j} \bmod qR$ and instead samples \mathbf{B}_j uniformly in $R_q^{d \times k-\ell}$, which is unbeknownst to \mathcal{A} under $\text{M-LWE}_{n,d,d,q,\mathcal{B}_1}^{k-\ell}$. $G_{j,5}$ adds the tag guess as $\mathbf{B}_j = \mathbf{B}'_j + \mathbf{t}^+ \mathbf{G}_{H,j}$ with \mathbf{B}'_j uniform. In $G_{j,6}$, it re-introduces a partial secret key to get $\mathbf{B}_j = \mathbf{R}_{1,j} + \mathbf{A}\mathbf{R}_{2,j} + \mathbf{t}^+ \mathbf{G}_{H,j}$ under $\text{M-LWE}_{n,d,d,q,\mathcal{B}_1}^{k-\ell}$. The trapdoor switching is used again in $G_{j,7}$ to use the partial trapdoor $\mathbf{R}_{1,j}, \mathbf{R}_{2,j}$ instead of $\mathbf{R}'_{1,j}, \mathbf{R}'_{2,j}$. In $G_{j,8}$, we replace $\mathbf{R}'_{1,j} + \mathbf{A}\mathbf{R}'_{2,j}$ by a uniform \mathbf{A}'_3 again under $\text{M-LWE}_{n,d,d,q,\mathcal{B}_1}^{k-\ell}$. Finally, \mathbf{A}_3 is again changed to be perfectly uniform in $G_{j,9}$ so that $G_{j,9} = G_{j+1,0}$.

Notice that in the hybrid games, the effective tag we need to used in the sampling procedure is $\mathbf{T} = \text{diag}(\mathbf{t} - \mathbf{t}^+, \dots, \mathbf{t} - \mathbf{t}^+, \mathbf{t}, \dots, \mathbf{t})$. As our sampler supports such tag matrices, we only need to make sure that \mathbf{T} is invertible in R_q , and that the parameters s_1, s_2 are set so as to consider a bound on $\|\mathbf{T}\|_2$ during each step of the hybrid. Regarding invertibility, we note that the sampling is only done for $\mathbf{t} \neq \mathbf{t}^+$ as the i^+ -th query is handled differently. We thus get that \mathbf{T} is diagonal with non-zero ternary polynomials on the diagonal. Using [LS18, Cor. 1.2], our choice of q , and in particular the splitting behavior of qR in R , entails that $\|\mathbf{t} - \mathbf{t}^+\|_\infty = 1 < q^{1/\kappa}/\sqrt{\kappa}$, thus proving $\mathbf{t} - \mathbf{t}^+ \in R_q^\times$. The same holds for \mathbf{t} which proves that $\mathbf{T} \in GL_d(R_q)$. Then, we have that $\|\mathbf{T}\|_2 \leq \max(\|\mathbf{t} - \mathbf{t}^+\|_1, \|\mathbf{t}\|_1) \leq 2w$ by definition of the tag space. This explains why s_2 is set with $2w$ and not just w , and that s_1^2 features a term in $4w^2$. All things considered, the analysis of the hybrid argument is exactly the same as in [AGJ⁺24], and it thus holds that by looping over $j \in [d]$, we have

$$\text{Adv}_{G_9}^\bullet[\mathcal{A}] \lesssim h^{od} \left(\text{Adv}_{G_{d,9}}^\bullet[\mathcal{A}] \right), \text{ and } \text{Adv}_{G_9}^\bullet[\mathcal{A}] \lesssim h^{od} \left(\text{Adv}_{G_{d,9}}^\bullet[\mathcal{A}] \right),$$

where h^{od} corresponds to the d -th composition power of the function h , which is itself defined by

$$h(x) = (k-\ell)\varepsilon_{\text{M-LWE}} + \delta \left(2(k-\ell)\varepsilon_{\text{M-LWE}} + \delta \left((k-\ell)\varepsilon_{\text{M-LWE}} + x \right)^{\frac{2\lambda-1}{2\lambda}} \right)^{\frac{2\lambda-1}{2\lambda}},$$

with

$$\delta = 1 + Q(\lambda - 1/2) \cdot \left(\left(\frac{1+\varepsilon}{1-\varepsilon} \right)^{2d(\ell+1)(n-1)+4d+4} \left(\frac{1+\varepsilon/ndk}{1-\varepsilon/ndk} \right)^{2ndk} - 1 \right)^2,$$

by the detailed loss of Algorithm 4.2 and the relative error lemma of [Pre17].

