Fully Asymmetric Anamorphic Homomorphic Encryption from LWE

Amit Deo and Benoît Libert

Zama, France

Abstract. As introduced by Persiano *et al.* (Eurocrypt'22), anamorphic encryption (AE) is a primitive enabling private communications against a dictator that forces users to surrender their decryption keys. In its fully asymmetric flavor (defined by Catalano *et al.*, Eurocrypt'24), anamorphic channels can work as hidden public-key mechanisms in the sense that anamorphic encryptors are not necessarily able to decrypt anamorphic ciphertexts. Unfortunately, fully asymmetric AE is hard to come by and even impossible to obtain from ordinary public-key encryption via black-box constructions. So far, only three schemes are known to rely on well-established assumptions. In this paper, we exhibit constructions from the standard LWE assumption based on Regev's cryptosystem and its dual version. In both cases, we retain the additive homomorphism of the schemes. We additionally show that dual Regev is public-key anamorphic in the sense of Persiano *et al.* (Crypto'24). In the FHE setting, we show that the dual GSW system provides fully asymmetric AE (while preserving its leveled homomorphism) when instantiated with binary/ternary secret keys. Along the way, we discuss the extent to which our schemes satisfy a generalization of Banfi *et al.*'s notion of robustness (Eurocrypt'24) to the case of homomorphically evaluated ciphertexts.

Keywords. Anamorphic encryption, homomorphic encryption, LWE.

1 Introduction

In normal deployments of cryptography, it is assumed that senders can encrypt any message of their choice and that only receivers know their decryption key. However, in a dictatorship, users' privacy may be compromised if they are forced to hand over their decryption keys. On the other side, senders may be coerced into sending particular messages and lose their ability to communicate freely.

In order to tackle these issues, Persiano, Phan and Yung [43] introduced the elegant paradigm of anamorphic encryption (AE), which aims at enabling covert communication in an environment where an overruling authority has the power of restricting users' privacy or their freedom of speech.

Persiano *et al.* [43] proposed two flavors of AE called *sender-anamorphic* and *receiver-anamorphic* encryption, which both allow covert private communication even in the presence of a dictator. Both flavors involve normal messages, which the dictator has access to, and *covert* messages that should remain hidden even under coercion. Informally, receiver-anamorphic encryption schemes

(which are the focus of this work) consist of an ordinary public key encryption (PKE) scheme (comprised of the usual $\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec}$ algorithms) and an anamorphic triplet consisting of algorithms $(\mathsf{aGen}, \mathsf{aEnc}, \mathsf{aDec})$, where aGen outputs an anamorphic key pair $(\mathsf{apk}, \mathsf{ask})$ that is indistinguishable from a regular key pair produced by KGen. In addition, aGen also outputs a *double key* dk which allows encrypting covert messages and a *trapdoor key* tk which allows decrypting covert messages. When coerced, the receiver only reveals ask to the dictator and can plausibly deny that $(\mathsf{dk}, \mathsf{tk})$ exists since $(\mathsf{ask}, \mathsf{apk})$ looks like a normal key pair. In fully asymmetric schemes [18], the pair $(\mathsf{dk}, \mathsf{tk})$ functions as the key pair of a covert asymmetric encryption mechanism while we have $\mathsf{tk} = \emptyset$ if the anamorphic mode is symmetric. Using apk and dk, aEnc produces an anamorphic ciphertext act containing both a normal plaintext μ and covert plaintext $\hat{\mu}$. The normal decryption procedure $\mathsf{Dec}(\mathsf{ask}, \cdot)$ operates on act exactly as on a normal ciphertext to recover μ whereas $\mathsf{aDec}(\mathsf{dk}, \mathsf{tk}, \mathsf{ask}, \mathsf{act})$ recovers $\hat{\mu}$. Security-wise, a dictator cannot tell apart anamorphic keys/ciphertexts $(\mathsf{ask}, \mathsf{apk}, \mathsf{act})$ and normal ones $(\mathsf{sk}, \mathsf{pk}, \mathsf{ct})$, meaning that $\hat{\mu}$ remains private even when ask is exposed.

As emphasized in [43], describing new cryptosystems that are intentionally designed to support an anamorphic channel is unlikely to be useful since a dictator can easily ban bespoke systems yielding the anamorphic security guarantees. The challenge is thus to demonstrate the existence of an anamorphic mode in *existing* encryption schemes that are already in use and were not initially designed for the purpose of enabling covert communication.

Several works [43,31,48,8,18] took significant steps forward in this direction. For example, Kutylowski *et al.* [31] showed that any randomness-recovering PKE scheme (such as Goldwasser-Micali [27] or Paillier [40]) has a symmetric anamorphic mode. Regarding generic constructions from any PKE scheme, Persiano *et al.* [43] gave a simple realization based on rejection-sampling,[1] where anamorphic ciphertexts can carry $O(\log \lambda)$ bits of covert messages if λ is the security parameter. On the negative side, Catalano *et al.* [19] also showed that no black-box AE construction admits a fully asymmetric anamorphic mode. This means that fully asymmetric AE systems can only be obtained by re-examining existing systems on a case-by-case basis, or by imposing conditions on the underlying PKE candidate, or via non-black-box techniques (like NIZK [11], obfuscation [9] or garbled circuits [50]) that unfortunately tend to suffer from a lack of efficiency.

On the constructive front, Catalano *et al.* [18] showed that some well-known existing schemes can be endowed with fully asymmetric anamorphic triplets. For example, the lite Cramer-Shoup CCA1 cryptosystem [21] and the Naor-Yung construction [39] fall into this category. Persiano *et al.* [44] showed that the Koppula-Waters cryptosystem [30] is in fact public-key anamorphic (which implies full asymmetry). The latter property means that the scheme can be set up in anamorphic mode without requiring any initial communication between the sender and anamorphic encryptors: anamorphic ciphertexts are indeed computable using only the anamorphic public key apk, without any double key dk.

[1] As shown in [20], this construction requires the underlying PKE system to have high min-entropy ciphertexts.

While encouraging, these positive results still leave us in a somewhat unsatisfactory situation where only a handful of PKE schemes are known to support fully asymmetric anamorphic channels. Moreover, they suffer from certain limitations in terms of efficiency and/or advanced functionalities. In particular, the problem of constructing fully asymmetric FHE systems remains open. As of today, we do not even have a fully asymmetric linearly homomorphic scheme where the expansion rate (i.e., the ratio between ciphertext and plaintext sizes) is at most polylogarithmic in the security parameter, let alone with post-quantum security.

1.1 Our Contributions

In this paper, we show that several well-known homomorphic PKE schemes based on the Learning-With-Errors (LWE) assumption [45,46] do have fully asymmetric anamorphic triplets. Moreover, we show that some of them can be made robust [8], meaning that normal ciphertexts are never mistakenly interpreted as conveying a covert plaintext by the anamorphic decryption algorithm.

As a first contribution, we prove that the dual Regev [24] cryptosystem admits a fully asymmetric anamorphic triplet, which is amenable to additionally satisfy the notion of robustness of Banfi *et al.* [8]. This yields the first plausibly quantum-safe linearly homomorphic AE scheme where the expansion rate is only $O(\log^2 \lambda)$, or even $O(\log \lambda)$ depending on the parameters. We first describe an anamorphic construction using lattice trapdoors [24,37] before providing an alternative construction that does not rely on such trapdoors. The former, trapdoor-based construction has the advantage of being public-key anamorphic [44], meaning that the double key dk is empty. As such, it allows anamorphic encryptors to send anamorphic messages without communicating with the receiver beforehand in order to obtain a covert encryption key. In conjunction with robustness, the latter public-key anamorphism solves an open question from [44].

A notable property of our alternative trapdoor-less construction is that its bandwidth rate [43] (i.e. the ratio between covert and normal message lengths) is naturally larger than 1. In the case where the normal message consists of a single plaintext slot, the bandwidth rate can be as large $\Theta(\lambda)$ when using ternary secret keys. If we use a Gaussian secret key, the *expected* bandwidth rate drops to ≈ 1. Previously, a bandwidth rate greater than 1 was also observed for the Goldwasser-Micali scheme in [31], but its anamorphic mode is symmetric. So far, the only known fully asymmetric AE system with bandwidth rate > 1 was obtained [44] from the Koppula-Waters PKE [30], which is significantly less efficient than dual Regev.

In a second contribution, we adapt the idea underlying our trapdoor-based anamorphic dual Regev construction to build a fully asymmetric anamorphic triplet for the primal Regev cryptosystem [45,46]. The resulting scheme retains its additive homomorphic property over the anamorphic message space while achieving an expansion rate that is poly-logarithmic in the security parameter.

As a third contribution, we improve on the results of Catalano *et al.* [18] who highlighted the existence of a (symmetric) anamorphic channel in the Gentry-Sahai-Waters FHE scheme [25]. We show that the dual GSW scheme admits a

fully asymmetric anamorphic triplet when its secret keys are binary or ternary (our proofs work in both cases) and sampled from a natural distribution.

Our final contribution is to provide a refined definition of robustness in the context of homomorphic encryption (HE). One caveat of the current definition [8] is that it only considers fresh ciphertexts and does not guarantee anything about homomorphically evaluated ciphertexts. We thus generalize the existing definition into a notion of homomorphic robustness for anamorphic HE. Our definition requires that, after homomorphically processing a set of normal ciphertexts, the resulting evaluated ciphertext does not anamorphically decrypt to a valid covert message, except with negligible probability. This captures applications where homomorphic evaluations are never performed on a mixture of normal and anamorphic ciphertexts. We prove that our dual Regev/GSW constructions achieve this robustness property when restricted to *additive* homomorphic operations. In dual GSW, we were unable to prove robustness under homomorphic multiplications and leave it as an interesting open question to do so without modifying the scheme. We also discuss issues that arise in attempts to strengthen our definition of homomorphic robustness when computing over a mix of normal and anamorphic ciphertexts.

1.2 Technical Overview

Anamorphic Triplets for dual Regev. Recall that the dual Regev cryptosystem [24] has a public key of the form $\mathsf{pk} = (\mathbf{A}, \mathbf{U} = \mathbf{A} \cdot \mathbf{E}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^{n \times n}$, where the secret key $\mathsf{sk} = \mathbf{E} \in \mathbb{Z}^{m \times n}$ is a small-norm matrix sampled from a high-entropy distribution (that can be binary, ternary or Gaussian). A ciphertext encrypting $\boldsymbol{\mu} \in \mathbb{Z}_p^n$ takes the form $(\mathbf{c}_0, \mathbf{c}_1) = (\mathbf{A}^\top \cdot \mathbf{s} + \mathbf{e}_0, \mathbf{U}^\top \cdot \mathbf{s} + \mathbf{e}_1 + \Delta \cdot \boldsymbol{\mu})$ where $\mathbf{s} \in \mathbb{Z}_q^n$ is uniform, $\mathbf{e}_0 \in \mathbb{Z}^m, \mathbf{e}_1 \in \mathbb{Z}^n$ are small noise vectors and $\Delta = q/p$. Decryption computes $\mathbf{c}_1 - \mathbf{E}^\top \mathbf{c}_0 \approx \Delta \cdot \boldsymbol{\mu}$ and decodes it to recover the message $\boldsymbol{\mu}$. In order to obtain an anamorphic triplet, we endow $\mathbf{A}$ with a lattice trapdoor [24,37] that serves as the trapdoor key tk. The anamorphic encryption algorithm builds on an idea suggested by Zhang *et al.* [51] and encodes the covert message in the high-order bits of $\mathbf{s}$. Namely, it replaces $\mathbf{s}$ by $\hat{\mathbf{s}} = \mathbf{s} + \Delta \cdot \hat{\boldsymbol{\mu}}$, where $\hat{\boldsymbol{\mu}} \in \mathbb{Z}_p^n$ is a covert message and $\mathbf{s}$ is a Gaussian vector sampled from the noise distribution. Using the lattice trapdoor tk, the anamorphic decryption algorithm aDec can recover $\hat{\mathbf{s}} \in \mathbb{Z}_q$ from $\mathbf{c}_0 = \mathbf{A} \cdot \hat{\mathbf{s}} + \mathbf{e}$ provided that $\mathbf{e}$ is small (as observed in [51], $\hat{\mathbf{s}}$ does not have to be small for this purpose). From $\hat{\mathbf{s}} = \mathbf{s} + \Delta \cdot \hat{\boldsymbol{\mu}}$, aDec can then compute $\hat{\boldsymbol{\mu}}$ using a standard decoding procedure. Interestingly, aEnc can encrypt using only the anamorphic public key $\mathsf{apk} = (\mathbf{A}, \mathbf{U})$, without obtaining any additional covert double key.

We prove (public-key) anamorphism under the HNF LWE assumption [6], which says that $\mathbf{A}^\top \mathbf{s} + \mathbf{e}$ is pseudorandom even when $(\mathbf{s}, \mathbf{e})$ are both sampled from the noise distribution. Even if the secret key $\mathbf{E}$ is exposed, we can rely on a noise randomization lemma due to Katsumata and Yamada [28] to properly simulate the encryption oracle and prove that a normal ciphertext is indistinguishable from one where $\mathbf{c}_0$ carries a covert message in the high-order bits of $\mathbf{s}$.

To achieve robustness, our solution is to have the normal encryption algo-

rithm Enc compute $\mathbf{c}_0 = \mathbf{A}^\top \cdot \mathbf{s} + \mathbf{e}_0$ for a uniformly chosen $\mathbf{s} \in \mathbb{Z}_q^n$ while aEnc computes $\mathbf{c}_0 = \mathbf{A}^\top \cdot (\mathbf{s} + \Delta \cdot \hat{\boldsymbol{\mu}}) + \mathbf{e}_0$ using a short LWE secret $\mathbf{s}$. This way, aDec can distinguish anamorphic ciphertexts from normal ones by inspecting the size of $\mathbf{s}$ when attempting to recover a covert message from a candidate $\hat{\mathbf{s}} = \mathbf{s} + \Delta \cdot \hat{\boldsymbol{\mu}}$.

In Supplementary Material C, we provide a different anamorphic triplet for dual Regev, which is not based on lattice trapdoors. Instead, we exploit the fact that a dual Regev secret key $\mathbf{e} \in \mathbb{Z}^m$ (which underlies the public key $\mathbf{u} = \mathbf{A} \cdot \mathbf{e} \in \mathbb{Z}_q^n$) sampled from a binary/ternary distribution is very likely to contain $\Omega(\lambda)$ zeroes. We can then give away a uniformly chosen proper subset $I \subset [m]$ of these zero positions in the double key dk without compromising the fully asymmetric property. This allows us to have aEnc embed the covert message $\hat{\boldsymbol{\mu}} \in \mathbb{Z}_p^k$ in ciphertext components indexed by I. When it comes to proving anamorphic security (against a dictator that knows $\mathbf{e}$), we can use the observation that, when the normal decryption algorithm Dec computes $c_1 - \mathbf{e}^\top \mathbf{c}_0$, $\mathbf{e}$ does not interfere with the covert message components since they are precisely located in components of $\mathbf{c}_0$ that correspond to zero positions of $\mathbf{e}$. This approach does not extend to a packed version with $\ell = \omega(1)$ regular message slots since this would require an expanded secret key $\mathbf{E} \in \mathbb{Z}^{m \times \ell}$ and we can only embed anamorphic messages in ciphertext positions that correspond to all-zeroes rows of $\mathbf{E}$. On the upside, we do not have this limitation in the length of anamorphic plaintexts, so that the bandwidth rate can be as large as $\Omega(\lambda)$.

While our trapdoor-based anamorphic triplet for dual Regev can easily be adapted to the ring setting [35], our trapdoor-less construction cannot as it inherently relies on the unstructured nature of the public matrix.

