GG-GSW: Chosen-Ciphertext Secure Leveled FHE From Gadget Trapdoors

Jérôme Nguyen

Universität der Bundeswehr München
jerome.nguyen@unibw.de

Abstract. We present a simple leveled fully homomorphic encryption (FHE) scheme that achieves IND-CCA1 security in the standard model without relying on SNARKs. Our construction combines the gadget trapdoors introduced by Micciancio and Peikert [Eurocrypt 2012] with the GSW leveled FHE scheme [Gentry, Sahai, Waters, Crypto 2013]. We make only minimal changes to GSW to obtain security: ciphertext size remains similar and the homomorphic operations are unchanged.
We prove IND-CCA1 security under a new delayed-target LWE assumption. We provide evidence for its hardness in two ways: under a lattice knowledge assumption, we give a reduction from standard LWE; for a restricted parameter regime, we give a quantum reduction from the one-more inhomogeneous short integer solution assumption.
Finally, using the randomized evaluation technique from Bourse et al. [Crypto 2016], we obtain IND-CPA^D security for branching programs using only polynomial noise flooding and without requiring statistical correctness.

Keywords: Fully Homomorphic Encryption · Chosen-Ciphertext Security · Gadget Trapdoors

1 Introduction

Fully homomorphic encryption (FHE) is a powerful primitive that allows computation over encrypted data without needing access to the secret key. Since Gentry's seminal paper [29], the field has seen dramatic improvements in the efficiency and functionality of FHE schemes. See, e.g. [9,8,10,18,19,26] for some of the key developments. However, as FHE approaches usable performance for practical applications, it becomes increasingly important to consider its security in more advanced models.

Historically, the security of FHE schemes has overwhelmingly been analyzed in the context of passive adversaries, i.e. in the indistinguishability under chosen-plaintext attack (IND-CPA) model. While this is a natural starting point, this may be overly optimistic in practice. Among other things, the security of an IND-CPA-secure FHE scheme depends on the honesty of the party performing the homomorphic evaluation. Yet, this party can mount a host of practical attacks, ranging from simple distinguishing attacks to catastrophic key-recovery

attacks. Furthermore, the difficulty of running such attacks is not necessarily correlated to their impact, e.g. a single leaked decryption of a ciphertext may allow an adversary to recover the secret-key in some cases [36].

Higher security models. Motivated by key-recovery attacks on approximate FHE schemes, Li and Micciancio [36] introduced the IND-CPA^D (IND-CPA with restricted decryptions) security notion. In addition to being able to observe encryptions and homomorphic evaluations as in the IND-CPA model, an adversary may query decryptions of honestly generated ciphertexts. While originally thought of as exclusively relevant for approximate FHE schemes, it was later shown that so-called "exact" FHE schemes suffer from similar attacks if their correctness guarantees are chosen too low [17,15].

Classically, the next target after IND-CPA security is indistinguishability under non-adaptive chosen-ciphertext attack (IND-CCA1). Here an adversary may see the decryption of arbitrarily chosen (possibly malformed) ciphertexts *before* receiving a distinguishing challenge. This is a well-studied notion, and security in this model already excludes key-recovery attacks against active adversaries.

Finally, beyond the CCA1 bound one may look at security against active adversaries with adaptive queries. It is well-known that the IND-CCA2 security model is not achievable for homomorphic encryption as it forbids any malleability. However, some intermediary models have been shown to be achievable. In this direction, Manulis and Nguyen [42] have introduced the vCCA notion, and recently Yang et al. [51] have introduced IV-CCA. Both models essentially allow an adversary to make adaptive decryption queries, but ensure that the ciphertexts do not depend on the challenge ciphertext.

Usual Techniques. Unfortunately, to date only a few techniques are known to provably achieve these higher levels of security, and each incur significant overhead.

A statistically correct CPA-secure FHE scheme achieves IND-CPA^D security [36]. While one may therefore consider that exact FHE schemes are beyond the scope of CPA^D style notions, we view achieving CPA^D security through correctness as its own method. Correctness is a delicate property for FHE schemes. There is often a mismatch between theoretical upper bounds governing correctness guarantees and practical efficiency considerations. Using correctness to argue CPA^D security implies that it becomes a security property which must hold with overwhelming probability. However, in practice correctness is often seen as a performance property that can be tuned to achieve better efficiency. Some applications may not require exponentially low decryption failure rates and a user may accept to set parameters such that they happen with noticeable probability e.g. 2^{-15}, to reduce the computational overhead. Additionally, libraries often use heuristic arguments to estimate correctness [4] in order to improve efficiency, which may be exploited by a malicious adversary [17].

Known CPA^D attacks rely on the dependence of the decryption algorithm on the LWE noise of a ciphertext. Li et al. [37] have shown that this can be mitigated through noise flooding at decryption. However this requires the noise flooding

term to come from an exponentially wide distribution. Notice that the security of both methods, correctness and noise flooding, require correct upper bounds on the noise of a ciphertext, and their performance scales with the tightness of these bounds.

To obtain security against active adversaries i.e. CCA1 or the recently introduced vCCA or IV-CCA security, the main technique is to prove that ciphertexts are well-formed via proofs of knowledge [13,42,51]. The evaluator proves that the evaluation algorithm has been correctly executed on valid inputs. This significantly increases the computational overhead of evaluation, often relies on strong knowledge assumptions and also hurts the compactness of the underlying FHE schemes. The latest iteration [51] of this technique is based on the LWE assumption, but the proof still grows poly-logarithmically with the circuit size.

Outside of FHE. If we consider existing techniques to achieve security against active adversaries outside of the scope of FHE, the landscape is different. The well-known Fujisaki-Okamoto transform [27] is a highly efficient technique to obtain IND-CCA2 security, but is fundamentally incompatible with homomorphism by design. However, the (light) Cramer-Shoup encryption scheme, and the hash proof system framework that followed it [20,21] achieve CCA1 security by embedding an implicit proof of knowledge within the ciphertext. Most importantly for our purposes, it is an additively homomorphic encryption (AHE) scheme. Unfortunately, to the best of our knowledge no translation of the technique to lattice-based schemes exists, not even if we only ask for AHE schemes. In light of this, we ask the following questions.

Can we construct an FHE scheme that achieves IND-CCA1 *security without relying on SNARKs? Can we achieve* IND-CPA^D *security without exponential noise flooding or statistical correctness?*

1.1 Our Contributions

We answer both questions positively. Our main result is a leveled FHE scheme that achieves IND-CCA1 security under a new variant of the LWE assumption, in the standard model and without requiring SNARKs. It is a variation of the (dual) GSW scheme [31] and requires only minimal changes: ciphertext size remains similar and all homomorphic operations are unchanged. It also yields the first CCA1-secure lattice-based AHE scheme that does not rely on proofs of knowledge.

In order to prove security, we introduce the delayed-target LWE (DT-LWE) assumption, a lattice analogue of delayed-target assumptions previously studied in the discrete-log setting [34,40]. Informally, DT-LWE splits the LWE problem into two phases. In the first, the adversary receives a uniformly random matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ and oracle access to a bounded-distance LWE inversion oracle: on input a vector $\mathbf{b}' = \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q \in \mathbb{Z}_q^m$, the oracle returns the vector $\mathbf{s}$ if and only if $\mathbf{e}$ is sufficiently short. In the second phase, the oracle is disabled, after which the adversary receives a target vector which is either an LWE sample

$\mathbf{b} = \mathbf{As} + \mathbf{e} \bmod q$ or uniformly random. The task is then to distinguish the cases.

We provide evidence for the hardness of DT-LWE in two ways. First, we show that any efficient distinguisher for DT-LWE yields either a distinguisher for standard LWE or an algorithm violating the lattice knowledge assumption LK-ϵ [41]. This mirrors an earlier result of Wu and Stinson, who derive delayed-target DDH from DDH together with a knowledge assumption [50]. Second, for a restricted parameter regime, we give a quantum reduction from the one-more inhomogeneous short integer solution (OM-ISIS) assumption to DT-LWE.

We extend the CPA^D model to the CCA1 setting. Brzuska et al. [12] show that CCA1 security does not imply CPA^D security. Conceptually, however, an active adversary should at least be as powerful as a passive one, since it should have access to any information available to the latter. Following [12], we introduce the IND-CCA1^D notion, a CPA^D version of the IND-CCA1 game by augmenting the latter with the CPA^D oracles (cf. Figure 4).

Finally, using the randomized evaluation technique of Bourse et al. [7] for branching programs, we upgrade our CCA1 scheme to achieve IND-CCA1^D security over NC^1, *almost for free*.

Notably, our security proof does not require statistical correctness. Indeed, one may chose parameters that do not guarantee correctness even for a single level of homomorphic evaluation, while security still holds. Hence, our construction is the first to achieve CPA^D security without deriving it from correctness or relying on exponential noise flooding.

1.2 Our Techniques

Our starting point is the CCA1-secure encryption scheme based on LWE given by Micciancio and Peikert [45], which we refer to as the MP scheme. It is a dual-style encryption scheme [30] that uses gadget trapdoors to recover the LWE noise underlying a ciphertext. This allows decryption to reject ciphertexts whose noise is too large, preventing malformed ciphertexts from leaking information about the secret key. However, the scheme is not homomorphic: proving CCA1 security requires sampling additional randomness in order to enable a puncturing argument. Our approach is to forgo this randomness and instead base security on the DT-LWE assumption.

The public key of our scheme is therefore a trapdoored matrix $\mathbf{A} \in \mathbb{Z}_q^{\bar{m} \times n}$, and ciphertexts have the form

$$\mathbf{AS} + \mathbf{E} + \mu\mathbf{G} \bmod q \in \mathbb{Z}_q^{\bar{m} \times k},$$

where $\mathbf{S}$ and $\mathbf{E}$ are the encryption randomness and $\mathbf{G}$ is the GSW gadget matrix. The secret key is a short matrix $\mathbf{T_R}$ satisfying

$$\mathbf{T_R A} = \hat{\mathbf{G}} \bmod q,$$

where $\hat{\mathbf{G}}$ is a second gadget matrix. This allows us to use the same integrity mechanism as in the MP scheme. Postponing the details of message recovery, we

have

$$\mathbf{T_R}(\mathbf{C} - \mu\mathbf{G}) = \hat{\mathbf{G}}\mathbf{S} + \mathbf{T_R}\mathbf{E} \bmod q.$$

Decryption then recovers $\mathbf{S}$ and $\mathbf{E}$ and, if and only if $\mathbf{E}$ satisfies some bound $\|\mathbf{E}\|_2 < D$, it outputs μ.

The scheme is clearly additive. Moreover, decryption applies the linear map $\mathbf{T_R}$ to the ciphertext. Together with the gadget-based message encoding from GSW, this yields the same approximate eigenvector phenomenon. Therefore, we inherit the GSW homomorphic additions and multiplications unchanged [31,5]. We thus obtain a leveled FHE scheme combining GSW homomorphism with the ciphertext integrity mechanism from MP. Since the construction uses two gadget matrices, $\mathbf{G}$ and $\hat{\mathbf{G}}$, we call it *double-gadget GSW* (GG-GSW).

Security from delayed-target LWE. At a high level, the known chosen-ciphertext attacks on lattice-based encryption schemes follow two broad avenues. An adversary may exploit the "almost" linearity of decryption to extract information on the secret key directly [41,38]; or, it may manipulate the noise of a ciphertext to induce decryption failures in a way that depends on the secret key [23,17]. In the MP scheme, decryption recovers and validates the noise. This integrity check already addresses both avenues. On the other hand, the additional randomness which obstructs homomorphism appears to only serve the puncturing argument in the security proof. This suggests retaining the noise recovery mechanism while replacing the puncturing argument with a different foundation.

The CCA1 security proof. We base the IND-CCA1 security proof of GG-GSW on two ideas. First, the output of the decryption algorithm only depends on the public key $\mathbf{A}$ and on the input ciphertext. Indeed, we show that except with negligible probability, its output is identical to an information-theoretical oracle parametrized by $\mathbf{A}$. Decryption accepts or rejects a ciphertext solely based on the distance of $\mathbf{C} - \mu\mathbf{G}$ to the q-ary lattice defined by $\mathbf{A}$. This allows us to replace the trapdoored public key by a uniformly random matrix by a leftover-hash argument.

Second, we show that this oracle may be simulated using the bounded-distance inversion oracle given by the DT-LWE problem. In the IND-CCA1 game, all decryption queries are issued before the challenge ciphertext is given, so the reduction is able to answer them while it has access to its oracle. Once this access is revoked, it receives the DT-LWE challenge, which it embed in the ciphertext challenge using standard techniques.

Evidence for the hardness of DT-LWE. Intuitively, the DT-LWE oracle is only useful to the adversary if it is able to produce at least one accepting query $\mathbf{b} = \mathbf{A}\mathbf{s}+\mathbf{e} \bmod q$ with $\mathbf{e}$ sufficiently small without already knowing $(\mathbf{s}, \mathbf{e})$. We call this task *obliviously sampling* an accepting vector. A randomly sampled vector $\mathbf{b}$ has negligible probability of being accepted, and in fact, to the best of our knowledge no (classical) algorithm is known to have non-negligible success probability. This is related to a lattice-based knowledge assumption LK-ϵ introduced by Loftus

et al. [41, Definition 2]. Under that assumption, if an adversary outputs an accepting vector $\mathbf{b}$, an extractor recovers the corresponding lattice point in the q-ary lattice, from which the reduction recovers $(\mathbf{s}, \mathbf{e})$. The reduction can therefore answer the queries to the DT-LWE oracle by using the extractor, and the problem reduces to standard LWE.

The LK-ϵ-based reduction supports the hardness of DT-LWE against classical adversaries: it shows that exploiting the additional oracle requires violating a natural knowledge assumption. However, since LK-ϵ is a non-falsifiable assumption, it is desirable to have evidence of a different kind. Furthermore, there are known quantum algorithms that, for some parameter regimes, enable oblivious sampling of LWE vectors, invalidating LK-ϵ [24,16]. However, these algorithms do not appear to break DT-LWE. Indeed, for a restricted parameter regime, we give a quantum reduction from the OM-ISIS assumption to DT-LWE.

The OM-ISIS problem provides a short preimage oracle for a matrix $\mathbf{A}$, and requires an adversary to return short preimages of *one more* random syndrome than the number of oracle queries it has made. In our reduction, we query the preimage oracle on the columns of a gadget matrix to sample a gadget trapdoor for the challenge matrix $\mathbf{A}$. This allows us to answer DT-LWE oracle queries using the MP noise-recovery procedure. We then use a search DT-LWE adversary as a search-LWE solver in the ISIS to LWE quantum reduction given by Chailloux and Hermouet [14]. This allows us to produce one more ISIS solution than the number of preimage queries we made. Combined with our search-to-decision reduction, this completes the reduction from OM-ISIS to decisional DT-LWE.

Obtaining IND-CCA1^D security over NC^1. To obtain IND-CCA1^D security, we adapt the randomized GSW evaluation algorithm of Bourse et al. [7] to dual GSW. Before evaluating a branching program, we add a small flooding term $\widetilde{\mathbf{Y}}_i$ to each input ciphertext, whose noise becomes $\mathbf{H}_i = \mathbf{E}_i + \widetilde{\mathbf{Y}}_i$. Generalizing the Bourse et al. analysis to matrix-valued noise, we show that the output-noise distribution depends on the input noise only through the values $(\mathbf{H}_i)_i$. We extend the DT-LWE problem to the hint setting, where the adversary obtains noisy linear combinations of the DT-LWE noise [33]. Using the "leaky" framework given by Lai et al. [35], we reduce this problem to DT-LWE. We then prove IND-CCA1^D security by treating the $\mathbf{H}_i = \mathbf{E}_i + \widetilde{\mathbf{Y}}_i$ as hints, which allows us to simulate the CPA^D oracles. This yields IND-CCA1^D security for branching programs using only polynomial noise flooding, without requiring the evaluation algorithm to be statistically correct.

1.3 Related Works

In a concurrent work, Ogilvie uses the HintLWE assumption to obtain IND-CPA^D and KR-CPA^D security [47]. Their work focuses on improving the analysis of noise-flooding techniques in the bit-security framework to reduce the overhead it generates. This is orthogonal to our work. They give an improved analysis for general LWE and RLWE schemes, which still requires exponential noise-flooding.

On the other hand, we give a tailored approach to the GSW scheme and branching programs, which allows us to use polynomial noise-flooding.

As mentioned above, the picture concerning IND-CCA1 security is more positive for AHE schemes, particularly for group-based constructions. Lipmaa [40] studied the security of ElGamal and Damgård's ElGamal under tailored Diffie-Hellman assumptions. More recently, Libert [39] showed that restricting the message space enables proofs of IND-CCA1 security from standard assumptions for several AHE schemes including Damgard's ElGamal [22]. Phan et al. [48] construct a new variant of Paillier-ElGamal encryption. It is a "true" linearly homomorphic encryption scheme, in the sense that it supports unlimited additions. They show its IND-CCA1 security under the decisional composite residuosity assumption. Moreover, under a non-falsifiable linear-only homomorphism assumption, they prove that it achieves vCCA security without relying on a SNARK.

2 Preliminaries

Notation. We write vectors in lower-case bold font (e.g. $\mathbf{x}$) and matrices in upper-case bold font (e.g. $\mathbf{A}$). We consider all vectors to be column vectors unless specified otherwise. For vectors $\mathbf{v}_1, \ldots \mathbf{v}_l$, we write $[\mathbf{v}_1|\mathbf{v}_2|\ldots|\mathbf{v}_l]$ to denote the matrix whose columns are the $\mathbf{v}_i$. For matrices $\mathbf{A}_1, \ldots, \mathbf{A}_l$, we write $\mathbf{A} = [\mathbf{A}_1; \ldots; \mathbf{A}_l]$ to denote the matrix $\mathbf{A}$ that is the vertical concatenation of the $\mathbf{A}_i$, i.e. $\mathbf{A} = \left[\begin{smallmatrix} \mathbf{A}_1 \\ \cdots \\ \mathbf{A}_l \end{smallmatrix}\right]$. For any set S, we write $x \xleftarrow{\$} S$ to denote that x is sampled from the uniform distribution over S. For any probability distribution $\mathcal{D}$, we write $x \leftarrow \mathcal{D}$ to denote that x is sampled according to the distribution $\mathcal{D}$.

2.1 Gaussian Distributions

For any $\sigma > 0$, the spherical Gaussian function on $\mathbb{R}^n$ with parameter σ is defined as $\rho_\sigma(\mathbf{x}) = \exp(-\pi\|\mathbf{x}\|^2/\sigma^2)$ for all $\mathbf{x} \in \mathbb{R}^n$. For a rank-$n$ matrix $\mathbf{S} \in \mathbb{R}^{m\times n}$ and a vector $\mathbf{c} \in \mathbb{R}^n$, the ellipsoidal Gaussian function on $\mathbb{R}^n$ centered at $\mathbf{c}$ with parameter $\mathbf{S}$, for all $\mathbf{x} \in \mathbb{R}^n$, is defined as:

$$\rho_{\mathbf{S},\mathbf{c}}(\mathbf{x}) = \exp(-\pi(\mathbf{x}-\mathbf{c})^\intercal(\mathbf{S}^\intercal\mathbf{S})^{-1}(\mathbf{x}-\mathbf{c})).$$

Note that equivalently, $\rho_{\mathbf{S}}$ is the function with covariance matrix $\mathbf{\Sigma} = \mathbf{S}^\intercal\mathbf{S}$. The spherical case is obtained when $\mathbf{S} = \sigma \cdot \mathbf{I}_n$ and we write $\rho_{\sigma,\mathbf{c}}$. If $\mathbf{c} = \mathbf{0}$, we write $\rho_{\mathbf{S}}(\cdot)$. For a set A, $\rho_{\mathbf{S}}(A)$ denotes $\sum_{\mathbf{x}\in A}\rho_{\mathbf{S}}(\mathbf{x})$.

Lemma 2.1 (Simplified generalized leftover hash lemma [25][7]). *Let* $\mathbf{e}$ *be any random variable over* $\mathbb{Z}_q^m$ *and* $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^k$. *Then,*

$$\Delta((\mathbf{X}\mathbf{e}, \mathbf{X}, f(e)), (r, \mathbf{X}, f(e))) \leq \frac{1}{2}\sqrt{q^{n+k}\cdot 2^{-H_\infty(\mathbf{e})}}$$

where $\mathbf{X} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$ *and* $\mathbf{r} \xleftarrow{\$} \mathbb{Z}_q^n$ *and* $H_\infty(\mathbf{e})$ *is the min entropy of* $\mathbf{e}$.

2.2 Lattices

An m-dimensional lattice Λ is a discrete additive subgroup of $\mathbb{R}^m$. For an integer $k \leq m$ and a rank k matrix $\mathbf{B} \in \mathbb{R}^{m \times k}$, $\Lambda(\mathbf{B}) = \{\mathbf{Bx} \in \mathbb{R}^m \mid \mathbf{x} \in \mathbb{Z}^k\}$ is the rank-k lattice generated by the columns of $\mathbf{B}$. We write Λ^* to denote the dual lattice $\{\mathbf{v} \in \mathbb{R}^n \mid \langle \mathbf{v}, \Lambda \rangle \in \mathbb{Z}\}$ and we write $\Lambda_{\mathbf{u}}^{\perp}(\mathbf{B}^\intercal)$ to denote $\{\mathbf{v} \in \mathbb{Z}^m \mid \mathbf{B}^\intercal \mathbf{v} = \mathbf{u} \bmod q\}$ for some vector $\mathbf{u} \in \mathbb{Z}_q^k$. Let $\lambda_\infty(\Lambda)$ denote the ℓ^∞ norm of a shortest vector in Λ. Furthermore, for $\mathbf{c} \in \mathbb{R}^m$, $d_\infty(\mathbf{c}, \Lambda) \coloneqq \min_{\mathbf{p} \in \Lambda} \|\mathbf{c} - \mathbf{p}\|_\infty$.

For a rank-n lattice Λ, $\mathbf{S} \in \mathbb{R}^{m \times n}$, and $\mathbf{c} \in \mathbb{R}^n$, the ellipsoidal Gaussian distribution with parameter $\mathbf{S}$ and support $\Lambda + \mathbf{c}$ is defined, for all $\mathbf{x} \in \Lambda + \mathbf{c}$, as:

$$\mathcal{D}_{\Lambda+\mathbf{c},\mathbf{S}}(\mathbf{x}) = \frac{\rho_{\mathbf{S}}(\mathbf{x})}{\rho_{\mathbf{S}}(\Lambda + \mathbf{c})}.$$

We recall the definition for the smoothing parameter [43].

Definition 2.1 (Smoothing parameter). *For an n-dimensional lattice Λ and a positive real $\epsilon > 0$, the smoothing parameter $\eta_\epsilon(\Lambda)$ is the smallest s such that $\rho_{1/s}(\Lambda^* \setminus \{0\}) \leq \epsilon$.*

Lemma 2.2 ([43], Lemma 3.3). *For any n-dimensional lattice Λ and $\epsilon > 0$,*

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

where $\lambda_n(\Lambda)$ is the smallest real number $r > 0$ such that $\dim(\mathrm{span}(\Lambda \cap r\mathcal{B})) = n$ and $\mathcal{B}$ is the n-dimensional unit ball centered at the origin.