Game G_9 . In this game, the challenger aborts if the tag guess is incorrect. More precisely, in branch $\rho = 1$, the adversary must return a type \bullet forgery which means that $\mathbf{t}^+ \notin \{\mathbf{t}^{(i)}; i \in [Q]\}$. Because \mathbf{t}^+ is hidden to the view of the adversary, the guess is correct with probability $1/(|\mathcal{T}_w| - Q)$. For the branch $\rho = 2$, the adversary must return a type \circ forgery meaning there exists $i^* \in [Q]$ such that $\mathbf{t}^* = \mathbf{t}^{(i^*)}$. The challenger thus aborts if $i^+ \neq i^*$, meaning the guess is correct with probability $1/Q$. We thus get

$$\text{Adv}_{G_{d,9}}^{\bullet}[\mathcal{A}] = (|\mathcal{T}_w| - Q)\text{Adv}_{G_9}^{\bullet}[\mathcal{A}], \text{ and } \text{Adv}_{G_{d,9}}^{\circ}[\mathcal{A}] = Q\text{Adv}_{G_9}^{\circ}[\mathcal{A}],$$

Exploiting the forgery. We now explain for each branch how to exploit the forgery outputted by \mathcal{A} to find a solution of a specific M-SIS instance. More precisely, we bound $\text{Adv}_{G_9}^{\bullet}[\mathcal{A}]$ and $\text{Adv}_{G_9}^{\circ}[\mathcal{A}]$ separately. The challenger indeed receives two instances $\overline{\mathbf{A}} = [\mathbf{I}_d | \mathbf{A} | \mathbf{A}_3 | \mathbf{d} | \mathbf{u}]$ and $\overline{\mathbf{A}}' = [\mathbf{I}_d | \mathbf{A} | \mathbf{A}_3]$ of $\text{M-SIS}_{n,d,2d+k-\ell+2,q,\beta_{\bullet}}$ and $\text{M-SIS}_{n,d,2d+k-\ell,q,\beta_{\bullet}}$ respectively. The first will be used to define the material when $\rho = 1$, while the other will serve for the branch where $\rho = 2$. Depending on the value ρ sampled at the outset, it discards one of these two instances. It then proceeds as in G_9 following the determined branch.

Branch $\rho = 1$. We start by bounding $\text{Adv}_{G_9}^{\bullet}[\mathcal{A}]$, i.e., in branch $\rho = 1$. It holds that $\mathbf{t}^* = \mathbf{t}^+$, which means that $\mathbf{t}^* \mathbf{G}_H - \mathbf{B} = \mathbf{t}^* \mathbf{G}_H - \mathbf{R}_1 - \mathbf{A} \mathbf{R}_2 - \mathbf{t}^+ \mathbf{G}_H = -\mathbf{R}_1 - \mathbf{A} \mathbf{R}_2 \pmod{qR}$. The challenger then aborts if $\| -[\mathbf{R}_1^T | \mathbf{R}_2^T]^T \mathbf{v}_2^* \|_2 > \sqrt{nd} \|\mathbf{v}_2^*\|_2$. Again, by the Johnson-Lindenstrauss bound of [AGJ+24, Lem. 2.4] (as $-\mathbf{R}_1^T | \mathbf{R}_2^T]^T$ is indeed drawn from \mathcal{B}_1), the challenger continues with probability at least $1/C$ for $C \approx 2$. We then re-write the equation recovering $\mathbf{v}_{1,1}^*$ as

$$(\mathbf{v}_{1,1}^* - \mathbf{R}_1 \mathbf{v}_2^*) + \mathbf{A}(\mathbf{v}_{1,2}^* - \mathbf{R}_2 \mathbf{v}_2^*) + \mathbf{A}_3 \mathbf{v}_3^* - m^* \mathbf{d} - \mathbf{u} = \mathbf{0} \pmod{qR},$$

i.e., $[\mathbf{I}_d | \mathbf{A}' | \mathbf{A}_3 | \mathbf{d} | \mathbf{u}] \mathbf{x}^* = \mathbf{0} \pmod{qR}$ with

$$\mathbf{x}^* = \begin{bmatrix} \begin{bmatrix} \mathbf{v}_{1,1}^* \\ \mathbf{v}_{1,2}^* \end{bmatrix} - \begin{bmatrix} \mathbf{R}_1 \\ \mathbf{R}_2 \end{bmatrix} \mathbf{v}_2^* \\ \mathbf{v}_3^* \\ -m^* \\ -1 \end{bmatrix},$$