Trapdoor-based primal Regev. The (packed) primal Regev system [45,46] involves public keys of the form $\mathsf{pk} = (\mathbf{A}, \mathbf{U} = \mathbf{S}^\top \cdot \mathbf{A} + \mathbf{E}^\top) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^{\ell \times m}$, where $\mathsf{sk} = \mathbf{S} \in \mathbb{Z}_q^{n \times \ell}$ is uniform and $\mathbf{E} \in \mathbb{Z}^{m \times \ell}$ is a Gaussian noise. An encryption of $\boldsymbol{\mu} \in \mathbb{Z}_p^\ell$ computes $\mathbf{c} = (\mathbf{c}_0, \mathbf{c}_1) = (\mathbf{A} \cdot \mathbf{r}, \mathbf{U} \cdot \mathbf{r} + \Delta \cdot \boldsymbol{\mu})$, where $\mathbf{r} \in \{0,1\}^m$ is uniformly random and $\Delta = q/p$. Decryption simply decodes $\mathbf{c}_1 - \mathbf{S}^\top \cdot \mathbf{c}_0 \approx \Delta \cdot \boldsymbol{\mu}$. Similarly to our trapdoor-based dual Regev, the aEnc algorithm embeds the covert message into the encryption randomness. In this case, things are slightly more complicated. In anamorphic mode, the uniform public matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ is replaced by a "close-to-low-rank" matrix $\mathbf{C}^\top \cdot \mathbf{B} + \mathbf{F}^\top$ (similarly to the lossy mode of LWE [26]). Here, $\mathbf{C} \in \mathbb{Z}_q^{\ell \times n}$ is a statistically uniform trapdoored matrix of rank $\ell \ll n$ whereas $\mathbf{B} \in \mathbb{Z}^{\ell \times m}$ and $\mathbf{F} \in \mathbb{Z}^{m \times n}$ are small-norm Gaussian matrices. The double key is $\mathsf{dk} = (\mathbf{C}, \mathbf{D} = \mathbf{C} \cdot \mathbf{S})$ while tk consists of a trapdoor [37] for $\mathbf{C}$. Given a covert message $\hat{\boldsymbol{\mu}} \in \mathbb{Z}_p^\ell$, an anamorphic encryption is obtained as $\mathbf{c}' = (\mathbf{c}'_0, \mathbf{c}'_1) = (\mathbf{A} \cdot \mathbf{r} + \mathbf{C}^\top \cdot (\Delta \cdot \hat{\boldsymbol{\mu}}), \mathbf{U} \cdot \mathbf{r} + \mathbf{D}^\top \cdot (\Delta \cdot \hat{\boldsymbol{\mu}}) + \Delta \cdot \boldsymbol{\mu})$. Since $\mathbf{c}'_0 \approx \mathbf{C}^\top \cdot (\mathbf{B} \cdot \mathbf{r} + \Delta \cdot \hat{\boldsymbol{\mu}})$, the trapdoor key tk can recover $\mathbf{B} \cdot \mathbf{r} + \Delta \cdot \hat{\boldsymbol{\mu}}$ and decode it to $\hat{\boldsymbol{\mu}}$ since $\mathbf{B} \cdot \mathbf{r}$ is small. Note that the $\mathbf{D}$-dependent term in $\mathbf{c}'_1$ ensures that $(\mathbf{c}'_0, \mathbf{c}'_1)$ decrypts to the normal message $\boldsymbol{\mu}$ under the secret key $\mathsf{ask} = \mathbf{S}$, which is necessary for the proof of anamorphism to work out. When it comes to proving the scheme fully asymmetric, we need to rely on a "first-are-errorless"

LWE assumption [22,15,3] (which has reductions from LWE) since the double key reveals ℓ noiseless LWE samples $\mathbf{D} = \mathbf{C} \cdot \mathbf{S}$.

Anamorphic dual GSW. In the dual GSW FHE, a natural choice is to use a ternary secret key $\mathbf{e} \in \{0, \pm 1\}^m$ and public key $\mathbf{B} = \begin{bmatrix} \mathbf{A}^\top \\ \mathbf{e}^\top \cdot \mathbf{A}^\top \end{bmatrix} \in \mathbb{Z}_q^{(m+1)\times n}$, where $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ is a uniform matrix. We will assume that $\mathbf{e}$ is sampled from the distribution $\mathcal{P}$ where $+1$ and -1 occur with probability $1/4$ and 0 occurs with probability $1/2$ (the construction and proofs also work for uniform binary secret keys). An encryption of $\mu \in \mathbb{Z}_p$ is a matrix $\mathbf{C} = \mathbf{B} \cdot \mathbf{S} + \mathbf{E} + \mu \cdot \mathbf{G}$, where $\mathbf{S} \in \mathbb{Z}_q^{n\times(m+1)\lceil \log q \rceil}$ is chosen uniformly, $\mathbf{E}$ is a Gaussian noise matrix and $\mathbf{G} = \mathbf{I}_{m+1} \otimes (1, 2, \ldots, 2^{\lfloor \log q \rfloor}) \in \mathbb{Z}_q^{(m+1)\times(m+1)\lceil \log q \rceil}$ is the gadget matrix [37]. Decryption entails computing $[-\mathbf{e}^\top \mid 1] \cdot \mathbf{C} \cdot \mathbf{v}$ for a short constant vector $\mathbf{v}$ such that $[-\mathbf{e}^\top \mid 1] \cdot \mathbf{C} \cdot \mathbf{v} = [-\mathbf{e}^\top \mid 1] \cdot \mathbf{E} \cdot \mathbf{v} + \Delta \cdot \mu$. Since $[-\mathbf{e}^\top \mid 1] \cdot \mathbf{E} \cdot \mathbf{v}$ is small, a standard decoding procedure allows computing μ.

The anamorphic key generation algorithm extends the idea of our trapdoorless dual Regev construction. For any matrix $\mathbf{M} \in \mathbb{Z}_q^{N\times M}$ and any set $I \subseteq [N]$, let $\mathbf{M}_I$ the matrix obtained from $\mathbf{M}$ by replacing the rows with indices outside of I by the vector $\mathbf{0}^M$. After sampling the secret key $\mathbf{e} = (e_1, \ldots, e_m)^\top \hookleftarrow \mathcal{P}^m$, aGen defines the double key dk as a uniformly chosen subset $I = (i_1, \ldots, i_{n_0}) \subset [m]$ of size $n_0 = \Theta(\lambda)$ such that $e_i = 0$ for all $i \in I$. The parameter n_0 is chosen carefully in such a way that a sufficiently large I exists with overwhelming probability. Then, aGen samples a short vector $\mathbf{t}' \in \mathbb{Z}^{n_0-1}$ with Gaussian entries and sets the trapdoor key $\mathsf{tk} = \mathbf{t} \in \mathbb{Z}^m$ to be the vector resulting from permuting the entries of $((\mathbf{t}')^\top \mid \mathbf{0}^{m-n_0} \mid -1)^\top$. The permutation is chosen so that $\mathbf{t}[i_j] = \mathbf{t}'[j]$ for $j \in [n_0 - 1]$, $\mathbf{t}[i_{n_0}] = -1$ and $\mathbf{t}[j] = 0$ for the remaining indices j. The matrix $\mathbf{A}$ is built by applying the same permutation to the rows of $\begin{bmatrix} \tilde{\mathbf{A}} \\ (\mathbf{t}'^\top \mid \mathbf{0}^{m-n_0}) \cdot \tilde{\mathbf{A}} + \bar{\mathbf{e}}_I^\top \end{bmatrix}$, where $\tilde{\mathbf{A}}$ is uniform in $\mathbb{Z}_q^{(m-1)\times n}$ and $\bar{\mathbf{e}}_I$ is a Gaussian noise.

The public key $\mathsf{apk} = \mathbf{B}$ is then obtained from $\mathbf{A}$ as in the normal key generation. Crucially, the trapdoor key $\mathsf{tk} = \mathbf{t}$ satisfies $\mathbf{t}_{[m]\setminus I} = \mathbf{0}$ and thus $[\mathbf{t}^\top \mid 0] \cdot \mathbf{B} = \bar{\mathbf{e}}_I^\top$ is small. By an HNF LWE assumption [6] in dimension $n_0 - 1$, the matrix $\mathbf{A}$ appears uniform even to any adversary that has $\mathsf{ask} = \mathbf{e}$.

Algorithm aEnc then encrypts regular/covert message pair $(\mu, \hat{\mu})$ as

$$\mathbf{C} = \mathbf{B} \cdot \mathbf{S} + \mathbf{E} + \mu \cdot \mathbf{G}_{[m+1]\setminus I} + \hat{\mu} \cdot \mathbf{G}_I, \tag{1}$$

where $\mathbf{S} \in \mathbb{Z}^{n\times(m+1)\lceil \log q \rceil}$ is small. The normal decryption process still outputs μ using $\mathsf{ask} = \mathbf{e}$ due to the fact that $[\mathbf{e}^\top \mid 1] \cdot \mathbf{G}_I = \mathbf{0}$. Since $[\mathbf{t}^\top \mid 0] \cdot \mathbf{G}_{[m+1]\setminus I} = \mathbf{0}$, aDec computes $[\mathbf{t}^\top \mid 0] \cdot \mathbf{C} \cdot \mathbf{v}' = \left(-\bar{\mathbf{e}}_I^\top \cdot \mathbf{S} + [\mathbf{t}^\top \mid 0] \cdot \mathbf{E}\right) \cdot \mathbf{v}' + \Delta \cdot \hat{\mu}$ for a short vector $\mathbf{v}'$ that depends only on i_{n_0}. In the latter equality, the first term of the right-hand-side member is small since $\mathbf{S}$ is small, so that aDec can decode $\hat{\mu}$.

We can also prove that multiplying two ciphertexts of the form (1) which encrypt $(\mu_1, \hat{\mu}_1)$ and $(\mu_2, \hat{\mu}_2)$ yields a product anamorphic ciphertext that normally (resp. anamorphically) decrypts to $\mu_1\mu_2$ (resp. $\hat{\mu}_1\hat{\mu}_2$).

The proof of anamorphism relies on an LWE assumption in dimension $n_0 - 1$ along with a statistical argument from [28] to show that anamorphic ciphertexts

are indistinguishable from normal ones, even when $\mathbf{e}$ is exposed. This is possible only after arguing that $\mathbf{A}$ appears uniform as mentioned above. The proof of full asymmetry involves a similar LWE assumption and the Leftover Hash Lemma that leverages the remaining min-entropy in the secret key $\mathbf{e}$ after the leakage of $\mathsf{dk} = I$. The key difference between our construction and the anamorphic GSW scheme of Catalano *et al.* [18, Section 4.3] is that the latter's double key reveals the entire secret key while we only leak a certain number of zero positions allowing to prove full asymmetry.

The scheme is shown robust with respect to homomorphic additions by exploiting the fact that aEnc samples $\mathbf{S}$ from a Gaussian distribution whereas Enc samples $\mathbf{S}$ uniformly.

1.3 Related Work

The concept of anamorphic encryption was introduced by Persiano, Phan and Yung [43]. They provided a simple rejection-sampling-based solution allowing to extend any standard PKE scheme with a receiver-anamorphic channel conveying up to $O(\log \lambda)$ bits of covert messages per ciphertext. In the same work, they showed that the well-known Naor-Yung paradigm [39] is receiver-anamorphic. They also pointed out that both the primal and dual Regev schemes are sender-anamorphic when all users share a common uniform matrix in their public keys.

Banfi *et al.* [8] gave a different generic construction from any PKE scheme where the sender and the receiver maintain a synchronized state.

Kutylowski *et al.* [31] proved that any PKE exhibiting a randomness recovery property is anamorphic. This covers the cases of RSA-OAEP [10], Goldwasser-Micali [27] and Paillier [40]. Further, [31] also introduced a distinction between multi-receiver and single-receiver anamorphism depending on whether dk allows decrypting regular messages or not (in both cases, anamorphic messages remain hidden without access to dk). Elgamal was shown [31] multi-receiver anamorphic whereas Cramer-Shoup [21] was proven single-receiver anamorphic.

The notion of fully asymmetric anamorphic encryption, where dk does not provide decryption access to either normal or covert message, was put forth by Catalano *et al.* [18]. They proved that the lite Cramer-Shoup system [21] and Naor-Yung constructions [39] can both be turned into fully asymmetric AE systems. They also showed [18, Appendix C] that full asymmetry is a stronger notion than single-receiver anamorphism. To our knowledge, their work is the only one describing AE schemes with homomorphic properties. Interestingly, they showed that lite Cramer-Shoup retains its linear homomorphism over anamorphic messages while their adaptation of Naor-Yung can be made fully homomorphic if the underlying NIZK proof is itself fully homomorphic [5]. In [18], they also proved that the GSW FHE scheme [25] is anamorphic (but not fully asymmetric).

In the context of symmetric anamorphic channels, Catalano *et al.* [18] further proved that the hybrid KEM/DEM encryption paradigm and the standard IBE-to-CCA transformation [12] both yield anamorphic encryption.

The robustness property of AE schemes was introduced in [48,8] where applying anamorphic decryption to normal ciphertexts results in an error message.

The notion of robustness in [48] goes one step further ensuring an error message when anamorphically decrypting anamorphic ciphertexts created using an incorrect double key. Wang *et al.* [48] then reformulate sender-anamorphism allowing a covert message to be sent across multiple anamorphic ciphertexts providing two constructions of robust anamorphic encryption: the first from a pseudorandom, robust PKE and the second from hybrid PKE.

Banfi *et al.* [8] also introduced anamorphic *extensions*, which allow multiple double keys to be chosen after the generation of the public key. They further showed that a property called selective randomness recovery (and satisfied by both Elgamal and Cramer-Shoup) enables robust anamorphic encryption. They separately proved that RSA-OAEP also has a robustly anamorphic extension.

Generic constructions of anamorphic encryption from standard PKE schemes are inherently limited in terms of efficiency. Catalano *et al.* [19] established that no black-box realization can have anamorphic message spaces of super-polynomial size. Recently [20], they strengthened their impossibility result by showing that it holds true for any stateless scheme regardless of the length of anamorphic messages.[2] They further showed that the rejection-sampling-based compiler of [43] actually requires the underlying PKE scheme to have high min-entropy ciphertexts unless one settles for a weaker notion of semi-adaptive security.[3] They also ruled out the existence of black-box fully asymmetric realizations even from PKE schemes with high min-entropy ciphertexts and even with semi-adaptive security. On the positive side, they proved that fully asymmetric AE with small anamorphic message space is achievable from indistinguishability obfuscation [9] if the underlying PKE has high min-entropy ciphertexts.

For the time being, very few fully asymmetric AE candidates are known and all of them suffer from certain limitations in terms of efficiency and/or additional functionalities. Indeed, the scheme of [44] is not particularly efficient and the one of [20] relies on heavy obfuscation machinery. Hence, the examples of [44,20] are mostly feasibility results and they "only" provide ordinary PKE schemes without enhanced functionalities. As for the homomorphic realizations of [18], their DDH-based scheme is restricted to polynomial-size message spaces (since the message is encoded in the exponent in a discrete-log-hard group) while their fully homomorphic extension of Naor-Yung resorts to expensive homomorphic NIZK proofs [5] for general NP statements. Finally, these constructions are only known to be instantiable under discrete-logarithm-related (and thus quantum-vulnerable) assumptions. Even if we disregard the homomorphic property, the only known post-quantum fully asymmetric scheme is obtained by instantiating the Naor-Yung-based AE scheme of [18] with LWE-based NIZK proofs [41,49].

Roadmap. Section 2 presents some background material. Our fully asymmetric dual/primal Regev constructions are described in Sections 3 and 4, respectively. Our anamorphic triplet for dual GSW is given in Section 5. The supplementary

[2] The black-box construction of [8] sidesteps the impossibility result by using a synchronized state.

[3] In this notion, the adversary is only given the secret key after the query phase in the anamorphic security game.

material is mostly dedicated to deferred proofs, except for our trapdoor-less constructions for dual Regev (in Section C), and our definitions of homomorphic robustness in Section E.

2 Background

For any $q \geq 2$, we let $\mathbb{Z}_q$ denote the ring of integers with addition and multiplication modulo q. We always set q as a prime integer. If $\mathbf{x}$ is a vector over $\mathbb{R}$, then $\|\mathbf{x}\|$ denotes its Euclidean norm, $\|\mathbf{x}\|_1$ denotes its 1-norm and $\|\mathbf{x}\|_\infty$ denotes its infinity norm. If $\mathbf{M} \in \mathbb{R}^{n \times m}$ is a matrix, then $\|\mathbf{M}\|$ denotes its operator norm $\|\mathbf{M}\| = \sup_{\|\mathbf{x}\|=1} \|\mathbf{M} \cdot \mathbf{x}\|$. For any positive integer n, we set $[n] = \{1, \ldots, n\}$.

If X and Y are distributions over the same domain, then $\Delta(X, Y)$ denotes their statistical distance. For a distribution D, we denote by $x \hookleftarrow D$ the action of sampling x from the distribution D. By $x \sim D$, we mean that x was sampled according to the distribution D. For finite set S, we denote by $U(S)$ the uniform distribution over S. We also define $\mathcal{P}$ to be the distribution over $\mathbb{Z}$ that outputs 0 with probability $1/2$ and ± 1 with probability $1/4$. As usual, we let λ denote the security parameter. Then, PPT indicates probabilistic polynomial time and we use the standard asymptotic notation $O, \omega, \Omega, \mathsf{poly}$, and negl, where $\widetilde{\cdot}$ indicates the omission of $O(\log \lambda)$ factors. Two distributions within $\mathsf{negl}(\lambda)$ statistical distance of each other are said to be statistically close.

2.1 Randomness Extraction

We first recall the Leftover Hash Lemma, as it was stated in [1].

Lemma 1 ([1]). *Let $\mathcal{H} = \{h : X \to Y\}_{h \in \mathcal{H}}$ be a family of universal hash functions and let $f : X \to Z$ be a function. Let $(T_i)_{i \leq k}$ be independent random variables over the set X, for some $k > 0$. Letting $\gamma = \max_{i \leq k} \gamma(T_i)$ where $\gamma(T_i) = \max_t \Pr[T_i = t]$, we have*

$$\Delta\big((h, (h(T_i), f(T_i))_{i \leq k}), (h, (U(Y))^{(i)}, f(T_i))_{i \leq k}))\big) \leq \frac{k}{2} \cdot \sqrt{\gamma \cdot |Y| \cdot |Z|}.$$

The Leftover Hash Lemma implies the following corollary, which is often used to re-randomize matrices over $\mathbb{Z}_q$ by multiplying them with small-norm matrices.