Lemma 2.3 ([45], Lemma 2.4). *Let $n, m, q \geq 2$ be positive integers. For $\mathbf{z} \in \mathbb{Z}_q^n$, let the subgroup $\mathbb{G}_{\mathbf{z}} = \{\langle \mathbf{a}, \mathbf{z} \rangle : \mathbf{a} \in \mathbb{Z}_q^n\} \subset \mathbb{Z}_q$, and let $g_{\mathbf{z}} = |\mathbb{G}_{\mathbf{z}}| = q/\gcd(z_1, \ldots, z_n, q)$. Let $\epsilon > 0, \eta \geq \eta_\epsilon(\mathbb{Z}^m)$ and $s > \eta$ be reals. Then for uniformly random $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$,*

$$\mathbb{E}\left[\rho_{1/s}(\Lambda^\perp(\mathbf{A})^*)\right] \leq (1 + \epsilon) \sum_{\mathbf{z} \in \mathbb{Z}_q^n} \max(1/g_{\mathbf{z}}, \eta/s)^m. \tag{1}$$

In particular, if $q = p^e$ is a power of a prime p, and

$$m \geq \max\left\{ n + \frac{\log(3 + 2/\epsilon)}{\log p}, \frac{n \log q + \log(2 + 2/\epsilon)}{\log(s/\eta)} \right\} \tag{2}$$

then $s \geq \eta_{2\epsilon/\delta}(\Lambda^\perp(\mathbf{A}))$ except with probability at most δ. In particular, letting $s = 2\eta_\epsilon(\mathbb{Z}^m)$ and $\epsilon = \mathsf{negl}(n)$, $m = n \log q + \omega(\log n)$ fulfills Equation (2).

2.3 Matrices

For $\mathbf{B} = [\mathbf{b}_1|\dots|\mathbf{b}_m] \in \mathbb{R}^{n\times m}$ and a vector norm $\|\cdot\|$, define $\|\mathbf{B}\| = \max_i(\|\mathbf{b}_i\|)$. There exists a *singular value decomposition* $\mathbf{B} = \mathbf{QDP}^\intercal$, where $\mathbf{Q} \in \mathbb{R}^{n\times n}$ and $\mathbf{P} \in \mathbb{R}^{m\times m}$ are orthogonal matrices and $\mathbf{D} \in \mathbb{R}^{n\times m}$ is a diagonal matrix. The diagonal entries $s_i \geq 0$ are called the singular values of $\mathbf{B}$. The largest singular value of $\mathbf{B}$ is written $s_1(\mathbf{B})$ and we have $s_1(\mathbf{B}) = \max_{\|\mathbf{v}\|_2=1}\|\mathbf{Bv}\|_2 \geq \|\mathbf{B}\|_2$. The following lemma follows from [45, Lemma 2.8, Lemma 2.9].

Lemma 2.4. *Let $\Lambda \subset \mathbb{R}^n$ be a lattice, $0 < \epsilon < 1$, and let $s \geq \eta_\epsilon(\Lambda)$. For $i \in [m]$, let $\mathbf{x}_i \leftarrow \mathcal{D}_{\Lambda+\mathbf{c}_i,s}$, where $\mathbf{c}_i \in \mathrm{span}(\Lambda)$, and set $\mathbf{X} = [\mathbf{x}_1|\dots|\mathbf{x}_m] \in \mathbb{R}^{n\times m}$. There exists a universal constant $C > 0$ such that for any $t \geq 0$, we have*

$$\Pr[s_1(\mathbf{X}) > Cs(\sqrt{m} + \sqrt{n} + t)] \leq 2\exp(\delta)\exp(-\pi t^2),$$

where $\delta = m\ln\left(\frac{1+\epsilon}{1-\epsilon}\right)$. If every $\mathbf{c}_i = \mathbf{0}$, the bound holds for any $s > 0$ and $\delta = 0$.

We also use the following ellipsoidal Gaussian tail bound. The first statement is [1, Lemma 3], and the centered case follows from [6, Lemma 1.5].

Lemma 2.5. *For a rank-n lattice Λ, vector $\mathbf{c} \in \mathbb{R}^n$, constant $0 < \epsilon < 1$, and matrix $\mathbf{S}$ such that $s_n(\mathbf{S}) \geq \eta_\epsilon(\Lambda)$, we have that for $\mathbf{v} \leftarrow \mathcal{D}_{\Lambda+\mathbf{c},\mathbf{S}}$,*

$$\Pr[\|\mathbf{v}\|_2 \geq s_1(\mathbf{S})\sqrt{n}] \leq \frac{1+\epsilon}{1-\epsilon}2^{-n}.$$

If $\mathbf{c} = \mathbf{0}$, then the probability is at most 2^{-n} for any rank-n matrix $\mathbf{S}$.

2.4 The LWE Problem

We recall the learning with error (LWE) assumption [49].

Definition 2.2 (LWE). *Let n, m, q be positive integers. Let χ_s, χ_e be probability distributions over $\mathbb{Z}_q^n$ and $\mathbb{Z}^m$ respectively. For $\mathsf{pp} = (n, m, q, \chi_s, \chi_e)$, the learning with error problem, written $\mathsf{LWE}_{\mathsf{pp}}$, is the task of distinguishing:*

$$(\mathbf{A}, \mathbf{As} + \mathbf{e} \bmod q) \qquad \textit{from} \qquad (\mathbf{A}, \mathbf{u})$$

where $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m\times n}, \mathbf{s} \leftarrow \chi_s$, $\mathbf{e} \leftarrow \chi_e$, $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$. The $\mathsf{LWE}_{\mathsf{pp}}$ assumption states that for any PPT adversary $\mathcal{A}$, the following advantage is negligible:

$$\mathsf{Adv}^{\mathsf{LWE}}_{\mathsf{pp},\mathcal{A}}(\lambda) \coloneqq |\Pr[\mathcal{A}(1^\lambda, \mathbf{A}, \mathbf{As} + \mathbf{e}) = 1] - \Pr[\mathcal{A}(1^\lambda, \mathbf{A}, \mathbf{u}) = 1]|.$$

In case χ_s is the uniform distribution and χ_e is the discrete spherical Gaussian with parameter σ, we may write $\mathsf{pp} = (n, m, q, \sigma)$.

2.5 Gadget Trapdoors

Let n, q be positive integers and set $l = \lceil \log q \rceil$. Define the gadget vector $\mathbf{g} \coloneqq (1, 2, \ldots, 2^{l-1})^\intercal \in \mathbb{Z}_q^l$. Gadget matrices are structured matrices of the form

$$\hat{\mathbf{G}} \coloneqq \mathbf{I}_n \otimes \mathbf{g} \in \mathbb{Z}_q^{nl \times n} \quad \text{or} \quad \mathbf{G} \coloneqq \mathbf{I}_n \otimes \mathbf{g}^\intercal = \mathrm{diag}(\mathbf{g}^\intercal, \ldots, \mathbf{g}^\intercal) \in \mathbb{Z}_q^{n \times nl}.$$

We recall the gadget LWE inversion algorithm introduced by Micciancio and Peikert [45], as well as the randomized gadget decomposition $\mathbf{G}_{\mathsf{rand}}^{-1}$ [45,5,7].

Lemma 2.6 (Gadget inversion). *Let m, k be positive integers, and let $\hat{\mathbf{G}} \in \mathbb{Z}_q^{nl \times n}$ be a gadget matrix. Let $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ and $\mathbf{T} \in \mathbb{Z}^{nl \times m}$ be matrices such that*

$$\mathbf{TA} = \hat{\mathbf{G}} \bmod q.$$

Let $d = 2$ if q is a power of two and $d = \sqrt{5}$ otherwise. Let $\mathbf{B} = \mathbf{AS} + \mathbf{E} \bmod q$ with $\mathbf{S} \in \mathbb{Z}_q^{n \times k}$ and $\mathbf{E} \in \mathbb{Z}_q^{m \times k}$ such that $\|\mathbf{E}\|_2 < q/2ds_1(\mathbf{T})$.

There exists a deterministic polynomial-time algorithm Invert that, on input $(\mathbf{T}, \mathbf{A}, \mathbf{B})$, returns $(\mathbf{S}, \mathbf{E})$.

Definition 2.3 (Gadget decomposition). *For $r = \tilde{O}(1)$, there is a PPT algorithm $\mathbf{G}_{\mathsf{rand}}^{-1} : \mathbb{Z}_q^n \longrightarrow \mathbb{Z}^{nl}$ such that $\mathbf{x} \leftarrow \mathbf{G}_{\mathsf{rand}}^{-1}(\mathbf{u})$ is statistically close to the discrete Guaussian $\mathcal{D}_{\Lambda_{\mathbf{0}}^{\perp}(\mathbf{G}) + \mathbf{G}_{\mathsf{det}}^{-1}(\mathbf{u}), r}$, where $\mathbf{G}_{\mathsf{det}}^{-1}$ denotes bit decomposition. For $\mathbf{V} \in \mathbb{Z}_q^{n \times p}$, we define $\mathbf{G}_{\mathsf{rand}}^{-1}(\mathbf{V}) \in \mathbb{Z}^{nl \times p}$ by applying $\mathbf{G}_{\mathsf{rand}}^{-1}$ independently to each column. We write $\mathbf{G}^{-1}$ in the case that either $\mathbf{G}_{\mathsf{det}}^{-1}$ or $\mathbf{G}_{\mathsf{rand}}^{-1}$ may be used.*

2.6 Homomorphic Encryption

Definition 2.4 ((Fully) Homomorphic Encryption). *A public-key homomorphic encryption (HE) scheme for a family of functions $\mathcal{F}$ with message space $\mathcal{M}$ is a quadruple of PPT algorithms $\mathsf{HE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ as follows:*

- $\mathsf{Gen}(1^\lambda)$*: Output a public key / secret key pair denoted* $(\mathsf{pk}, \mathsf{sk})$.
- $\mathsf{Enc}_{\mathsf{pk}}(m)$*: On input a message $m \in \mathcal{M}$ and the public-key* pk*, output a ciphertext c. We may write $\mathsf{Enc}_{\mathsf{pk}}(m; r)$ to denote the use of randomness r.*
- $\mathsf{Dec}_{\mathsf{sk}}(c)$*: On input a ciphertext c and the secret key* sk*, output $m \in \mathcal{M} \cup \{\perp\}$.*
- $\mathsf{Eval}_{\mathsf{pk}}(f, c_1, \ldots, c_l)$*: On input a function $f \in \mathcal{F}$, and ciphertexts $c_1, \ldots, c_l$, output a ciphertext $\hat{c}$. The public key may be omitted when it is not required.*

The scheme is a *fully homomorphic encryption* (FHE) scheme if $\mathcal{F}$ contains all efficiently computable functions. Let $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^\lambda)$, $m_1, \ldots, m_l \in \mathcal{M}$ and $c_i \leftarrow \mathsf{Enc}_{\mathsf{pk}}(m_i)$. The scheme is statistically correct if, for every $f \in \mathcal{F}$,

$$\Pr[\mathsf{Dec}_{\mathsf{sk}}(\mathsf{Eval}_{\mathsf{pk}}(f, c_1, \ldots, c_l)) \neq f(m_1, \ldots, m_l)] \leq \mathsf{negl}(\lambda).$$

We depart from the convention of treating so-called "exact" FHE schemes as statistically correct. We call an FHE scheme *exact* if decryption is considered as successful only when it returns $f(m_1, \ldots, m_l)$, rather than an approximation. Exactness is therefore orthogonal to the correctness probability. In particular, an exact scheme may have arbitrarily low correctness probability.

We recall the standard definitions of IND-CPA and IND-CCA1, as well as the definition of IND-CPA^D from [36], in Appendix A.

3 The GG-GSW Scheme

Our CCA1-secure leveled FHE scheme combines the key generation and decryption from the CCA1-secure scheme given by Micciancio and Peikert [45] with the gadget-based message encoding and homomorphic multiplication from the GSW scheme [31,5].

3.1 A leveled FHE scheme

Let n, m, q be positive integers with $n = n(\lambda)$, $q = 2^l$ and $m = nl$. Set $\bar{m} = 2m$ and $k = \bar{m} \cdot l$. Let $s, \sigma > 0$ be Gaussian parameters for key generation and encryption respectively with $s = \Theta(1)$ chosen sufficiently large for the leftover hash lemma 2.1 and let $D > 0$ with $\sigma\sqrt{\bar{m}} \ll D$. Define the gadget vector $\mathbf{g} = (1, 2, \ldots, 2^{l-1})$ and the gadget matrices

$$\mathbf{G} = \mathbf{I}_{\bar{m}} \otimes \mathbf{g}^\intercal \in \mathbb{Z}_q^{\bar{m} \times k}, \qquad \hat{\mathbf{G}} = \mathbf{I}_n \otimes \mathbf{g} \in \mathbb{Z}_q^{m \times n}.$$

We use $\mathbf{G}$ for message encoding and $\hat{\mathbf{G}}$ for the gadget trapdoor which ensures CCA1 security. Our leveled FHE scheme GG-GSW is defined as follows:

- $\mathsf{GG\text{-}GSW.Gen}(1^\lambda)$: Sample $\mathbf{A}_1 \xleftarrow{\$} \mathbb{Z}_q^{m \times n}$, $\mathbf{R} \leftarrow \mathcal{D}_{\mathbb{Z}^{m \times m}, s}$ and compute $\mathbf{A}_2 = \hat{\mathbf{G}} - \boldsymbol{R}\boldsymbol{A_1} \bmod q$. Return the public key/secret key pair,

$$\mathsf{pk} = \mathbf{A} \coloneqq \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix} \in \mathbb{Z}_q^{\bar{m} \times n} \quad \mathsf{sk} \coloneqq \mathbf{R}.$$

- $\mathsf{GG\text{-}GSW.Enc}(\mathsf{pk} = \mathbf{A}, \mu \in \{0,1\})$: Sample $\mathbf{S} \xleftarrow{\$} \mathbb{Z}_q^{n \times k}$ and $\mathbf{E} \leftarrow \mathcal{D}^k_{\mathbb{Z}^{\bar{m}}, \sigma}$. Return,

$$\mathbf{C} = \mathbf{AS} + \mathbf{E} + \mu\mathbf{G} \mod q.$$

- $\mathsf{GG\text{-}GSW.Dec}(\mathsf{sk}, \mathbf{C})$:
  1. Set $\mathbf{T}_\mathbf{R} \coloneqq [\mathbf{R} \mid \mathbf{I}_m] \in \mathbb{Z}_q^{m \times \bar{m}}$.
  2. For each $\mu \in \{0,1\}$, do:
     - Compute $(\mathbf{S}', \mathbf{E}') \leftarrow \mathsf{Invert}(\mathbf{T}_\mathbf{R}, \mathbf{A}, \mathbf{C} - \mu\mathbf{G})$.
     - If the call returns $\perp$, or if $\mathbf{C} \neq \mathbf{AS}' + \mathbf{E}' + \mu\mathbf{G} \bmod q$, continue with the next value of μ.
     - If $\|\mathbf{E}'\| \geq D$, continue with the next value of μ.
     - Return μ.
  3. Return $\perp$.

Similarly to GSW, our scheme enjoys better noise growth properties with a binary message space $\{0,1\}$. However, GG-GSW readily supports any polynomially sized message space $\mathcal{M} \subset \mathbb{Z}_q$, simply by extending Step 2 in the natural way. For large message spaces, decryption can avoid looping over $\mathcal{M}$ by augmenting the secret key with $\mathbf{x} \in \mathbb{Z}_q^m$ and the public key with the row $\mathbf{a}_3^\intercal = -\mathbf{x}^\intercal\mathbf{A}_1$. One then decrypts by first recovering μ using $\mathbf{x}$, and then running the integrity check using $\mathbf{R}$.

3.2 Correctness

The decryption algorithm tests both possible values of μ. Thus, to show correctness, we must exclude that both values may be accepted. This can only happen if there exists $\mathbf{S}, \mathbf{S}' \in \mathbb{Z}_q^{n\times k}$ and sufficiently short $\mathbf{E}, \mathbf{E}' \in \mathbb{Z}^{\bar{m}\times k}$ such that

$$\mathbf{AS} + \mathbf{E} = \mathbf{AS}' + \mathbf{E}' + \mathbf{G} \bmod q. \tag{3}$$

To this end, we extend the minimum-distance argument on q-ary lattices given in [30, Lemma 5.3] to shifted cosets. The proof appears in Appendix C.

Lemma 3.1. *Let n, m, q be positive integers. For $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m\times n}$, let $\mathcal{L}_{\mathbf{A}} \coloneqq \mathbf{A}\mathbb{Z}^n + q\mathbb{Z}^m$. For every r satisfying $1 \le r < q$, we have*

$$\Pr_{\mathbf{A}}[\lambda_\infty(\mathcal{L}_{\mathbf{A}}) < r] \le q^n \left(\frac{3r}{q}\right)^m.$$

Moreover, suppose that q is a power of two. Then, for every unit vector $\mathbf{u} \in \mathbb{Z}^m$,

$$\Pr_{\mathbf{A}}\left[d_\infty\left(\frac{q}{2}\mathbf{u}, \mathcal{L}_{\mathbf{A}}\right) < \frac{q}{6}\right] \le q^n 2^{-m}.$$

Correctness now follows from two arguments. Since the secret key $\mathbf{R}$ is sufficiently short, whenever $\mathbf{C}$ has short noise, the Invert algorithm succeeds when called on the correct μ. The previous lemma then ensures that this is the case for no other message.

Lemma 3.2 (Correctness). *Let $\mathbf{C} = \mathbf{AS} + \mathbf{E} + \mu\mathbf{G} \bmod q$ for any $\mathbf{S} \in \mathbb{Z}_q^{n\times k}$ and $\mu \in \{0, 1\}$. Suppose that*

$$D < \frac{q}{4(C \cdot s \cdot (2\sqrt{m} + \sqrt{\lambda}) + 1)},$$

where $C \ge 1$ is the constant from Lemma 2.4. Then, the decryption algorithm returns μ if $\|\mathbf{E}\| < D$, except with negligible probability in the randomness of key generation. In particular, if $\sigma\sqrt{\bar{m}} < D$, $\mathsf{GG\text{-}GSW}$ is statistically correct.

Proof. Let $(\mathsf{pk}, \mathsf{sk}) = (\mathbf{A}, \mathbf{R}) \leftarrow \mathsf{GG\text{-}GSW.Gen}(1^\lambda)$. From Lemma 2.4, we have

$$s_1(\mathbf{R}) \le Cs\left(2\sqrt{m} + \sqrt{\lambda}\right), \tag{4}$$

except with probability $2^{-\Omega(\lambda)}$. Let $\mathbf{C} \in \mathbb{Z}_q^{\bar{m}\times k}$ be as in the statement of the lemma. Step 2 is successful if the Invert algorithm returns $\mathbf{S}, \mathbf{E}$. By Lemma 2.6, this is the case if $\|\mathbf{E}\|_2 < q/4\sqrt{s_1(\mathbf{R})^2 + 1}$. By assumption,

$$\|\mathbf{E}\|_2 < D < \frac{q}{4(C \cdot s \cdot (2\sqrt{m} + \sqrt{\lambda}) + 1)} \le \frac{q}{4\sqrt{s_1(\mathbf{R})^2 + 1}} \tag{5}$$

where the last inequality uses Equation (4). Hence, conditioned on Equation (4) holding, Step 2 accepts μ.

From the leftover hash lemma 2.1, the distribution of $\mathbf{A}$ is statistically close to uniform. Then, for any unit vector $\mathbf{u}$, Lemma 3.1 gives

$$d_\infty\left(\frac{q}{2}\mathbf{u}, \mathcal{L}_\mathbf{A}\right) \geq \frac{q}{6}$$

except with probability at most $2^{-n\log q}$. Suppose that another message $\mu' \neq \mu$ also passes the checks. Then there exist $\mathbf{S}', \mathbf{E}'$ such that $\mathbf{A}(\mathbf{S}-\mathbf{S}') = \mathbf{E}' - \mathbf{E} + (\mu'-\mu)\mathbf{G} \bmod q$. Let $\mathbf{u}$ be a unit vector such that $(q/2)\mathbf{u}$ is a column of $\mathbf{G}$, and let $\mathbf{e}, \mathbf{e}'$ be the corresponding columns of $\mathbf{E}, \mathbf{E}'$. By assumption, we have $D < q/12$, and

$$d_\infty\left(\frac{q}{2}\mathbf{u}, \mathcal{L}_\mathbf{A}\right) \leq \|\mathbf{e}-\mathbf{e}'\|_\infty < 2D < \frac{q}{6},$$

contradicting the lower bound. Hence, the decryption algorithm returns μ.

For the second statement, let $\mathbf{C} \leftarrow \mathsf{GG\text{-}GSW.Enc}(\mathsf{pk}, \mu)$. It has noise $\mathbf{E} \leftarrow \mathcal{D}^k_{\mathbb{Z}^{\bar{m}},\sigma}$. By Lemma 2.5, $\|\mathbf{E}\|_2 < \sigma\sqrt{\bar{m}} < D$ except with negligible probability, and the claim follows. □

Homomorphic Operations. The main difference between our scheme and the regular GSW scheme is in the key generation. Hence, once our trapdoored matrix $\mathbf{A}$ is generated, our scheme behaves as the GSW scheme. In particular, it inherits all the homomorphic properties from it. We define the homomorphic operations in exactly the same way as in [5]:

- $\mathsf{Add}(\mathbf{C}_1, \mathbf{C}_2)$: $\mathbf{C}_1 \boxplus \mathbf{C}_2 = \mathbf{C}_1 + \mathbf{C}_2$,
- $\mathsf{Mult}(\mathbf{C}_1, \mathbf{C}_2)$: $\mathbf{C}_1 \boxdot \mathbf{C}_2 \leftarrow \mathbf{C}_1 \cdot \mathbf{G}^{-1}(\mathbf{C}_2)$. As in the original GSW scheme, this operation is non-commutative and right associative. A scalar-ciphertext product between a and $\mathbf{C}$ is computed as $\mathbf{C} \boxdot a \leftarrow \mathbf{C} \cdot \mathbf{G}^{-1}(a\mathbf{G})$.

The following lemma shows that GG-GSW has the same noise-growth behavior as GSW. The proof is essentially unchanged from [5, Lemma 3.5], and is therefore omitted.

Lemma 3.3. *Let $\mathbf{C}_1, \mathbf{C}_2$ be two ciphertexts encrypting μ_1, μ_2 with noise $\mathbf{E}^{(1)}$, $\mathbf{E}^{(2)}$ respectively. Then $\mathbf{C}_1 \boxplus \mathbf{C}_2$ is an encryption of $\mu_1 + \mu_2$ and has noise $\mathbf{E}^{(1)} + \mathbf{E}^{(2)}$; $\mathbf{C}_1 \boxdot \mathbf{C}_2$ is an encryption of $\mu_1 \cdot \mu_2$ and has noise $\mathbf{E}^{(1)}\mathbf{X} + \mu_1\mathbf{E}^{(2)}$, where $\mathbf{X} \leftarrow \mathbf{G}^{-1}(\mathbf{C}_2)$.*

Having established that GG-GSW is a leveled FHE scheme, we now turn to its security properties. The next section introduces our new delayed-target LWE assumption and provides evidence for its hardness. Finally, we use it to prove the IND-CCA1 security of GG-GSW.