The last coefficient is non zero which ensures $\mathbf{x}^* \neq \mathbf{0}$. We can then directly bound

$$\|\mathbf{x}^*\|_2 \leq \sqrt{\left(\sqrt{B_{1,1}^2 + B_{1,2}^2} + \sqrt{nd} B_2\right)^2 + B_3^2 + n + 1} = \beta_{\bullet},$$

thus proving that \mathbf{x}^* is a solution of $\text{M-SIS}_{n,d,2d+k-\ell+2,q,\beta_{\bullet}}$. We get $\text{Adv}_{\text{M-SIS}}[\mathcal{A}] \geq \text{Adv}_{G_9}^{\bullet}[\mathcal{A}]/C$ which leads to

$$\text{Adv}_{G_9}^{\bullet}[\mathcal{A}] \leq C \varepsilon_{\text{M-SIS}}^{\bullet}.$$

Branch $\rho = 2$. We now bound $\text{Adv}_{G_{13}}^{\bullet}[\mathcal{A}]$. Recall we are also in the case where $\mathbf{t}^* = \mathbf{t}^+$ so that $\mathbf{t}^* \mathbf{G}_H - \mathbf{B} = -\mathbf{R}_1 - \mathbf{A} \mathbf{R}_2 \bmod qR$. We then define $\Delta \mathbf{v}_1 = [\mathbf{v}_{1,1}^* - \mathbf{v}_{1,1}^{(i^+)}, \mathbf{v}_{1,2}^* - \mathbf{v}_{1,2}^{(i^+)}]$ where $\mathbf{v}_{1,i}^{(i^+)}$ was part of the signature in the i^+ -th query. We also define $\Delta \mathbf{v}_{2m} = [\mathbf{v}_2^* - \mathbf{v}_2^{(i^+)}, m^* - m^+]$. The challenger then aborts if $\|[-\mathbf{R}_1^T | \mathbf{R}_2^T]^T \mathbf{s} | \Delta \mathbf{v}_{2m}\|_2 > \sqrt{nd} \|\Delta \mathbf{v}_{2m}\|_2$ which happens with probability at most $1 - 1/C$ for $C \approx 2$ using [AGJ⁺24, Lem. 2.3] again. Then, because of how \mathbf{u} is set, we have

$$[\mathbf{I}_d | \mathbf{A} | \mathbf{A}_3] \begin{bmatrix} \Delta \mathbf{v}_1 - [[\mathbf{R}_1^T | \mathbf{R}_2^T]^T \mathbf{s} | \Delta \mathbf{v}_{2m}] \\ \mathbf{v}_3^* - \mathbf{v}_3^{(i^+)} \end{bmatrix} = \mathbf{0} \bmod qR.$$

Because of the change in G_2 , it holds that $m^* \neq m^+$. Using the same argument as in previous works, e.g., [LLM⁺16, LNPS21, LNP22, JRS23, AGJ⁺24], the unpredictability of \mathbf{s} ensures that \mathbf{x}^* is non-zero except with negligible probability. We now bound \mathbf{x}^* . Because of the bounds enforced in previous games, we have

$$\|\mathbf{x}^*\|_2 \leq \sqrt{\left(2\sqrt{B_{1,1}^2 + B_{1,2}^2} + \sqrt{nd}\sqrt{4B_2^2 + n}\right)^2 + 4B_3^2} = \beta_{\bullet},$$

thus proving that \mathbf{x}^* is a solution of $\text{M-SIS}_{n,d,2d+k-\ell,q,\beta_{\bullet}}$. We then obtain that $\text{Adv}_{\text{M-SIS}}[\mathcal{A}] \geq \text{Adv}_{G_9}^{\bullet}[\mathcal{A}]/C - \text{negl}(\lambda)$ which leads to

$$\text{Adv}_{G_9}^{\bullet}[\mathcal{A}] \leq C\varepsilon_{\text{M-SIS}}^{\bullet} + \text{negl}(\lambda).$$

Advantage Bound. We can now combine all of the advantage bounds from each game hop with the derived bounds on the advantages in G_9 . We obtain the following equations.

$$\begin{aligned} \text{Adv}_{G_3}^{\bullet}[\mathcal{A}] &\lesssim h^{\circ d} (C(|\mathcal{T}_w| - Q)\varepsilon_{\text{M-SIS}}^{\bullet}), \\ \text{Adv}_{G_3}^{\bullet}[\mathcal{A}] &\lesssim \varepsilon_{\text{M-LWE}} + \frac{1 + \varepsilon}{1 - \varepsilon} \frac{4C^2 M_{1,1} M_{1,2}}{(1 - p)^4} h^{\circ d} (CQ\varepsilon_{\text{M-SIS}}^{\bullet}). \end{aligned}$$

Combining these inequalities gives the claimed security. \square

5.2 Performance

We now compare the performance of our standard model signature scheme to the one from [AGJ⁺24, Sec. 3]. As the parameters of the latter were selected towards being plugged in an anonymous credentials system, we review the parameter selection to be tailored to a standalone signature. In particular, the signed message which previously contained a user secret key and a vector of several attributes is now a single ring element $m = \mathcal{H}(\mathbf{m})$ as in our signature presented above. This allows us to slightly reduce the parameters. For clarity of comparison, we still refer to this parameter-optimized version as [AGJ⁺24]. The latter relies on the elliptic sampler with full gadget, recalled in Section 4.1.