Lemma 2. *Take integers $m, n, k, \bar{n}$ such that $m > (n + \bar{n}) \log q + 2\lambda$, for some prime $q > 2$. Let $\mathbf{A} \hookleftarrow U(\mathbb{Z}_q^{n \times m}), \widetilde{\mathbf{A}} \hookleftarrow U(\mathbb{Z}_q^{n \times k})$ and $\mathbf{R} \hookleftarrow \mathcal{P}^{m \times k}$. For any matrix $\mathbf{F} \in \mathbb{Z}_q^{\bar{n} \times m}$, the distributions $(\mathbf{A}, \mathbf{A} \cdot \mathbf{R}, \mathbf{F} \cdot \mathbf{R})$ and $(\mathbf{A}, \widetilde{\mathbf{A}}, \mathbf{F} \cdot \mathbf{R})$ are within statistical distance $k \cdot 2^{-\lambda}$.*

2.2 Lattices and Discrete Gaussian Distributions

An n-dimensional lattice $\Lambda \subseteq \mathbb{R}^n$ is the set $\Lambda = \{\sum_{i=1}^n z_i \cdot \mathbf{b}_i \mid \mathbf{z} \in \mathbb{Z}^n\}$ of all integer linear combinations of a set of linearly independent basis vectors

$\mathbf{B} = \{\mathbf{b}_1, \ldots, \mathbf{b}_n\} \subseteq \mathbb{R}^n$. The dual of a lattice Λ is defined to be $\hat{\Lambda} = \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{y}^\top \cdot \mathbf{x} \in \mathbb{Z} \ \forall\ \mathbf{y} \in \Lambda\}$. For a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, we define the lattices $\Lambda_q^\perp(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{x} = \mathbf{0} \bmod q\}$ and $\Lambda(\mathbf{A}) = \mathbf{A}^\top \cdot \mathbb{Z}^n + q\mathbb{Z}^m$.

Discrete Gaussian Distributions. Let $\mathbf{\Sigma} \in \mathbb{R}^{n \times n}$ be a symmetric positive definite matrix, and let $\mathbf{c} \in \mathbb{R}^n$. We define the Gaussian function on $\mathbf{x} \in \mathbb{R}^n$ by $\rho_{\mathbf{\Sigma},\mathbf{c}}(\mathbf{x}) = \exp(-\pi(\mathbf{x}-\mathbf{c})^\top \mathbf{\Sigma}^{-1}(\mathbf{x}-\mathbf{c}))$. In the case $\mathbf{\Sigma} = \sigma^2 \cdot \mathbf{I}_n$ for a real value $\sigma > 0$ and $\mathbf{c} = \mathbf{0}$, we denote by ρ_σ the Gaussian function $\rho_\sigma(\mathbf{x}) = \exp(-\pi\|\mathbf{x}\|^2/\sigma^2)$ with standard deviation σ. For any lattice $\Lambda \subset \mathbb{R}^n$, the discrete Gaussian distribution $D_{\Lambda,\sigma}$ has probability mass $\Pr_{X \sim D_{\Lambda,\sigma}}[X = \mathbf{x}] = \frac{\rho_\sigma(\mathbf{x})}{\rho_\sigma(\Lambda)}$ for any $\mathbf{x} \in \Lambda$, where we define $\rho_\sigma(\Lambda) = \sum_{\mathbf{y} \in \Lambda} \rho_\sigma(\mathbf{y})$. For a lattice Λ, the smoothing parameter $\eta_\varepsilon(\Lambda)$ [38] is defined to be the smallest $r > 0$ such that $\rho_{1/r}(\hat{\Lambda} \setminus \mathbf{0}) \leq \varepsilon$ with $\hat{\Lambda}$ denoting the dual of Λ, for any $\varepsilon \in (0,1)$. By [38, Lemma 3.3], we have $\eta_{2^{-n}}(\mathbb{Z}^n) \leq \sqrt{(\ln(2n(1+2^n)))/\pi} < \sqrt{n}$.

Lemma 3 ([7, Lemma 1.5],[34, Lemma 4.4]).

1. *For any $k > 0$, $\Pr[|z| > k \cdot \sigma \mid z \hookleftarrow D_{\mathbb{Z},\sigma}] \leq 2 \cdot \exp(-k^2/2)$.*
2. *For any $k > 1$, $\Pr[\|\mathbf{z}\| > k \cdot \sigma\sqrt{m} \mid z \hookleftarrow D_{\mathbb{Z}^m,\sigma}] \leq k^m \cdot \exp(\frac{m}{2}(1-k^2))$.*

In our security proofs, we rely on a noise randomization technique introduced by Katsumata and Yamada [28].

Lemma 4 ([28, Lemma 1]). *Let q, m, t be positive integers and r a positive real satisfying $r > \max(\omega(\sqrt{\log m}), \omega(\sqrt{\log t}))$. Let $\mathbf{b} \in \mathbb{Z}_q^m$ and $\mathbf{x} \hookleftarrow D_{\mathbb{Z}^m,r}$. Then, for any $\mathbf{V} \in \mathbb{Z}^{m \times t}$ and any positive real $s > \|\mathbf{V}\|$, there exists a PPT algorithm $\mathsf{ReRand}(\mathbf{V}, \mathbf{b}+\mathbf{x}, r, s)$ that outputs $\mathbf{b}' = \mathbf{V}^\top \mathbf{b} + \mathbf{x}' \in \mathbb{Z}_q^t$ where $\mathbf{x}'$ is distributed statistically close to $D_{\mathbb{Z}^t,2rs}$.*

We also rely on the following lemma proven by Boneh and Freeman [13].

Lemma 5 ([13, Lemma 4.12]). *Let $\Lambda_1, \Lambda_2 \subseteq \mathbb{Z}^m$ be full-rank lattices, $\sigma_1, \sigma_2 \in \mathbb{R}$ and take independent random variables $X \sim D_{\Lambda_1,\sigma_1}$ and $Y \sim D_{\Lambda_2,\sigma_2}$. Define $\tau = \frac{\sigma_1\sigma_2}{\sqrt{\sigma_1^2+\sigma_2^2}} = \left(\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2}\right)^{-1/2}$ and suppose that $\tau \geq \eta_\epsilon(\Lambda_1 \cap \Lambda_2)$ for some negligible ϵ. Then, the random variable $Z = X+Y$ is a sample from a distribution statistically close to $D_{\Lambda_1+\Lambda_2,\sqrt{\sigma_1^2+\sigma_2^2}}$.*

The LWE Assumption. We now recall the Learning With Errors problem [45].

Definition 1. *Let $\lambda \in \mathbb{N}$ be a security parameter and take integers $n = n(\lambda)$, $m = m(\lambda)$, $q = q(\lambda)$. Let $\chi = \chi(\lambda)$ be an efficiently sampleable distribution over $\mathbb{Z}_q$. The $\mathsf{LWE}_{n,m,q,\chi}$ assumption posits that the following distance is a negligible function for any PPT algorithm $\mathcal{A}$:*

$$\begin{aligned}\mathbf{Adv}^{\mathcal{A},\mathsf{LWE}}_{n,m,q,\chi}(\lambda) := \big|\Pr[\mathcal{A}(1^\lambda, \mathbf{A}, \mathbf{u}) = 1 \mid \mathbf{A} \hookleftarrow U(\mathbb{Z}_q^{m \times n}), \mathbf{u} \hookleftarrow U(\mathbb{Z}_q^m)] \\ - \Pr[\mathcal{A}(1^\lambda, \mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e}) = 1 \mid \mathbf{A} \hookleftarrow U(\mathbb{Z}_q^{m \times n}), \mathbf{s} \hookleftarrow U(\mathbb{Z}_q^n), \mathbf{e} \hookleftarrow \chi^m]\big|.\end{aligned}$$

A typical choice for χ is the integer Gaussian distribution $D_{\mathbb{Z},\alpha q}$ for some parameter $\alpha \in (\sqrt{n}/q, 1)$. In particular, choosing $\alpha q > 2\sqrt{n}$ allows for quantum reductions from standard lattice problems with approximation factor $\gamma = \widetilde{O}(n/\alpha) = \widetilde{O}(\sqrt{n}q)$ to LWE (see, e.g., [45,15]). The best lattice algorithms for approximation factor γ run in time at least $2^{\widetilde{\Omega}(n/\log\gamma)}$ [47]. In the following, we sometimes rely on the hardness of LWE in HNF form denoted $\mathsf{HNF\text{-}LWE}_{n,m,q,\chi}$, where the secret $\mathbf{s}$ is sampled from the distribution χ^n. As shown in [6], this variant is as hard as the standard LWE problem. When using alternative secret distributions $\mathcal{S}$, we denote the corresponding LWE problems by $\mathsf{LWE}^{\mathcal{S}}_{n,m,q,\chi}$.

Lattice Trapdoors. Micciancio and Peikert [37] described a lattice trapdoor mechanism that simplifies [24]. They use a "gadget" matrix $\mathbf{G} = \mathbf{I}_n \otimes (1, 2, 4, \ldots, 2^{\lfloor \log q \rfloor}) \in \mathbb{Z}_q^{n \times n\lceil \log q \rceil}$ for which anyone can publicly sample short vectors $\mathbf{x} \in \mathbb{Z}^m$ such that $\mathbf{G} \cdot \mathbf{x} = \mathbf{0}$. The function $\mathbf{G}^{-1} : \mathbb{Z}_q^n \to \{0,1\}^{n \cdot \lceil \log q \rceil}$ maps a vector $\mathbf{v} \in \mathbb{Z}_q^n$ to a binary vector $\mathbf{w} = \mathbf{G}^{-1}(\mathbf{v})$ such that $\mathbf{G} \cdot \mathbf{w} = \boldsymbol{v}$ using binary decomposition. This definition can be extended to matrices by applying $\mathbf{G}^{-1}$ to each column in turn. As in [37], we call $\mathbf{R} \in \mathbb{Z}^{\bar{m} \times n\lceil \log q \rceil}$ a $\mathbf{G}$-trapdoor for a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times (\bar{m} + n\lceil \log q \rceil)}$ if $\mathbf{A} \cdot [-\mathbf{R}^\top \mid \mathbf{I}_{n \cdot \lceil \log q \rceil}]^\top = \mathbf{H} \cdot \mathbf{G}$ for some invertible matrix $\mathbf{H} \in \mathbb{Z}_q^{n \times n}$. Moreover, if $\mathbf{H} \in \mathbb{Z}_q^{n \times n}$ is invertible, then $\mathbf{R}$ allows inverting the LWE function $(\mathbf{s}, \mathbf{e}) \mapsto \mathbf{A}^\top \cdot \mathbf{s} + \mathbf{e}$, for any $\mathbf{s} \in \mathbb{Z}_q^n$ and any sufficiently short $\mathbf{e} \in \mathbb{Z}^{\bar{m}+nk}$, where $k = \lceil \log q \rceil$.

Lemma 6 ([37, Section 5]). *Let $\bar{m} \geq n \log q + 2\lambda$ and $k = \lceil \log q \rceil$. There exists a PPT algorithm $\mathsf{GenTrap}$ that inputs matrices $\bar{\mathbf{A}} \sim U(\mathbb{Z}_q^{n \times \bar{m}})$, $\mathbf{H} \in \mathbb{Z}_q^{n \times n}$ and outputs matrices $\mathbf{R} \sim \mathcal{P}^{\bar{m} \times nk}$ and $\mathbf{A} = \left[\bar{\mathbf{A}} \mid \bar{\mathbf{A}}\mathbf{R} + \mathbf{H}\mathbf{G}\right] \in \mathbb{Z}_q^{n \times (\bar{m}+nk)}$ such that the distribution of $\mathbf{A}$ is within statistical distance $2^{-\Omega(\lambda)}$ from the uniform distribution $U(\mathbb{Z}_q^{n \times (\bar{m}+nk)})$. Moreover, if $\mathbf{H} \in \mathbb{Z}_q^{n \times n}$ is invertible, there exists a deterministic polynomial time algorithm Invert that takes as inputs $\mathbf{R} \sim \mathcal{P}^{\bar{m} \times nk}$, $\mathbf{H} \in \mathbb{Z}_q^{n \times n}$ and a vector*

$$\begin{pmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \end{pmatrix} = \mathbf{A}^\top \cdot \mathbf{s} + \begin{pmatrix} \mathbf{e}_1 \\ \mathbf{e}_2 \end{pmatrix}, \tag{2}$$

where $\mathbf{s} \in \mathbb{Z}_q^n$, $\mathbf{e}_1 \in \mathbb{Z}^{\bar{m}}$ and $\mathbf{e}_2 \in \mathbb{Z}^{nk}$, and outputs $\mathbf{s}$ and $\mathbf{e} = (\mathbf{e}_1^\top \mid \mathbf{e}_2^\top)^\top$ as long as $\|\mathbf{e}_2 - \mathbf{R}^\top \mathbf{e}_1\| \leq q/(2\sqrt{k})$.

The Invert algorithm of Lemma 6 builds on the observation [37] that the lattice $\Lambda_q^\perp(\mathbf{G})$ has a public trapdoor [24] (i.e., a matrix $\mathbf{S} \in \mathbb{Z}^{nk \times nk}$ of norm $\|\mathbf{S}\| \leq \max(\sqrt{5}, \sqrt{k})$ which is invertible over $\mathbb{Q}$ and such that $\mathbf{G} \cdot \mathbf{S} = 0 \bmod q$) that makes it possible to compute $(\mathbf{s}, \mathbf{e})$ from $\mathbf{y} = \mathbf{G}^\top \mathbf{s} + \mathbf{e}$ as long as $\|\mathbf{e}\| < q/(2\sqrt{k})$. In short, given a vector $\mathbf{y} = (\mathbf{y}_1^\top \mid \mathbf{y}_2^\top)^\top$ of the form (2), Invert first computes

$$\mathbf{y}_2 - \mathbf{R}^\top \mathbf{y}_1 = \mathbf{G}^\top \mathbf{H}^\top \mathbf{s} + (\mathbf{e}_2 - \mathbf{R}^\top \mathbf{e}_1) \bmod q$$

and then $\mathbf{S}^\top \cdot (\mathbf{y}_2 - \mathbf{R}^\top \mathbf{y}_1) \bmod q = \mathbf{S}^\top \cdot (\mathbf{e}_2 - \mathbf{R}^\top \mathbf{e}_1) \bmod q$. Since $\|\mathbf{S}\| \leq \sqrt{k}$ and $\|\mathbf{e}_2 - \mathbf{R}^\top \mathbf{e}_1\| \leq q/(2\sqrt{k})$, the right-hand-side member of the latter equality

is actually $\mathbf{S}^\top \cdot (\mathbf{e}_2 - \mathbf{R}^\top \mathbf{e}_1)$ over $\mathbb{Z}$, which allows computing $(\mathbf{e}_2 - \mathbf{R}^\top \mathbf{e}_1)$ since $\mathbf{S}$ has full rank over $\mathbb{Q}$. In turn, this allows computing $\mathbf{s}$ from $\mathbf{G}^\top \mathbf{H}^\top \mathbf{s}$.

2.3 Other Useful Lemmas

Message Encoding. As in [51], we use a pair of algorithms $(\mathsf{encode}_d, \mathsf{decode}_d)$ parameterized by integers (n, q, d) such that $\mathsf{encode}_d : \mathbb{Z}_d^n \to \mathbb{Z}_q^n$ maps any $\mathbf{v} \in \mathbb{Z}_d^n$ to $\mathsf{encode}_d(\mathbf{v}) = (v_1 \cdot \lfloor \frac{q}{d} \rceil, \ldots, v_n \cdot \lfloor \frac{q}{d} \rceil)$ while $\mathsf{decode}_d : \mathbb{Z}_q^n \to \mathbb{Z}_d^n$ maps $\mathbf{u} \in \mathbb{Z}_q^n$ to $\mathsf{decode}_d(\mathbf{u}) = (\lfloor u_1 \cdot \frac{d}{q} \rceil, \ldots, \lfloor u_n \cdot \frac{d}{q} \rceil)$. As shown by Lemma 7, decode_d undoes the encoding of encode_d for a suitable parameter choice.

Lemma 7 ([51, Lemma 7]). *Let n, q, d positive integers such that $2 \leq d \leq \sqrt{q}$. Then, for any $\mathbf{v} \in \mathbb{Z}_d^n$ and any $\mathbf{e} \in \mathbb{Z}^n$ such that $\|\mathbf{e}\|_\infty < \frac{q-(d-1)d}{2d}$, we have $\mathbf{v} = \mathsf{decode}_d(\mathsf{encode}_d(\mathbf{v}) + \mathbf{e})$.*

Norm of a Random Matrix. We need the following upper bound on the norm of a random matrix over $\{-1, 1\}$.