4 Delayed-Target LWE

In the the delayed-target LWE (DT-LWE) problem, the adversary first receives a uniformly random matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m\times n}$, as well as access to an oracle $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$ that on input $\mathbf{b} \in \mathbb{Z}_q^m$ returns $\mathbf{s} \in \mathbb{Z}_q^n$ if and only if

$$\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q, \quad \text{with } \|\mathbf{e}\|_2 < \beta.$$

$\mathrm{Exp}^{\mathsf{DT\text{-}LWE}}_{\mathsf{pp},\mathcal{A}}(\lambda)$:

$\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m\times n}$
$\mathsf{st} \leftarrow \mathcal{A}_1^{\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{A},\cdot)}(\mathbf{A})$
$b \xleftarrow{\$} \{0,1\}$
$\mathbf{s} \leftarrow \chi_s, \mathbf{e} \leftarrow \chi_e$
if $b = 0$ **then**
  $\mathbf{b} := \mathbf{As} + \mathbf{e} \bmod q$
else $\mathbf{b} \xleftarrow{\$} \mathbb{Z}_q^m$
$b' \leftarrow \mathcal{A}_2(\mathsf{st}, \mathbf{b})$
return $(b \stackrel{?}{=} b')$

$\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{A}, \mathbf{b})$:

if $\exists (\mathbf{s}, \mathbf{e}) \in \mathbb{Z}_q^n \times \mathbb{Z}_q^m$ such that
  $\mathbf{b} = \mathbf{As} + \mathbf{e} \bmod q \ \wedge \ \|\mathbf{e}\|_2 < \beta$ **then**
  return $\mathbf{s}$
return $\perp$

Fig. 1. The delayed-target LWE experiment. The parameter β is chosen such that the tuple $(\mathbf{s}, \mathbf{e})$ is unique if it exists.

In a second phase, oracle access is disabled and the adversary must distinguish the following distributions

$$\mathbf{b} \xleftarrow{\$} \mathbb{Z}_q^m \quad \text{and} \quad \mathbf{b} = \mathbf{As} + \mathbf{e} \bmod q, \quad \text{where } \mathbf{s} \leftarrow \chi_s, \mathbf{e} \leftarrow \chi_e.$$

Throughout this section, n, m, q are positive integers with $n = n(\lambda)$ and $m = n + \omega(\log n)$ unless specified otherwise. Furthermore, $\sigma, \beta > 0$ denote the Gaussian error parameter and decoding radius, respectively.

Definition 4.1 (Delayed-target LWE). *Let χ_s, χ_e be probability distributions over $\mathbb{Z}^n$ and $\mathbb{Z}^m$ respectively and set $\mathsf{pp} = (n, m, q, \chi_s, \chi_e, \beta)$. The delayed-target LWE (DT-LWE) problem is defined by the experiment in Figure 1. For an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, define its advantage as*

$$\mathsf{Adv}^{\mathsf{DT\text{-}LWE}}_{\mathsf{pp},\mathcal{A}}(\lambda) := \left| \Pr\left[\mathrm{Exp}^{\mathsf{DT\text{-}LWE}}_{\mathsf{pp},\mathcal{A}}(\lambda) = 1\right] - \frac{1}{2} \right|.$$

The delayed-target LWE assumption states that for any PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, $\mathsf{Adv}^{\mathsf{DT\text{-}LWE}}_{\mathsf{pp},\mathcal{A}}(\lambda)$ is negligible.

If χ_s is the uniform distribution over $\mathbb{Z}_q^n$ and χ_e is a discrete Gaussian over $\mathbb{Z}^m$ with parameter σ, then we write $\mathsf{pp} = (n, m, q, \sigma, \beta)$.

The DT-LWE problem also admits a search variant, denoted s-DT-LWE. We give a search to decision reduction for DT-LWE based on the sample-preserving reduction for LWE by Micciancio and Mol [44]. The sample-preserving property allows us to simulate $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$, as the oracle is defined for a fixed matrix $\mathbf{A}$. We state the definition of s-DT-LWE and prove the following proposition in Appendix D.1.

$\mathcal{G}_{\mathsf{LWE}}(n, m, q)$:

1 : Sample $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m \times n}$
2 : Compute a (canonical) basis $\mathbf{B}$ for the lattice $\mathcal{L}_{\mathbf{A}} \coloneqq \mathbf{A}\mathbb{Z}^n + q\mathbb{Z}^m$ by computing the hermite normal form of $[\mathbf{A}|q\mathbf{I}_m]$
3 : Return $\mathbf{B}$

Fig. 2. The q-ary lattice basis generation algorithm $\mathcal{G}_{\mathsf{LWE}}(n, m, q)$.

Proposition 4.1 (Search to decision). *Let* $\mathsf{pp} = (n, m, q, \sigma, \beta)$ *be parameters for the* DT-LWE *problem. Suppose that* $m \geq n + \omega(\log n)$*, that* $q = \mathsf{poly}(n)$*, and that one of the following conditions hold:*

1. *The modulus* q *is prime.*
2. *The modulus* $q = p^w$ *is a prime power for prime* $p = \mathsf{poly}(n)$ *and* $\sigma = o(p/\sqrt{n})$.

Then, for every PPT *adversary* $\mathcal{A}$ *against* $\mathsf{DT\text{-}LWE}_{\mathsf{pp}}$ *with non-negligible advantage, there exists a* PPT *adversary* $\mathcal{B}$ *against* $\mathsf{s\text{-}DT\text{-}LWE}_{\mathsf{pp}}$ *with non-negligible advantage.*

On the falsifiability of the assumption. In general, the DT-LWE assumption is not known to be falsifiable in the sense of [46,32]. Indeed, assuming the hardness of the corresponding LWE problem, the challenger cannot implement the oracle $\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{b})$ for a matrix $\mathbf{A}$ sampled from the uniform distribution in polynomial-time. However, in some parameter regimes, lattice-trapdoor algorithms such as those of [30,45] allows the challenger to sample a tuple $(\mathbf{A}, \tau)$ such that the distribution of $\mathbf{A}$ is statistically close to uniform and, given τ, efficiently implement $\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{b})$. In this paper, we will focus on that regime.

4.1 On the Hardness of DT-LWE

Recall from our discussion in the introduction that the $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$ oracle is only useful if the adversary is able to obliviously sample $\mathbf{b} = \mathbf{As} + \mathbf{e} \bmod q$ with sufficiently short $\mathbf{e}$. We show that under the LK-ϵ assumption [41, Definition 2], the DT-LWE problem is no easier than the LWE problem. Let $\mathbf{B}$ be a basis of a random lattice generated by some algorithm $\mathcal{G}$. Informally, the assumption states that any adversary that takes $\mathbf{B}$ as input and outputs a point sufficiently close to the lattice, must know the corresponding lattice point. We recall the formal definition in Appendix B.

In this paper, we will choose $\mathcal{G}$ to generate a basis for random q-ary lattices that correspond to the LWE problem. More precisely, let n, m, q be positive integers. We define the basis generation algorithm $\mathcal{G}_{\mathsf{LWE}}(n, m, q)$ in Figure 2.

Remark 4.1. As alluded to before, this assumption is believed to hold for classical adversaries. However, there exists quantum algorithms [24,16] that obliviously sample LWE vectors, contradicting LK-ϵ for the $\mathcal{G}_{\mathsf{LWE}}$ algorithm against quantum adversaries - at least in *some* parameter regimes.

Proposition 4.2. *Let* $\mathsf{pp} = (n, m, q, \chi_s, \chi_e, \beta)$ *be parameters for the* $\mathsf{DT\text{-}LWE}$ *problem, let* $\epsilon \in (0, 1/2)$ *and set* $\mathsf{pp}' = (n, m, q, \chi_s, \chi_e)$*. Suppose that*

$$m \geq n \log q \quad \textit{and} \quad \beta \leq \frac{\epsilon q}{12}.$$

For any PPT *adversary* $\mathcal{A}$ *against* $\mathsf{DT\text{-}LWE}_{\mathsf{pp}}$*, there exist* PPT *adversaries* $\mathcal{B}$ *and* $\mathcal{C}$ *with extractor* $\mathcal{C}^*$ *such that*

$$\mathsf{Adv}^{\mathsf{DT\text{-}LWE}}_{\mathsf{pp},\mathcal{A}}(\lambda) \leq \mathsf{Adv}^{\mathsf{LWE}}_{\mathsf{pp}',\mathcal{B}}(\lambda) + \mathsf{Adv}^{\text{LK-}\epsilon}_{\mathcal{G}_{\mathsf{LWE}},\mathcal{C},\mathcal{C}^*}(\lambda) + \mathsf{negl}(\lambda).$$

Proof. We fix a (deterministic) Smith normal form algorithm (SNF). For a basis $\mathbf{B}$ of $\mathcal{L}(\mathbf{B}) = \mathbf{A}\mathbb{Z}^n + q\mathbb{Z}^m$ for some $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m \times n}$, define $\mathrm{Can}(\mathbf{B})$ as:

1. Compute the SNF of $\mathbf{B}$ to obtain matrices $\mathbf{U}, \mathbf{V} \in \mathrm{GL}_m(\mathbb{Z})$ such that $\mathbf{UBV} = \mathbf{D}$, with diagonal $\mathbf{D} = \mathrm{diag}(d_1, \ldots, d_m) \in \mathbb{Z}^{m \times m}$.
2. Return $\mathbf{A}_0 \coloneqq \mathbf{U}^{-1} \begin{bmatrix} \mathbf{I}_n \\ \mathbf{0} \end{bmatrix} \bmod q \in \mathbb{Z}_q^{m \times n}$.

Notice that $\mathbf{A}_0$ is left-invertible. Moreover, since $m \geq n + \omega(\log n)$, $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ is left invertible with overwhelming probability, and conditioned on this

$$\mathbb{Z}^m / \mathcal{L}(\mathbf{B}) \cong \mathbb{Z}_q^m / \operatorname{im}(\mathbf{A}) \cong \mathbb{Z}_q^{m-n} \cong \bigoplus_{i=1}^{m} \mathbb{Z}/d_i\mathbb{Z}.$$

From the properties of the SNF, we get $\mathbf{D} = \mathrm{diag}(\mathbf{I}_n, q\mathbf{I}_{m-n})$. Hence,

$$\begin{aligned} \mathcal{L}(\mathbf{B}) \bmod q &= \mathbf{U}^{-1} \mathrm{diag}(\mathbf{I}_n, q\mathbf{I}_{m-n}) \mathbf{V}^{-1} \mathbb{Z}^m \bmod q \\ &= \mathbf{U}^{-1} (\mathbb{Z}^n \times q\mathbb{Z}^{m-n}) \bmod q = \operatorname{im}(\mathbf{A}_0). \end{aligned}$$

Let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a PPT adversary against $\mathsf{DT\text{-}LWE}_{\mathsf{pp}}$. We first construct an LK-ϵ adversary $\mathcal{C}$ against $\mathcal{G}_{\mathsf{LWE}}$ from $\mathcal{A}$.

The LK-ϵ *adversary* $\mathcal{C}(\mathbf{B}; r)$. The random tape r includes $\mathbf{X} \xleftarrow{\$} \mathrm{GL}_n(\mathbb{Z}_q)$ and random coins ρ for $\mathcal{A}$.

1. Compute $\mathbf{A}_0 = \mathrm{Can}(\mathbf{B})$ and set $\mathbf{A}' \coloneqq \mathbf{A}_0\mathbf{X} \bmod q$.
2. Run $\mathsf{st} \leftarrow \mathcal{A}_1^{\mathcal{O}_{\mathsf{DT\text{-}LWE}}}(\mathbf{A}'; \rho)$. Answer a query $\mathbf{b} \in \mathbb{Z}_q^m$ to $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$ with:
   - Query $\mathbf{p} \leftarrow \mathcal{O}(\mathbf{b}')$, where $\mathbf{b}'$ is the integer lift of $\mathbf{b}$.
   - If $\mathbf{p} \notin \mathcal{L}(\mathbf{B})$ or $\|\mathbf{p} - \mathbf{b}'\|_2 \geq \beta$, return $\perp$.
   - Otherwise, compute $\mathbf{s} \in \mathbb{Z}_q^n$ satisfying $\mathbf{A}'\mathbf{s} = \mathbf{p} \bmod q$.
   - Return $\mathbf{s}$.
3. Halt when $\mathcal{A}_1$ halts.

By the LK-ϵ assumption, there exists a PPT extractor $\mathcal{C}^*$ corresponding to $\mathcal{G}_{\mathsf{LWE}}$ and $\mathcal{C}$. We give an adversary $\mathcal{B}$ against $\mathsf{LWE}_{\mathsf{pp}'}$ that uses $\mathcal{A}$ and $\mathcal{C}^*$ as a subroutine.

The LWE *adversary* $\mathcal{B}(\mathbf{A}, \mathbf{b})$. On input the LWE challenge $(\mathbf{A}, \mathbf{b}) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$. If $\mathbf{A}$ is not left-invertible, abort. Otherwise:

1. Compute the basis $\mathbf{B}$ from $\mathbf{A}$ as $\mathcal{G}_{\mathsf{LWE}}$ does and set $\mathbf{A}_0 = \mathrm{Can}(\mathbf{B})$.
2. Compute $\mathbf{L}_0$ the left inverse of $\mathbf{A}_0$ and set $\mathbf{X} \coloneqq \mathbf{L}_0\mathbf{A} \bmod q \in \mathbb{Z}_q^{n \times n}$.
3. Sample random coins ρ for $\mathcal{A}$ and set $\mathsf{st}' \coloneqq (\mathbf{B}, r = (\mathbf{X}, \rho))$.
4. Run $\mathsf{st} \leftarrow \mathcal{A}_1^{\mathcal{O}_{\mathsf{DT\text{-}LWE}}}(\mathbf{A}; \rho)$. Answer its oracle queries exactly as $\mathcal{C}$ does above, using $\mathcal{C}^*$ with state st' in place of $\mathcal{O}$.
5. Return $b' \leftarrow \mathcal{A}_2(\mathsf{st}, \mathbf{b})$.

Since $\mathbf{B}$ only depends on $\mathrm{im}(\mathbf{A})$, conditioned on $\mathbf{B}$ and $\mathbf{A}$ being left-invertible, $\mathbf{X} = \mathbf{L}_0\mathbf{A}$ is uniform over $\mathrm{GL}_n(\mathbb{Z}_q)$. In that case, $(\mathbf{B}, (\mathbf{X}, \rho))$ supplied to $\mathcal{C}^*$ has the expected distribution.

It remains to verify that the $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$ oracle is correctly simulated. Suppose that $\mathbf{b}$ is a valid query, i.e. $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q$ with $\|\mathbf{e}\|_2 < \beta$, and let $\mathbf{p}* = \mathbf{b}' - \mathbf{e} \in \mathcal{L}(\mathbf{B})$. Then $d_\infty(\mathbf{b}', \mathcal{L}(\mathbf{B})) < \beta$. From Lemma 3.1 with $r = q/12$, except with probability at most q^{-n}, we have

$$\beta \leq \epsilon\lambda_\infty(\mathcal{L}(\mathbf{B})). \tag{6}$$

Therefore, unless the bad flag is set in the LK-ϵ experiment, $\mathcal{C}^*$ returns some $\mathbf{p} \in \mathcal{L}(\mathbf{B})$ satisfying

$$\|\mathbf{p} - \mathbf{b}'\|_\infty \leq \epsilon\lambda_\infty(\mathcal{L}(\mathbf{B})).$$

Since $\mathbf{p} - \mathbf{p}^* \in \mathcal{L}(\mathbf{B})$ and $\|\mathbf{p} - \mathbf{p}^*\|_\infty < 2\epsilon\lambda_\infty(\mathcal{L}(\mathbf{B}))$ with $\epsilon < 1/2$, we have $\mathbf{p} = \mathbf{p}^*$. Thus every valid query receives the correct response. Conversely, for any query $\mathbf{b} \in \mathbb{Z}_q^m$ accepted by the simulator

$$\mathbf{b} = \mathbf{A}\mathbf{s} + (\mathbf{b}' - \mathbf{p}) \bmod q$$

with error norm less than β. Hence, conditioned on Equation (6) holding and $\overline{\mathsf{bad}}$, $\mathcal{B}$ perfectly simulates the $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$.

Putting it together, conditioned on $\mathbf{A}$ being left invertible, Equation (6) holding, and $\overline{\mathsf{bad}}$, the adversary $\mathcal{B}$ has advantage at least $\mathsf{Adv}^{\mathsf{DT\text{-}LWE}}_{\mathsf{pp},\mathcal{A}}(\lambda)$, and the claimed bound follows. □

Remark 4.2. Proposition 4.2 shows that it is necessary to be able to obliviously sample an accepting vector $\mathbf{b}$ in order for the $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$ oracle to be useful. However, it is not sufficient to solve the DT-LWE problem. For example, consider an adversary that has access to an oracle that outputs accepting vectors $\mathbf{b}$. Then submitting these vectors to the DT-LWE oracle always outputs 1, which is of no use to the adversary. This can be further generalized to any oracle that outputs vectors $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q$ where $(\mathbf{s}, \mathbf{e})$ are taken from publicly samplable distributions. In particular, the quantum oblivious sampling algorithms from [24,16] invalidate the LK-ϵ assumption for $\mathcal{G}_{\mathsf{LWE}}$, but to the best of our knowledge, do not invalidate the DT-LWE assumption.

Quantum reduction from OM-ISIS. Recently, Chailloux and Hermouet gave a quantum reduction from the inhomogeneous short integer solution (ISIS) problem to the search LWE problem [14]. Based on this, we give a quantum reduction from the one-more ISIS (OM-ISIS) problem [2] to DT-LWE. We stress that this reduction does *not* cover the parameter regime used by the construction given in Section 3. We recall the definition of the one-more ISIS problem in Appendix B.

Proposition 4.3. *Let* $\mathsf{pp} = (n, m, q, \sigma_{\mathsf{OM}}, \beta_{\mathsf{OM}})$ *and* $\mathsf{pp}' = (n, m, q, \sigma_{\mathsf{DT}}, \beta_{\mathsf{DT}})$ *be parameters for the* OM-ISIS *and* DT-LWE *problems respectively. Suppose that* $q \geq 3$ *and*

$$m \geq n \log q, \qquad \sigma_{\mathsf{OM}} = \Omega(\sqrt{m}), \qquad \min(\sigma_{\mathsf{DT}}, \frac{q}{\sigma_{\mathsf{DT}}}) = \omega(\sqrt{\log \lambda})$$

$$\beta_{\mathsf{OM}} \geq \sqrt{m}\frac{q}{2\sigma_{\mathsf{DT}}} \quad \textit{and} \quad \beta_{\mathsf{DT}} \leq \frac{q}{c \cdot \sigma_{\mathsf{OM}}(\sqrt{\kappa} + \sqrt{m} + \sqrt{\lambda})},$$

where $\kappa = n\lceil \log q \rceil$ *and* $c > 0$ *is a universal constant. If there exists an efficient, possibly quantum, algorithm for* $\mathsf{s\text{-}DT\text{-}LWE}_{\mathsf{pp}'}$ *that succeeds with non-negligible probability* ϵ_{DT}*, then there exists an efficient quantum algorithm for* $\mathsf{OM\text{-}ISIS}_{\mathsf{pp}}$ *that succeeds with probability* $\epsilon_{\mathsf{OM}} = \Omega(\epsilon_{\mathsf{DT}})$*.*

Note that known attacks on OM-ISIS impose $\sigma_{\mathsf{DT}} = \Omega(q/(\sigma_{\mathsf{OM}}\sqrt{m}))$ for the corresponding OM-ISIS problem to be plausibly hard [2]. We prove the proposition in Appendix D.2.

Corollary 4.1. *Let* pp *and* pp′ *be as given in Proposition 4.3 and* q *fulfill the conditions given in Proposition 4.1. If there exists an efficient (quantum) algorithm for* $\mathsf{DT\text{-}LWE}_{\mathsf{pp}'}$*, then there exists an efficient quantum algorithm for* $\mathsf{OM\text{-}ISIS}_{\mathsf{pp}}$*.*

Proof. This follows directly from Proposition 4.3 and Proposition 4.1. □

4.2 IND-CCA1 Security of GG-GSW

Finally, we are ready to prove the IND-CCA1 security of the GG-GSW scheme under the DT-LWE assumption.

Theorem 4.1 (IND-CCA1 Security). *Let* $\mathsf{pp} = (n, \bar{m}, q, \sigma, D)$*. For any* PPT *adversary* $\mathcal{A}$ *against the* IND-CCA1 *security of the* GG-GSW *scheme, there exist a* DT-LWE *distinguisher* $\mathcal{B}$ *such that*

$$\mathsf{Adv}^{\text{IND-CCA1}}_{\mathsf{GG\text{-}GSW},\mathcal{A}}(\lambda) \leq 2k \cdot \mathsf{Adv}^{\mathsf{DT\text{-}LWE}}_{\mathcal{B}}(\lambda) + \mathsf{negl}(\lambda).$$

Proof. In Figure 3, we describe an alternate decryption algorithm $\mathsf{Dec}^{\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{A},\cdot)}$ that uses the DT-LWE oracle $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$ as a subroutine. We show in Lemma 4.1 that for every ciphertext $\mathbf{C} \in \mathbb{Z}_q^{\bar{m} \times k}$, both decryption algorithms behave the same except with negligible probability over the key generation randomness.

$\mathsf{Dec}^{\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{A},\cdot)}(\mathbf{C})$:

1 : **for** $\mu \in \{0,1\}$ **do**
2 : $\quad [\mathbf{b}_{1,\mu}|\ldots|\mathbf{b}_{k,\mu}] \coloneqq \mathbf{C} - \mu\mathbf{G}$
3 : $\quad$ **for** $i \in [k]$ **do**
4 : $\quad\quad \mathbf{s}_{i,\mu} \leftarrow \mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{b}_{i,\mu})$
5 : $\quad$ **if** $\perp \notin \{\mathbf{s}_{i,\mu}\}_{i\in[k]}$ **then**
6 : $\quad\quad$ **return** μ
7 : **return** $\perp$

Fig. 3. The oracle-aided decryption algorithm $\mathsf{Dec}^{\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{A},\cdot)}(\cdot)$.

We consider a sequence of hybrid games. For each index i, let W_i be the event that the adversary $\mathcal{A}$ wins in game G_i.

Game G_0: This is the normal IND-CCA1 game.

Game G_1: In this game, the challenger now uses the oracle-aided decryption algorithm $\mathsf{Dec}^{\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{A},\cdot)}$ instead of $\mathsf{GG\text{-}GSW.Dec}$ to answer decryption queries. It follows from Lemma 4.1 that,

$$|\Pr[W_0] - \Pr[W_1]| \leq \mathsf{negl}(\lambda).$$

Game G_2: We now change how the public key is computed. Sample $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{\bar{m}\times n}$ directly from the uniform distribution. Since decryption queries are now answered through the statistical oracle, it is not an issue that $\mathbf{A}_2$ does not contain the $\hat{\mathbf{G}}$ term. Hence, it follows from the leftover hash lemma that,

$$|\Pr[W_1] - \Pr[W_2]| \leq \mathsf{negl}(\lambda).$$

Game G_3: In this game, we sample $\mathbf{C} \xleftarrow{\$} \mathbb{Z}_q^{\bar{m}\times k}$, and set the challenge ciphertext to be $\mathbf{C}^* = \mathbf{C} + \mu_b\mathbf{G} \bmod q$.