5.2.1 Size. The resulting construction of [AGJ⁺24] yields a standard model signature of 6.72 KB and a public key of 47.5 KB for 125 bits of security in the Core-SVP model. Thanks to our truncated sampler, the scheme described in Section 5.1 yields signatures of 4.82 KB and public keys of 28.5 KB for 120 bits of security (NIST-II level) in the Core-SVP model, which represents an improvement of respectively 28 % and 40 %. In Table 5.1, we report the sizes and security for different values of ℓ , $\ell = 0$ corresponding to [AGJ⁺24]. These sizes are obtained by setting the parameters according to the security reductions to M-SIS and M-LWE and by taking into account the reduction loss. More aggressive parameters (and hence better sizes) could be obtained by only considering the state-of-the-art cryptanalysis. An example parameter set is given in Table A.1.

	pk	t	v _{1,2}	v ₂	v ₃	Tot.	Sec.
$\ell = 0$	47.50 KB	0.03	1.63	4.05	1.01	6.72 KB	125
$\ell = 1$	38.00 KB	0.03	1.63	3.39	0.85	5.90 KB	121
$\ell = 2$	28.50 KB	0.03	1.61	2.54	0.64	4.82 KB	120

Table 5.1. Comparison of the signature sizes for different values of ℓ ($k = 5$). All sizes are in KB. The case $\ell = 0$ corresponds to [AGJ⁺24].

5.2.2 Computational Complexity. We now compare the computational complexity of both approaches for standard model digital signatures. Note that this actually boils down to compare the underlying samplers as they actually represent 95 % of the signature generation process. To be more specific, we compare the performance of our sampler `RingTruncatedSampler` when used to generate signatures in the standard model (as described in Section 5.1) with the most efficient alternative relying on full gadget, namely [AGJ⁺24] implementing the improved elliptic sampler from Section 4.1.

We provide a full implementation in C⁴ of said signatures and benchmark our new samplers. Our implementation is built upon that of [AGJ⁺24] for a fair comparison. Our benchmarks were performed on a laptop with an Intel Core i7 12800H CPU running at 4.6 GHz. We use the same compilation options for both constructions, that is `-O3 -march=native` using `gcc 11.4.0` with `pthread` disabled when building FLINT. The resulting timings are reported in Table 5.2. The main components of both `RingTruncatedSampler` and the elliptic sampler are the `SamplePerturb` algorithm and the Klein sampler whose performance are also indicated in the table. While these timings are already practical, we note that both implementations could be optimized much further. They are however sufficient for comparing both options.

⁴ <https://github.com/truncatedsampler/truncated-sampler>

Logically, the simpler structure of the perturbation sampler in the case of full gadget leads to better performance (about 30 % faster) when generating a signature than in the case of truncated gadgets. Conversely, the signature verification is 30% faster in the latter case thanks to the reduced dimension. However, one must keep in mind that for current applications of standard model signatures (group signatures, anonymous credentials, etc), the main bottleneck is the generation and the verification of the zero-knowledge proof of knowledge of the signature, not the issuance of the latter. For example, in the case of anonymous credentials, the analysis in [AGJ⁺24] shows that the former step is around 10 times slower than the signature issuance in their case. As lattice zero-knowledge proofs are very sensitive to the witness dimension, we expect timing improvements thanks to our truncated sampler, which should compensate the issuance overhead. Moreover, we recall that in such contexts, signature issuance occurs once per user whereas credential showing can be performed many times by the same user, increasing further the performance gains. Evaluating precisely the improvement resulting from our shorter signatures would however require to fully implement the corresponding zero-knowledge proof systems, which we leave for future works.