Lemma 8 ([1, Lemma 15],[33, Fact 2.4]). *Let $\mathbf{R} \sim U(\{-1, 1\})^{m \times n}$ be a random matrix. There exists a universal constant $C > 0$ such that we have $\Pr_{\mathbf{R}}[\|\mathbf{R}\| > C\sqrt{m+n}] \leq \exp(-(m+n))$*

As pointed out in [1, Lemma 15], the constant C is at most 12. As a corollary, the same bound holds when the matrix $\mathbf{R}$ is sampled from the distribution $\mathcal{P}$, as shown in Supplementary Material A.1.

2.4 Definitions for Anamorphic Encryption

We now recall the syntax and security definitions of anamorphic encryption [43]. In particular, we consider the fully asymmetric property defined by Catalano *et al.* [18] and the robustness property introduced by Banfi *et al.* [8].

An anamorphic encryption scheme consists of a public-key encryption (PKE) scheme $\Pi = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ where the key generation, encryption and decryption algorithms proceed in the usual way. In addition, it admits an anamorphic triplet $\Sigma = (\mathsf{aGen}, \mathsf{aEnc}, \mathsf{aDec})$ with the following syntax.

Definition 2. *A triplet $\Sigma = (\mathsf{aGen}, \mathsf{aEnc}, \mathsf{aDec})$ is* **anamorphic** *if:*

- aGen *is a PPT algorithm that takes as input a security parameter $\lambda \in \mathbb{N}$ and outputs an anamorphic public key apk, an anamorphic secret key ask, a (possibly empty) trapdoor key tk, and a double key dk*
- aEnc *is a PPT algorithm that takes as input apk and dk, a real message $\mu \in \mathcal{M}$, and a covert message $\hat{\mu} \in \widehat{\mathcal{M}}$ and outputs a ciphertext act.*
- aDec *is a deterministic algorithm that inputs dk, tk, ask and an anamorphic ciphertext act produced by aEnc. It outputs the anamorphic message $\hat{\mu} \in \widehat{\mathcal{M}}$ or a special symbol $\perp \notin \widehat{\mathcal{M}}$.*

Definition 3. *A PKE scheme* $\Pi = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ *is* **anamorphic** *if it provides IND-CPA security and there exists an anamorphic triplet* Σ *such that, for any PPT dictator* $\mathcal{D}$*, there exists a negligible function* $\nu(\lambda)$ *such that*

$$\mathbf{Adv}^{\mathsf{Anamorphism}}_{\mathcal{D},\Pi,\Sigma}(\lambda) = |\Pr[\mathsf{RealG}_{\Pi}(\lambda,\mathcal{D}) = 1] - \Pr[\mathsf{AnamorphicG}_{\Sigma}(\lambda,\mathcal{D}) = 1]| \leq \nu(\lambda)$$

where the real and ideal experiments are defined as follows:

$\mathsf{RealG}_{\Pi}(\lambda,\mathcal{D})$:
1. $(\mathsf{pk},\mathsf{sk}) \leftarrow \mathsf{KGen}(1^{\lambda})$;
2. *Return* $\mathcal{D}^{\mathcal{O}_e(\mathsf{pk},\cdot,\cdot)}(\mathsf{pk},\mathsf{sk})$

where $\mathcal{O}_e(\mathsf{pk},\mu,\hat{\mu}) = \mathsf{Enc}(\mathsf{pk},\mu)$

$\mathsf{AnamorphicG}_{\Sigma}(\lambda,\mathcal{D})$
1. $((\mathsf{apk},\mathsf{ask}),\mathsf{tk},\mathsf{dk}) \leftarrow \mathsf{aGen}(1^{\lambda})$;
2. *Return* $\mathcal{D}^{\mathcal{O}_a(\mathsf{apk},\cdot,\cdot)}(\mathsf{apk},\mathsf{ask})$

where $\mathcal{O}_a(\mathsf{apk},\mu,\hat{\mu}) = \mathsf{aEnc}(\mathsf{apk},\mathsf{dk},\mu,\hat{\mu})$

Catalano *et al.* [18] introduced a property called *fully asymmetric* which ensures that, in the anamorphic mode, the double key dk can be used as an asymmetric encryption key that allows encrypting without necessarily being able to decrypt. The pair $(\mathsf{dk},\mathsf{tk})$ can thus serve as an anamorphic counterpart of the asymmetric key pair $(\mathsf{apk},\mathsf{ask})$.

Definition 4. *An anamorphic PKE scheme* $\Pi = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ *equipped with an anamorphic triplet* $\Sigma = (\mathsf{aGen}, \mathsf{aEnc}, \mathsf{aDec})$ *is* **fully asymmetric** *if, for any PPT adversary* $\mathcal{A}$*, there exists a negligible function* $\nu(\lambda)$ *such that*

$$\mathbf{Adv}^{\mathsf{FAsym}}_{\mathcal{A},\Pi,\Sigma}(\lambda) = |\Pr[\mathsf{Fasym}^0_{\mathcal{A},\Pi,\Sigma}(\lambda) = 1] - \Pr[\mathsf{Fasym}^1_{\mathcal{A},\Pi,\Sigma}(\lambda) = 1]| \leq \nu(\lambda)$$

where the experiments $\mathsf{Fasym}^0_{\mathcal{A},\Pi,\Sigma}(\lambda)$ *and* $\mathsf{Fasym}^1_{\mathcal{A},\Pi,\Sigma}(\lambda)$ *are defined as follows.*

$\mathsf{Fasym}^b_{\mathcal{A},\Pi,\Sigma}(\lambda)$:
1. $((\mathsf{apk},\mathsf{ask}),\mathsf{tk},\mathsf{dk}) \leftarrow \mathsf{aGen}(1^{\lambda})$
2. $(\mu_0,\mu_1,\hat{\mu}_0,\hat{\mu}_1) \leftarrow \mathcal{A}(\mathsf{apk},\mathsf{dk})$
3. $\mathsf{act} \leftarrow \mathsf{aEnc}(\mathsf{apk},\mathsf{dk},\mu_b,\hat{\mu}_b)$
4. *Return* $\mathcal{A}(\mathsf{act})$

Banfi *et al.* [8] introduced a notion of *robustness* for anamorphic encryption. Informally speaking, it captures the infeasibility of finding a message that anamorphically decrypts to $\hat{\mu} \neq \perp$ when encrypted normally. This property is formalized by means of an indistinguishability requirement between two oracles. The first oracle always runs the anamorphic decryption algorithm on a normal encryption of an input plaintext μ. The second one always returns $\perp$.

Definition 5. *An anamorphic PKE scheme* $\Pi = (\mathsf{ParGen}, \mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ *endowed with an anamorphic triplet* $\Sigma = (\mathsf{aGen}, \mathsf{aEnc}, \mathsf{aDec})$ *is* **robust** *if, for any PPT adversary* $\mathcal{A}$*, there exists a negligible function* $\nu(\lambda)$ *such that*

$$\mathbf{Adv}^{\mathsf{rob}}_{\mathcal{A},\Pi,\Sigma}(\lambda) = |\Pr[\mathsf{Robust}^0_{\mathcal{A},\Pi,\Sigma}(\lambda) = 1] - \Pr[\mathsf{Robust}^1_{\mathcal{A},\Pi,\Sigma}(\lambda) = 1]| \leq \nu(\lambda)$$

where $\mathsf{Robust}^0_{\mathcal{A},\Pi,\Sigma}(\lambda)$ *and* $\mathsf{Robust}^1_{\mathcal{A},\Pi,\Sigma}(\lambda)$ *are defined as follows.*

$\mathsf{Robust}^b_{\mathcal{A},\Pi,\Sigma}(\lambda)$:
1. $((\mathsf{apk},\mathsf{ask}),\mathsf{tk},\mathsf{dk}) \leftarrow \mathsf{aGen}(1^{\lambda})$;
2. *Return* $\mathcal{A}^{\mathcal{O}_b(\mathsf{apk},\mathsf{ask},\mathsf{dk},\mathsf{tk},\cdot)}(\mathsf{apk},\mathsf{ask})$
   where $\mathcal{O}_0(\mathsf{apk},\mathsf{ask},\mathsf{dk},\mathsf{tk},\mu) = \mathsf{aDec}(\mathsf{dk},\mathsf{tk},\mathsf{ask},\mathsf{Enc}(\mathsf{apk},\mu))$
   and $\mathcal{O}_1(\mathsf{apk},\mathsf{ask},\mathsf{dk},\mathsf{tk},\mu) = \perp$

We note that the above definition only considers robustness for fresh ciphertexts. If the underlying encryption scheme is homomorphic, Definition 5 does not say anything about ciphertexts obtained by e.g. adding normal ciphertexts. In Supplementary Material E, we generalize the definition of robustness to ciphertexts that have undergone homomorphic operations.

2.5 (Anamorphic) Homomorphic Encryption

For completeness, we recall the notion of (somewhat) homomorphic public-key encryption. We use the definition from [16], but assume that the evaluation key is part of the public key output by KGen for simplicity.

Definition 6. *Let $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda\in\mathcal{N}}$ be a class of circuits and take a PKE scheme $\Pi = (\mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$. The tuple of algorithms*

$$\mathsf{HE} = (\mathsf{HE.KGen} = \mathsf{KGen}, \mathsf{HE.Enc} = \mathsf{Enc}, \mathsf{HE.Dec} = \mathsf{Dec}, \mathsf{HE.Eval})$$

is **$\mathcal{C}$-homomorphic** *if for any sequence of circuits $\mathsf{C}_\lambda \in \mathcal{C}_\lambda$, and respective inputs $\mu_1, \ldots, \mu_\ell \in \mathcal{M}$, there exists a negligible function $\nu(\lambda)$ such that*

$$\Pr[\mathsf{HE.Dec}(\mathsf{sk}, \mathsf{HE.Eval}(\mathsf{pk}, \mathsf{C}_\lambda, (c_1, \ldots, c_\ell))) \neq \mathsf{C}_\lambda(\mu_1, \ldots, \mu_\ell)] \leq \nu(\lambda)$$

where $\mathsf{pk}, \mathsf{sk} \hookleftarrow \mathsf{HE.KGen}(1^\lambda)$ and $c_i \hookleftarrow \mathsf{HE.Enc}(\mathsf{pk}, \mu_i)$.

Finally, as in [18], we extend the anamorphism definition to capture homomorphic encryption schemes, by ensuring that the homomorphic properties are retained over the covert message space $\widehat{\mathcal{M}}$.

Definition 7. *A $\mathcal{C}$-homomorphic encryption scheme $\mathsf{HE} = (\mathsf{HE.KGen}, \mathsf{HE.Enc}, \mathsf{HE.Dec}, \mathsf{HE.Eval})$ is* **anamorphic** *if:*

- $(\mathsf{HE.KGen}, \mathsf{HE.Enc}, \mathsf{HE.Dec})$ *is an anamorphic PKE scheme with anamorphic triplet* $(\mathsf{HE.aGen}, \mathsf{HE.Enc}, \mathsf{HE.Dec})$
- HE *has* **$\mathcal{C}$-homomorphic correctness for covert messages**. *That is, for any sequence of circuits $\mathsf{C}_\lambda \in \mathcal{C}_\lambda$, and respective inputs $\mu_1, \ldots, \mu_\ell \in \mathcal{M}$, $\hat{\mu}_1, \ldots, \hat{\mu}_\ell \in \widehat{\mathcal{M}}$, there exists a negligible function $\nu(\lambda)$ such that*

  $$\Pr\left[\mathsf{HE.aDec}\big(\mathsf{dk}, \mathsf{tk}, \mathsf{ask}, \mathsf{HE.Eval}(\mathsf{apk}, \mathsf{C}_\lambda, (c_1, \ldots, c_\ell))\big) \neq \mathsf{C}_\lambda(\hat{\mu}_1, \ldots, \hat{\mu}_\ell)\right] \leq \nu(\lambda)$$

  where $\mathsf{apk}, \mathsf{ask}, \mathsf{tk}, \mathsf{dk} \hookleftarrow \mathsf{HE.aGen}(1^\lambda)$ and $c_i \hookleftarrow \mathsf{HE.aEnc}(\mathsf{pk}, \mathsf{dk}, \mu_i, \hat{\mu}_i)$.

3 Public-Key Anamorphic Linearly Homomorphic Encryption from Dual Regev

The dual Regev system [24] was shown to be sender-anamorphic [43] when all users share a common matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$. We show that it is receiver-anamorphic and fully asymmetric when public keys contain a user-specific matrix $\mathbf{A}$.

To do this, our idea is to embed the anamorphic message in the high-order

bits of the LWE secret $\mathbf{s}$ as suggested by [51] in a different context. The anamorphic trapdoor key tk consists of a lattice trapdoor for the matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, which makes it possible to recover the anamorphic message. At the same time, we can prove that the anamorphic message remains invisible to a dictator even when the normal decryption key is exposed.

To our knowledge, we thus obtain the first fully asymmetric linearly homomorphic AE scheme based on a post-quantum assumption. It is also the first one where the expansion rate (i.e., the ratio between the ciphertext size and the length of regular/anamorphic plaintexts) is only poly-logarithmic in λ. In the description hereunder, anamorphic messages are as long as regular plaintexts.[4]

Note that the scheme is *public-key anamorphic*, as defined in [44]. Indeed, the double key dk is empty and anyone can run aEnc using only apk. The scheme can thus be set up in anamorphic mode without requiring any initial secret communication where the key owner would discreetly provide encryptors with a double key enabling covert communication. Since the scheme can be made robust (see Supplementary Material E.2), it solves an open question by Persiano *et al.* [44], which is to simultaneously provide public-key anamorphism and robustness.

$\mathsf{KGen}(1^\lambda)$: Given a security parameter 1^λ,

1. Choose dimensions $n, \bar{m} \in \mathsf{poly}(\lambda)$, a plaintext modulus $p \in \mathsf{poly}(\lambda)$, a ciphertext modulus $q \in \mathsf{poly}(\lambda)$ such that $q > p$, $\bar{m} \geq n \log q + 2\lambda$ and $k = \lceil \log q \rceil$. Choose an error rate $\alpha \in (0,1)$ and a standard deviation $\sigma > \alpha q$. Define parameters $\mathsf{par} := \Big(q, p, n, \bar{m}, \alpha, \sigma\Big)$.
2. Choose a random matrix $\mathbf{A} \hookleftarrow \mathbb{Z}_q^{n \times m}$, where $m = \bar{m} + nk$.
3. Sample[5] $\mathbf{E} \hookleftarrow \mathcal{P}^{m \times n}$ and compute $\mathbf{U} = \mathbf{A} \cdot \mathbf{E} \in \mathbb{Z}_q^{n \times n}$.

Return $\mathsf{sk} = \mathbf{E} \in \{-1, 0, 1\}^{m \times n}$ and $\mathsf{pk} = (\mathsf{par}, \mathbf{A}, \mathbf{U})$.

$\mathsf{aGen}(1^\lambda)$: Given a security parameter 1^λ,

1. Generate parameters $\mathsf{par} := \Big(q, p, n, \bar{m}, \alpha, \sigma\Big)$ as in step 1 of KGen.
2. Generate a statistically uniform matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, where $m = \bar{m} + nk$, together with a trapdoor $\mathbf{R}$ by applying Lemma 6. Namely, choose a random matrix $\bar{\mathbf{A}} \hookleftarrow U(\mathbb{Z}_q^{n \times \bar{m}})$. Then, sample $\mathbf{R} \hookleftarrow \mathcal{P}^{\bar{m} \times nk}$ and compute
$$\mathbf{A} = [\bar{\mathbf{A}} \mid \bar{\mathbf{A}} \cdot \mathbf{R} + \mathbf{G}].$$
3. Sample $\mathbf{E} \hookleftarrow \mathcal{P}^{m \times n}$ and compute $\mathbf{U} = \mathbf{A} \cdot \mathbf{E} \bmod q$.

Return $\mathsf{ask} = \mathbf{E} \in \{-1, 0, 1\}^{m \times n}$, $\mathsf{apk} = (\mathsf{par}, \mathbf{A}, \mathbf{U})$, $\mathsf{dk} = \emptyset$ and $\mathsf{tk} = \mathbf{R} \in \{-1, 0, 1\}^{\bar{m} \times k}$.

[4] The length of normal plaintexts can nevertheless be increased as in the packed dual Regev scheme of [24]. Then, the bandwidth rate becomes smaller than 1.

[5] The scheme and proofs work just as well if $\mathbf{E} \sim U(\{-1, 1\}^{m \times n})$ but, in the trapdoorless variant and in the dual GSW case of Section 5, we need a distribution containing zero in its support. The proof of Theorem 1 carries over to the case of binary secret keys $\mathbf{E} \sim U(\{0, 1\}^{m \times n})$ but requires a larger σ due to a larger upper bound on $\|\mathbf{E}\|$.