We show that game G_2 and G_3 are indistinguishable under the $\mathsf{DT\text{-}LWE}$ assumption. For $i \in \{0,\ldots,k\}$, let H_i be the hybrid game in which key generation and decryption are handled as in game G_2, but the challenge ciphertext is computed as $\mathbf{C}^* = [\mathbf{c}_1|\cdots|\mathbf{c}_k] + \mu_b\mathbf{G} \bmod q$, where

$$\mathbf{c}_j = \begin{cases} \mathbf{u}_j, & j \leq i, \\ \mathbf{A}\mathbf{s}_j + \mathbf{e}_j \bmod q, & j > i, \end{cases}$$

where $\mathbf{u}_j \xleftarrow{\$} \mathbb{Z}_q^{\bar{m}}$, $\mathbf{s}_j \xleftarrow{\$} \mathbb{Z}_q^n$, and $\mathbf{e}_j \xleftarrow{\$} D_{\mathbb{Z}^{\bar{m}},\sigma}$ are sampled independently. Thus, $\mathsf{H}_0 = G_2$ and $\mathsf{H}_k = G_3$.

We construct a $\mathsf{DT\text{-}LWE}_{\mathsf{pp}}$ distinguisher $\mathcal{B}$ as follows. Given $\mathbf{A}$ and oracle access to $\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{A}, \cdot)$, $\mathcal{B}$ first samples an index $i^* \xleftarrow{\$} [k]$. It then runs the first stage of $\mathcal{A}$, and answers its decryption queries using $\mathsf{Dec}^{\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{A},\cdot)}$. When $\mathcal{A}$ outputs messages (μ_0, μ_1), $\mathcal{B}$ enters the second phase of the game. After the oracle is disabled, $\mathcal{B}$ receives its $\mathsf{DT\text{-}LWE}$ target $\mathbf{b}$ and samples a challenge bit b. It then computes the challenge ciphertext by sampling columns $\mathbf{c}_j$ from the uniform distribution for $j < i^*$, setting $\mathbf{c}_{i^*} = \mathbf{b}$ and sampling the remaining columns for $i^* < j \leq k$ as described above. It then sets $\mathbf{C}^* = [\mathbf{c}_1 | \cdots | \mathbf{c}_k] + \mu_b \mathbf{G} \bmod q$ and runs the second stage of $\mathcal{A}$ with input $\mathbf{C}^*$. Finally, $\mathcal{B}$ outputs 1 if $\mathcal{A}$ correctly guesses b, and outputs 0 otherwise.

If $\mathbf{b}$ is an LWE sample, then $\mathcal{B}$ perfectly simulates H_{i^*-1}; if $\mathbf{b}$ is uniform, it perfectly simulates H_{i^*}. It follows from a standard hybrid argument that

$$|\Pr[W_2] - \Pr[W_3]| \leq 2k \cdot \mathsf{Adv}^{\mathsf{DT\text{-}LWE}}_{\mathsf{pp},\mathcal{B}}(\lambda).$$

Finally, in game G_3 the adversary's view is statistically independent from the challenge bit b, and we have

$$\Pr[W_3] = \frac{1}{2},$$

which concludes the proof. □

Lemma 4.1. *Let* $(\mathsf{pk} = \mathbf{A}, \mathsf{sk} = \mathbf{R}) \leftarrow \mathsf{GG\text{-}GSW.Gen}(1^\lambda)$. *Except with negligible probability over the randomness of key generation, for every* $\mathbf{C} \in \mathbb{Z}_q^{\bar{m} \times k}$,

$$\mathsf{GG\text{-}GSW.Dec}(\mathsf{sk}, \mathbf{C}) = \mathsf{Dec}^{\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{A},\cdot)}(\mathbf{C}).$$

Proof. Suppose that there exists $(\mathbf{S}, \mathbf{E}, \mu) \in \mathbb{Z}_q^{n \times k} \times \mathbb{Z}_q^{\bar{m} \times k} \times \{0, 1\}$ such that $\mathbf{C} = \mathbf{AS} + \mathbf{E} + \mu\mathbf{G} \bmod q$ and $\|\mathbf{E}\|_2 < D$. From Lemma 3.2, except with negligible probability over the randomness of key generation, $\mathsf{GG\text{-}GSW.Dec}(\mathsf{sk}, \mathbf{C})$ returns μ. Furthermore, let $[\mathbf{b}_{1,\mu} | \ldots | \mathbf{b}_{k,\mu}] \coloneqq \mathbf{C} - \mu\mathbf{G} = \mathbf{AS} + \mathbf{E}$. By definition of $\mathsf{DT\text{-}LWE}_{\mathsf{pp}}$, the oracle $\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{b}_{i,\mu})$ returns $\mathbf{s}_{i,\mu} \neq \perp$, the i-th column of $\mathbf{S}$, since $\|\mathbf{E}\|_2 < D$. Hence $\mathsf{Dec}^{\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{A},\cdot)}(\mathbf{C})$ also returns μ with probability 1.

Conversely, suppose that no such tuple $(\mathbf{S}, \mathbf{E}, \mu)$ exists. By the assertion in Step 2 of $\mathsf{GG\text{-}GSW.Dec}$, it returns $\perp$. Similarly, $\mathsf{Dec}^{\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{A},\cdot)}$ returns $\perp$ as there must exists at least one column $\mathbf{b}_{i,\mu}$ such that $\perp \leftarrow \mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{b}_{i,\mu})$. □

5 Revisiting IND-CCA1 Security

Brzuska et al. [12] gave a separation between CCA1 and CPA^D security. An active adversary, such as the one modeled by the CCA1 game, should be at least as powerful as a passive adversary. This separation is therefore rather counterintuitive and highlights that CCA1 security may be insufficient in settings where CPA^D security is relevant. In order to bridge this gap, we introduce the IND-CCA1^D game, which combines both games. We show that this restores the expected hierarchy between active and passive security notions. Finally, we show how to make GG-GSW IND-CCA1^D-secure without relying on correctness.

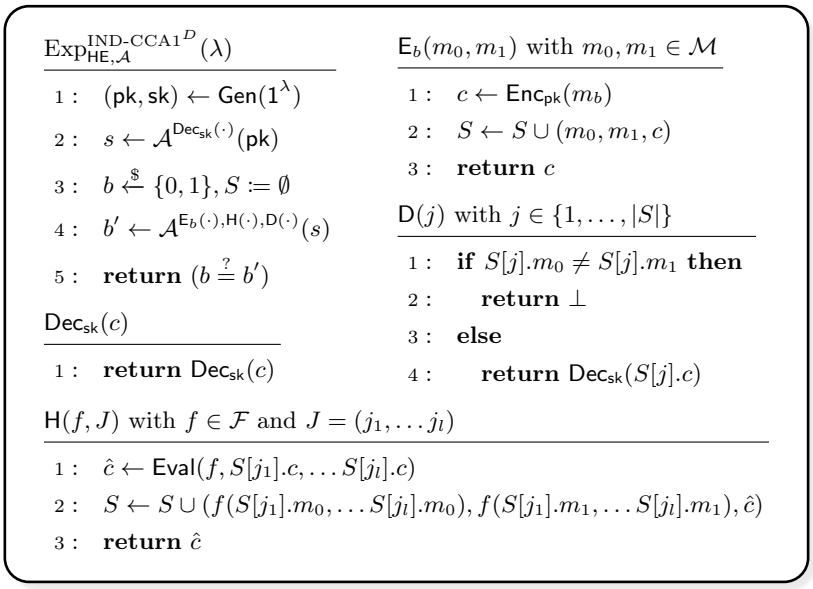

Fig. 4. The IND-CCA1^D experiment $\mathrm{Exp}^{\text{IND-CCA1}^D}_{\mathsf{HE},\mathcal{A}}(\lambda)$, with the unrestricted decryption oracle $\mathsf{Dec}_{\mathsf{sk}}(\cdot)$ and the CPA^D oracles $\mathsf{E}_b(m_0,m_1)$, $\mathsf{H}(f,J)$, and $\mathsf{D}(j)$.

5.1 A Unified Security Model

Informally, the IND-CCA1^D game proceeds in two phases. The adversary first receives unrestricted access to the decryption oracle. In the second phase, it now interacts with the CPA^D encryption, evaluation, and decryption oracles.

Definition 5.1 (IND-CCA1^D)**.** *Let* $\mathsf{HE} = (\mathsf{Gen},\mathsf{Enc},\mathsf{Dec},\mathsf{Eval})$ *be a public-key HE scheme. Consider the* IND-CCA1^D *experiment given in Figure 4. The advantage of an adversary* $\mathcal{A}$ *is given by:*

$$\mathsf{Adv}^{\text{IND-CCA1}^{\text{D}}}_{\mathsf{HE},\mathcal{A}}(\lambda) = |\Pr[\mathrm{Exp}^{\text{IND-CCA1}^D}_{\mathsf{HE},\mathcal{A}}(\lambda) = 1] - \frac{1}{2}|.$$

We say that HE *is* indistinguishable against non-adaptive chosen-ciphertext attack and passive decryption (IND-CCA1^D) *if for all* PPT *adversaries* $\mathcal{A}$*, the advantage* $\mathsf{Adv}^{\text{IND-CCA1}^{\text{D}}}_{\mathsf{HE},\mathcal{A}}(\lambda)$ *is negligible in* λ*.*

The following proposition shows that IND-CCA1 and IND-CPA^D security do not imply IND-CCA1^D security. Indeed, the CCA1 oracle may uncover a partial hint to an attacker that requires the CPA^D decryption oracle to complete. We give the proof in Appendix E.

Proposition 5.1. *If there exists a statistically correct,* IND-CCA1 *secure HE scheme, then there exists an HE scheme that is both* IND-CCA1 *secure and* IND-CPA^D *secure, but not* IND-CCA1^D *secure.*

Lemma 5.1. *Any public-key FHE scheme that is* IND-CCA1 *secure and statistically correct achieves* IND-CCA1^D *security.*

Proof. The proof follows the reduction from IND-CPA security and statistical correctness to IND-CPA^D security [36], using the CCA1 oracle for first-phase decryption before applying the CPA^D reduction.

Since GG-GSW is an exact FHE scheme, we can choose its parameters to obtain statistical correctness, in which case IND-CCA1^D security follows from Lemma 5.1 and Theorem 4.1. However, this is somewhat unsatisfactory. Most applications do not require such high correctness guarantees in practice, and may benefit more from trading off correctness for better performance. The remainder of this section is dedicated to obtaining IND-CCA1^D security under parameters that do not guarantee correctness.

5.2 Delayed-target LWE with Hints

We introduce the delayed-target hint LWE (DT-HintLWE) problem, which we will use in Section 5.4 to prove IND-CCA1^D security. This is a delayed-target analogue of the HintLWE problem. We show that DT-HintLWE reduces to DT-LWE.

Definition 5.2 (Delayed-target HintLWE). *Let n, m, q, t be positive integers and $\beta > 0$. Let χ_s, χ_e, χ_f be probability distributions over $\mathbb{Z}^n, \mathbb{Z}^m$ and $\mathbb{Z}^t$ respectively. Let $\mathcal{L} \subseteq \mathbb{Z}_q^{t\times n} \times \mathbb{Z}^{t\times m}$ be a polynomial-time decidable family of matrices. Set $\mathsf{pp} = (n, m, q, t, \chi_s, \chi_e, \chi_f, \mathcal{L}, \beta)$. The delayed-target HintLWE (*DT-HintLWE*) problem is defined by the experiment in Figure 5. For an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, define its advantage as*

$$\mathsf{Adv}^{\mathsf{DT\text{-}HintLWE}}_{\mathsf{pp},\mathcal{A}}(\lambda) := \left|\Pr\left[\mathrm{Exp}^{\mathsf{DT\text{-}HintLWE}}_{\mathsf{pp},\mathcal{A}}(\lambda) = 1\right] - \frac{1}{2}\right|$$

The delayed-target HintLWE assumption states that for any PPT *adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, $\mathsf{Adv}^{\mathsf{DT\text{-}HintLWE}}_{\mathsf{pp},\mathcal{A}}(\lambda)$ is negligible.*

If χ_s is the uniform distribution over $\mathbb{Z}_q^n$ and χ_e and χ_f are discrete Gaussians with parameter σ and σ_f respectively, then we write $\mathsf{pp} = (n, m, q, t, \sigma, \sigma_f, \mathcal{L}, \beta)$.

Lai, Swarnakar, and Woo establish the hardness of "leaky" LWE problems from the corresponding LWE problem. The following proposition is essentially a special case of their framework [35, Theorem 3, Condition 2]. Their result supports arbitrary secret distributions and nonspherical Gaussian errors. However, we only state the specialization to uniform secrets, spherical discrete-Gaussian errors, and error-only hints, which matches our setting in Section 5.4.

Proposition 5.2. *Let $\mathsf{pp} = (n, m, q, t, \sigma, \sigma_f, \mathcal{L}, \beta)$ be a parameter set for the* DT-HintLWE *problem. Let $\epsilon \in (0, 1)$, $\sigma' > 0$ and let $\mathsf{pp}' = (n, m, q, \sigma', \beta)$. Suppose that $\sigma' \geq \sqrt{2}\eta_\epsilon(\mathbb{Z}^m)$, and that every $(\mathbf{L}_s, \mathbf{L}_e) \in \mathcal{L}$ satisfies*

$$\mathbf{L}_s = \mathbf{0} \quad \textit{and} \quad s_1(\mathbf{L}_e)^2 \leq \sigma_f^2 \left(\frac{1}{2\eta_\epsilon(\mathbb{Z}^m)^2 + \sigma'^2} - \frac{1}{\sigma^2}\right).$$

$\mathrm{Exp}^{\mathsf{DT\text{-}HintLWE}}_{\mathsf{pp},\mathcal{A}}(\lambda)$:

1 : $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m\times n}$

2 : $(\mathsf{st}, \mathbf{L}_s, \mathbf{L}_e) \leftarrow \mathcal{A}_1^{\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{A},\cdot)}(\mathbf{A})$; **assert** $(\mathbf{L}_s, \mathbf{L}_e) \in \mathcal{L}$

3 : $b \xleftarrow{\$} \{0,1\}$

4 : $\mathbf{s} \leftarrow \chi_s, \mathbf{e} \leftarrow \chi_e, \mathbf{f} \leftarrow \chi_f$

5 : **if** $b = 0$ **then** $\mathbf{b} \coloneqq \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q$

6 : **if** $b = 1$ **then** $\mathbf{b} \xleftarrow{\$} \mathbb{Z}_q^m$

7 : $\mathbf{h} \coloneqq [\mathbf{L}_s | \mathbf{L}_e] \begin{bmatrix} \mathbf{s} \\ \mathbf{e} \end{bmatrix} + \mathbf{f}$

8 : $b' \leftarrow \mathcal{A}_2(\mathsf{st}, \mathbf{b}, \mathbf{h})$

9 : **return** $(b \stackrel{?}{=} b')$

Fig. 5. The delayed-target HintLWE experiment. The $\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{A}, \cdot)$ oracle is defined in Figure 1.

For any PPT *adversary* $\mathcal{A}$ *against* $\mathsf{DT\text{-}HintLWE}_{\mathsf{pp}}$, *there exists a* PPT *adversary* $\mathcal{B}$ *against* $\mathsf{DT\text{-}LWE}_{\mathsf{pp}'}$ *such that*

$$\mathsf{Adv}^{\mathsf{DT\text{-}HintLWE}}_{\mathsf{pp},\mathcal{A}}(\lambda) \leq \mathsf{Adv}^{\mathsf{DT\text{-}LWE}}_{\mathsf{pp}',\mathcal{B}}(\lambda) + \frac{2\epsilon}{1-\epsilon}.$$

5.3 Randomized Evaluation For Branching Programs

We recall branching programs in Appendix F. Building on techniques introduced by Brakerski and Vaikuntanathan [11] as well as Bourse et al. [7], we define a randomized evaluation algorithm $\mathsf{GG\text{-}GSW.Eval}$ for branching programs. Throughout this section, we use the notation and parameters from Section 3. We overload notation and write $\mathsf{GG\text{-}GSW.Enc}_{\Gamma}(\cdot)$ to denote the encryption algorithm with ellipsoidal discrete Gaussian noise distribution with parameter Γ.

Remark 5.1. The results of this section can also be used to obtain a CPA^D-secure scheme from the CPA-secure dual-GSW scheme. The rerandomization argument does not depend on any special property of the public-key $\mathbf{A}$ or the decryption procedure. Indeed, we expect similar techniques to apply to the GSW scheme itself.

$\mathsf{GG\text{-}GSW.Eval}(\Pi, (\mathbf{C}_1, \dots \mathbf{C}_N))$: On input a branching program Π of length L and width W and ciphertexts $(\mathbf{C}_1, \dots \mathbf{C}_N)$, the algorithm defines the initial state $\mathbf{V}_0 = (\mathbf{G}, \mathbf{0}, \dots, \mathbf{0})$, a tuple of noiseless encryptions of $\mathbf{v}_0 = (1, 0, \dots, 0)$. It then rerandomizes each input ciphertext by computing, for $i \in [N]$

$$\mathbf{C}'_i = \mathbf{C}_i + \widetilde{\mathbf{Y}}_i \quad \text{where } \widetilde{\mathbf{Y}}_i \leftarrow \mathcal{D}^k_{\mathbb{Z}^{\bar{m}},\sigma_f}. \tag{7}$$

The algorithm then evaluates the L levels of the program iteratively. For every $t \in [L]$ and $w \in [W]$, it samples $\mathbf{Y}_{t,w} \leftarrow \mathcal{D}^k_{\mathbb{Z}^{\bar{m}}, r\sqrt{2}}$, and for $b \in \{0,1\}$

$$\mathbf{X}_{t,w,b} \leftarrow \mathbf{G}^{-1}_{\mathsf{rand}}(\mathbf{V}_{t-1}[\pi^{-1}_{t,b}(w)]).$$

Then, it sets

$$\mathbf{V}_t[w] \coloneqq \mathbf{C}'_{\mathsf{var}(t)} \cdot \mathbf{X}_{t,w,1} + (\mathbf{G} - \mathbf{C}'_{\mathsf{var}(t)}) \cdot \mathbf{X}_{t,w,0} + \mathbf{Y}_{t,w}. \tag{8}$$

Finally, the algorithm outputs the ciphertext $\mathbf{V}_L[1]$.

The following lemma can be seen as a matrix dual-GSW generalization of [7, Lemma 5.4]. It gives the noise distribution of $\mathbf{V}_t[w]$ at every level. The proof appears in Appendix G.

Lemma 5.2 (Noise distribution). *Let $\sigma_f > 0$ be a Gaussian parameter and let $r = \Omega\big((\sigma + \sigma_f)\sqrt{k\lambda \log(k)}\big)$. For a branching program $\Pi : \{0,1\}^N \to \{0,1\}$ of length L, let $\tau_{t,i}$ be the number of times the i-th variable x_i has been called after t steps of the evaluation. Let $(\mathsf{pk} = \mathbf{A}, \mathsf{sk}) \leftarrow \mathsf{GG\text{-}GSW.Gen}(1^\lambda)$. For any $(x_1, \dots x_N) \in \{0,1\}^N$, let $\mathbf{C}_i = \mathbf{A}\mathbf{S}_i + \mathbf{E}_i + x_i\mathbf{G} \leftarrow \mathsf{GG\text{-}GSW.Enc}(\mathsf{pk}, x_i)$. For all $t \in [L]$ and $w \in [W]$, let $\mathbf{V}_t[w]$ be as defined in Equation (8). We have,*

$$\left(\mathbf{V}_t[w], (\mathbf{C}_i)_{i=1}^N\right) \overset{s}{\approx} \left(\mathbf{V}'_{t,w}, (\mathbf{C}_i)_{i=1}^N\right),$$

where $\mathbf{V}'_{t,w} \leftarrow \mathsf{GG\text{-}GSW.Enc}_{\Gamma_t}(\mathsf{pk}, v_t[w])$ with $\Gamma_t = r\left(2\sum_{i=1}^N \tau_{t,i}\hat{\mathbf{E}}_i\hat{\mathbf{E}}_i^\intercal\right)^{1/2}$, $\hat{\mathbf{E}}_i = [\mathbf{E}_i + \widetilde{\mathbf{Y}}_i | \mathbf{I}_{\bar{m}}]$ and $\widetilde{\mathbf{Y}}_i \leftarrow \mathcal{D}^k_{\mathbb{Z}^{\bar{m}}, \sigma_f}$.

Using this lemma, we can now bound the noise magnitude of an evaluated ciphertext. We focus on one-hop evaluation. The multihop case can be derived analogously.

Lemma 5.3 (Noise magnitude). *Given the same setting as Lemma 5.2, let $\delta = O((\sigma + \sigma_f)\sqrt{k})$ and define*

$$\rho_L \coloneqq r\sqrt{2L(\delta^2 + 1)\bar{m}}.$$

Then the output noise $\widehat{\mathbf{E}}$ of the evaluation of a branching program of length L satisfies $\|\widehat{\mathbf{E}}\|_2 < \rho_L$ except with negligible probability. In particular, the evaluated ciphertext decrypts correctly whenever $\rho_L < D$.

Proof. Let $\hat{\mathbf{E}}$ be the noise matrix of $\hat{\mathbf{C}}$. From Lemma 5.2, it follows a discrete Gaussian distribution with parameter matrix Γ_L. We have

$$s_1(\Gamma_L) \le r\sqrt{2\sum_{i=1}^N \tau_{L,i}(s_1(\hat{\mathbf{E}}_i\hat{\mathbf{E}}_i^\intercal) + 1)}.$$

We have $s_1(\hat{\mathbf{E}}_i\hat{\mathbf{E}}_i^\intercal) = s_1(\mathbf{E}_i + \widetilde{\mathbf{Y}}_i)^2 \le \delta^2$ for all i, except with negligible probability. Since $\sum_i \tau_{L,i} = L$, we have, $s_1(\Gamma_L) \le r\sqrt{2L(\delta^2+1)}$. It follows from Lemma 2.5 and a union bound that with overwhelming probability, $\|\mathbf{E}\| < s_1(\Gamma_L)\sqrt{\bar{m}} \le r\sqrt{2L(\delta^2+1)\bar{m}} = \rho_L$. The second claims follows from Lemma 3.2. □

5.4 IND-CCA1^D Security Without Correctness Guarantees

For ease of exposition, we state our result for the one-hop variant of the randomized evaluation algorithm, in which evaluation may only be applied to fresh ciphertexts. This can be enforced by augmenting each ciphertext with a freshness flag, which is set to true by encryption and to false by evaluation. Decryption ignores the flag and the evaluation algorithm simply returns $\perp$ if any input is not fresh. We suppress the flag from the notation. We write L_{max} to denote the maximum length of the branching programs we accept.

Let $\mathsf{GG\text{-}GSW}^*$ denote the $\mathsf{GG\text{-}GSW}$ with the one-hop randomized evaluation algorithm. We show that $\mathsf{GG\text{-}GSW}^*$ is an IND-CCA1^D secure leveled FHE scheme for branching programs.