Procedure	Time (ms)			
	mean	med	min	max
SamplePerturb (full)	51.983	51.838	50.883	53.636
SamplePerturb (truncated)	80.234	80.280	77.287	82.914
Gadget sampler (Klein)	1.821	1.818	1.763	1.906
Gadget sampler (Klein)	1.822	1.820	1.761	1.897
Elliptic sampler	56.526	56.534	55.349	57.931
RingTruncatedSampler	83.931	83.891	79.302	85.769
Sign	56.949	56.953	55.651	58.583
Sign	84.255	84.168	79.898	86.821
Verify	1.127	1.126	1.104	1.191
Verify	0.771	0.771	0.746	0.802

Table 5.2. Benchmark results. Statistics over 100 executions. The highlighted rows correspond to our implementation based on the truncated sampler, while the clear rows correspond to our parameter-optimized version of [AGJ⁺24]. Where applicable, the key and message were randomized.

5.3 Application to Advanced Signatures

Our sampler is designed to retain the main features of the sampler from [MP12] while supporting truncated gadgets. It can therefore be readily adapted to all the constructions that were relying on the MP sampler so far, leading to direct improvements that we describe for some prominent primitives of this area,

for different choices of ℓ (the truncation parameter). Echoing the result in Section 4.1, we can actually consider that existing constructions use our sampler with $\ell = 0$, which allows to unify notations. In other words, our sampler adds a parameter to these systems where $\ell = 0$ corresponds to the state-of-the-art and where any value $\ell > 0$ corresponds to our new sampler. The scripts used to derive the size and security estimations of these advanced primitives are provided alongside our implementation⁵ of the signature of Section 5.1.

5.3.1 Static Group Signature. We start by plugging our sampler in the group signature of [LNPS21][LNP22, Sec. 6.4]. We only give a high level description of the scheme and how our sampler fits into it, and refer to the latter works for a more detailed description. Each user possess an identity defined by a tag t (in a similar space to the one defined in the standard model signature), and gets a secret key $(\mathbf{v}_{1,1}, \mathbf{v}_{1,2}, \mathbf{v}_2, \mathbf{v}_3)$ from the group manager (signer). This user secret key is sampled by the manager as in our standard model signature with the exception that $\mathbf{d} = \mathbf{0}$ (no signed message), and $\mathbf{u} = \mathbf{v}_{1,1}^{\text{gm}} + \mathbf{A}\mathbf{v}_{1,2}^{\text{gm}} + \mathbf{B}\mathbf{v}_2^{\text{gm}} + \mathbf{A}_3\mathbf{v}_3^{\text{gm}}$. Then, the group signature consists of a zero-knowledge proof of $(t, \mathbf{v}_{1,1}, \mathbf{v}_{1,2}, \mathbf{v}_2, \mathbf{v}_3)$ using the framework of [LNP22], as well as an encryption of t under the tracing authority’s encryption key. The construction, much like our standard model signature, is therefore agnostic to the method used to sample these user secret keys. Plugging a different sampler will just result in different parameters and performance (adjusted to keep the same security). In particular, our truncated sampler produces user secret keys of smaller dimension which leads to shorter zero-knowledge proofs and thus shorter group signatures. In Table 5.3, we report the sizes of the group manager’s public key pk , user secret keys usk , and group signatures (including the ciphertext) with (gsig_v) and without (gsig) verifiable encryption proof. Also, we report sizes with and without the recent partial trapdoor switching technique of [AGJ⁺24] which also applies to this group signature, in order to reduce further the dimension of \mathbf{A}_3 .

We observe that using our sampler procures a gain between 14 and 17% on the group signature size, 40% on the manager’s public key, and roughly 40-45% on the user secret key size. These figures were obtained by adapting the parameter selection script from [LNP22]. The gain is slightly smaller when using the partial trapdoor switching technique because the latter diminishes the impact of the gadget dimension on the overall group signature.

5.3.2 Anonymous Credentials. Let us now study the impact of our sampler on the anonymous credentials construction of [AGJ⁺24]. As their scheme is an adaptation of their standard model signature (to add a hiding part $[\mathbf{L}_d|\mathbf{A}]\mathbf{r}$ to the commitment) which is then plugged into zero-knowledge proofs, we can readily replace it by the one presented in Section 5.1 with the same minor adjustments. In Table 5.4, we report the sizes of the issuance transcript components (commitment \mathbf{c} , commitment opening proof π_1 , signature sig), as well as that

⁵ <https://github.com/truncatedsampler/truncated-sampler>

	Key Material		Group Signature	
	$ \text{pk} $	$ \text{usk} $	$ \text{gsig} $	$ \text{gsig}_v $
$\ell = 0$	47.5 KB	26.98 KB	91.55 KB	98.02 KB
$\ell = 1$	38.0 KB	16.11 KB	80.95 KB	87.36 KB
$\ell = 2$	28.5 KB	13.82 KB	76.03 KB	82.65 KB
with partial trapdoor switching [AGJ ⁺ 24]				
$\ell = 0$	47.5 KB	21.59 KB	80.30 KB	86.77 KB
$\ell = 1$	38.0 KB	13.44 KB	72.20 KB	78.61 KB
$\ell = 2$	28.5 KB	11.82 KB	69.09 KB	75.71 KB