Enc$(\mathsf{pk}, \boldsymbol{\mu})$**:** Given the public key $\mathsf{pk} = (\mathsf{par}, \mathbf{A}, \mathbf{U})$ and a message $\boldsymbol{\mu} \in \mathbb{Z}_p^n$, sample $\mathbf{s} \hookleftarrow U(\mathbb{Z}_q^n)$, $\mathbf{e}_0 \hookleftarrow D_{\mathbb{Z}^m,\sigma}$ and $\mathbf{e}_1 \hookleftarrow D_{\mathbb{Z}^n,\sigma}$. Set

$$\mathbf{c}_0 = \mathbf{A}^\top \cdot \mathbf{s} + \mathbf{e}_0 \in \mathbb{Z}_q^m, \qquad \mathbf{c}_1 = \mathbf{U}^\top \cdot \mathbf{s} + \mathbf{e}_1 + \Delta \cdot \boldsymbol{\mu} \in \mathbb{Z}_q^n,$$

where $\Delta = \lfloor q/p \rceil$. Output the ciphertext $\mathsf{ct} = (\mathbf{c}_0, \mathbf{c}_1) \in \mathbb{Z}_q^{(m+n)}$.

aEnc$(\mathsf{apk}, \mathsf{dk}, \boldsymbol{\mu}, \hat{\boldsymbol{\mu}})$**:** Given $\mathsf{apk} = (\mathsf{par}, \mathbf{A}, \mathbf{U})$, $\mathsf{dk} = \emptyset$, and messages $\boldsymbol{\mu} \in \mathbb{Z}_p^n$ and $\hat{\boldsymbol{\mu}} \in \mathbb{Z}_p^n$, sample $\mathbf{s} \hookleftarrow D_{\mathbb{Z}^m,\alpha q}$, $\mathbf{e}_0 \hookleftarrow D_{\mathbb{Z}^m,\sigma}$ and $\mathbf{e}_1 \hookleftarrow D_{\mathbb{Z}^n,\sigma}$. Then,

1. Compute $\hat{\mathbf{s}} = \mathbf{s} + \mathsf{encode}_p(\hat{\boldsymbol{\mu}}) = \mathbf{s} + \Delta \cdot \hat{\boldsymbol{\mu}} \in \mathbb{Z}_q^n$ where $\mathsf{encode}_p : \mathbb{Z}_p^n \to \mathbb{Z}_q^n$ is the encoding algorithm of Section 2.3 and $\Delta = \lfloor q/p \rfloor$.
2. Compute

$$\mathbf{c}_0 = \mathbf{A}^\top \cdot \hat{\mathbf{s}} + \mathbf{e}_0 \in \mathbb{Z}_q^m, \qquad \mathbf{c}_1 = \mathbf{U}^\top \cdot \hat{\mathbf{s}} + \mathbf{e}_1 + \Delta \cdot \boldsymbol{\mu} \in \mathbb{Z}_q^n,$$

Output the anamorphic ciphertext $\mathsf{act} = (\mathbf{c}_0, \mathbf{c}_1) \in \mathbb{Z}_q^{(m+n)}$.

Dec$(\mathsf{pk}, \mathsf{sk}, \mathsf{ct})$**:** Given a ciphertext $\mathsf{ct} = (\mathbf{c}_0, \mathbf{c}_1) \in \mathbb{Z}_q^{(m+n)}$ and the secret key $\mathsf{sk} = \mathbf{E} \in \{-1, 0, 1\}^{m \times n}$, compute and return $\boldsymbol{\mu} = \lfloor (\mathbf{c}_1 - \mathbf{E}^\top \cdot \mathbf{c}_0)/\Delta \rceil$.

aDec$(\mathsf{dk}, \mathsf{tk}, \mathsf{ask}, \mathsf{act})$**:** Given a ciphertext $\mathsf{act} = (\mathbf{c}_0, \mathbf{c}_1) \in \mathbb{Z}_q^{(m+n)}$ and the trapdoor key $\mathsf{tk} = \mathbf{R} \in \{-1, 0, 1\}^{\bar{m} \times k}$,

1. Compute $(\hat{\mathbf{s}}, \mathbf{e}_0) \leftarrow \mathsf{Invert}(\mathbf{R}, \mathbf{c}_0)$ using the Invert algorithm of Lemma 6. If $(\hat{\mathbf{s}}, \mathbf{e}_0) = (\perp, \perp)$ (meaning there exists no $\mathbf{e}_0 = (\mathbf{e}_{0,1}^\top \mid \mathbf{e}_{0,2}^\top)^\top \in \mathbb{Z}^m$ such that $\|\mathbf{e}_{0,2} - \mathbf{R}^\top \mathbf{e}_{0,1}\| \leq q/(2\sqrt{k})$ and $\mathbf{c}_0 = \mathbf{A}^\top \hat{\mathbf{s}} + \mathbf{e}_0$ for some $\hat{\mathbf{s}} \in \mathbb{Z}_q^n$), return $\perp$.
2. Output $\hat{\boldsymbol{\mu}} = \mathsf{decode}_p(\hat{\mathbf{s}}) \in \mathbb{Z}_p^n$ using the decode algorithm of Section 2.3.

PARAMETERS. Lemma 4 requires $\alpha q = \omega(\sqrt{\log m})$ and we choose $\alpha q = \Omega(\sqrt{n})$ so as to also satisfy the condition $\alpha q > 2\sqrt{n}$ for the hardness of LWE. We choose $\sigma = 26 \cdot \alpha q \sqrt{m+n} = O(m)$ to satisfy the hypothesis of Theorem 1.

In order for aDec to decrypt properly, we need to make sure that $\|\mathbf{e}_0\|_\infty < \frac{q-(p-1)p}{2p}$ due to the constraint of Lemma 7. Also, in order to apply the Invert algorithm of Lemma 6, the noise term $\mathbf{e}_0 = (\mathbf{e}_{0,1}^\top \mid \mathbf{e}_{0,2}^\top)^\top \in \mathbb{Z}^{\bar{m}} \times \mathbb{Z}^{nk}$ of $\mathbf{c}_0$ must satisfy $\|\mathbf{e}_{0,2} - \mathbf{R}^\top \mathbf{e}_{0,1}\| \leq q/(2\sqrt{k})$, where $k = \lceil \log q \rceil$. Lemma 9 will show that w.h.p. $\|\mathbf{R}^\top \mathbf{e}_{0,1}\|_2 \leq 2\sigma\sqrt{\lambda \bar{m} n k}$, so that

$$\|\mathbf{e}_{0,2} - \mathbf{R}^\top \mathbf{e}_{0,1}\| \leq O(\sigma\sqrt{\lambda \bar{m} n k}) \leq O(m^2\sqrt{\lambda}) \leq O(m^2 n^{1/2}) \tag{3}$$

In order to satisfy the condition $\|\mathbf{e}_{0,2} - \mathbf{R}^\top \mathbf{e}_{0,1}\| \leq q/(2\sqrt{k})$, we can thus choose $q = \Theta(m^{2.5}) = \tilde{\Theta}(n^{2.5})$. Then, if we set $p = O(1)$, we also have $\|\mathbf{s}\|_\infty < \frac{q-(p-1)p}{2p}$ in aEnc (since $\|\mathbf{s}\|_\infty \leq \alpha q\sqrt{2\lambda} = O(m^{1/2} n^{1/2})$ with overwhelming probability by Lemma 3) and $q \geq 2p \cdot \alpha q \sqrt{2\lambda}$.

Lemma 9. *The scheme is correct with probability at least* $1 - 2^{-\Omega(\lambda)}$ *over the randomness of* KGen, aGen, Enc *and* aEnc *if* KGen *and* aGen *choose parameters such that* $q = \Theta(m^{2.5})$, $p = O(1)$, $\alpha q = \Theta(\sqrt{m})$, *and* $\sigma = 2\alpha q s$, *where* $s \geq 13\sqrt{m+n}$. (The proof is available in Supplementary Material A.2).

We remark that the proof of Lemma 9 still works if we set $p = \Theta(n)$ without increasing q or other parameters. In this case, the expansion rate is only $O(\log \lambda)$.

Theorem 1. *If $\sigma = 2\alpha q s$ where $s \geq 13\sqrt{m+n}$, the scheme is anamorphic under the $\mathsf{LWE}_{n,m,q,\chi}$ and $\mathsf{HNF\text{-}LWE}_{n,m,q,\chi}$ assumptions with $\chi = D_{\mathbb{Z},\alpha q}$.*

Proof. The proof considers a sequence of games. For each i, W_i denotes the event that the adversary outputs 1 in Game_i. The first game is identical to experiment RealG_{Π} of Definition 3 and the last game is identical to $\mathsf{AnamorphicG}_{\Sigma}$.

Game_0**:** This is the real experiment RealG_{Π}. At the beginning of the experiment, the adversary $\mathcal{A}$ is given $(\mathsf{pk}, \mathsf{sk}) = ((\mathsf{par}, \mathbf{A}, \mathbf{U}), \mathbf{E})$. At each query $\mathcal{O}_e(\mathsf{pk}, \boldsymbol{\mu}, \hat{\boldsymbol{\mu}})$, the challenger returns a ciphertext $(\mathbf{c}_0, \mathbf{c}_1) \leftarrow \mathsf{Enc}(\mathsf{pk}, \boldsymbol{\mu})$. We call W_0 the event that $\mathcal{A}$ outputs 1 at the end of the game.

Game_1**:** This game is like Game_0 except that the challenger generates the matrix $\mathbf{A}$ with a trapdoor, by running algorithm $\mathsf{GenTrap}$ from Lemma 6. By Lemma 6, the distribution of the public key $(\mathbf{A}, \mathbf{U})$ is statistically close to that of Game_0 and $|\Pr[W_1] - \Pr[W_0]| \leq 2^{-\Omega(\lambda)}$.

Game_2**:** We modify the encryption oracle of RealG_{Π}. At each query $\mathcal{O}_e(\mathsf{pk}, \boldsymbol{\mu}, \hat{\boldsymbol{\mu}})$, instead of running the real Enc algorithm, the challenger samples $\mathbf{s} \hookleftarrow U(\mathbb{Z}_q^n)$, $\mathbf{x} \hookleftarrow D_{\mathbb{Z}^m,\alpha q}$ and uses the ReRand algorithm of Lemma 4 to compute

$$\begin{aligned} \mathbf{c}'_0 &= \mathbf{A}^\top \cdot \mathbf{s} + \mathbf{x} \in \mathbb{Z}_q^m \\ \begin{bmatrix} \mathbf{c}_0 \\ \mathbf{c}_1 \end{bmatrix} &= \mathsf{ReRand}([\mathbf{I}_m \mid \mathbf{E}], \mathbf{c}'_0, \alpha q, s) + \begin{bmatrix} \mathbf{0}^m \\ \Delta \cdot \boldsymbol{\mu} \end{bmatrix} \in \mathbb{Z}_q^{m+n}, \end{aligned} \tag{4}$$

where $s \geq \sqrt{1 + 144(m+n)}$. Since Lemma 12 ensures that $\|\mathbf{E}\| \leq 12\sqrt{m+n}$ with overwhelming probability $\geq 1 - 2 \cdot \exp(-(m+n))$, we can apply Lemma 4 with $\mathbf{V} = [\mathbf{I}_m \mid \mathbf{E}]$ (for which we have $\|\mathbf{V}\| \leq s$), $\mathbf{b} = \mathbf{A}^\top \cdot \mathbf{s} \in \mathbb{Z}_q^m$, and $r = \alpha q$ to obtain that the above distribution of $(\mathbf{c}_0, \mathbf{c}_1)$ is statistically close to the one obtained by computing

$$\begin{bmatrix} \mathbf{c}_0 \\ \mathbf{c}_1 \end{bmatrix} = \begin{bmatrix} \mathbf{A}^\top \\ \mathbf{U}^\top \end{bmatrix} \cdot \mathbf{s} + \begin{bmatrix} \mathbf{e}_0 \\ \mathbf{e}_1 \end{bmatrix} + \begin{bmatrix} \mathbf{0}^m \\ \Delta \cdot \boldsymbol{\mu} \end{bmatrix},$$

where $\mathbf{s} \hookleftarrow U(\mathbb{Z}_q^n)$, $\mathbf{e}_0 \sim D_{\mathbb{Z}^m,2\alpha q s}$, $\mathbf{e}_1 \sim D_{\mathbb{Z}^n,2\alpha q s}$. This shows that $|\Pr[W_2] - \Pr[W_1]| \leq Q \cdot 2^{-\Omega(\lambda)}$, where Q is the number of queries to the oracle $\mathcal{O}_e$.

Game_3**:** We change again the encryption oracle. At each query $\mathcal{O}_e(\mathsf{pk}, \boldsymbol{\mu}, \hat{\boldsymbol{\mu}})$, instead of computing a ciphertext as per (4), the challenger computes

$$\begin{aligned} \mathbf{c}'_0 &\hookleftarrow U(\mathbb{Z}_q^m) \\ \begin{bmatrix} \mathbf{c}_0 \\ \mathbf{c}_1 \end{bmatrix} &= \mathsf{ReRand}([\mathbf{I}_m \mid \mathbf{E}], \mathbf{c}'_0, \alpha q, s) + \begin{bmatrix} \mathbf{0}^m \\ \Delta \cdot \boldsymbol{\mu} \end{bmatrix}, \end{aligned} \tag{5}$$

and returns $(\mathbf{c}_0, \mathbf{c}_1)$. Under the LWE assumption, this change goes unnoticed. By a standard hybrid argument over all queries to the encryption oracle $\mathcal{O}_e$, we obtain $|\Pr[W_3] - \Pr[W_2]| \leq Q \cdot \mathbf{Adv}^{\mathsf{LWE}}_{n,m,q,\chi}(\lambda)$.

Game_4**:** We change again the encryption oracle. At each query $\mathcal{O}_e(\mathsf{pk}, \boldsymbol{\mu}, \hat{\boldsymbol{\mu}})$, the challenger now samples $\mathbf{u} \hookleftarrow U(\mathbb{Z}_q^m)$ uniformly and computes

$$\begin{aligned} \mathbf{c}_0' &= \mathbf{u} + \mathbf{A}^\top \cdot \mathsf{encode}_p(\hat{\boldsymbol{\mu}}) \\ \begin{bmatrix} \mathbf{c}_0 \\ \mathbf{c}_1 \end{bmatrix} &= \mathsf{ReRand}([\mathbf{I}_m \mid \mathbf{E}], \mathbf{c}_0', \alpha q, s) + \begin{bmatrix} \mathbf{0}^m \\ \Delta \cdot \boldsymbol{\mu} \end{bmatrix}, \end{aligned} \tag{6}$$

Clearly, this change has no impact on the distribution of $(\mathbf{c}_0, \mathbf{c}_1)$ since we have $\mathbf{c}_0' \sim U(\mathbb{Z}_q^m)$ exactly as in (5). Therefore we have $\Pr[W_4] = \Pr[W_3]$.

Game_5**:** We change the output distribution of $\mathcal{O}_e$. At each query $\mathcal{O}_e(\mathsf{pk}, \boldsymbol{\mu}, \hat{\boldsymbol{\mu}})$, the challenger now samples $\mathbf{s} \hookleftarrow D_{\mathbb{Z}^n, \alpha q}$, $\mathbf{x} \hookleftarrow D_{\mathbb{Z}^m, \alpha q}$ and computes

$$\begin{aligned} \mathbf{c}_0' &= (\mathbf{A}^\top \cdot \mathbf{s} + \mathbf{x}) + \mathbf{A}^\top \cdot \mathsf{encode}_p(\hat{\boldsymbol{\mu}}) \\ \begin{bmatrix} \mathbf{c}_0 \\ \mathbf{c}_1 \end{bmatrix} &= \mathsf{ReRand}([\mathbf{I}_m \mid \mathbf{E}], \mathbf{c}_0', \alpha q, s) + \begin{bmatrix} \mathbf{0}^m \\ \Delta \cdot \boldsymbol{\mu} \end{bmatrix}, \end{aligned} \tag{7}$$

which amounts to replacing the uniform $\mathbf{u}$ of (6) by a pseudorandom vector $\mathbf{u} = \mathbf{A}^\top \cdot \mathbf{s} + \mathbf{x}$ at each query. Under the LWE assumption in HNF form [6], this change does not affect $\mathcal{A}$'s view and a standard hybrid argument over all queries to $\mathcal{O}_e$ implies $|\Pr[W_5] - \Pr[W_4]| \leq Q \cdot \mathbf{Adv}^{\mathsf{HNF\text{-}LWE}}_{n,m,q,\chi}(\lambda)$.