Theorem 5.1 (IND-CCA1^D **Security).** *Let $L_{\max} = \mathsf{poly}(\lambda)$, and let $\rho_{L_{\max}}$ be as defined in Lemma 5.3. Suppose that there exists $B > D$ such that*

$$\rho_{L_{\max}} \leq B \qquad \text{and} \qquad B + D < q/6.$$

For any PPT adversary $\mathcal{A}$ that makes $Q_{\mathsf{enc}} = Q_{\mathsf{enc}}(\lambda)$ encryption queries and $Q_{\mathsf{eval}} = Q_{\mathsf{eval}}(\lambda)$ homomorphic evaluation queries, let

$$d \coloneqq \bar{m} Q_{\mathsf{eval}}, \qquad \mathbf{L}_e \coloneqq \begin{bmatrix} \mathbf{I}_{\bar{m}} \\ \vdots \\ \mathbf{I}_{\bar{m}} \end{bmatrix} \in \mathbb{Z}^{d \times \bar{m}}, \qquad \mathcal{L} \coloneqq \{(\mathbf{0}_{d \times n}, \mathbf{L}_e)\},$$

and set $\mathsf{pp} = (n, \bar{m}, q, d, \sigma, \sigma_f, \mathcal{L}, D)$. Then, there exists a PPT distinguisher $\mathcal{B}$ such that,

$$\mathsf{Adv}^{\text{IND-CCA1}^{\mathrm{D}}}_{\mathsf{GG\text{-}GSW}^*,\mathcal{A}}(\lambda) \leq 2 Q_{\mathsf{enc}} k \cdot \mathsf{Adv}^{\mathsf{DT\text{-}HintLWE}}_{\mathsf{pp},\mathcal{B}}(\lambda) + \mathsf{negl}(\lambda).$$

Proof. We consider a sequence of games G_i. For every index i, let W_i be the event that the adversary wins the game G_i.

Game G_0: This is the normal IND-CCA1^D game.

Game G_1: In this game, queries to the unrestricted decryption oracle are answered using $\mathsf{Dec}^{\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{A},\cdot)}$, the oracle defined in Figure 3. As in the proof of Theorem 4.1, we have $|\Pr[W_0] - \Pr[W_1]| \leq \mathsf{negl}(\lambda)$.

Game G_2: In this game, we augment the state that the challenger maintains to also keep track of ciphertext noise, i.e., entries have the form $S[u] = (m_0, m_1, \mathbf{C}, \mathbf{E})$, where $\mathbf{E}$ is the noise associated with $\mathbf{C}$. The CPA^D oracles are implemented as follows.

- $\mathsf{E}(m_0, m_1)$: The oracle returns $\mathbf{C} = \mathbf{AS} + \mathbf{E} + m_b \mathbf{G} \bmod q$ as usual. It then updates the state with $S \leftarrow S \cup (m_0, m_1, \mathbf{C}, \mathbf{E})$.
- $\mathsf{H}(\Pi, J)$: Similarly, for every homomorphic evaluation query, the oracle keeps track of the noise in the ciphertext, i.e.

  $$S \leftarrow S \cup (\Pi(S[j_1].m_0, \ldots S[j_N].m_0), \Pi(S[j_1].m_1, \ldots S[j_N].m_1), \hat{\mathbf{C}}, \hat{\mathbf{E}}),$$

  where $\hat{\mathbf{E}}$ is computed from the corresponding $S[j_i].\mathbf{E}$ and the evaluation randomness.

- $\mathsf{D}(j)$: If $S[j].m_0 \neq S[j].m_1$, return $\perp$. Otherwise, the restricted decryption oracle returns $S[j].m_0$ if $\|S[j].\mathbf{E}\|_2 < D$, and returns $\perp$ if $\|S[j].\mathbf{E}\|_2 \geq D$.

We show in Lemma 5.4 that the restricted decryption oracles $\mathcal{D}$ behave identically in game G_1 and G_2 except with negligible probability,

$$|\Pr[W_1] - \Pr[W_2]| \leq \mathsf{negl}(\lambda).$$

Game G_3: In this game, both the public-key $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{\bar{m} \times n}$ and the ciphertexts $\mathbf{C} \xleftarrow{\$} \mathbb{Z}_q^{\bar{m} \times k}$ output by the encryption oracle E are sampled from the uniform distribution. The oracle still samples $\mathbf{E} \in \mathbb{Z}^{\bar{m} \times k}$ and updates the state with $\mathbf{C}$ and $\mathbf{E}$ as in game G_2. All other oracles are unchanged. Using the leftover hash lemma and a hybrid argument, we show in Lemma H.1 that

$$|\Pr[W_2] - \Pr[W_3]| \leq 2Q_{\mathsf{enc}} k \cdot \mathsf{Adv}^{\mathsf{DT\text{-}HintLWE}}_{\mathsf{pp},\mathcal{B}}(\lambda) + \mathsf{negl}(\lambda).$$

The proof appear in Appendix H. The key observation is that the DT-HintLWE hints correspond to the flooded noise $\mathbf{E}_i + \widetilde{\mathbf{Y}}_i$. As we have shown in Lemma 5.2 that the noise distribution of an evaluated ciphertext only depend on these values, this allows the evaluation algorithm to update the state without knowing the actual challenge noise $\mathbf{E}_i$.

Finally, notice that in game G_3, the view of the adversary is independent of the challenge bit. Hence, we have, $\Pr[W_3] = 1/2$, which concludes the proof. □

Lemma 5.4. *Game G_1 and G_2 are statistically indistinguishable.*

The proof is given in Appendix H. At a high level, except with negligible probability every error recorded in the state has norm smaller than B. When the recorded error has norm smaller than D, the real and simulated decryption oracles agree by Lemma 3.2. Otherwise, the simulated oracle rejects, whereas acceptance by the real oracle would yield another representation with error smaller than D. If this representation has the same message, it contradicts the minimum-distance bound for $\mathcal{L}_{\mathbf{A}}$; if it has the other message, it contradicts the shifted-coset bound. Both follow from Lemma 3.1, since $B + D < q/6$.

Corollary 5.1. *Let $\sigma' = \Omega(\sqrt{\lambda})$, $\sigma = 2\sigma'$ and $\sigma_f = 2\sigma'\sqrt{Q_{\mathsf{eval}}}$ and let $\mathsf{pp}' = (n, m, q, \sigma', D)$. Then, for any PPT adversary $\mathcal{A}$ that makes $Q_{\mathsf{enc}} = Q_{\mathsf{enc}}(\lambda)$ encryption queries and $Q_{\mathsf{eval}} = Q_{\mathsf{eval}}(\lambda)$ homomorphic evaluation queries, there exists a PPT DT-LWE distinguisher $\mathcal{B}$ such that,*

$$\mathsf{Adv}^{\text{IND-CCA1}^{\mathrm{D}}}_{\mathsf{GG\text{-}GSW}^*,\mathcal{A}}(\lambda) \leq 2Q_{\mathsf{enc}} k \cdot \mathsf{Adv}^{\mathsf{DT\text{-}LWE}}_{\mathsf{pp}',\mathcal{B}}(\lambda) + \mathsf{negl}(\lambda)$$

Proof. This directly follows from the parameter choice, Lemma 2.2, Proposition 5.2 and Theorem 5.1. □

5.5 Parameters

We conclude with an example parameter comparison. Take $n = \Theta(\lambda)$, $m = n \log q$, $\sigma' = \Theta(\sqrt{n})$, $\bar{m} = 2m$, and $k = \bar{m} \log q$. Choose $s = \Theta(1)$ sufficiently large for the leftover hash lemma. Following Corollary 5.1, set $\sigma = 2\sigma'$, $\sigma_f = 2\sigma'\sqrt{Q_{\mathsf{eval}}}$, where $Q_{\mathsf{eval}} = \mathsf{poly}(n)$ bounds the number of homomorphic evaluations, and choose $r = \Theta\big((\sigma + \sigma_f)\sqrt{k\lambda \log(k)}\big)$.

Let $T \coloneqq Cs(2\sqrt{m} + \sqrt{\lambda}) + 1$ be the factor appearing in Lemma 3.2. Let $L_{\max}$ bound the length of the branching programs to be evaluated, and let $\rho_{L_{\max}}$ be the corresponding noise bound from Lemma 5.3. We compare two ways of choosing the modulus q.

As discussed above, the standard way to obtain CPA^D or IND-CCA1^D security for an exact scheme is through statistical correctness. From Lemma 5.3, this approach requires $\rho_{L_{\max}} < D < q/(4T)$, and hence $q_{\text{correct}} = \Omega(T\rho_{L_{\max}})$.

Our result provides a second route. Say we decide to only guarantee correctness for fresh ciphertexts and allow evaluated ciphertexts to have noticeable decryption-failure probability. We may then set $D = \Theta(\sigma\sqrt{\bar{m}})$ independently of $\rho_{L_{\mathsf{max}}}$. Using Theorem 5.1, security only requires a parameter B such that

$$\rho_{L_{\mathsf{max}}} < B \quad \text{and} \quad B + D < q/6.$$

Taking $B = \Theta(\rho_{L_{\mathsf{max}}})$ therefore gives $q_{\text{secure}} = \Omega(\max(DT, \rho_{L_{\mathsf{max}}}))$. For the parameters above, $T\sigma\sqrt{\bar{m}} = \widetilde{\Theta}(n^{3/2})$, whereas $\rho_{L_{\max}} = \widetilde{\Theta}\left(n^3 Q_{\mathsf{eval}}\sqrt{L_{\max}}\right)$. Finally, the two choices become

$$q_{\text{correct}} = \widetilde{\Omega}\Big(n^{7/2} Q_{\mathsf{eval}}\sqrt{L_{\max}}\Big), \quad q_{\text{secure}} = \widetilde{\Omega}\Big(n^3 Q_{\mathsf{eval}}\sqrt{L_{\max}}\Big).$$

From a performance point of view, this saves a modest $T = \widetilde{\Theta}(\sqrt{n})$ factor in the modulus. However, the theoretical result is more striking. Using only polynomial noise flooding, we obtain IND-CCA1^D security for a nearly maximally incorrect FHE scheme: fresh ciphertexts decrypt correctly, but the scheme need not support even one level of correct evaluation.

Acknowledgments. We thank an anonymous CRYPTO reviewer for pointing out an error in an earlier version of the IND-CCA1 security proof, and Renaud Sirdey for helpful discussions.

References

1. Agrawal, S., Gentry, C., Halevi, S., Sahai, A.: Discrete gaussian leftover hash lemma over infinite domains. In: Sako, K., Sarkar, P. (eds.) Advances in Cryptology - ASIACRYPT 2013. pp. 97–116. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
2. Agrawal, S., Kirshanova, E., Stehlé, D., Yadav, A.: Practical, round-optimal lattice-based blind signatures. In: Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security. p. 39–53. CCS '22, Association for Computing Machinery, New York, NY, USA (2022). https://doi.org/10.1145/3548606.3560650

3. Akavia, A.: Learning noisy characters, multiplication codes, and cryptographic hardcore predicates (2008), http://hdl.handle.net/1721.1/43032
4. Al Badawi, A., Bates, J., Bergamaschi, F., Cousins, D.B., Erabelli, S., Genise, N., Halevi, S., Hunt, H., Kim, A., Lee, Y., Liu, Z., Micciancio, D., Quah, I., Polyakov, Y., R.V., S., Rohloff, K., Saylor, J., Suponitsky, D., Triplett, M., Vaikuntanathan, V., Zucca, V.: Openfhe: Open-source fully homomorphic encryption library. In: Proceedings of the 10th Workshop on Encrypted Computing & Applied Homomorphic Cryptography. p. 53–63. WAHC'22, Association for Computing Machinery, New York, NY, USA (2022). https://doi.org/10.1145/3560827.3563379
5. Alperin-Sheriff, J., Peikert, C.: Faster bootstrapping with polynomial error. In: Garay, J.A., Gennaro, R. (eds.) Advances in Cryptology – CRYPTO 2014. pp. 297–314. Springer Berlin Heidelberg, Berlin, Heidelberg (2014)
6. Banaszczyk, W.: New bounds in some transference theorems in the geometry of numbers. Mathematische Annalen **296**(1), 625–635 (1993)
7. Bourse, F., Del Pino, R., Minelli, M., Wee, H.: Fhe circuit privacy almost for free. In: Robshaw, M., Katz, J. (eds.) Advances in Cryptology – CRYPTO 2016. pp. 62–89. Springer Berlin Heidelberg, Berlin, Heidelberg (2016)
8. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (leveled) fully homomorphic encryption without bootstrapping. In: Proceedings of the 3rd Innovations in Theoretical Computer Science Conference. p. 309–325. ITCS '12, Association for Computing Machinery, New York, NY, USA (2012). https://doi.org/10.1145/2090236.2090262
9. Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) lwe. In: Proceedings of the 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science. p. 97–106. FOCS '11, IEEE Computer Society, USA (2011). https://doi.org/10.1109/FOCS.2011.12
10. Brakerski, Z., Vaikuntanathan, V.: Fully homomorphic encryption from ring-lwe and security for key dependent messages. In: Rogaway, P. (ed.) Advances in Cryptology – CRYPTO 2011. pp. 505–524. Springer Berlin Heidelberg, Berlin, Heidelberg (2011)
11. Brakerski, Z., Vaikuntanathan, V.: Lattice-based fhe as secure as pke. In: Proceedings of the 5th Conference on Innovations in Theoretical Computer Science. p. 1–12. ITCS '14, Association for Computing Machinery, New York, NY, USA (2014). https://doi.org/10.1145/2554797.2554799
12. Brzuska, C., Canard, S., Fontaine, C., Phan, D.H., Pointcheval, D., Renard, M., Sirdey, R.: Relations among new CCA security notions for approximate FHE. IACR Communications in Cryptology **2**(1) (2025). https://doi.org/10.62056/aee0iv7sf
13. Canetti, R., Raghuraman, S., Richelson, S., Vaikuntanathan, V.: Chosen-ciphertext secure fully homomorphic encryption. In: Fehr, S. (ed.) Public-Key Cryptography – PKC 2017. pp. 213–240. Springer Berlin Heidelberg, Berlin, Heidelberg (2017)
14. Chailloux, A., Hermouet, P.: On the quantum equivalence between $S|\mathrm{LWE}\rangle$ and ISIS. In: Heninger, N., Rosulek, M. (eds.) Advances in Cryptology – CRYPTO 2026. pp. 391–428. Springer Nature Switzerland, Cham (2026)
15. Checri, M., Sirdey, R., Boudguiga, A., Bultel, J.P.: On the practical cpa-d security of "exact" and threshold fhe schemes and libraries. In: Reyzin, L., Stebila, D. (eds.) Advances in Cryptology – CRYPTO 2024. pp. 3–33. Springer Nature Switzerland, Cham (2024)
16. Chen, Y., Hu, Z., Liu, Q., Luo, H., Tu, Y.: Lwe with quantum amplitudes: Algorithm, hardness, and oblivious sampling. In: Tauman Kalai, Y., Kamara, S.F.

(eds.) Advances in Cryptology – CRYPTO 2025. pp. 513–544. Springer Nature Switzerland, Cham (2025)
17. Cheon, J.H., Choe, H., Passelègue, A., Stehlé, D., Suvanto, E.: Attacks against the ind-cpad security of exact fhe schemes. In: Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security. p. 2505–2519. CCS '24, Association for Computing Machinery, New York, NY, USA (2024). https://doi.org/10.1145/3658644.3690341
18. Cheon, J.H., Kim, A., Kim, M., Song, Y.: Homomorphic encryption for arithmetic of approximate numbers. In: Takagi, T., Peyrin, T. (eds.) Advances in Cryptology – ASIACRYPT 2017. pp. 409–437. Springer International Publishing, Cham (2017)
19. Chillotti, I., Gama, N., Georgieva, M., Izabachène, M.: Tfhe: fast fully homomorphic encryption over the torus. Journal of Cryptology **33**(1), 34–91 (2020)
20. Cramer, R., Shoup, V.: A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In: Krawczyk, H. (ed.) Advances in Cryptology — CRYPTO '98. pp. 13–25. Springer Berlin Heidelberg, Berlin, Heidelberg (1998)
21. Cramer, R., Shoup, V.: Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In: Knudsen, L.R. (ed.) Advances in Cryptology — EUROCRYPT 2002. pp. 45–64. Springer Berlin Heidelberg, Berlin, Heidelberg (2002)
22. Damgård, I.: Towards practical public key systems secure against chosen ciphertext attacks. In: Feigenbaum, J. (ed.) Advances in Cryptology — CRYPTO '91. pp. 445–456. Springer Berlin Heidelberg, Berlin, Heidelberg (1992)
23. D'Anvers, J.P., Guo, Q., Johansson, T., Nilsson, A., Vercauteren, F., Verbauwhede, I.: Decryption failure attacks on ind-cca secure lattice-based schemes. In: Lin, D., Sako, K. (eds.) Public-Key Cryptography – PKC 2019. pp. 565–598. Springer International Publishing, Cham (2019)
24. Debris-Alazard, T., Fallahpour, P., Stehlé, D.: Quantum oblivious lwe sampling and insecurity of standard model lattice-based snarks. In: Proceedings of the 56th Annual ACM Symposium on Theory of Computing. p. 423–434. STOC 2024, Association for Computing Machinery, New York, NY, USA (2024). https://doi.org/10.1145/3618260.3649766
25. Dodis, Y., Reyzin, L., Smith, A.: Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In: Cachin, C., Camenisch, J.L. (eds.) Advances in Cryptology - EUROCRYPT 2004. pp. 523–540. Springer Berlin Heidelberg, Berlin, Heidelberg (2004)
26. Ducas, L., Micciancio, D.: Fhew: Bootstrapping homomorphic encryption in less than a second. In: Oswald, E., Fischlin, M. (eds.) Advances in Cryptology – EUROCRYPT 2015. pp. 617–640. Springer Berlin Heidelberg, Berlin, Heidelberg (2015)
27. Fujisaki, E., Okamoto, T.: Secure integration of asymmetric and symmetric encryption schemes. In: Wiener, M. (ed.) Advances in Cryptology — CRYPTO' 99. pp. 537–554. Springer Berlin Heidelberg, Berlin, Heidelberg (1999)
28. Genise, N., Micciancio, D., Peikert, C., Walter, M.: Improved discrete gaussian and subgaussian analysis for lattice cryptography. In: Kiayias, A., Kohlweiss, M., Wallden, P., Zikas, V. (eds.) Public-Key Cryptography - PKC 2020 - 23rd IACR International Conference on Practice and Theory of Public-Key Cryptography, Edinburgh, UK, May 4-7, 2020, Proceedings, Part I. Lecture Notes in Computer Science, vol. 12110, pp. 623–651. Springer (2020). https://doi.org/10.1007/978-3-030-45374-9_21
29. Gentry, C.: Fully homomorphic encryption using ideal lattices. In: Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing. p. 169–178.

STOC '09, Association for Computing Machinery, New York, NY, USA (2009). https://doi.org/10.1145/1536414.1536440
30. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing. p. 197–206. STOC '08, Association for Computing Machinery, New York, NY, USA (2008). https://doi.org/10.1145/1374376.1374407
31. Gentry, C., Sahai, A., Waters, B.: Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In: Canetti, R., Garay, J.A. (eds.) Advances in Cryptology – CRYPTO 2013. pp. 75–92. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
32. Gentry, C., Wichs, D.: Separating succinct non-interactive arguments from all falsifiable assumptions. In: Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing. p. 99–108. STOC '11, Association for Computing Machinery, New York, NY, USA (2011). https://doi.org/10.1145/1993636.1993651
33. Kim, D., Lee, D., Seo, J., Song, Y.: Toward practical lattice-based proof of knowledge from hint-mlwe. In: Handschuh, H., Lysyanskaya, A. (eds.) Advances in Cryptology – CRYPTO 2023. pp. 549–580. Springer Nature Switzerland, Cham (2023)
34. Koblitz, N., Menezes, A.: Another look at non-standard discrete log and diffie-hellman problems. Cryptology ePrint Archive, Paper 2007/442 (2007), https://eprint.iacr.org/2007/442
35. Lai, R.W.F., Swarnakar, M., Woo, I.K.Y.: Leaky LWE: Learning with errors with semi-adaptive secret- and error-leakage. IACR Communications in Cryptology **2**(3) (2025). https://doi.org/10.62056/ah89ksuc2
36. Li, B., Micciancio, D.: On the security of homomorphic encryption on approximate numbers. In: Canteaut, A., Standaert, F.X. (eds.) Advances in Cryptology – EUROCRYPT 2021. pp. 648–677. Springer International Publishing, Cham (2021)
37. Li, B., Micciancio, D., Schultz-Wu, M., Sorrell, J.: Securing approximate homomorphic encryption using differential privacy. In: Dodis, Y., Shrimpton, T. (eds.) Advances in Cryptology – CRYPTO 2022. pp. 560–589. Springer Nature Switzerland, Cham (2022)
38. Li, Z., Galbraith, S.D., Ma, C.: Preventing adaptive key recovery attacks on the gsw levelled homomorphic encryption scheme. In: Chen, L., Han, J. (eds.) Provable Security. pp. 373–383. Springer International Publishing, Cham (2016)
39. Libert, B.: Leveraging small message spaces for cca1 security in additively homomorphic and bgn-type encryption. In: Advances in Cryptology – EUROCRYPT 2025: 44th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Madrid, Spain, May 4–8, 2025, Proceedings, Part II. p. 34–63. Springer-Verlag, Berlin, Heidelberg (2025). https://doi.org/10.1007/978-3-031-91124-8_2
40. Lipmaa, H.: On the cca1-security of elgamal and damgård's elgamal. In: Lai, X., Yung, M., Lin, D. (eds.) Information Security and Cryptology. pp. 18–35. Springer Berlin Heidelberg, Berlin, Heidelberg (2011)
41. Loftus, J., May, A., Smart, N.P., Vercauteren, F.: On cca-secure somewhat homomorphic encryption. In: Miri, A., Vaudenay, S. (eds.) Selected Areas in Cryptography. pp. 55–72. Springer Berlin Heidelberg, Berlin, Heidelberg (2012)
42. Manulis, M., Nguyen, J.: Fully homomorphic encryption beyond ind-cca1 security: Integrity through verifiability. In: Advances in Cryptology – EUROCRYPT 2024: 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26–30, 2024, Proceedings, Part II.

p. 63–93. Springer-Verlag, Berlin, Heidelberg (2024). https://doi.org/10.1007/978-3-031-58723-8_3
43. Micciancio, D., Regev, O.: Worst-case to average-case reductions based on gaussian measures. In: 45th Annual IEEE Symposium on Foundations of Computer Science. pp. 372–381 (2004). https://doi.org/10.1109/FOCS.2004.72
44. Micciancio, D., Mol, P.: Pseudorandom knapsacks and the sample complexity of lwe search-to-decision reductions. In: Rogaway, P. (ed.) Advances in Cryptology – CRYPTO 2011. pp. 465–484. Springer Berlin Heidelberg, Berlin, Heidelberg (2011)
45. Micciancio, D., Peikert, C.: Trapdoors for lattices: Simpler, tighter, faster, smaller. In: Pointcheval, D., Johansson, T. (eds.) Advances in Cryptology – EUROCRYPT 2012. pp. 700–718. Springer Berlin Heidelberg, Berlin, Heidelberg (2012), full version available at Cryptology ePrint Archive, Paper 2011/501
46. Naor, M.: On cryptographic assumptions and challenges. In: Boneh, D. (ed.) Advances in Cryptology - CRYPTO 2003. pp. 96–109. Springer Berlin Heidelberg, Berlin, Heidelberg (2003)
47. Ogilvie, T.: Ind-cpa-d and kr-d security with reduced noise from the hintlwe problem. In: Hanaoka, G., Yang, B.Y. (eds.) Advances in Cryptology – ASIACRYPT 2025. pp. 263–295. Springer Nature Singapore, Singapore (2026)
48. Phan, D.H., Sirdey, R., Vacher, J.: On the CCA security properties (and more) of a new variant of paillier-ElGamal. Cryptology ePrint Archive, Paper 2026/452 (2026), https://eprint.iacr.org/2026/452
49. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. J. ACM **56**(6) (Sep 2009). https://doi.org/10.1145/1568318.1568324
50. Wu, J., Stinson, D.R.: On the security of the ElGamal encryption scheme and damgard's variant. Cryptology ePrint Archive, Paper 2008/200 (2008), https://eprint.iacr.org/2008/200
51. Yang, R., Yu, Z., Susilo, W.: Fully homomorphic encryption with chosen-ciphertext security from lwe. In: Advances in Cryptology – CRYPTO 2025: 45th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17–21, 2025, Proceedings, Part III. p. 371–405. Springer-Verlag, Berlin, Heidelberg (2025). https://doi.org/10.1007/978-3-032-01881-6_12

A Standard Security Notions

We recall the definitions of IND-CPA^D and IND-CCA1 security.