Table 5.3. Comparison of the group manager’s public key, user secret keys, and group signatures with and without verifiable encryption proof for different values of ℓ ($k = 5$). All sizes are in KB. The case $\ell = 0$ corresponds to the construction of [LNPS21,LNP22].

of the show proof π_2 . For a fair comparison, we keep almost all the parameters equal (n, d, q, k, b , etc.), except for ℓ and the Gaussian widths which depend on the sampler. We also give the achieved security for anonymity and unforgeability using the proofs and formulae present in [Jeu24]. The parameters are obtained by adapting the parameter selection script of [AGJ⁺24].

	Issuance Transcript				Show	Sec.
	$ \text{c} $	$ \pi_1 $	$ \text{sig} $	Tot.	$ \pi_2 $	$\lambda_{\text{anon}}, \lambda_{\text{uf}}$
$\ell = 0$	2.37	35.99	6.81	45.17 KB	79.58 KB	125, 123
$\ell = 1$	2.37	35.99	5.97	44.33 KB	75.73 KB	123, 122
$\ell = 2$	2.37	35.99	4.87	43.23 KB	71.46 KB	123, 122
with zero-knowledge optimizations [LNP22, Sec. 4.4, App. A][LN22]						
$\ell = 0$	2.37	24.91	6.81	34.09 KB	60.76 KB	137, 125
$\ell = 1$	2.37	24.91	5.97	33.25 KB	57.53 KB	135, 125
$\ell = 2$	2.37	24.91	4.87	32.15 KB	53.98 KB	135, 125

Table 5.4. Comparison of the issuance transcript sizes between the user and signer, the credential proof sizes, and the security for different values of ℓ ($k = 5$). All sizes are in KB. The case $\ell = 0$ corresponds to the construction of [AGJ⁺24] with the elliptic sampler akin that of Section 4.1. The λ_{anon} and λ_{uf} are the bit-security for the anonymity and unforgeability of the anonymous credentials system.

Most of the issuance transcript does not change as it is independent of the sampler used for the signature afterwards. However, the showing proof π_2 is improved for $\ell > 0$ because the witness dimension is smaller. As for the group signature, the intricacies and overhead involved by the zero-knowledge proof

framework of [LNP22], even with further optimizations like that of [LN22] which were not considered in [AGJ+24], dilute the improvement to around 10% for $\ell = 2$ (compared to the 28% for the standalone signature), but it comes for free as it does not require significant changes in the overall construction.

5.3.3 Blind Signature. We can apply the same change to the recent blind signature of [JS24]. In their case, the commitment uses the whole matrix \mathbf{A}_t so as to inject excess randomness in the commitment to be recycled later to mask part of the blind signature. We can however adapt their construction and security proof so as to rely on our new truncated sampler as well. As before, in Table 5.5, we report the issuance transcript sizes and blind signature size for our new preimage sampler while keeping all other parameters the same. The parameters are obtained by adapting the parameter selection script of [JS24]. Our change of sampler incurs no security loss, with 126 bits of one-more unforgeability and 125 bits of anonymity. We again obtain a 10% gain on the size of the blind signature, but also on the issuance transcript as the commitment and commitment proof both depend on $k - \ell$.

	Issuance Transcript					Blind Signature			
	$ t $	$ c $	$ ct $	$ \pi_1 $	$ v $	Tot.	$ w_L $	$ \pi_2 $	$ bsig $
$\ell = 0$	0.03	3.59	1.62	45.68	8.70	59.63 KB	5.38	35.74	41.12 KB
$\ell = 1$	0.03	3.59	1.62	41.32	6.67	53.21 KB	4.62	31.66	36.28 KB

Table 5.5. Comparison of the issuance transcript sizes between the user and signer and the blind signature ($bsig$) sizes for different values of ℓ ($k = 3$). All sizes are in KB. The case $\ell = 0$ corresponds to the construction of [JS24] with the elliptic sampler akin that of Section 4.1.

References

- AGJ+24. S. Argo, T. Güneysu, C. Jeudy, G. Land, A. Roux-Langlois, and O. Sanders. Practical Post-Quantum Signatures for Privacy. In *CCS*, 2024.
- Ajt96. M. Ajtai. Generating Hard Instances of Lattice Problems (Extended Abstract). In *STOC*, 1996.
- Ban93. W. Banaszczyk. New Bounds in Some Transference Theorems in the Geometry of Numbers. *Math. Ann.*, 1993.
- BDGL16. A. Becker, L. Ducas, N. Gama, and T. Laarhoven. New Directions in Nearest Neighbor Searching with Applications to Lattice Sieving. In *SODA*, 2016.
- BEP+21. P. Bert, G. Eberhart, L. Prabel, A. Roux-Langlois, and M. Sabt. Implementation of Lattice Trapdoors on Modules and Applications. In *PQCrypto*, 2021.