Game_6**:** In this game, the challenger answers all encryption queries $\mathcal{O}_e(\mathsf{pk}, \boldsymbol{\mu}, \hat{\boldsymbol{\mu}})$ by sampling $\mathbf{s} \hookleftarrow D_{\mathbb{Z}^n, \alpha q}$, $\mathbf{e}_0 \hookleftarrow D_{\mathbb{Z}^m, \sigma}$, $\mathbf{e}_1 \hookleftarrow D_{\mathbb{Z}^n, \sigma}$ and computing

$$\begin{aligned} \mathbf{c}_0 &= \mathbf{A}^\top \cdot (\mathbf{s} + \mathsf{encode}_p(\hat{\boldsymbol{\mu}})) + \mathbf{e}_0 \\ \mathbf{c}_1 &= \mathbf{U}^\top \cdot (\mathbf{s} + \mathsf{encode}_p(\hat{\boldsymbol{\mu}})) + \mathbf{e}_1 + \Delta \cdot \boldsymbol{\mu}, \in \mathbb{Z}_q^n, \end{aligned} \tag{8}$$

By applying Lemma 4 again with $\mathbf{V} = [\mathbf{I}_m \mid \mathbf{E}]$, $\mathbf{b} = \mathbf{A}^\top \cdot \mathbf{s} \in \mathbb{Z}_q^m$, and $r = \alpha q$, the distribution of $(\mathbf{c}_0, \mathbf{c}_1)$ is statistically close to that of Game_5. We have $|\Pr[W_6] - \Pr[W_5]| \leq Q \cdot 2^{-\Omega(\lambda)}$, where Q is the number of $\mathcal{O}_e$-queries.

In Game_6, the challenger is running $\mathsf{AnamorphicG}_\Sigma$ with the adversary. By the triangle inequality, we obtain the following inequality which proves the result:

$$\mathbf{Adv}^{\mathsf{Anamorphism}}_{\mathcal{D},\Pi,\Sigma}(\lambda) \leq Q \cdot \left(\mathbf{Adv}^{\mathsf{LWE}}_{n,m,q,\chi}(\lambda) + \mathbf{Adv}^{\mathsf{HNF\text{-}LWE}}_{n,m,q,\chi}(\lambda)\right) + (2Q+1) \cdot 2^{-\Omega(\lambda)}.$$

□

Remark. Since the scheme is public-key anamorphic [44] (and $\mathsf{dk} = \emptyset$), it is also fully asymmetric in the sense of Definition 4. Indeed, by the anamorphic property, the challenge ciphertext $\mathsf{act} \leftarrow \mathsf{aEnc}(\mathsf{apk}, \mathsf{dk}, \mu_b, \hat{\mu}_b)$ and the anamorphic public key apk are indistinguishable from a normal ciphertext $\mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{pk}, \mu_b)$ encrypted under a normal public key pk. Then, the standard IND-CPA security property ensures that b is computationally hidden.

Achieving Robustness. As described above, the scheme does not provide robustness. It can easily be made robust for fresh ciphertexts if we modify the

second step of aDec and have it first compute $\mathbf{s}' = \hat{\mathbf{s}} - \Delta \cdot \hat{\boldsymbol{\mu}} \bmod q$ and return $\perp$ if $\|\mathbf{s}'\|_\infty > \alpha q\sqrt{2\lambda}$. Since $\mathbf{s}$ is chosen uniformly in $\mathbb{Z}_q^n$ in the normal encryption algorithm, aDec can only obtain $\mathbf{s}'$ such that $\|\mathbf{s}'\|_\infty \leq \alpha q\sqrt{2\lambda}$ with negligible probability. However, this creates a problem with Definition 7 since the sum of two anamorphic ciphertexts can anamorphically decrypt to $\perp$. In Supplementary Material E.2, we modify aDec to achieve robustness, even for ciphertexts obtained from homomorphic additions, and preserve homomorphic correctness.

Construction Without Lattice Trapdoors. In Supplementary Material C, we provide an alternative fully asymmetric anamorphic variant of dual Regev. While this construction is no longer public-key anamorphic (since $\mathsf{dk} \neq \emptyset$), it is more efficient and does not rely on lattice trapdoors.

4 Fully Asymmetric AE from Primal Regev

We now show that the primal Regev cryptosystem [45,46] can also be endowed with a fully asymmetric anamorphic mechanism.

Our proof relies on the lossy mode of LWE [26], which was used in several works [4,32,29]. This lossy mode relies on the pseudorandomness of matrices of the form $\mathbf{A}^\top = \mathbf{B}^\top\mathbf{C} + \mathbf{F}$, where $\mathbf{C} \hookleftarrow U(\mathbb{Z}_q^{\ell\times n})$, $\mathbf{B} \hookleftarrow U(\mathbb{Z}_q^{\ell\times m})$, $\mathbf{F} \hookleftarrow \chi^{m\times n}$ for $\ell \ll n$. Under the $\mathsf{LWE}_{\ell,m,q,\chi}$ assumption, such matrices are known to be computationally indistinguishable from uniformly random matrices in $\mathbb{Z}_q^{n\times m}$.

When $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ is produced by the lossy sampler, a secret $\mathbf{s} \in \mathbb{Z}^n$ where each entry is sampled from a narrow interval $[-\gamma, \gamma]$ has high entropy conditionally on $\mathbf{A}^\top\mathbf{s} + \mathbf{e}$, when $\mathbf{e}$ is sampled from a suitable noise distribution (as shown in, e.g., [4,29]). However, we do not rely on the entropy of $\mathbf{s}$ in lossy mode here. We only use the property that lossy matrices are "close" to matrices with smaller rank $\ell \ll n$. This allows us to instantiate the scheme with uniform secret keys $\mathbf{s} \sim U(\mathbb{Z}_q^n)$, which is necessary to prove that the scheme is fully asymmetric.[6]

Moreover, we also need to modify the generation of lossy matrices in such a way that the matrix $\mathbf{B}$ is sampled from the noise distribution χ. Under the $\mathsf{HNF\text{-}LWE}_{\ell,n,q,\chi}$ assumption, a matrix produced by our modified $\mathsf{SampleLossy}$ for a random $\mathbf{C} \sim U(\mathbb{Z}_q^{\ell\times n})$ is indistinguishable from a uniform $\mathbf{A} \sim U(\mathbb{Z}_q^{n\times m})$.

Our anamorphic variant is identical to the original scheme (in its multi-bit version [42]) with uniformly random secret keys $\mathbf{S} \sim U(\mathbb{Z}_q^{n\times\ell})$.

$\mathsf{KGen}(1^\lambda)$: Given a security parameter 1^λ,

1. Choose dimensions $n, \ell, m, k \in \mathsf{poly}(\lambda)$, a plaintext modulus $p \in \mathsf{poly}(\lambda)$ and a ciphertext modulus $q \in \mathsf{poly}(\lambda)$ such that $q > p$, $n = 2\ell k + 2\lambda$ and $m \geq (n+\ell)k + 2\lambda$ where $k = \lceil \log q \rceil$. Choose an error rate $\alpha \in (0,1)$ and output the common public parameters $\mathsf{par} := \Big(q, p, n, m, \ell, \alpha\Big)$.
2. Choose a random matrix $\mathbf{A} \hookleftarrow \mathbb{Z}_q^{n\times m}$.

[6] The reason is that we use a "first-are-errorless" variant of LWE [22,15,3], which was only shown as hard as LWE for uniform secrets.

3. Sample $\mathbf{S} \hookleftarrow U(\mathbb{Z}_q^{n\times\ell})$, $\mathbf{E} \hookleftarrow D_{\mathbb{Z}^m,\alpha q}^{\ell}$ and compute

$$\mathbf{U}^\top = \mathbf{A}^\top\mathbf{S} + \mathbf{E} \in \mathbb{Z}_q^{m\times\ell}$$

Return $\mathsf{sk} = \mathbf{S} \in \mathbb{Z}_q^{n\times\ell}$ and $\mathsf{pk} = (\mathsf{par}, \mathbf{A}, \mathbf{U})$.

aGen(1^λ)**:** Given a security parameter 1^λ,

1. Run step 1 of KGen to generate parameters $\mathsf{par} = (q, p, n, m, \ell, \alpha)$.
2. Let $\mathbf{G}_\ell = \mathbf{I}_\ell \otimes (1, 2, \ldots, 2^{k-1})$. Generate a statistically uniform matrix $\mathbf{C} = [\bar{\mathbf{C}} \mid \bar{\mathbf{C}} \cdot \mathbf{R}_C + \mathbf{G}_\ell] \in \mathbb{Z}_q^{\ell\times n}$ together with a trapdoor $\mathbf{R}_C \sim \mathcal{P}^{(\ell k+2\lambda)\times\ell k}$ for $\Lambda_q^\perp(\mathbf{C})$ using the GenTrap algorithm of Lemma 6.
3. Compute a matrix $\mathbf{A}^\top = \mathbf{B}^\top\mathbf{C} + \mathbf{F}$, where $\mathbf{B} \hookleftarrow \chi^{\ell\times m}$, $\mathbf{F} \hookleftarrow \chi^{m\times n}$.
4. Sample $\mathbf{S} \hookleftarrow U(\mathbb{Z}_q^{n\times\ell})$, $\mathbf{E} \hookleftarrow D_{\mathbb{Z}^m,\alpha q}^{\ell}$ and compute $\mathbf{U}^\top = \mathbf{A}^\top\mathbf{S} + \mathbf{E}$.
5. Compute $\mathbf{D} = \mathbf{C} \cdot \mathbf{S} \in \mathbb{Z}_q^{\ell\times\ell}$.

Return $\mathsf{ask} = \mathbf{S} \in \mathbb{Z}_q^{n\times\ell}$, $\mathsf{apk} = (\mathsf{par}, \mathbf{A}, \mathbf{U})$, $\mathsf{dk} = (\mathbf{C}, \mathbf{D}) \in \mathbb{Z}_q^{\ell\times n} \times \mathbb{Z}_q^{\ell\times\ell}$ and $\mathsf{tk} = \mathbf{R}_C \in \{-1, 0, 1\}^{(\ell k+2\lambda)\times\ell k}$.

Enc$(\mathsf{pk}, \boldsymbol{\mu})$**:** Given the public key $\mathsf{pk} = (\mathsf{par}, \mathbf{A}, \mathbf{U})$ and a message $\boldsymbol{\mu} \in \mathbb{Z}_p^\ell$, sample $\mathbf{r} \hookleftarrow U(\{0,1\}^m)$ and compute

$$\mathbf{c}_0 = \mathbf{A} \cdot \mathbf{r} \in \mathbb{Z}_q^n, \qquad \mathbf{c}_1 = \mathbf{U} \cdot \mathbf{r} + \Delta \cdot \boldsymbol{\mu} \in \mathbb{Z}_q^\ell,$$

where $\Delta = \lfloor q/p \rceil$. Output the ciphertext $\mathsf{ct} = (\mathbf{c}_0, \mathbf{c}_1) \in \mathbb{Z}_q^{(n+\ell)}$.

aEnc$(\mathsf{apk}, \mathsf{dk}, \boldsymbol{\mu}, \hat{\boldsymbol{\mu}})$**:** Given $\mathsf{apk} = (\mathsf{par}, \mathbf{A}, \mathbf{U})$, $\mathsf{dk} = (\mathbf{C}, \mathbf{D}) \in \mathbb{Z}_q^{\ell\times n} \times \mathbb{Z}_q^{\ell\times\ell}$, and messages $\boldsymbol{\mu} \in \mathbb{Z}_p^\ell$ and $\hat{\boldsymbol{\mu}} \in \mathbb{Z}_p^\ell$, sample $\mathbf{r} \hookleftarrow U(\{0,1\}^m)$ and

1. Compute $\hat{\mathbf{s}} = \mathsf{encode}_p(\hat{\boldsymbol{\mu}}) = \Delta \cdot \hat{\boldsymbol{\mu}} \in \mathbb{Z}_q^\ell$ where $\mathsf{encode}_p : \mathbb{Z}_p^\ell \to \mathbb{Z}_q^\ell$ is the encoding algorithm of Section 2.3 and $\Delta = \lfloor q/p \rceil$.
2. Compute

$$\mathbf{c}_0 = \mathbf{A} \cdot \mathbf{r} + \mathbf{C}^\top \cdot \hat{\mathbf{s}}, \qquad \mathbf{c}_1 = \mathbf{U} \cdot \mathbf{r} + \mathbf{D}^\top \cdot \hat{\mathbf{s}} + \Delta \cdot \boldsymbol{\mu},$$

Output the anamorphic ciphertext $\mathsf{act} = (\mathbf{c}_0, \mathbf{c}_1) \in \mathbb{Z}_q^{(n+\ell)}$.

Dec$(\mathsf{pk}, \mathsf{sk}, \mathsf{ct})$**:** Given a ciphertext $\mathsf{ct} = (\mathbf{c}_0, \mathbf{c}_1) \in \mathbb{Z}_q^{(n+\ell)}$ and the secret key $\mathsf{sk} = \mathbf{S} \in \mathbb{Z}_q^{n\times\ell}$, compute and return $\boldsymbol{\mu} = \lfloor (\mathbf{c}_1 - \mathbf{S}^\top \cdot \mathbf{c}_0)/\Delta \rceil$.

aDec$(\mathsf{dk}, \mathsf{tk}, \mathsf{ask}, \mathsf{act})$**:** Given a ciphertext $\mathsf{act} = (\mathbf{c}_0, \mathbf{c}_1) \in \mathbb{Z}_q^{(m+n)}$ and the trapdoor key $\mathsf{tk} = \mathbf{R}_C \in \{-1, 0, 1\}^{(\ell k+2\lambda)\times\ell k}$,

1. Compute $(\hat{\mathbf{s}}', \mathbf{e}_0) \leftarrow \mathsf{Invert}(\mathbf{R}_C, \mathbf{c}_0)$ using the Invert algorithm of Lemma 6. If $(\hat{\mathbf{s}}', \mathbf{e}_0) = (\perp, \perp)$, return $\perp$.
2. Compute $\hat{\boldsymbol{\mu}} = \mathsf{decode}_p(\hat{\mathbf{s}}') \in \mathbb{Z}_p^\ell$ using the decoding algorithm of Section 2.3 and return $\hat{\boldsymbol{\mu}} \in \mathbb{Z}_p^\ell$.

Parameters. Lemma 10 requires $q > 18 \cdot \alpha q m \cdot \ell^{3/2} k^{3/2}$ to ensure correctness. We can set $\alpha q = \Omega(\sqrt{\ell})$ to guarantee the hardness of LWE instances used in Theorem 2 and Theorem 3. Since $n = \Theta(\ell \log q)$ and $m = \Omega(n \log q)$, we can thus choose $q = \Theta(\ell^3 \log^{7/2} \ell) = \tilde{\Theta}(\ell^3)$ and $p = O(1)$ or even $p = \Theta(\ell)$.

Lemma 10. *The scheme provides correctness with overwhelming probability over the randomness of* KGen, aGen, Enc *and* aEnc *for* $p = O(1)$ *and* $p = \Theta(\ell)$, *if* $q > 18 \cdot \alpha q m \cdot \ell^{3/2} k^{3/2}$. (The proof is given in Supplementary Material B.1.)

Compared to our scheme based on dual Regev, the above construction offers a better concrete security in its proof of anamorphism as the bound on the adversary's advantage only loses a factor $O(m)$ (regardless of the number of encryption queries) with respect to the LWE assumption.

Theorem 2. *The scheme is anamorphic in the sense of Definition 3 under the* $\mathsf{HNF\text{-}LWE}_{\ell,n,q,\chi}$ *assumption with* $\chi = D_{\mathbb{Z},\alpha q}$.

Proof. The proof considers a sequence of games. The first game is identical to experiment RealG_{Π} of Definition 3 and the last game is identical to $\mathsf{AnamorphicG}_{\Sigma}$. We let W_i denote the event that $\mathcal{A}$ outputs 1 at the end of the Game_i.

Game_0**:** This is the real experiment RealG_{Π}, where the adversary $\mathcal{A}$ is initially given $(\mathsf{pk}, \mathsf{sk}) = ((\mathsf{par}, \mathbf{A}, \mathbf{U}), \mathbf{S})$. At each encryption query $\mathcal{O}_e(\mathsf{pk}, \boldsymbol{\mu}, \hat{\boldsymbol{\mu}})$, the challenger returns a ciphertext $(\mathbf{c}_0, \mathbf{c}_1) \leftarrow \mathsf{Enc}(\mathsf{pk}, \boldsymbol{\mu})$ of the form

$$\mathbf{c}_0 = \mathbf{A} \cdot \mathbf{r}, \qquad \mathbf{c}_1 = \mathbf{U} \cdot \mathbf{r} + \Delta \cdot \boldsymbol{\mu},$$

Note that $\mathbf{c}_1$ can equivalently be written $\mathbf{c}_1 = \mathbf{S}^\top \cdot \mathbf{c}_0 + \mathbf{E}^\top \cdot \mathbf{r} + \Delta \cdot \boldsymbol{\mu}$.