Definition A.1 (IND-CPA^D**).** *Let $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ be a public-key HE scheme. Consider the* IND-CPA^D *experiment given in Figure 6. The advantage of an adversary $\mathcal{A}$ is given by:*

$$\mathsf{Adv}^{\text{IND-CPA}^{\text{D}}}_{\Pi,\mathcal{A}}(\lambda) = |\Pr[\mathrm{Exp}^{\text{IND-CPA}^{D}}_{\Pi,\mathcal{A}}(\lambda) = 1] - \frac{1}{2}|.$$

We say that Π is indistinguishable against chosen-plaintext attacks with passive decryption (IND-CPA^D) *if for all* PPT *adversaries $\mathcal{A}$, $\mathsf{Adv}^{\text{IND-CPA}^{\text{D}}}_{\Pi,\mathcal{A}}(\lambda)$ is negligible in λ.*

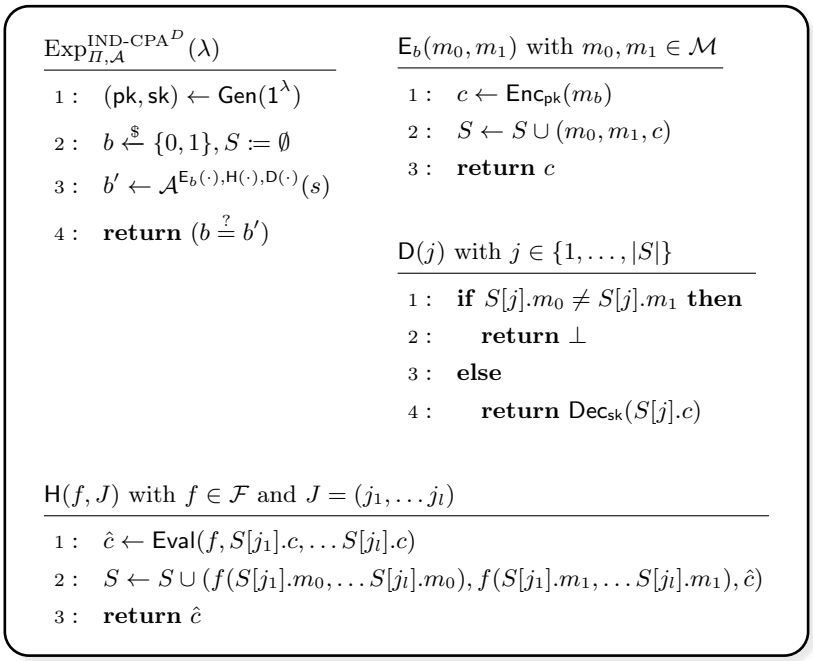

Fig. 6. The IND-CPA^D experiment $\mathrm{Exp}^{\text{IND-CPA}^D}_{\Pi,\mathcal{A}}(\lambda)$.

Definition A.2 (IND-CCA1)**.** *Let $\Pi = (\mathsf{Gen},\mathsf{Enc},\mathsf{Dec})$ be a public-key encryption scheme. Consider the* IND-CCA1 *experiment given in Figure 7. The advantage of an adversary $\mathcal{A}$ is given by:*

$$\mathsf{Adv}^{\text{IND-CCA1}}_{\Pi,\mathcal{A}}(\lambda) = |\Pr[\mathrm{Exp}^{\text{IND-CCA1}}_{\Pi,\mathcal{A}}(\lambda) = 1] - \frac{1}{2}|.$$

We say that Π is indistinguishable against non-adaptive chosen-ciphertext attacks (IND-CCA1) *if for all* PPT *adversaries $\mathcal{A}$, $\mathsf{Adv}^{\text{IND-CCA1}}_{\Pi,\mathcal{A}}(\lambda)$ is negligible in λ.*

B More lattice problems

Definition B.1 (One-more ISIS [2]). *Let n,m,q be positive integers and let $\sigma,\beta > 0$. The $\mathsf{OM\text{-}ISIS}_{n,m,q,\sigma,\beta}$ experiment is given in Figure 8. For an adversary $\mathcal{A}$, define its advantage as*

$$\mathsf{Adv}^{\mathsf{OM\text{-}ISIS}}_{\mathsf{pp},\mathcal{A}}(\lambda) := \Pr\left[\mathrm{Exp}^{\mathsf{OM\text{-}ISIS}}_{\mathsf{pp},\mathcal{A}}(\lambda) = 1\right]$$

The one-more-ISIS$_{n,m,q,\sigma,\beta}$ assumption states that for any PPT *adversary $\mathcal{A}$, $\mathsf{Adv}^{\mathsf{OM\text{-}ISIS}}_{\mathsf{pp},\mathcal{A}}(\lambda)$ is negligible.*

$\underline{\mathrm{Exp}^{\text{IND-CCA1}}_{\Pi,\mathcal{A}}(\lambda)}$

1 : $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^\lambda)$
2 : $(m_0, m_1, s) \leftarrow \mathcal{A}^{\mathsf{Dec}_{\mathsf{sk}}(\cdot)}(\mathsf{pk})$
3 : $b \xleftarrow{\$} \{0,1\}$
4 : $c^* \leftarrow \mathsf{Enc}(\mathsf{pk}, m_b)$
5 : $b' \leftarrow \mathcal{A}(s, c^*)$
6 : **return** $(b \stackrel{?}{=} b')$

Fig. 7. The IND-CCA1 experiment $\mathrm{Exp}^{\text{IND-CCA1}}_{\Pi,\mathcal{A}}(\lambda)$.

$\underline{\mathrm{Exp}^{\mathsf{OM\text{-}ISIS}}_{n,m,q,\sigma,\beta}(\lambda)}$

1 : $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$
2 : $S \coloneqq \emptyset$
3 : $k \coloneqq 0$
4 : $\{(\mathbf{y}_j, \mathbf{t}_j)\}_{j\in l} \leftarrow \mathcal{A}^{\mathcal{O}_{\mathsf{syn}}(), \mathcal{O}_{\mathsf{pre}}(\cdot)}(\mathbf{A})$
5 : **if** $l = k+1$ **then**
6 : **return** $\bigwedge_{j\in l} (\mathbf{A}\mathbf{y}_j \bmod q \in S \wedge \|\mathbf{y}_j\|_2 \le \beta)$
7 : **return** 0

$\underline{\mathcal{O}_{\mathsf{syn}}()}$

1 : $\mathbf{t} \xleftarrow{\$} \mathbb{Z}_q^n$
2 : $S \leftarrow S \cup \{\mathbf{t}\}$
3 : **return t**

$\underline{\mathcal{O}_{\mathsf{pre}}(\mathbf{t}')}$

1 : $\mathbf{y}' \xleftarrow{\$} \mathcal{D}_{\Lambda^\perp_{\mathbf{t}'}(\mathbf{A}),\sigma}$
2 : $k \leftarrow k+1$
3 : **return** $\mathbf{y}'$

Fig. 8. The OM-ISIS experiment.

Definition B.2 (LK-ϵ)**.** *Let $\epsilon \in (0, 1/2)$ be a fixed constant. Let $\mathcal{G}$ denote an algorithm that on input 1^λ outputs a basis $\mathbf{B}$ of an $m = m(\lambda)$ dimensional lattice $\mathcal{L}(\mathbf{B})$. Let $\mathcal{C}$ be an algorithm that takes $\mathbf{B}$ as input and returns nothing. Let $\mathcal{C}^*$ denote an extraction algorithm which takes as input a vector $\mathbf{c} \in \mathbb{R}^m$ and some state, and returns a vector $\mathbf{p} \in \mathbb{R}^m$ and updated state. The* LK-ϵ *experiment is*

$\mathrm{Exp}^{\text{LK-}\epsilon}_{\mathcal{G},\mathcal{C},\mathcal{C}^*}(\lambda)$:

1 : $\mathbf{B} \leftarrow \mathcal{G}(1^\lambda)$
2 : Choose random tape $r$for $\mathcal{C}$
3 : $\Lambda \coloneqq \mathcal{L}(\mathbf{B}); \gamma \coloneqq \epsilon\lambda_\infty(\Lambda)$
4 : $\mathsf{st} \coloneqq (\mathbf{B}, r)$
5 : $\mathsf{bad} \coloneqq 0$
6 : Run $\mathcal{C}^{\mathcal{O}(\cdot)}(\mathbf{B}; r)$ until it halts.
7 : **return** bad

$\mathcal{O}(\mathbf{c})$:

1 : $(\mathbf{p}, \mathsf{st}) \leftarrow \mathcal{C}^*(\mathbf{c}, \mathsf{st})$
2 : $\mathsf{close} \coloneqq \left[d_\infty(\mathbf{c}, \Lambda) \le \gamma\right]$
3 : $\mathsf{valid} \coloneqq \left[\mathbf{p} \in \Lambda \wedge \|\mathbf{p} - \mathbf{c}\|_\infty \le \gamma\right]$
4 : **if** close $\wedge \neg$valid **then**
5 : bad $\leftarrow 1$
6 : **return p**

Fig. 9. The LK-ϵ experiment.

defined in Figure 9 [1]*. Define the advantage*

$$\mathsf{Adv}^{\text{LK-}\epsilon}_{\mathcal{G},\mathcal{C},\mathcal{C}^*}(\lambda) \coloneqq \Pr[\mathrm{Exp}^{\text{LK-}\epsilon}_{\mathcal{G},\mathcal{C},\mathcal{C}^*} = 1].$$

We say that $\mathcal{G}$ satisfies the LK-ϵ *assumption if for every* PPT $\mathcal{C}$*, there exists a* PPT *extractor $\mathcal{C}^*$ such that $\mathsf{Adv}^{\text{LK-}\epsilon}_{\mathcal{G},\mathcal{C},\mathcal{C}^*}(\lambda)$ is negligible.*

C Proof of Lemma 3.1

Lemma 3.1. *Let n, m, q be positive integers. For $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m\times n}$, let $\mathcal{L}_\mathbf{A} \coloneqq \mathbf{A}\mathbb{Z}^n + q\mathbb{Z}^m$. For every r satisfying $1 \le r < q$, we have*

$$\Pr_\mathbf{A}[\lambda_\infty(\mathcal{L}_\mathbf{A}) < r] \le q^n \left(\frac{3r}{q}\right)^m.$$

Moreover, suppose that q is a power of two. Then, for every unit vector $\mathbf{u} \in \mathbb{Z}^m$,

$$\Pr_\mathbf{A}\left[d_\infty\left(\frac{q}{2}\mathbf{u}, \mathcal{L}_\mathbf{A}\right) < \frac{q}{6}\right] \le q^n 2^{-m}.$$

Proof. Fix $\mathbf{s} \in \mathbb{Z}_q^n \setminus \{\mathbf{0}\}$, set $d \coloneqq \gcd(s_1, \ldots, s_n, q)$ and let $\mathbb{G} = d\mathbb{Z}_q$. As shown in the proof of [45, Lemma 2.4, full version], the vector $\mathbf{As}$ is uniform over $\mathbb{G}^m$. If $d \ge r$, then $\mathbb{G}^m$ contains no nonzero vectors with norm less than r. Otherwise, there are at most $(3r/d)^m$ vectors $\mathbf{v} \in (d\mathbb{Z}_q)^m$ with $\|\mathbf{v}\|_\infty < r$, since $2\left\lceil \frac{r}{d} \right\rceil - 1 \le \frac{3r}{d}$. Therefore,

$$\Pr_\mathbf{A}\left[0 < \|\mathbf{As}\|_\infty < r\right] \le \frac{(3r/d)^m}{(q/d)^m} = \left(\frac{3r}{q}\right)^m.$$

[1] The close condition is stated in the discussion following Figure 2 of [41], but is omitted from the displayed experiment itself. It is necessary for the assumption to not be vacuous.

Finally, since $r < q$, any nonzero vector in $\mathcal{L}_{\mathbf{A}}$ that has norm less than r corresponds to a nonzero $\mathbf{s}$. Taking a union bound over the nonzero $\mathbf{s} \in \mathbb{Z}_q^n$ proves the first claim.

For the second claim, suppose that q is a power of two. Since $\mathbf{s} \neq 0$, we have $d < q$. Since q is a power of two, it follows that $d \mid q/2$ and $d \leq q/2$. Hence, $(q/2)\mathbf{u} \in \mathbb{G}^m$, so $\mathbf{As} - (q/2)\mathbf{u}$ is also uniform over $\mathbb{G}^m$. Taking $r = q/6$, we have $r/d \geq 1/3$, so $2\lceil r/d \rceil - 1 \leq 3r/d$ again. Since $r = q/6$, a vector in $\mathcal{L}_{\mathbf{A}}$ that has distance to $(q/2)\mathbf{u}$ less than r must correspond to a nonzero $\mathbf{s}$. We may therefore apply the same argument and obtain the bound $q^n(3r/q)^m = q^n 2^{-m}$. □

D Hardness of DT-LWE

D.1 A search to decision reduction

Definition D.1 (Search Delayed-target LWE). *Let n, m, q be positive integers and $\beta > 0$. Let χ_s, χ_e be probability distributions over $\mathbb{Z}^n$ and $\mathbb{Z}^m$ respectively and set $\mathsf{pp} = (n, m, q, \chi_s, \chi_e, \beta)$. We define the search delayed-target LWE (*$\mathsf{s\text{-}DT\text{-}LWE}$*) experiment in Figure 10. For an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, define its advantage as*

$$\mathsf{Adv}^{\mathsf{s\text{-}DT\text{-}LWE}}_{\mathsf{pp},\mathcal{A}}(\lambda) \coloneqq \Pr\left[\mathrm{Exp}^{\mathsf{s\text{-}DT\text{-}LWE}}_{\mathsf{pp},\mathcal{A}}(\lambda) = 1\right].$$

The search delayed-target LWE assumption states that for any PPT *adversary $\mathcal{A}$, $\mathsf{Adv}^{\mathsf{s\text{-}DT\text{-}LWE}}_{\mathsf{pp},\mathcal{A}}(\lambda)$ is negligible.*

$\mathrm{Exp}^{\mathsf{s\text{-}DT\text{-}LWE}}_{\mathsf{pp},\mathcal{A}}(\lambda)$

1 : $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m \times n}$
2 : $\mathsf{st} \leftarrow \mathcal{A}_1^{\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{A},\cdot)}(\mathbf{A})$
3 : $\mathbf{s} \leftarrow \chi_s, \mathbf{e} \leftarrow \chi_e$
4 : $\mathbf{b} \coloneqq \mathbf{As} + \mathbf{e} \bmod q$
5 : $\mathbf{s}' \leftarrow \mathcal{A}_2(\mathsf{st}, \mathbf{b})$
6 : **return** $(\mathbf{s} \stackrel{?}{=} \mathbf{s}')$

$\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{A}, \mathbf{b})$:

1 : **if** $\exists\, (\mathbf{s}, \mathbf{e}) \in \mathbb{Z}_q^n \times \mathbb{Z}^m$ such that
$\mathbf{b} = \mathbf{As} + \mathbf{e} \bmod q \ \wedge\ \|\mathbf{e}\|_2 < \beta$ **then**
2 : **return** $\mathbf{s}$
3 : **return** $\perp$

Fig. 10. The search delayed-target LWE experiment. The parameter β is chosen such that the tuple $(\mathbf{s}, \mathbf{e})$ is unique if it exists.

Proposition 4.1 (Search to decision). *Let $\mathsf{pp} = (n, m, q, \sigma, \beta)$ be parameters for the* $\mathsf{DT\text{-}LWE}$ *problem. Suppose that $m \geq n + \omega(\log n)$, that $q = \mathsf{poly}(n)$, and that one of the following conditions hold:*

1. *The modulus q is prime.*

2. *The modulus* $q = p^w$ *is a prime power for prime* $p = \mathsf{poly}(n)$ *and* $\sigma = o(p/\sqrt{n})$.

Then, for every PPT *adversary* $\mathcal{A}$ *against* $\mathsf{DT\text{-}LWE}_{\mathsf{pp}}$ *with non-negligible advantage, there exists a* PPT *adversary* $\mathcal{B}$ *against* $\mathsf{s\text{-}DT\text{-}LWE}_{\mathsf{pp}}$ *with non-negligible advantage.*

The following two lemmas shows that the search to decision reduction by Micciancio and Mol [44] is compatible with the delayed-target setting. The first lemma ensures 1) that the matrices $\mathbf{A}_i$ derived from the search-LWE challenge matrix $\mathbf{A}$ on which the decision-LWE adversary is run are chosen independently of the target $\mathbf{b}$; 2) that these matrices have a compatible form. The second lemma shows that given oracle access to $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$ with respect to $\mathbf{A}$, it is possible to efficiently simulate the $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$ oracle with respect to $\mathbf{A}_i$.

Lemma D.1. *Let* n, m, q, σ *be as in Proposition 4.1, and let* $(\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e})$ *be a search-*$\mathsf{LWE}_{n,m,q,\sigma}$ *challenge. For every* PPT $\mathsf{LWE}_{n,m,q,\sigma}$ *distinguisher* $\mathcal{D}$*, there exists a* PPT *algorithm* $R = (R_1, R_2)$ *solving search-*LWE *such that*

$$(\mathbf{A}_1, \ldots, \mathbf{A}_L, \mathbf{v}_1, \ldots, \mathbf{v}_L, \mathsf{st}) \leftarrow R_1(\mathbf{A})$$
$$\mathbf{s}' \leftarrow R_2^{\mathcal{D}(\mathbf{A}_1,\cdot),\ldots,\mathcal{D}(\mathbf{A}_L,\cdot)}(\mathsf{st}, \mathbf{b}).$$

Moreover,

- *the tuple* $(\mathbf{A}_1, \ldots, \mathbf{A}_L, \mathbf{v}_1, \ldots, \mathbf{v}_L, \mathsf{st})$ *is independent of the target* $\mathbf{b}$*, and for every* $i \in [L]$ *satisfies*

$$\mathrm{im}(\mathbf{A}_i) \subseteq \mathrm{im}(\mathbf{A}) + \mathrm{span}(\mathbf{v}_i) \quad \textit{where } \mathbf{v}_i \in \mathbb{Z}_q^m;$$

- *if* $\mathcal{D}$ *has non-negligible advantage,* $\mathbf{s} = \mathbf{s}'$ *with non-negligible probability.*

Proof. Let R be the search-to-decision reduction from [44, Proposition 3]. We show how to decompose its execution into the two algorithms R_1, R_2.

At a high level, R first turns an LWE instance $(\mathbf{A}, \mathbf{b}) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$ into a knapsack instance

$$(\mathbf{H}, \mathbf{y}) \in \mathbb{Z}_q^{(m-n)\times m} \times \mathbb{Z}_q^{m-n}$$

satisfying $\mathbf{H}\mathbb{Z}_q^m = \mathbb{Z}_q^{m-n}$ and $\ker(\mathbf{H}) = \mathrm{im}(\mathbf{A})$, see [44, Lemma 9]. Then, through queries from the significant Fourier transform (SFT) algorithm, R obtains $\mathbf{r}_1, \ldots, \mathbf{r}_L \in \mathbb{Z}_l^m$ and defines instances

$$\mathbf{H}_i = \mathbf{H} - \mathbf{u}_i \mathbf{r}_i^{\intercal}$$

with $\mathbf{u}_i \in \mathbb{Z}_q^{m-n}$, see[44, Lemma 5, Proposition 2]. Finally, each knapsack matrix $\mathbf{H}_i$ is transformed into an LWE matrix $\mathbf{A}_i$ such that $\ker(\mathbf{H}_i) = \mathrm{im}(\mathbf{A}_i)$), see [44, Lemma 10].

The SFT algorithm is non adaptive. It generates the tuple $(\mathbf{r}_1, \ldots, \mathbf{r}_L)$ solely based on its own randomness [3, Remark following Theorem 3.3]. Since $\mathbf{H}\mathbb{Z}_q^m = \mathbb{Z}_q^{m-n}$, one can efficiently compute $\mathbf{v}_i$ such that $\mathbf{H}\mathbf{v}_i = \mathbf{u}_i$. It is therefore possible

to generate the matrix $\mathbf{H}$, the vectors $\mathbf{r}_i, \mathbf{u}_i$ and hence the LWE matrix $\mathbf{A}_i$ and the vector $\mathbf{v}_i$ before receiving $\mathbf{b}$. This defines R_1. The second stage R_2 simply completes R.

Furthermore, recall that $\mathrm{im}(\mathbf{A}_i) = \ker(\mathbf{H}_i)$. For any $\mathbf{x} \in \mathrm{im}(\mathbf{A}_i)$, we have

$$\begin{aligned}\mathbf{0} &= \mathbf{H}_i\mathbf{x} \\ &= (\mathbf{H} - \mathbf{u}\mathbf{r}_i^\intercal)\mathbf{x} \\ &= \mathbf{H}(\mathbf{x} - \mathbf{v}_i\mathbf{r}_i^\intercal\mathbf{x}),\end{aligned}$$

and $\mathbf{x} - \mathbf{v}_i\mathbf{r}_i^\intercal\mathbf{x} \in \ker(\mathbf{H})$. Since $\ker(\mathbf{H}) = \mathrm{im}(\mathbf{A})$,

$$\mathbf{x} \in \mathrm{im}(\mathbf{A}) + \mathbf{v}_i(\mathbf{r}_i^\intercal\mathbf{x})$$

which proves the first claim. The second claim follows from [44, Proposition 3].

Lemma D.2. *Let $\mathsf{pp} = (n, m, q, \sigma, \beta)$ be as in Proposition 4.1. Let $\mathbf{A}, \mathbf{A}' \in \mathbb{Z}_q^{m\times n}$ such that*

$$\mathrm{im}(\mathbf{A}') \subseteq \mathrm{im}(\mathbf{A}) + \mathrm{span}(\mathbf{v}) \quad \textit{for some known } \mathbf{v} \in \mathbb{Z}_q^m.$$

Then, a query to the $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$ oracle associated with $\mathbf{A}'$ can efficiently be simulated using at most q queries to the $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$ oracle associated with $\mathbf{A}$.