- Ber11. D. S. Bernstein. *Matrix Mathematics: Theory, Facts, and Formulas*. Princeton University Press, second edition, 2011.
- CGM19. Y. Chen, N. Genise, and P. Mukherjee. Approximate Trapdoors for Lattices and Smaller Hash-and-Sign Signatures. In *ASIACRYPT*, 2019.
- Cha82. D. Chaum. Blind Signatures for Untraceable Payments. In *CRYPTO*, 1982.
- Cha85. D. Chaum. Security Without Identification: Transaction Systems to Make Big Brother Obsolete. *Commun. ACM*, 1985.
- CvH91. D. Chaum and E. van Heyst. Group Signatures. In *EUROCRYPT*, 1991.
- DKL⁺18. L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler, and D. Stehlé. CRYSTALS-Dilithium: A Lattice-Based Digital Signature Scheme. *TCHES*, 2018.
- dPLS18. R. del Pino, V. Lyubashevsky, and G. Seiler. Lattice-Based Group Signatures and Zero-Knowledge Proofs of Automorphism Stability. In *CCS*, 2018.
- EWY23. T. Espitau, A. Wallet, and Y. Yu. On Gaussian Sampling, Smoothing Parameter and Application to Signatures. In *ASIACRYPT*, 2023.
- GM18. Nicholas Genise and Daniele Micciancio. Faster Gaussian Sampling for Trapdoor Lattices with Arbitrary Modulus. In *EUROCRYPT*, 2018.
- GMPW20. N. Genise, D. Micciancio, C. Peikert, and M. Walter. Improved Discrete Gaussian and Subgaussian Analysis for Lattice Cryptography. In *PKC*, 2020.
- GPV08. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for Hard Lattices and New Cryptographic Constructions. In *STOC*, 2008.
- Jeu24. C. Jeudy. *Design of Advanced Post-Quantum Signatures Schemes*. PhD thesis, Université de Rennes 1, Rennes, France, 2024. Accessible at <https://tel.archives-ouvertes.fr/tel-04727543v1>.
- JRS23. C. Jeudy, A. Roux-Langlois, and O. Sanders. Lattice Signature with Efficient Protocols, Application to Anonymous Credentials. In *CRYPTO*, 2023.
- JRS24. C. Jeudy, A. Roux-Langlois, and O. Sanders. Phoenix: Hash-and-Sign with Aborts from Lattice Gadgets. In *PQCrypto*, 2024.
- JS24. C. Jeudy and O. Sanders. Improved Lattice Blind Signatures from Recycled Entropy. *IACR Cryptol. ePrint Arch.*, page 1289, 2024.
- Kle00. P. N. Klein. Finding the closest lattice vector when it's unusually close. In *SODA*, 2000.
- LLLW23. Q. Lai, F.-H. Liu, A. Lysyanskaya, and Z. Wang. Lattice-based Commit-Transferrable Signatures and Applications to Anonymous Credentials. *IACR Cryptol. ePrint Arch.*, page 766, 2023.
- LLM⁺16. B. Libert, S. Ling, F. Mouhartem, K. Nguyen, and H. Wang. Signature Schemes with Efficient Protocols and Dynamic Group Signatures from Lattice Assumptions. In *ASIACRYPT*, 2016.
- LN22. V. Lyubashevsky and N. K. Nguyen. BLOOM: Bimodal Lattice One-Out-of-Many Proofs and Applications. *ASIACRYPT*, 2022.
- LNP22. V. Lyubashevsky, N. K. Nguyen, and M. Plançon. Lattice-Based Zero-Knowledge Proofs and Applications: Shorter, Simpler, and More General. *CRYPTO*, 2022.
- LNPS21. V. Lyubashevsky, N. K. Nguyen, M. Plançon, and G. Seiler. Shorter Lattice-Based Group Signatures via "Almost Free" Encryption and Other Optimizations. In *ASIACRYPT*, 2021.