Game_1**:** We modify the encryption oracle. At each query $\mathcal{O}_e(\mathsf{pk}, \boldsymbol{\mu}, \hat{\boldsymbol{\mu}})$, instead of running Enc, the challenger samples $\mathbf{r} \hookleftarrow U(\{0,1\}^m)$ and computes

$$\mathbf{c}_0 \hookleftarrow U(\mathbb{Z}_q^n), \qquad \mathbf{c}_1 = \mathbf{S}^\top \cdot \mathbf{c}_0 + \mathbf{E}^\top \cdot \mathbf{r} + \Delta \cdot \boldsymbol{\mu} \in \mathbb{Z}_q^\ell, \tag{9}$$

where $\mathbf{E} \in \mathbb{Z}^{m \times \ell}$ is the noise matrix of the public key $\mathbf{U}^\top = \mathbf{A}^\top \mathbf{S} + \mathbf{E}$. The only change is that we replace $\mathbf{c}_0 = \mathbf{A} \cdot \mathbf{r}$ by a truly uniform $\mathbf{c}_0$. We claim that Game_1 is statistically indistinguishable from Game_0. Indeed, by Lemma 1, the distribution $\{(\mathbf{A} \cdot \mathbf{r} \bmod q, \mathbf{E}^\top \cdot \mathbf{r}) \mid \mathbf{r} \hookleftarrow U(\{0,1\}^m)\}$ is within statistical distance $\frac{1}{2}\sqrt{2^{-m} \cdot q^n \cdot q^\ell} < 2^{-\lambda}$ from

$$\{(\mathbf{c}_0, \mathbf{E}^\top \cdot \mathbf{r}) \mid \mathbf{r} \hookleftarrow U(\{0,1\}^m), \mathbf{c}_0 \hookleftarrow U(\mathbb{Z}_q^n)\}$$

since the function $f(\mathbf{r}) = \mathbf{E}^\top \mathbf{r}$ has image size $\ll q^\ell$ and the choice of parameters implies $m > (n+\ell)\log q + 2\lambda$. This implies $|\Pr[W_2] - \Pr[W_1]| \leq Q \cdot 2^{-\lambda}$, where Q is the number of queries to the oracle $\mathcal{O}_e$.

Game_2**:** In this game, at each query $\mathcal{O}_e(\mathsf{pk}, \boldsymbol{\mu}, \hat{\boldsymbol{\mu}})$, the challenger generates a ciphertext by sampling $\mathbf{c}_0' \hookleftarrow U(\mathbb{Z}_q^n)$, $\mathbf{r} \hookleftarrow U(\{0,1\}^m)$ and computing

$$\begin{aligned} \mathbf{c}_0 &= \mathbf{c}_0' + \mathbf{C}^\top \cdot \mathsf{encode}_p(\hat{\boldsymbol{\mu}}) \\ \mathbf{c}_1 &= \mathbf{S}^\top \cdot \mathbf{c}_0 + \mathbf{E}^\top \cdot \mathbf{r} + \Delta \cdot \boldsymbol{\mu} \\ &= \mathbf{S}^\top \cdot \mathbf{c}_0' + \mathbf{E}^\top \cdot \mathbf{r} + \Delta \cdot \boldsymbol{\mu} + \mathbf{D}^\top \cdot \mathsf{encode}_p(\hat{\boldsymbol{\mu}}) \end{aligned} \tag{10}$$

where $\mathbf{E} \in \mathbb{Z}^{m \times \ell}$ is the noise matrix contained in $\mathbf{U}^\top = \mathbf{A}^\top \mathbf{S} + \mathbf{E}$. Clearly, the distribution of $(\mathbf{c}_0, \mathbf{c}_1)$ is exactly the same as in Game_1 since $\mathbf{c}_0 \sim U(\mathbb{Z}_q^n)$ in both games. We have $\Pr[W_2] = \Pr[W_1]$.

Game_3**:** This game is identical to Game_2 except that, at each encryption query $\mathcal{O}_e(\mathsf{pk}, \boldsymbol{\mu}, \hat{\boldsymbol{\mu}})$, the challenger replaces the truly uniform $\mathbf{c}'_0$ by $\mathbf{c}'_0 = \mathbf{A} \cdot \mathbf{r}$ with $\mathbf{r} \hookleftarrow U(\{0,1\}^m)$ in (10). Namely, the challenger computes $(\mathbf{c}_0, \mathbf{c}_1)$ by sampling $\mathbf{r} \hookleftarrow U(\{0,1\}^m)$ and computing

$$\mathbf{c}_0 = \mathbf{A} \cdot \mathbf{r} + \mathbf{C}^\top \cdot \mathsf{encode}_p(\hat{\boldsymbol{\mu}}), \qquad \mathbf{c}_1 = \mathbf{S}^\top \cdot \mathbf{c}_0 + \mathbf{E}^\top \cdot \mathbf{r} + \Delta \cdot \boldsymbol{\mu} \tag{11}$$

Note that $\mathbf{c}_1$ is now equal to

$$\begin{aligned}\mathbf{c}_1 &= \mathbf{S}^\top \cdot (\mathbf{A} \cdot \mathbf{r}) + \mathbf{E}^\top \cdot \mathbf{r} + \Delta \cdot \boldsymbol{\mu} + \mathbf{D}^\top \cdot \mathsf{encode}_p(\hat{\boldsymbol{\mu}}) \\ &= \mathbf{U} \cdot \mathbf{r} + \Delta \cdot \boldsymbol{\mu} + \mathbf{D}^\top \cdot \mathsf{encode}_p(\hat{\boldsymbol{\mu}})\end{aligned}$$

By Lemma 1 and the same argument as in the transition from Game_0 to Game_1, the distribution of each ciphertext is within statistical distance $2^{-\lambda}$ from that of Game_2. Consequently, we have $|\Pr[W_3] - \Pr[W_2]| \leq Q \cdot 2^{-\lambda}$, where Q is the number of queries to $\mathcal{O}_e$.

Game_4**:** This game is identical to Game_3 except that, in the key generation phase, the challenger replaces the uniform matrix $\mathbf{A} \sim U(\mathbb{Z}_q^{n \times m})$ by a lossy matrix of the form $\mathbf{A}^\top = \mathbf{B}^\top \mathbf{C} + \mathbf{F}$, where $\mathbf{C} \hookleftarrow U(\mathbb{Z}_q^{\ell \times n})$, $\mathbf{B} \hookleftarrow \chi^{\ell \times m}$, $\mathbf{F} \hookleftarrow \chi^{m \times n}$. Under the $\mathsf{HNF\text{-}LWE}_{\ell,n,q,\chi}$ assumption, this change is not noticeable to $\mathcal{A}$. By a standard hybrid argument over the rows of $\mathbf{A}^\top$, we concretely obtain $|\Pr[W_4] - \Pr[W_3]| \leq m \cdot \mathbf{Adv}^{\mathsf{HNF\text{-}LWE}}_{\ell,n,q,\chi}(\lambda)$.

Game_5**:** This game is as Game_4 but the challenger now generates the matrix $\mathbf{C} \in \mathbb{Z}_q^{\ell \times n}$ as a statistically uniform matrix with a trapdoor, by running algorithm $\mathsf{GenTrap}$ from Lemma 6. By Lemma 6, the distribution of $\mathbf{C}$ is statistically close to that of Game_4 and so is the distribution of $(\mathbf{A}, \mathbf{U})$. Therefore, we have $|\Pr[W_5] - \Pr[W_4]| \leq 2^{-\Omega(\lambda)}$.

In Game_5, the challenger is running $\mathsf{AnamorphicG}_\Sigma$ with the adversary. By the triangle inequality, we obtain the following bound which proves the result:

$$\mathbf{Adv}^{\mathsf{Anamorphism}}_{\mathcal{D},\Pi,\Sigma}(\lambda) = |\Pr[W_0] - \Pr[W_5]| \leq m \cdot \mathbf{Adv}^{\mathsf{HNF\text{-}LWE}}_{\ell,n,q,\chi}(\lambda) + \frac{2Q+1}{2^{\Omega(\lambda)}}.$$

□

Theorem 3. *Under the* $\mathsf{HNF\text{-}LWE}_{\ell,n,q,\chi}$ *and* $\mathsf{LWE}_{n-\ell,m+\ell,q,\chi}$ *assumptions with noise distribution* $\chi = D_{\mathbb{Z},\alpha q}$*, the scheme is fully asymmetric.*

Proof. The proof proceeds with a sequence of hybrid games where W_i^b denotes the event that the adversary outputs 1 in Game_i^b. The first game is the experiment $\mathsf{Fasym}^b_{\mathcal{A},\Pi,\Sigma}$ whereas the final hybrid game is completely independent of b.

Game_0^b**:** This is the real experiment $\mathsf{Fasym}^b_{\mathcal{A},\Pi,\Sigma}$. The challenger initially runs aGen to generate a statistically uniform matrix $\mathbf{C} \in \mathbb{Z}_q^{\ell \times n}$ with a trapdoor $\mathbf{R}_C$ by running the $\mathsf{GenTrap}$ algorithm of Lemma 6. It then uses $\mathbf{C}$ to generate a lossy matrix $\mathbf{A}^\top = \mathbf{B}^\top \cdot \mathbf{C} + \mathbf{F}$, where $\mathbf{B} \hookleftarrow D^\ell_{\mathbb{Z}^m,\alpha q}$, $\mathbf{F} \hookleftarrow D^n_{\mathbb{Z}^m,\alpha q}$. The

adversary $\mathcal{A}$ is run on $\mathsf{apk} = (\mathsf{par}, \mathbf{A}, \mathbf{U})$, $\mathsf{dk} = (\mathbf{C}, \mathbf{D} = \mathbf{C} \cdot \mathbf{S})$. In the challenge phase, $\mathcal{A}$ chooses messages $(\hat{\boldsymbol{\mu}}_0, \hat{\boldsymbol{\mu}}_1, \boldsymbol{\mu}_0, \boldsymbol{\mu}_1)$ and obtains a challenge $\mathsf{act} = (\mathbf{c}_0, \mathbf{c}_1) \leftarrow \mathsf{aEnc}(\mathsf{apk}, \mathsf{dk}, \hat{\boldsymbol{\mu}}_b, \boldsymbol{\mu}_b)$ of the form

$$\mathbf{c}_0 = \mathbf{A} \cdot \mathbf{r} + \mathbf{C}^\top \cdot \mathsf{encode}_p(\hat{\boldsymbol{\mu}}_b), \qquad \mathbf{c}_1 = \mathbf{U} \cdot \mathbf{r} + \mathbf{D}^\top \cdot \mathsf{encode}_p(\hat{\boldsymbol{\mu}}_b) + \Delta \cdot \boldsymbol{\mu}_b.$$

Game_1^b**:** This game is identical to Game_0^b except that the challenger chooses $\mathbf{C} \hookleftarrow U(\mathbb{Z}_q^{n \times m})$ truly uniformly instead of generating it with a trapdoor. By Lemma 6, the distributions of $\mathsf{dk} = (\mathbf{C}, \mathbf{D})$ and $\mathsf{apk} = (\mathsf{par}, \mathbf{A}, \mathbf{U})$ are statistically close to those of Game_0^b and $|\Pr[W_1^b] - \Pr[W_0^b]| \leq 2^{-\Omega(\lambda)}$.

Game_2^b**:** This game is like Game_1^b with the difference that, in the key generation phase, the challenger replaces the lossy matrix $\mathbf{A}^\top = \mathbf{B}^\top \mathbf{C} + \mathbf{F}$ by a uniform matrix $\mathbf{A} \hookleftarrow U(\mathbb{Z}_q^{n \times m})$. Under the $\mathsf{HNF\text{-}LWE}_{\ell,n,q,\chi}$ assumption, this change has no noticeable impact on $\mathcal{A}$. A standard hybrid argument over the rows of $\mathbf{A}$ leads to the inequality $|\Pr[W_2^b] - \Pr[W_1^b]| \leq m \cdot \mathbf{Adv}^{\mathsf{HNF\text{-}LWE}}_{\ell,n,q,\chi}(\lambda)$.

Game_3^b**:** We modify the distribution of $(\mathsf{apk}, \mathsf{dk})$, replacing $\mathbf{U}^\top = \mathbf{A}^\top \mathbf{S} + \mathbf{E}$ and $\mathbf{D} = \mathbf{C} \cdot \mathbf{S}$ by uniformly random matrices $\mathbf{U}^\top \hookleftarrow U(\mathbb{Z}_q^{m \times \ell})$ and $\mathbf{D} \hookleftarrow U(\mathbb{Z}_q^{\ell \times \ell})$. We claim that, under the $\mathsf{LWE}_{n-\ell,m+\ell,q,\chi}$ assumption, Game_3^b is indistinguishable from Game_2^b. The reason is that, under the "first-are-errorless" LWE assumption [22,15,3] in dimension n, the distribution

$$\big\{(\mathbf{A}^\top, \mathbf{A}^\top \mathbf{s} + \mathbf{e}, \mathbf{C}, \mathbf{C} \cdot \mathbf{s}) \mid \mathbf{A} \hookleftarrow U(\mathbb{Z}_q^{n \times m}),\ \mathbf{s} \hookleftarrow U(\mathbb{Z}_q^n),\ \mathbf{e} \hookleftarrow D_{\mathbb{Z}^m, \alpha q},\ \mathbf{C} \hookleftarrow U(\mathbb{Z}_q^{\ell \times n})\big\}$$

is indistinguishable from

$$\big\{(\mathbf{A}^\top, \mathbf{u}, \mathbf{C}, \mathbf{d}) \mid \mathbf{A} \hookleftarrow U(\mathbb{Z}_q^{n \times m}),\ \mathbf{u} \hookleftarrow U(\mathbb{Z}_q^m),\ \mathbf{C} \hookleftarrow U(\mathbb{Z}_q^{\ell \times n}),\ \mathbf{d} \hookleftarrow U(\mathbb{Z}_q^\ell)\big\}.$$

As shown in [22, Lemma C.6] and [3, Lemma 3], the "first-are-errorless" LWE problem in dimension n (with m errorful and ℓ errorless samples) is as hard as the LWE problem in dimension $n - \ell$ (with $m + \ell$ samples) when the secret is sampled from $U(\mathbb{Z}_q^n)$. By a standard hybrid argument over the columns of $\mathbf{U}^\top$ and $\mathbf{D}$, we get $|\Pr[W_3^b] - \Pr[W_2^b]| \leq \ell \cdot \mathbf{Adv}^{\mathsf{LWE}}_{n-\ell,m+\ell,q,\chi}(\lambda)$.

Game_4^b**:** In this game, we now compute the challenge ciphertext as

$$\mathbf{c}_0 = \mathbf{c}'_0 + \mathbf{C}^\top \cdot \mathsf{encode}_p(\hat{\boldsymbol{\mu}}_b), \qquad \mathbf{c}_1 = \mathbf{c}'_1 + \mathbf{D}^\top \cdot \mathsf{encode}_p(\hat{\boldsymbol{\mu}}_b) + \Delta \cdot \boldsymbol{\mu}_b,$$

where $\mathbf{c}'_0 \hookleftarrow U(\mathbb{Z}_q^n)$ and $\mathbf{c}'_1 \hookleftarrow U(\mathbb{Z}_q^\ell)$. By Lemma 2, the statistical distance between the distributions of the challenge ciphertext in Game_4^b and Game_3^b is at most $\frac{1}{2}\sqrt{2^{-m} \cdot q^{n+\ell}} < 2^{-\lambda+1}$ since the parameters are chosen so that $m > (n + \ell) \cdot \log q + 2\lambda$. Therefore $|\Pr[W_4^b] - \Pr[W_3^b]| \leq 2^{-\lambda+1}$.

In Game_4^b, the challenge ciphertext perfectly hides $b \in \{0, 1\}$ since the pair $(\mathbf{c}'_0, \mathbf{c}'_1)$ acts as a one-time pad. Consequently, we have $\Pr[W_4^0] = \Pr[W_4^1]$ and, by the triangle inequality, we obtain

$$\mathbf{Adv}^{\mathsf{FAsym}}_{\mathcal{A},\Pi,\Sigma}(\lambda) \leq 2^{-\Omega(\lambda)} + 2m \cdot \mathbf{Adv}^{\mathsf{HNF\text{-}LWE}}_{\ell,n,q,\chi}(\lambda) + 2\ell \cdot \mathbf{Adv}^{\mathsf{LWE}}_{n-\ell,m,q,\chi}(\lambda).$$

□

5 Fully Asymmetric Anamorphic FHE from Dual GSW

We now show that the dual version of the leveled GSW FHE admits a fully asymmetric anamorphic mode. In contrast with the constructions described in earlier sections, we do not rely on lattice trapdoors here.

For simplicity, we assume that normal secret keys $\mathbf{s}$ are sampled from the ternary distribution $\mathcal{P}$. However, the proof extends to the variant of [14] where secret keys are sampled from a uniform binary distribution since the expected number of zeroes in $\mathbf{s}$ is exactly the same.

In the description hereafter, if $\mathbf{M} \in \mathbb{Z}_q^{N\times M}$ is a matrix and $I \subset [N]$ is a subset of indexes, $\mathbf{M}_I \in \mathbb{Z}_q^{N\times M}$ denotes the matrix obtained by replacing the rows with indexes outside I with all-zeroes rows. For vector $\mathbf{v} \in \mathbb{Z}_q^N$, we similarly define $\mathbf{v}_I \in \mathbb{Z}_q^N$ as the vector where entries in $[N] \setminus I$ are replaced by zeroes. For simplicity, we only specify a multiplicative depth L in unary for KGen and aGen as additions contribute less significantly to parameters. However, one can be more specific about the circuit class by e.g. specifying a number of additions between multiplications in arithmetic circuits (see $(L_{\mathsf{add}}, L_{\mathsf{mult}})$ in Theorem 4).