Proof. Consider the oracle simulator in Figure 11. Suppose that $\mathbf{b} = \mathbf{A}'\mathbf{s}' + \mathbf{e}'$ with $\|\mathbf{e}'\|_2 < \beta$. Then, by hypothesis there exist $\mathbf{s} \in \mathbb{Z}_q^n$ and $z \in \mathbb{Z}_q$ such that

$$\mathbf{A}'\mathbf{s}' = \mathbf{A}\mathbf{s} + z\mathbf{v}.$$

It follows that $\mathbf{b} - z\mathbf{v} = \mathbf{A}\mathbf{s} + \mathbf{e}'$, so $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$ returns $\mathbf{s}$. Since $\mathbf{b} - \mathbf{e}' \in \mathrm{im}(\mathbf{A}')$, $\mathcal{O}_{\mathsf{sim}}^{\mathbf{A},\mathbf{A}',\mathbf{v}}(\mathbf{b})$ returns $\mathbf{s}'$.

Conversely, suppose that $\mathcal{O}_{\mathsf{sim}}^{\mathbf{A},\mathbf{A}',\mathbf{v}}(\mathbf{b})$ returns $\mathbf{s}'$ for some $\mathbf{b} \in \mathbb{Z}_q^m$. Then $\mathbf{b} = \mathbf{A}'\mathbf{s}' + \mathbf{e}' \bmod q$ and $\|\mathbf{e}'\|_2 < \beta$. Hence, $\mathcal{O}_{\mathsf{sim}}^{\mathbf{A},\mathbf{A}',\mathbf{v}}(\mathbf{b})$ perfectly simulates a query to the $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$ oracle associated with $\mathbf{A}'$ in at most q queries. □

Based on the two previous lemmas, we can now give a short proof of Proposition 4.1.

Proof. Let $R = (R_1, R_2)$ be the algorithm from Lemma D.1, and let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a $\mathsf{DT\text{-}LWE}_{\mathsf{pp}}$ adversary. We give a $\mathsf{s\text{-}DT\text{-}LWE}_{\mathsf{pp}}$ algorithm $\mathcal{B} = (\mathcal{B}_1, \mathcal{B}_2)$ that uses R and $\mathcal{A}$ as a subroutine.

The first stage $\mathcal{B}_1^{\mathcal{O}_{\mathsf{DT\text{-}LWE}}}(\mathbf{A})$. On input $\mathbf{A} \in \mathbb{Z}_q^{m\times n}$ and given access to the $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$ oracle associated with $\mathbf{A}$:

1. Run $((\mathbf{A}_1, \ldots, \mathbf{A}_L, \mathbf{v}_1, \ldots, \mathbf{v}_L, \mathsf{st}_R)) \leftarrow R_1(\mathbf{A})$.
2. For $i \in [L]$, run $\mathsf{st}_i \leftarrow \mathcal{A}_1^{\mathcal{O}_{\mathsf{sim}}^{\mathbf{A},\mathbf{A}',\mathbf{v}}(\cdot)}(\mathbf{A}_i)$, where $\mathcal{O}_{\mathsf{sim}}^{\mathbf{A},\mathbf{A}',\mathbf{v}}(\mathbf{b})$ simulates the $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$ oracle associated to $\mathbf{A}_i$ as described in Lemma D.2.
3. Output $\mathsf{st} = (\mathsf{st}_R, \{\mathsf{st}_i, \mathbf{A}_i\}_{i\in[L]})$

$\mathcal{O}_{\mathsf{sim}}^{\mathbf{A},\mathbf{A}',\mathbf{v}}(\mathbf{b})$:

for $z \in \mathbb{Z}_q$ **do**
  $\mathbf{s} \leftarrow \mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{b} - z\mathbf{v})$
  if $\mathbf{s} \neq \perp$ **then**
    $\mathbf{e}' \coloneqq \mathbf{b} - z\mathbf{v} - \mathbf{A}\mathbf{s} \bmod q$
    if $\|\mathbf{e}'\|_2 < \beta \wedge (\mathbf{b} - \mathbf{e}') \in \mathrm{im}(\mathbf{A}')$ **then**
      Compute $\mathbf{s}'$such that $\mathbf{A}'\mathbf{s}' = \mathbf{b} - \mathbf{e}' \bmod q$
      return $\mathbf{s}'$
return $\perp$

Fig. 11. The oracle simulates $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$ associated with $\mathbf{A}'$ given oracle access to $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$ associated with $\mathbf{A}'$.

The second stage $\mathcal{B}_2(\mathsf{st}, \mathbf{b})$. On input st and $\mathbf{b} \in \mathbb{Z}_q^m$:

1. For $i \in [L]$, set $\mathcal{D}(\mathbf{A}_i, \cdot) = \mathcal{A}_2(\mathsf{st}_i, \cdot)$.
2. Run $\mathbf{s}' \leftarrow R_2^{\mathcal{D}(\mathbf{A}_1,\cdot),\ldots,\mathcal{D}(\mathbf{A}_L,\cdot)}(\mathsf{st}_R, \mathbf{b})$.
3. Return $\mathbf{s}'$.

By *Lemma D.*2, $\mathcal{B}_1$ perfectly simulates the oracle queries to $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$ associated to $\mathbf{A}_i$, and the st_i are well distributed. Therefore, $\mathcal{D}(\mathbf{A}_i, \cdot)$ behaves as the LWE adversary expected by R. From Lemma D.1, if $\mathcal{A}$ has non-negligible advantage against $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$, then $\mathcal{B}$ has non-negligible advantage against $\mathsf{s\text{-}DT\text{-}LWE}$. □

D.2 A quantum reduction from OM-ISIS

We start by slightly strengthening the guarantees of [14, Proposition 17] by using a larger ISIS norm bound. This will be required by our amplification argument in the main proof. For a fixed matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, we write search-$\mathsf{LWE}_\sigma(\mathbf{A}^\intercal)$ and $\mathsf{ISIS}_\beta(\mathbf{A})$ for the corresponding problems with public matrix fixed to $\mathbf{A}^\intercal$ and $\mathbf{A}$, respectively.

Lemma D.3. *Let n, m, q be positive integers, let $\beta, \sigma > 0$ and let $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$. Suppose that*

$$\min(\sigma, \frac{q}{\sigma}) = \omega(\sqrt{\log \lambda}) \quad \textit{and} \quad \beta \geq \sqrt{m}\frac{q}{2\sigma}.$$

Given an efficient search-$\mathsf{LWE}_\sigma(\mathbf{A}^\intercal)$ algorithm with non-negligible success probability ϵ_{LWE}, there exists an efficient quantum algorithm R for $\mathsf{ISIS}_\beta(\mathbf{A})$ that succeeds with probability at least

$$\epsilon_{\mathsf{ISIS}} \coloneqq \epsilon_{\mathsf{LWE}} - \mathsf{negl}(\lambda). \tag{9}$$

Proof. From [14, Theorem 1 and Proposition 16], there exists a quantum algorithm R with success probability at least

$$\epsilon_{\mathsf{ISIS}} = \epsilon_{\mathsf{LWE}}\gamma - 2\sqrt{\epsilon_{\mathsf{LWE}}(1-\epsilon_{\mathsf{LWE}})(1-\gamma)},$$

where, since $\min(\sigma, \frac{q}{\sigma}) = \omega(\sqrt{\log\lambda})$, $\gamma = \Pr[\|\mathcal{D}_{\mathbb{Z}^m,r}\|_2 \leq \beta] + \mathsf{negl}(\lambda)$ with $r = q/2\sigma$. It follows from Lemma 2.5 that $\gamma = 1 - \mathsf{negl}(\lambda)$. Substituting γ in the above equation gives the claimed bound. □

Proposition 4.3. *Let* $\mathsf{pp} = (n, m, q, \sigma_{\mathsf{OM}}, \beta_{\mathsf{OM}})$ *and* $\mathsf{pp}' = (n, m, q, \sigma_{\mathsf{DT}}, \beta_{\mathsf{DT}})$ *be parameters for the* OM-ISIS *and* DT-LWE *problems respectively. Suppose that* $q \geq 3$ *and*

$$m \geq n\log q, \qquad \sigma_{\mathsf{OM}} = \Omega(\sqrt{m}), \qquad \min(\sigma_{\mathsf{DT}}, \frac{q}{\sigma_{\mathsf{DT}}}) = \omega(\sqrt{\log\lambda})$$

$$\beta_{\mathsf{OM}} \geq \sqrt{m}\frac{q}{2\sigma_{\mathsf{DT}}} \quad \textit{and} \quad \beta_{\mathsf{DT}} \leq \frac{q}{c\cdot\sigma_{\mathsf{OM}}(\sqrt{\kappa}+\sqrt{m}+\sqrt{\lambda})},$$

where $\kappa = n\lceil\log q\rceil$ *and* $c > 0$ *is a universal constant. If there exists an efficient, possibly quantum, algorithm for* $\mathsf{s\text{-}DT\text{-}LWE}_{\mathsf{pp}'}$ *that succeeds with non-negligible probability* ϵ_{DT}*, then there exists an efficient quantum algorithm for* $\mathsf{OM\text{-}ISIS}_{\mathsf{pp}}$ *that succeeds with probability* $\epsilon_{\mathsf{OM}} = \Omega(\epsilon_{\mathsf{DT}})$*.*

Proof. Let R be the quantum algorithm from Lemma D.3 and let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be an efficient (possibly quantum) adversary against the s-DT-LWE problem that succeeds with probability ϵ_{DT}. We give a quantum reduction $\mathcal{B}$ against the OM-ISIS problem that makes $\kappa = n\lceil\log q\rceil$ classical queries to the preimage oracle $\mathcal{O}_{\mathsf{pre}}$ and at most $(\kappa+1)L$ calls to the quantum reduction R, where $L = L(\lambda)$ is polynomially bounded.

We start by showing how to implement a search-LWE adversary given a s-DT-LWE adversary and advice $\mathbf{R} \in \mathbb{Z}^{m\times\kappa}$ such that $\mathbf{AR} = \mathbf{G} \bmod q$, where $\mathbf{G}$ is the gadget matrix.

*The search-*LWE *adversary* $\mathcal{A}'_{\mathbf{R}}(\mathbf{A}^\intercal, \mathbf{b}^*)$. On input $(\mathbf{A}^\intercal, \mathbf{b}^*) \in \mathbb{Z}_q^{m\times n} \times \mathbb{Z}_q^m$:

1. Run $\mathsf{st} \leftarrow \mathcal{A}_1^{\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\mathbf{A}^\intercal,\cdot)}(\mathbf{A}^\intercal)$. Answer queries $\mathbf{b} \in \mathbb{Z}_q^m$ to $\mathcal{O}_{\mathsf{DT\text{-}LWE}}(\cdot)$ with:
   - Compute $(\mathbf{s}, \mathbf{e}) \leftarrow \mathsf{Invert}(\mathbf{R}^\intercal, \mathbf{A}^\intercal, \mathbf{b})$ such that $\mathbf{b} = \mathbf{A}^\intercal\mathbf{s} + \mathbf{e} \bmod q$. If this step fails, return $\perp$ to $\mathcal{A}_1$.
   - If $\|\mathbf{e}\|_2 \geq \beta_{\mathsf{DT}}$, return $\perp$.
   - Return $\mathbf{s}$.
2. Run $\mathbf{s} \leftarrow \mathcal{A}_2(\mathsf{st}, \mathbf{b}^*)$.
3. Return $\mathbf{s}$.

We write $R[\mathcal{A}'_{\mathbf{R}}](\cdot,\cdot)$ to denote the instantiation of R with the search-LWE adversary $\mathcal{A}'_{\mathbf{R}}$.

The OM-ISIS *adversary* $\mathcal{B}(\mathbf{A})$. On input a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ from the OM-ISIS challenger:

1. For $i \in [\kappa]$, let $\mathbf{g}_i$ be the i-th column of the gadget matrix $\mathbf{G} \in \mathbb{Z}_q^{n \times \kappa}$. Query $\mathbf{g}_i$ to the preimage oracle $\mathcal{O}_{\mathsf{pre}}$ to obtain $\mathbf{r}_i \in \mathbb{Z}^m$ such that $\mathbf{A}\mathbf{r}_i = \mathbf{g}_i \bmod q$. Set
$$\mathbf{R} \coloneqq [\mathbf{r}_1 | \mathbf{r}_2 | \dots | \mathbf{r}_k].$$
2. For $i \in [\kappa + 1]$, do:
   (a) For $j \in [L]$:
   - Query the syndrome oracle $\mathcal{O}_{\mathsf{syn}}$ to obtain $\mathbf{t}'$.
   - Run the quantum reduction $\mathbf{y}' \leftarrow R[\mathcal{A}'_{\mathbf{R}}](\mathbf{A}, \mathbf{t}')$.
   - If $\mathbf{t}' = \mathbf{A}\mathbf{y}' \bmod q$ and $\|\mathbf{y}'\|_2 \leq \beta_{\mathsf{OM}}$, set $(\mathbf{y}_i, \mathbf{t}_i) \coloneqq (\mathbf{y}', \mathbf{t}')$ and exit the loop.

   (b) If after L iterations no valid pair was found, return a failure symbol $\perp$.
3. Return $\{(\mathbf{y}_i, \mathbf{t}_i)\}_{i \in [\kappa+1]}$.

Recall that the columns $\mathbf{r}_i$ of $\mathbf{R}$ are sampled from $\mathcal{D}_{\Lambda^{\perp}_{\mathbf{g}_i}(\mathbf{A}), \sigma_{\mathsf{OM}}}$. Let C be the constant from Lemma 2.4 and let bad denote the event

$$s_1(\mathbf{R}) > C\sigma_{\mathsf{OM}}(\sqrt{\kappa} + \sqrt{m} + \sqrt{\lambda}). \tag{10}$$

Since $m \geq n \log q$ and $\sigma_{\mathsf{OM}} = \Omega(\sqrt{m})$, it follows from Lemma 2.3 and Lemma 2.4 that

$$\Pr[\mathsf{bad}] < \mathsf{negl}(\lambda).$$

Conditioned on $\overline{\mathsf{bad}}$ and by assumption on β_{DT}, with $c \geq 2C\sqrt{5}$ we have

$$\beta_{\mathsf{DT}} \leq \frac{q}{c \cdot \sigma_{\mathsf{OM}}(\sqrt{\kappa} + \sqrt{m} + \sqrt{\lambda})} \leq \frac{q}{2\sqrt{5} s_1(\mathbf{R})}.$$

It follows from Lemma 2.6 that $\mathcal{A}'_{\mathbf{R}}$ perfectly simulates the $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$ oracle. It returns $\mathbf{s}$ in Step 1 if and only if

$$\mathbf{b} = \mathbf{A}^{\intercal}\mathbf{s} + \mathbf{e} \bmod q \quad \text{with } \|\mathbf{e}\|_2 < \beta_{\mathsf{DT}}.$$

For $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, let $\epsilon_{\mathsf{DT}}(\mathbf{A})$ denote the conditional success probability of $\mathcal{A}$ on input $\mathbf{A}^{\intercal}$. Define

$$\mathcal{G} \coloneqq \{\mathbf{A} \in \mathbb{Z}_q^{n \times m} \mid \epsilon_{\mathsf{DT}}(\mathbf{A}) \geq \epsilon_{\mathsf{DT}}/2\}.$$

Let $\mathcal{G}_{\mathbf{A}}$ denote the event $\mathbf{A} \in \mathcal{G}$ and $p \coloneqq \Pr[\mathcal{G}_{\mathbf{A}}]$. We have

$$\epsilon_{\mathsf{DT}} = \mathbb{E}_{\mathbf{A}}[\epsilon_{\mathsf{DT}}(\mathbf{A})] \leq p + (1-p)\frac{\epsilon_{\mathsf{DT}}}{2}.$$

From this we obtain

$$\Pr_{\mathbf{A}}[\mathbf{A} \in \mathcal{G}] \geq \frac{\epsilon_{\mathsf{DT}}}{2 - \epsilon_{\mathsf{DT}}} \geq \frac{\epsilon_{\mathsf{DT}}}{2}.$$

Conditioned on $\mathcal{G}_{\mathbf{A}}$ and $\overline{\mathsf{bad}}$, the search-LWE adversary $\mathcal{A}'_{\mathbf{R}}$ has advantage at least $\epsilon_{\mathsf{LWE}} \coloneqq \epsilon_{\mathsf{DT}}/2$. By Lemma D.3, each call to $R[\mathcal{A}'_{\mathbf{R}}]$ therefore returns a valid ISIS solution with probability at least

$$\frac{\epsilon_{\mathsf{DT}}}{2} - \mathsf{negl}(\lambda) \geq \frac{\epsilon_{\mathsf{DT}}}{3},$$

for sufficiently big λ. Every call to $R[\mathcal{A}'_{\mathbf{R}}](\mathbf{A}, \mathbf{t}')$ uses a freshly sampled syndrome $\mathbf{t}'$. Hence, the probability that all L attempts fail to produce one solution is at most

$$\left(1 - \frac{\epsilon_{\mathsf{DT}}}{3}\right)^L \leq \exp\left(-\frac{L\epsilon_{\mathsf{DT}}}{3}\right).$$

By assumption, ϵ_{DT} is non-negligible. It follows that there exists polynomially bounded $L \geq \left\lceil \frac{3}{\epsilon_{\mathsf{DT}}}(\lambda + \ln(\kappa+1)) \right\rceil$ and applying a union bound over the $\kappa+1$ required solutions gives a conditional success probability at least $1-e^{-\lambda}$. Finally, Let $\mathcal{E} \coloneqq \mathcal{G}_{\mathbf{A}} \wedge \overline{\mathsf{bad}}$. Then

$$\begin{aligned}
\Pr\left[\mathrm{Exp}^{\mathsf{OM\text{-}ISIS}}_{\mathsf{pp},\mathcal{B}}(\lambda) = 1\right] &\geq \Pr\left[\mathrm{Exp}^{\mathsf{OM\text{-}ISIS}}_{\mathsf{pp},\mathcal{B}}(\lambda) = 1 \middle| \mathcal{E}\right] \Pr\left[\mathcal{E}\right] \\
&\geq (1 - e^{-\lambda})\left(\frac{\epsilon_{\mathsf{DT}}}{2} - \mathsf{negl}(\lambda)\right) \\
&= \Omega(\epsilon_{\mathsf{DT}}),
\end{aligned}$$

where we used $\Pr[\mathbf{A} \in \mathcal{G} \wedge \mathsf{bad}] \leq \Pr[\mathsf{bad}]$. This concludes the proof. □

E Proof of Proposition 5.1

Proof. Let $\Pi' = (\mathsf{Gen}', \mathsf{Enc}', \mathsf{Dec}', \mathsf{Eval}')$ be a statistically correct, IND-CCA1-secure HE scheme. We build $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$, an *approximate* HE scheme, as follows:

- $\mathsf{Gen}(1^\lambda)$: Sample $(\mathsf{pk}', \mathsf{sk}') \leftarrow \mathsf{Gen}'(1^\lambda)$, a random message $m_w \xleftarrow{\$} \mathcal{M}$ and let $c_w \leftarrow \mathsf{Enc}'_{\mathsf{pk}'}(m_w)$. Output the key pair $\mathsf{pk} = (\mathsf{pk}', c_w)$ and $\mathsf{sk} = (\mathsf{sk}', m_w)$.
- $\mathsf{Enc}_{\mathsf{pk}}(m)$: Output $c = (\mathsf{Enc}'_{\mathsf{pk}'}(m), \mathsf{Enc}'_{\mathsf{pk}'}(m))$.
- $\mathsf{Eval}(f, c)$: Parse c as (c_0, c_1). Output $\hat{c} = (\mathsf{Eval}'(f, c_0), c_1)$.
- $\mathsf{Dec}(c)$: Parse c as (c_0, c_1). Compute $m = \mathsf{Dec}(c_0)$ and $d = \mathsf{Dec}(c_1)$. If either m or d is the $\perp$ symbol, return $\perp$. If $d = m_w$ return $m + \alpha$, where $\alpha \neq 0$ is a fixed public element. Otherwise, output m.

Without loss of generality, suppose that $|\mathcal{M}|$ is exponential in the security parameter. If not, c_w can be the bitwise encryption of m_w, adapting the encryption, evaluation, and decryption algorithm accordingly. We only give the case of univariate functions f, but the general case follows naturally.

The scheme remains IND-CCA1 secure because α is independent of the secret key and public. The reduction chooses m_w itself, so it has no issues answering all decryption queries. Furthermore, from the point of view of an IND-CPA^D adversary, the scheme is statistically correct. An adversary cannot get an arbitrary

ciphertext in its state, as this would contradict IND-CPA security. In particular, it cannot have c_w, or any homomorphic evaluation of it as part of its state. Recall that for statistically correct schemes, IND-CPA^D security is equivalent to CPA security.

However, Π is not IND-CCA1^D secure. Using the unrestricted decryption oracle, the adversary queries $(c_w, \mathsf{Enc}'_{\mathsf{pk}'}(m'))$ for a random $m' \in \mathcal{M}$. Except with negligible probability, $m' \neq m_w$, so the oracle returns m_w. The adversary then submits $(0, m_w)$ as its challenge messages and applies $f(x) = x - x$ to the challenge ciphertext. Decrypting the resulting ciphertext yields 0 when $b = 0$ and α when $b = 1$, except with negligible probability which allows to distinguish and win the IND-CCA1^D game. □

F Branching Programs

A permutation branching program Π of length L and width W with input space $\{0,1\}^N$ is a sequence of L tuples of the form $(\mathsf{var}(t), \pi_{t,0}, \pi_{t,1})$, where

- $\mathsf{var}(t) : [L] \to [N]$ is a function that associates the tuple for the t-th step to an input bit $x_{\mathsf{var}(t)}$.
- $\pi_{t,0}, \pi_{t,1} : [W] \to [W]$ are permutations associated to the t-th step of the program that dictate the change of the state as a function of an input bit $x_{\mathsf{var}(t)}$.

On input a tuple $(x_1, \ldots, x_N) \in \{0,1\}^N$, the program Π proceeds as follows. At every period $t \in [L]$, the program has a state $\mathbf{v}_t \in \{0,1\}^W$, with the initial state $\mathbf{v}_0 = (1, 0, \ldots, 0)$. Then, for every $t \in L$ and $w \in [W]$, the state is computed recursively as

$$\mathbf{v}_t[w] = x_{\mathsf{var}(t)} \cdot v_{t-1}[\pi_{t,1}^{-1}(w)] + (1 - x_{\mathsf{var}(t)}) \cdot v_{t-1}[\pi_{t,0}^{-1}(w)]$$

where $\mathbf{v}_t[w]$ is the w-th coefficient of $\mathbf{v}_t$. The output of the program is 1 if $\mathbf{v}_L[1] = 1$ and 0 otherwise.

G Proof of Lemma 5.2

We give a proof of Lemma 5.2 that largely follow the methodology from Bourse et al. [7], simply generalizing to support matrix noise.

G.1 Ciphertext rerandomization

We start with the matrix noise generalization of [7, Lemma 3.1].