- LNS21. V. Lyubashevsky, N. K. Nguyen, and G. Seiler. Shorter Lattice-Based Zero-Knowledge Proofs via One-Time Commitments. In *PKC*, 2021.
- LS02. T.-T. Lu and S.-H. Shiou. Inverses of 2×2 Block Matrices. *Computers & Mathematics With Applications*, 43:119–129, 2002.
- LS15. A. Langlois and D. Stehlé. Worst-case to Average-case Reductions for Module Lattices. *DCC*, 2015.
- LS18. V. Lyubashevsky and G. Seiler. Short, Invertible Elements in Partially Splitting Cyclotomic Rings and Applications to Lattice-Based Zero-Knowledge Proofs. In *EUROCRYPT*, 2018.
- MP12. D. Micciancio and C. Peikert. Trapdoors for Lattices: Simpler, Tighter, Faster, Smaller. In *EUROCRYPT*, 2012.
- MR07. D. Micciancio and O. Regev. Worst-Case to Average-Case Reductions Based on Gaussian Measures. *SIAM J. Comput.*, 2007.
- Pei10. C. Peikert. An Efficient and Parallel Gaussian Sampler for Lattices. In *CRYPTO*, 2010.
- PFH⁺20. T. Prest, P.-A. Fouque, J. Hoffstein, P. Kirchner, V. Lyubashevsky, T. Pornin, T. Ricosset, G. Seiler, W. Whyte, and Z. Zhang. *FALCON*. *Tech. rep.*, 2020. Available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- Pre17. T. Prest. Sharper Bounds in Lattice-Based Cryptography Using the Rényi Divergence. In *ASIACRYPT*, 2017.
- RSA78. R. L. Rivest, A. Shamir, and L. Adleman. A Method for Obtaining Digital Signatures and Public-Key Cryptosystems. *Commun. ACM*, 1978.
- YJW23. Y. Yu, H. Jia, and X. Wang. Compact Lattice Gadget and Its Applications to Hash-and-Sign Signatures. In *CRYPTO*, 2023.

A Parameters for the Standard Model Signature

Symbol	Description	Value
Signature Parameters		
λ	Security parameter	128
n	Signature ring degree	256
d	Module rank	4
q	Modulus	$370673 \approx 2^{18.5}$
k	Gadget length	5
ℓ	Number of truncated gadget entries	2
b	Gadget base	13
ε	Smoothing loss for samplers	2^{-40}
s_1	Gaussian width 1 of sampler	5462.023
s_2	Gaussian width 2 of sampler	448.535
s_3	Gaussian width 3 of sampler ($\mathbf{v}_{1,2}$)	5443.575
s_4	Gaussian width 4 of sampler ($\mathbf{v}_2, \mathbf{v}_3$)	77.689
$s_{1,1}$	Final Gaussian width of $\mathbf{v}_{1,1}$	7989.601
w	Hamming weight of tags	5
κ	Number of splitting factors of q	8
Q	Maximal number of signature queries	2^{32}
α_1, α_2	Rejection sampling slack (sec. proof)	22.07, 15.03
$M_{1,1}, M_{1,2}$	Rejection sampling repetition rate (sec. proof)	1.006, 1.014
$B_{1,1}$	Verification bound of $\mathbf{v}_{1,1}$	114085.50
$B_{1,2}$	Verification bound of $\mathbf{v}_{1,2}$	77730.16
B_2	Verification bound of \mathbf{v}_2	1834.48
B_3	Verification bound of \mathbf{v}_3	976.78
Security Estimates		
BKZ [Ⓛ]	Required BKZ blocksize for M-SIS [Ⓛ]	645
BKZ [Ⓢ]	Required BKZ blocksize for M-SIS [Ⓢ]	554
BKZ	Required BKZ blocksize for M-LWE	473
$\varepsilon_{\text{M-SIS}}^{\bullet}$	Hardness bound for M-SIS [Ⓛ]	$2^{-188.65}$
$\varepsilon_{\text{M-SIS}}^{\circ}$	Hardness bound for M-SIS [Ⓢ]	$2^{-162.03}$
$\varepsilon_{\text{M-LWE}}$	Hardness bound for M-LWE	$2^{-138.34}$
λ_{uf}	Reached unforgeability bit security (Thm. 5.1)	120
Efficiency Estimates		
$ \text{pk} $	Size of public key B	28.5 KB
$ \text{sk} $	Size of secret key R	6.0 KB
$ \text{sig} $	Size of signature ($\mathbf{t}, \mathbf{v}_{1,2}, \mathbf{v}_2, \mathbf{v}_3$)	4939 B

Table A.1. Suggested parameter set for the standard model signature. The hardness bounds are estimated in the Core-SVP model, i.e., obtained from the BKZ blocksize B with sieving SVP oracle as $2^{-B \log_2(\sqrt{3/2})} \approx 2^{-0.292B}$ [BDGL16].