Lemma G.1. *Let $\varepsilon, \varepsilon' > 0$, and let $r > \eta_\varepsilon(\Lambda^\perp(\mathbf{G}))$ be a Gaussian parameter. For any matrix $\mathbf{E} \in \mathbb{Z}_q^{\bar{m} \times k}$ and $\mathbf{v} \in \mathbb{Z}_q^{\bar{m}}$, if $r \geq \max\Big(8\big((1-\varepsilon)(2\varepsilon')^2\big)^{-\frac{1}{k}}, \sqrt{5}(1 + s_1(\mathbf{E}))\sqrt{\frac{\ln(2k(1+1/\varepsilon))}{\pi}}\Big)$*

$$\Delta((\mathbf{S}, \mathbf{Sx}, \mathbf{Ex} + \mathbf{y}), (\mathbf{S}, \mathbf{u}, \mathbf{e}')) < \varepsilon' + \varepsilon,$$

where $\mathbf{x} \leftarrow \mathbf{G}_{\mathsf{rand}}^{-1}(\mathbf{v})$ *has parameter* r, $\mathbf{S} \xleftarrow{\$} \mathbb{Z}_q^{n\times k}, \mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^k, \mathbf{y} \leftarrow \mathcal{D}_{\mathbb{Z}^{\bar{m}},r}$ *and* $\mathbf{e}' \leftarrow \mathcal{D}_{\mathbb{Z}^{\bar{m}},r\hat{\mathbf{E}}^\intercal}$ *with* $\hat{\mathbf{E}} = [\mathbf{E}|\mathbf{I}_{\bar{m}}]$.

Proof. The proof is unchanged from [7, Lemma 3.1]. We simply replace the use of [7, Lemma 3.6] with its matrix generalization Lemma G.3. Using matrix $\mathbf{E}$ instead of vector $\mathbf{e}$ also introduces a factor 2 in the first term of the bound through the generalized leftover hash lemma 2.1. □

The following lemma is a special case of the linear transformation theorem from [28][Theorem 1]. Note that we replaced the max-log distance bound by a statistical distance bound.

Lemma G.2. *Let* $\Lambda \subset \mathbb{R}^n$ *be a lattice , let* $\mathbf{a} \in \mathbb{R}^n$*, and let* $A \coloneqq \Lambda + \mathbf{a}$ *be a lattice coset. Let* $\epsilon \in [0,1)$*, and* $r > 0$*. Let* $\mathbf{T} \in \mathbb{R}^{d\times n}$ *have full row rank. Suppose that* $\ker(\mathbf{T}) = \mathrm{span}(\Lambda \cap \ker(\mathbf{T}))$ *and* $r \geq \eta_\epsilon(\Lambda \cap \ker(\mathbf{T}))$*. Then,*

$$\Delta\left(\mathbf{T} \cdot D_{A,r}, D_{\mathbf{T}A,r\mathbf{T}^\intercal}\right) \leq \epsilon.$$

Lemma G.3. *Let* $\varepsilon \in [0,1)$ *and* $r > 0$*. Let* $\mathbf{E} \in \mathbb{Z}^{\bar{m}\times k}$ *be a matrix and* $\mathbf{c} \in \mathbb{Z}^k$*. Suppose that* $r \geq \sqrt{5}(1 + s_1(\mathbf{E}))\sqrt{\frac{\ln(2k(1+1/\varepsilon))}{\pi}}$*. Then*

$$\Delta(\mathbf{E}\mathbf{x} + \mathbf{y}, \mathbf{e}') < \varepsilon,$$

where $\mathbf{x} \leftarrow \mathcal{D}_{\Lambda^\perp(\mathbf{G})+\mathbf{c},r}, \mathbf{y} \leftarrow \mathcal{D}_{\mathbb{Z}^{\bar{m}},r}$*, and* $\mathbf{e}' \leftarrow \mathcal{D}_{\mathbb{Z}^{\bar{m}},r\widehat{\mathbf{E}}^\intercal}$ *with* $\widehat{\mathbf{E}} = [\mathbf{E}|\mathbf{I}_{\bar{m}}]$.

Proof. Let $\mathcal{L} = \Lambda^\perp(\mathbf{G}) \times \mathbb{Z}^{\bar{m}}$, $\widetilde{\mathbf{c}} \coloneqq (\mathbf{c}, \mathbf{0})$ and set $A \coloneqq \mathcal{L} + \widetilde{\mathbf{c}}$. Notice that $(\mathbf{x}, \mathbf{y}) \sim D_{A,r}$. From *Lemma G.4*, we have $r \geq \eta_\varepsilon\big(\mathcal{L} \cap \ker(\widehat{\mathbf{E}})\big)$. Furthermore,

$$\mathcal{L} \cap \ker(\widehat{\mathbf{E}}) = \left\{ \begin{bmatrix} \mathbf{z} \\ -\mathbf{E}\mathbf{z} \end{bmatrix} : \mathbf{z} \in \Lambda^\perp(\mathbf{G}) \right\}.$$

Since $\Lambda^\perp(\mathbf{G})$ is full rank, we have $\mathrm{span}(\mathcal{L} \cap \ker(\widehat{\mathbf{E}})) = \ker(\widehat{\mathbf{E}})$. From Lemma G.2, we have

$$\Delta\left(\mathbf{E}\mathbf{x} + \mathbf{y}, \mathcal{D}_{\widehat{\mathbf{E}}A,r\widehat{\mathbf{E}}^\top}\right) \leq \varepsilon.$$

Finally, $\widehat{\mathbf{E}}A = \mathbf{E}\big(\Lambda^\perp(\mathbf{G}) + \mathbf{c}\big) + \mathbb{Z}^{\bar{m}} = \mathbb{Z}^{\bar{m}}$ and the claims follows. □

Lemma G.4. *Let* $\varepsilon > 0$ *and* $\mathcal{L} = \Lambda^\perp(\mathbf{G}) \times \mathbb{Z}^{\bar{m}} \subseteq \mathbb{Z}^{k+\bar{m}}$*. For* $\mathbf{E} \in \mathbb{Z}^{\bar{m}\times k}$*, let* $\widehat{\mathbf{E}} = [\mathbf{E}|\mathbf{I}_{\bar{m}}]$ *we have*

$$\eta_\varepsilon\big(\mathcal{L} \cap \ker(\widehat{\mathbf{E}})\big) \leq \sqrt{5}(1 + s_1(\mathbf{E}))\sqrt{\frac{\ln(2k(1+1/\varepsilon))}{\pi}}.$$

The proof follows directly from the proof of [7, Lemma 3.7], simply replacing $\hat{\mathbf{e}}$ with $\hat{\mathbf{E}}$ when defining the shortest vectors $\mathbf{u}_k$.

G.2 Randomized homomorphic evaluation

Lemma G.5. *Let $r = \Omega\big(\sigma\sqrt{k\lambda \log k}\big)$ with $\sigma > 0$ the Gaussian parameter used for encryption. Let $(\mathsf{pk} = \mathbf{A}, \mathsf{sk}) \leftarrow \mathsf{GG\text{-}GSW.Gen}(1^\lambda)$ and let $\mathbf{C} = \mathbf{AS} + \mathbf{E} \leftarrow \mathsf{GG\text{-}GSW.Enc}(\mathsf{pk}, 0)$. For any $\mathbf{V} \in \mathbb{Z}_q^{\bar{m}\times k}$, we have*

$$\Big(\mathbf{C} \cdot \mathbf{G}^{-1}_{\mathsf{rand}}(\mathbf{V}) + \mathbf{Y}, \mathbf{C}\Big) \overset{s}{\approx} (\mathbf{C}', \mathbf{C}),$$

where $\mathbf{G}^{-1}_{\mathsf{rand}}(\cdot)$ has parameter r, $\mathbf{Y} \leftarrow \mathcal{D}_{\mathbb{Z}^{\bar{m}\times k}, r}$, and $\mathbf{C}' \leftarrow \mathsf{GG\text{-}GSW.Enc}_\Gamma(\mathsf{pk}, 0)$ with $\Gamma = r[\mathbf{E}|\mathbf{I}_{\bar{m}}]^\intercal$.

Proof. From Lemma 2.4, we have $s_1(\mathbf{E}) = O(\sigma\sqrt{k})$ with overwhelming probability. Hence $r = \Omega(\sigma\sqrt{k\lambda \log k})$ is sufficient to apply Lemma G.1. Applying that lemma columnwise proves the claim, exactly as in [7, Lemma 4.1]. □

Note that for $\mathbf{E}, \widetilde{\mathbf{Y}}$, sampled with parameter σ and σ_f respectively, $s_1(\mathbf{E} + \widetilde{\mathbf{Y}}) \leq s_1(\mathbf{E}) + s_1(\widetilde{\mathbf{Y}}) \leq O((\sigma + \sigma_f)\sqrt{k})$ with overwhelming probability. Hence, to rerandomize a ciphertext with noise $\mathbf{E} + \widetilde{\mathbf{Y}}$ for $\widetilde{\mathbf{Y}}_i \leftarrow \mathcal{D}^k_{\mathbb{Z}^{\bar{m}}, \sigma_f}$, it is sufficient to choose rerandomization parameter $r = \Omega\big((\sigma + \sigma_f)\sqrt{k\lambda \log(k)}\big)$.

The following lemma then directly follows from Lemma G.5 using the same proof as [7, Lemma 5.3].

Lemma G.6 (adapted from [7], Lemma 5.3). *Let $\sigma_f > 0$ and let $r = \Omega\big((\sigma+\sigma_f)\sqrt{k\lambda \log(k)}\big)$. Let $\mu, v_0, v_1 \in \{0,1\}$, $(\mathsf{pk} = \mathbf{A}, \mathsf{sk}) \leftarrow \mathsf{GG\text{-}GSW.Gen}(1^\lambda)$, $\mathbf{C} = \mathbf{AS}+\mathbf{E}+\mu\mathbf{G} \leftarrow \mathsf{GG\text{-}GSW.Enc}(\mathsf{pk}, \mu)$ and let $\mathbf{C}' = \mathbf{C}+\widetilde{\mathbf{Y}}$ with $\widetilde{\mathbf{Y}} \leftarrow \mathcal{D}^k_{\mathbb{Z}^{\bar{m}}, \sigma_f}$. We have,*

$$\Big(\mathbf{C}' \cdot \mathbf{G}^{-1}_{\mathsf{rand}}(\mathbf{V}_1) + (\mathbf{G} - \mathbf{C}') \cdot \mathbf{G}^{-1}_{\mathsf{rand}}(\mathbf{V}_0) + \mathbf{Y}, \mathbf{C}\Big) \overset{s}{\approx} (\mathbf{V}'_\mu, \mathbf{C}),$$

where $\mathbf{V}_b \leftarrow \mathsf{GG\text{-}GSW.Enc}_\Gamma(\mathsf{pk}, v_b)$ for $b \in \{0,1\}$, $\mathbf{Y} \leftarrow \mathcal{D}^k_{\mathbb{Z}^{\bar{m}}, r\sqrt{2}}$, and $\mathbf{V}'_\mu \leftarrow \mathsf{GG\text{-}GSW.Enc}_\Xi(\mathsf{pk}, v_\mu)$ with $\Xi = (\Gamma^\intercal\Gamma + 2r^2\hat{\mathbf{E}}\hat{\mathbf{E}}^\intercal)^{\frac{1}{2}}$ and $\hat{\mathbf{E}} = [\mathbf{E} + \widetilde{\mathbf{Y}}|\mathbf{I}_{\bar{m}}]$.

Finally, we consider Lemma 5.2. Its proof follows essentially unchanged from the induction of [7, Lemma 5.4]. The first level follows from Lemma G.5, while the remaining levels follow from Lemma G.6. We therefore omit the proof.

Lemma 5.2 (Noise distribution). *Let $\sigma_f > 0$ be a Gaussian parameter and let $r = \Omega\big((\sigma + \sigma_f)\sqrt{k\lambda \log(k)}\big)$. For a branching program $\Pi : \{0,1\}^N \to \{0,1\}$ of length L, let $\tau_{t,i}$ be the number of times the i-th variable x_i has been called after t steps of the evaluation. Let $(\mathsf{pk} = \mathbf{A}, \mathsf{sk}) \leftarrow \mathsf{GG\text{-}GSW.Gen}(1^\lambda)$. For any $(x_1, \dots x_N) \in \{0,1\}^N$, let $\mathbf{C}_i = \mathbf{AS}_i + \mathbf{E}_i + x_i\mathbf{G} \leftarrow \mathsf{GG\text{-}GSW.Enc}(\mathsf{pk}, x_i)$. For all $t \in [L]$ and $w \in [W]$, let $\mathbf{V}_t[w]$ be as defined in Equation (8). We have,*

$$\big(\mathbf{V}_t[w], (\mathbf{C}_i)_{i=1}^N\big) \overset{s}{\approx} \big(\mathbf{V}'_{t,w}, (\mathbf{C}_i)_{i=1}^N\big),$$

where $\mathbf{V}'_{t,w} \leftarrow \mathsf{GG\text{-}GSW.Enc}_{\Gamma_t}(\mathsf{pk}, v_t[w])$ with $\Gamma_t = r\left(2\sum_{i=1}^N \tau_{t,i}\hat{\mathbf{E}}_i\hat{\mathbf{E}}_i^\intercal\right)^{1/2}$, $\hat{\mathbf{E}}_i = [\mathbf{E}_i + \widetilde{\mathbf{Y}}_i|\mathbf{I}_{\bar{m}}]$ and $\widetilde{\mathbf{Y}}_i \leftarrow \mathcal{D}^k_{\mathbb{Z}^{\bar{m}}, \sigma_f}$.

H Proof of Lemma H.1

Lemma 5.4. *Game G_1 and G_2 are statistically indistinguishable.*

Proof. Except with negligible probability, it follows from Lemmas 2.5 and 5.3, that each ciphertext recorded in the state has noise smaller than $\rho_{L_{\max}} \leq B$. Let $\Lambda_{\mathbf{A}} = \mathbf{A}\mathbb{Z}^n + q\mathbb{Z}^{\bar{m}}$. From Lemma 3.1, Lemma 2.1, we have

$$\lambda_\infty(\mathcal{L}_{\mathbf{A}}) \geq \frac{q}{6} \quad \text{and} \quad d_\infty\left(\frac{q}{2}\mathbf{u}, \mathcal{L}_{\mathbf{A}}\right) \geq q/6, \tag{11}$$

for any unit vector $\mathbf{u}$, except with negligible probability. Conditioned on these events, consider a query j to $\mathcal{D}(\cdot)$. For readability, we drop the index and write the corresponding state entry as $(m_0, m_1, \mathbf{C}, \mathbf{E})$. If $m_0 \neq m_1$, the oracle returns $\perp$ in both game G_1 and G_2. Otherwise suppose that $\mu = m_0 = m_1$. For some $\mathbf{S} \in \mathbb{Z}_q^{n\times k}$, we have

$$\mathbf{C} = \mathbf{AS} + \mathbf{E} + \mu\mathbf{G} \bmod q, \quad \text{with} \quad \|\mathbf{E}\|_2 < B.$$

From Lemma 3.2, if $\|\mathbf{E}\|_2 < D$ the $\mathcal{D}$ oracle returns μ in both games. It remains to consider the case where $D \leq \|\mathbf{E}\|_2 < B$. By definition, $\mathcal{D}(j)$ returns $\perp$ in game G_2. In game G_1, the oracle could return $\mu' \in \{0,1\}$ only if there exists $(\mathbf{S}', \mathbf{E}')$ such that

$$\mathbf{C} = \mathbf{AS}' + \mathbf{E}' + \mu'\mathbf{G} \quad \text{with} \quad \|\mathbf{E}'\| < D.$$

We have $\|\mathbf{E} - \mathbf{E}'\| < B + D$. Since by assumption $B + D < q/6$, this contradicts the first part of Equation (11) if $\mu = \mu'$, and using the same arguments as in Lemma 3.2, it contradicts the second if $\mu \neq \mu'$. Hence, both oracles agree on every query except with negligible probability. □

Lemma H.1. *For any* PPT *adversary $\mathcal{A}$ that makes $Q_{\mathsf{enc}} = Q_{\mathsf{enc}}(\lambda)$ queries to the encryption oracle and $Q_{\mathsf{eval}} = Q_{\mathsf{eval}}(\lambda)$ queries to the homomorphic evaluation oracle, there exists a* PPT DT-HintLWE *distinguisher $\mathcal{B}$ such that*

$$|\Pr[W_2] - \Pr[W_3]| \leq 2Q_{\mathsf{enc}}k \cdot \mathsf{Adv}_{\mathcal{B}}^{\mathsf{DT\text{-}HintLWE}_{\mathsf{pp}}}(\lambda) + \mathsf{negl}(\lambda).$$

For convenience, we write $\mathsf{DT\text{-}HintLWE}^{(k)}$ to denote the matrix variant of DT-HintLWE. In this variant, the challenger samples $\mathbf{S} \xleftarrow{\$} \mathbb{Z}_q^{n\times k}$ and error matrices $\mathbf{E}, \mathbf{F}$ columnwise from their respective distributions, and returns

$$\mathbf{B} = \begin{cases} \mathbf{AS} + \mathbf{E} \bmod q, & b = 0, \\ \mathbf{U} \xleftarrow{\$} \mathbb{Z}_q^{m\times k}, & b = 1, \end{cases} \qquad \mathbf{H} = \mathbf{L}_s\mathbf{S} + \mathbf{L}_e\mathbf{E} + \mathbf{F}.$$

The same column-by-column hybrid as in the proof of Theorem 4.1 shows that, for every distinguisher $\mathcal{B}'$, there exists a distinguisher $\mathcal{B}$ such that

$$\mathsf{Adv}_{\mathsf{pp},\mathcal{B}'}^{\mathsf{DT\text{-}HintLWE}^{(k)}}(\lambda) \leq k \cdot \mathsf{Adv}_{\mathsf{pp},\mathcal{B}}^{\mathsf{DT\text{-}HintLWE}}(\lambda).$$

Proof. Let $\overline{G_2}$ be defined as G_2, except that $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{\bar{m} \times n}$ is sampled uniformly. Since the decryption oracles in G_2 no longer use $\mathbf{R}$, the leftover hash lemma gives

$$\left|\Pr[W_2] - \Pr[\overline{W_2}]\right| \leq \mathsf{negl}(\lambda).$$

We introduce a series of hybrid games. For $j \in 0, \ldots, Q_{\mathsf{enc}}$, we define H_j to be the same game as $\overline{G_2}$ except the first j encryption queries are answered as in G_3, and the last $Q_{\mathsf{enc}} - j$ queries are answered as in game $\overline{G_2}$. Thus, $\mathsf{H}_0 = \overline{G_2}$ and $\mathsf{H}_{Q_{\mathsf{enc}}} = G_3$.

We build a $\mathsf{DT\text{-}HintLWE}_{\mathsf{pp}}^{(k)}$ distinguisher $\mathcal{B}$ that uses $\mathcal{A}$ as a subroutine. On input the challenge matrix $\mathbf{A}$, $\mathcal{B}$ runs $\mathsf{st} \leftarrow \mathcal{A}_1^{\mathsf{GG\text{-}GSW.Dec}}(\mathbf{A})$. It answers queries to the decryption oracle using its $\mathcal{O}_{\mathsf{DT\text{-}LWE}}$ oracle as described in the proof of Theorem 4.1.

Then $\mathcal{B}$ outputs $(\mathbf{L}_s, \mathbf{L}_e)$ with $\mathbf{L}_s = \mathbf{0}$ and receives challenge $\mathbf{B} \in \mathbb{Z}_q^{\bar{m} \times k}$ and hint

$$\mathbf{H} = \begin{bmatrix} \mathbf{H}^{(1)} \\ \vdots \\ \mathbf{H}^{Q_{\mathsf{eval}}} \end{bmatrix}, \quad \text{with } \mathbf{H}^{(\psi)} = \mathbf{E} + \widetilde{\mathbf{Y}}^{(\psi)} \text{ for } \psi \in [Q_{\mathsf{eval}}].$$

First, $\mathcal{B}$ samples an index $j^* \xleftarrow{\$} [Q_{\mathsf{enc}}]$ and a challenge bit b'. Then, it runs the second phase $\mathcal{A}_2(\mathsf{st})$, initializes the state S and answers queries as follows:

- $\mathsf{E}(m_0, m_1)$: The i-th encryption query is answered as follows.
  - If $i > j^*$: Sample $\mathbf{S}, \mathbf{E}$ and compute $\mathbf{C}$ as in game G_2. Output $\mathbf{C}$ and update the state with $S \cup (m_0, m_1, \mathbf{C}, \mathbf{E})$.
  - If $i < j^*$: Return $\mathbf{C} \xleftarrow{\$} \mathbb{Z}_q^{\bar{m} \times k}$ and update the state as in game G_3.
  - If $i = j^*$: Output $\mathbf{C} = \mathbf{B} + m_{b'}\mathbf{G}$ and update the state with $S \leftarrow S \cup (m_0, m_1, \mathbf{C}, \perp)$. Let u^* be the index of this query in the state.
- $\mathsf{H}(\Pi, J)$: Given input (Π, J) the ψ-th homomorphic evaluation query for $\psi \in [Q_{\mathsf{eval}}]$, let $\mathbf{C}_1, \ldots \mathbf{C}_N$ be the ciphertexts corresponding to the tuple of indices J. Let $x_1, \ldots, x_N \in \{0,1\}$ be the corresponding messages they encrypt and for $i \in J$, let $\mathbf{E}_i$ be the corresponding ciphertext noise i.e. $x_i = S[i].m_{b'}$ and $\mathbf{E}_i = S[i].\mathbf{E}$. Let $\hat{x} = \Pi(x_1, \ldots, x_N)$. Set,

$$\hat{\mathbf{E}}_i = \begin{cases} [\mathbf{E}_i + \widetilde{\mathbf{Y}}_i | \mathbf{I}_{\bar{m}}] \text{ with } \widetilde{\mathbf{Y}}_i \leftarrow \mathcal{D}_{\mathbb{Z}^{\bar{m} \times k}, r_1}, & \text{if } i \neq u^* \\ [\mathbf{H}^{(\psi)} | \mathbf{I}_{\bar{m}}] & \text{if } i = u^*. \end{cases}$$

  The evaluated ciphertext is computed as $\hat{\mathbf{C}} \leftarrow \mathsf{GG\text{-}GSW.Enc}_{\Gamma_L}(\mathsf{pk}, \hat{x})$ where $\Gamma_L = r\sqrt{2\sum_{i=1}^{N} \tau_{L,i} \hat{\mathbf{E}}_i \hat{\mathbf{E}}_i^\intercal}$. The oracle returns $\hat{\mathbf{C}}$ and updates the state as in game G_2, recording the sampled output error.
- $\mathsf{D}(j)$: Queries to the restricted decryption oracle are answered as in game G_3.

Finally, $\mathcal{B}$ outputs 0 if $\mathcal{A}$ guesses b' and 1 otherwise.

It follows from Lemma 5.2 that the output of H is correctly distributed. Moreover, if $\mathbf{B} = \mathbf{AS}' + \mathbf{E}' \bmod q$, then $\mathcal{B}$ simulates hybrid H_{j^*-1} whereas if $\mathbf{B}$ is uniformly distributed, $\mathcal{B}$ simulates H_{j^*}. The claim follows from a standard hybrid over the Q_{enc} encryption queries. □