$\mathsf{KGen}(1^\lambda, 1^L)$: Given security parameter 1^λ and multiplicative depth 1^L,

1. Choose dimensions $m, n, n_0 \in \mathsf{poly}(\lambda)$, a constant plaintext modulus p and a ciphertext modulus q that depends on (λ, L) such that $n_0 \leq \frac{m}{2} - \lceil\sqrt{\lambda m/2}\rceil$, $m - n_0 > nk + 2\lambda$ with $k = \lceil \log q\rceil = \mathsf{poly}(\log\lambda)$. Further, choose error rate $\alpha \in (0,1)$, a standard deviation $\sigma > \alpha q$ and set $\chi = D_{\mathbb{Z},\alpha q}$. Define parameters $\mathsf{par} = (q, p, m, n, \alpha, \sigma, L)$.
2. Sample $\mathbf{A} \hookleftarrow U(\mathbb{Z}_q^{n\times m})$, and $\mathbf{s} \hookleftarrow \mathcal{P}^m$ from the distribution $\mathcal{P}$ that outputs 0 with probability 1/2 and ± 1 with probability 1/4 each. Compute

$$\mathbf{B} = \left[\frac{\mathbf{A}^\top}{\mathbf{s}^\top\mathbf{A}^\top}\right] \in \mathbb{Z}_q^{(m+1)\times n}$$

Output $\mathsf{pk} = (\mathsf{par}, \mathbf{B})$ and $\mathsf{sk} = \mathbf{s} \in \{-1, 0, 1\}^m$.

$\mathsf{aGen}(1^\lambda, 1^L)$: Given security parameter 1^λ and multiplicative depth 1^L,

1. Run step 1 of KGen. This generates parameters $m, n, n_0 \in \mathsf{poly}(\lambda)$ and moduli p, q such that $n_0 \leq m/2 - \lceil\sqrt{\lambda m/2}\rceil$ and $m - n_0 > nk + 2\lambda$, where $k = \lceil\log q\rceil = \mathsf{poly}(\log\lambda)$. Define $\mathsf{par} = (q, p, m, n, n_0, \alpha, \sigma, L)$ where α and σ are also the same as in KGen.
2. Sample a secret key $\mathbf{s} = (s_1, \ldots, s_m)^\top \hookleftarrow \mathcal{P}^m$ from the distribution $\mathcal{P}$ that outputs 0 with probability 1/2 and ± 1 with probability 1/4 each. Let $J = \{i \in [m] : s_i = 0\} \subseteq [m]$. If $|J| < n_0$, abort and output $\perp$. Otherwise, choose a uniformly random subset $I \subseteq J$ such that $|I| = n_0$ and parse it as $I = \{i_1, \ldots, i_{n_0}\}$.
3. Choose a short vector $\mathbf{t} \in \mathbb{Z}^m$ such that $\mathbf{t}_{[m]\setminus I} = \mathbf{0}^m$ and a matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ such that the distribution of $\mathbf{t}^\top\mathbf{A}^\top \bmod q = \mathbf{t}^\top\mathbf{A}_I^\top \bmod q = \mathbf{e}_I^\top$ is χ^n.

This is done by sampling $\bar{\mathbf{A}}_0 \hookleftarrow U(\mathbb{Z}_q^{n\times(n_0-1)})$, $\mathbf{t}' \hookleftarrow \chi^{n_0-1}$, $\mathbf{e}_I \hookleftarrow \chi^n$ and computing

$$\bar{\mathbf{A}}^\top = \left[\frac{\bar{\mathbf{A}}_0^\top}{\mathbf{t}'^\top \bar{\mathbf{A}}_0^\top + \mathbf{e}_I^\top}\right] \in \mathbb{Z}_q^{n_0\times n}. \tag{12}$$

Then, the matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ is obtained by setting

$$\mathbf{A}[\cdot, i_j] = \bar{\mathbf{A}}[\cdot, j] \qquad\qquad \forall j \in [n_0]$$

and choosing other columns of $\mathbf{A}$ uniformly in $\mathbb{Z}_q^n$. Then, $\mathbf{t} \in \mathbb{Z}^m$ is defined by setting $\mathbf{t}[i_j] = -\mathbf{t}'[j]$ for each $j \in [n_0 - 1]$, $\mathbf{t}[i_{n_0}] = 1$ and $\mathbf{t}[i] = 0$ for all $i \in [m] \setminus I$.

4. Compute

$$\mathbf{B} = \left[\frac{\mathbf{A}^\top}{\mathbf{s}^\top \mathbf{A}^\top}\right] \in \mathbb{Z}_q^{(m+1)\times n}.$$

Finally, output $\mathsf{apk} = (\mathsf{par}, \mathbf{B} \in \mathbb{Z}_q^{(m+1)\times n})$, $\mathsf{ask} = \mathbf{s} \in \{-1, 0, 1\}^m$, $\mathsf{dk} = I$ and $\mathsf{tk} = \mathbf{t} \in \mathbb{Z}^m$.

Enc(pk, μ)**:** Given $\mathsf{pk} = (\mathsf{par}, \mathbf{B})$ and a plaintext $\mu \in \mathbb{Z}_p$, sample $\mathbf{S} \hookleftarrow U(\mathbb{Z}_q^{n\times M})$, $\mathbf{E} \hookleftarrow (D_{\mathbb{Z}^{m+1},\sigma})^M$, and compute

$$\mathbf{C} = \mathbf{B} \cdot \mathbf{S} + \mu \cdot \mathbf{G} + \mathbf{E}$$

where $\mathbf{G} = \mathbf{I}_{m+1} \otimes \mathbf{g}^\top \in \mathbb{Z}_q^{(m+1)\times M}$ is the gadget matrix of rank $m+1$ where $M = k(m+1)$. Then output $\mathsf{ct} = \mathbf{C} \in \mathbb{Z}_q^{(m+1)\times M}$.

aEnc$(\mathsf{apk}, \mathsf{dk}, \mu, \hat{\mu})$**:** Given $\mathsf{apk} = (\mathsf{par}, \mathbf{B} \in \mathbb{Z}_q^{n\times m})$, $\mathsf{dk} = I \subseteq [m]$ and plaintexts $\mu \in \mathbb{Z}_p$, $\hat{\mu} \in \mathbb{Z}_p$, define the matrix $\mathbf{J}_{\mu,\hat{\mu}} \in \mathbb{Z}_q^{(m+1)\times(m+1)}$ as

$$\mathbf{J}_{\mu,\hat{\mu}}[i,j] = \begin{cases} \mu & \text{if } \ i = j \ \wedge \ i \notin I \\ \hat{\mu} & \text{if } \ i = j \ \wedge \ i \in I \\ 0 & \text{if } \ i \neq j \end{cases}$$

Then, sample $\mathbf{S} \hookleftarrow \chi^{n\times M}$, $\mathbf{E} \hookleftarrow (D_{\mathbb{Z}^{m+1},\sigma})^M$ and compute

$$\begin{aligned} \mathbf{C} &= \mathbf{B} \cdot \mathbf{S} + \mathbf{J}_{\mu,\hat{\mu}} \otimes \mathbf{g}^\top + \mathbf{E} \\ &= \mathbf{B} \cdot \mathbf{S} + \mu \cdot \mathbf{G}_{[m+1]\setminus I} + \hat{\mu} \cdot \mathbf{G}_I + \mathbf{E} \end{aligned}$$

where $\mathbf{G} = \mathbf{I}_{m+1} \otimes \mathbf{g}^\top$ and $M = k(m+1)$. Output $\mathsf{act} = \mathbf{C} \in \mathbb{Z}_q^{(m+1)\times M}$.

Dec$(\mathsf{sk}, \mathsf{ct})$**:** On input of a secret key $\mathsf{sk} = \mathbf{s} \in \{-1, 0, 1\}^m$ and a ciphertext $\mathsf{ct} = \mathbf{C} \in \mathbb{Z}_q^{(m+1)\times M}$, define $\hat{\mathbf{e}}_{m+1} = (0, \ldots, 0, 1)^\top$ and compute

$$\nu = [-\mathbf{s}^\top \mid 1] \cdot \mathbf{C} \cdot \mathbf{G}^{-1}(\Delta \cdot \hat{\mathbf{e}}_{m+1}) \in \mathbb{Z}_q$$

where $\Delta = \lfloor q/p \rceil$. Then, return $\mu = \lfloor \nu/\Delta \rceil \in \mathbb{Z}_p$.

aDec$(\mathsf{dk}, \mathsf{tk}, \mathsf{ask}, \mathsf{act})$**:** Given a ciphertext $\mathsf{act} = \mathbf{C} \in \mathbb{Z}_q^{(m+1)\times M}$ and the trapdoor key $\mathbf{t} \in \mathbb{Z}^m$, define $\hat{\mathbf{e}}_{i_{n_0}} = (0, \ldots, 0, 1, 0, \ldots, 0)^\top \in \{0,1\}^{m+1}$ as the i_{n_0}-th unit vector and compute

$$\nu = [\ \mathbf{t}^\top \mid 0] \cdot \mathbf{C} \cdot \mathbf{G}^{-1}(\Delta \cdot \hat{\mathbf{e}}_{i_{n_0}}) \in \mathbb{Z}_q$$

where $\Delta = \lfloor q/p \rceil$. Then output $\hat{\mu} = \lfloor \nu/\Delta \rceil \in \mathbb{Z}_p$.

Homomorphism. We show that the above scheme is homomorphic. In particular, we may write two initial ciphertexts encrypting $(\mu_1, \hat{\mu}_1)$ and $(\mu_2, \hat{\mu}_2)$ as

$$\mathbf{C}_i \approx \mathbf{B} \cdot \mathbf{S}_i + \mu_i \cdot \mathbf{G}_{[m+1]\setminus I} + \hat{\mu}_i \cdot \mathbf{G}_I, \qquad \forall i \in \{1,2\}$$

where the approximations hide $\mathbf{E}_i$. Then, multiplying the two ciphertexts yields

$$\begin{aligned}\mathbf{C}^\times \triangleq \mathbf{C}_1 \cdot \mathbf{G}^{-1}(\mathbf{C}_2) &\approx \mathbf{B} \cdot \mathbf{S}_1 \cdot \mathbf{G}^{-1}(\mathbf{C}_2) + \mu_1 \cdot \mathbf{G}_{[m+1]\setminus I} \cdot \mathbf{G}^{-1}(\mathbf{C}_2) \\ &\qquad + \hat{\mu}_1 \cdot \mathbf{G}_I \cdot \mathbf{G}^{-1}(\mathbf{C}_2) \\ &\approx \mathbf{B} \cdot \mathbf{S}_1 \cdot \mathbf{G}^{-1}(\mathbf{C}_2) + \mathbf{B}_{[m+1]\setminus I} \cdot \big(\mu_1 \cdot \mathbf{S}_2\big) + \mu_1 \cdot \mu_2 \cdot \mathbf{G}_{[m+1]\setminus I} \\ &\qquad + \mathbf{B}_I \cdot \big(\hat{\mu}_1 \cdot \mathbf{S}_2\big) + \hat{\mu}_1 \cdot \hat{\mu}_2 \cdot \mathbf{G}_I,\end{aligned}$$

where the final approximation hides $\mathbf{E}_1 \cdot \mathbf{G}^{-1}(\mathbf{C_2}) + \mu_1 \cdot (\mathbf{E}_2^\top)_{[m+1]\setminus I} + \hat{\mu}_1 (\mathbf{E}_2^\top)_I$ which can be considered as an error term when p is small. From the above expression, it can be observed that $[-\mathbf{s}^\top \mid 1] \cdot \mathbf{C}^\times \approx (\mu_1 \cdot \mu_2) \cdot [-\mathbf{s}^\top \mid 1] \cdot \mathbf{G}$ and $[\mathbf{t}^\top \mid 0] \cdot \mathbf{C}^\times \approx (\hat{\mu}_1 \cdot \hat{\mu}_2) \cdot [\mathbf{t}^\top \mid 0] \cdot \mathbf{G}$ as required. This follows from the fact that

$$[-\mathbf{s}^\top \mid 1] \cdot \mathbf{B} = [-\mathbf{s}^\top \mid 1] \cdot \mathbf{B}_{[m+1]\setminus I} = [-\mathbf{s}^\top \mid 1] \cdot \mathbf{B}_I = \mathbf{0}^{1\times n},$$

and

$$[\mathbf{t}^\top \mid 0] \cdot \mathbf{B} = [\mathbf{t}^\top \mid 0] \cdot \mathbf{B}_I = \mathbf{e}_I^\top \approx \mathbf{0}^{1\times n}, \qquad [\mathbf{t}^\top \mid 0] \cdot \mathbf{B}_{[m+1]\setminus I} = \mathbf{0}^{1\times n}.$$

Although $\mathbf{C}^\times$ does not exactly follow the distribution of a fresh ciphertext, one can still show that repeated multiplications are possible. To do so, simply note that one can write $\mathbf{C}^\times \approx \mathbf{B}' + \mu_1 \cdot \mu_2 \cdot \mathbf{G}_{[m+1]\setminus I} + \hat{\mu}_1 \cdot \hat{\mu}_2 \cdot \mathbf{G}_I$ where $[-\mathbf{s}^\top \mid 1] \cdot \mathbf{B}' = \mathbf{0}$ *and* $[\mathbf{t}^\top \mid 0] \cdot \mathbf{B}'$ is small. Now, suppose we have two such ciphertexts $\mathbf{C}_1^\times \approx \mathbf{B}' + \mathbf{J}_{\mu,\hat{\mu}} \otimes \mathbf{g}^\top$ and $\mathbf{C}_2^\times \approx \mathbf{B}'' + \mathbf{J}_{\mu',\hat{\mu}'} \otimes \mathbf{g}^\top$ where we have $[-\mathbf{s}^\top \mid 1] \cdot \mathbf{B}' = [-\mathbf{s}^\top \mid 1] \cdot \mathbf{B}'' = \mathbf{0}$ and the products $[\mathbf{t}^\top \mid 0] \cdot \mathbf{B}'$ and $[\mathbf{t}^\top \mid 0] \cdot \mathbf{B}''$ are both small. Then, we can write

$$\mathbf{C}^{\times\times} \triangleq \mathbf{C}_1^\times \cdot \mathbf{G}^{-1}(\mathbf{C}_2^\times) \approx \overbrace{\mathbf{B}' \cdot \mathbf{G}^{-1}(\mathbf{C}_2^\times) + \mu' \cdot \mathbf{B}''_{[m+1]\setminus I} + \hat{\mu}' \cdot \mathbf{B}''_I}^{\triangleq \mathbf{B}'''} + \mu' \cdot \mu'' \cdot \mathbf{G}_{[m+1]\setminus I} + \hat{\mu}' \cdot \hat{\mu}'' \cdot \mathbf{G}_I$$

where $[-\mathbf{s}^\top \mid 1] \cdot \mathbf{B}''' = \mathbf{0}$ and $[\ \mathbf{t}^\top \mid 0] \cdot \mathbf{B}'''$ are small as required for correct decryption. This intuition can be extended to show correctness for more general computation (as shown by Theorem 4 and Supplementary Material D.1).

PARAMETERS. We consider the case of arithmetic circuits with $L_{\mathsf{add}} = \mathsf{poly}(\lambda)$ and $L_{\mathsf{mult}} = \widetilde{O}(1)$ as in Theorem 4. Assuming that $k = \lceil \log q \rceil = O(n^{\epsilon})$ for constant $\epsilon \in (0, 1/2)$, we have $M = \Theta(km)$, $n_0 = \frac{m}{2} - \lceil \frac{\sqrt{\lambda m}}{2} \rceil = \Theta(m)$ for the Hoeffding bound in Theorem 5, and $m = \Theta(kn)$ for the Leftover Hash Lemma. Also, to apply Lemma 4 in Theorem 5, we take $\sigma \geq 2\alpha q \cdot \sqrt{m+1}$, requiring $\alpha q = \omega(\sqrt{\log m})$. For the $\mathsf{LWE}_{n,\cdot,q,\chi}$ and $\mathsf{LWE}_{n_0-1,\cdot,q,\chi}$ assumptions, we consider the former as $n_0 > n$. Taking $\alpha q \geq 2\sqrt{n}$, we rely on lattice problems in dimension n with approximation factor $\gamma = \widetilde{O}(q \cdot \sqrt{n})$. If we take $q = \Omega(k \cdot (L_{\mathsf{add}} \cdot k^2 n)^{L_{\mathsf{mult}}} \cdot (\alpha q)^2 \cdot \sqrt{kn} \cdot kn)$ for Theorem 4, we can set $n = O(\lambda \cdot \log \gamma) = \widetilde{O}(L_{\mathsf{mult}} \cdot \lambda)$. We then conclude that a modulus of $q = \Omega(k \cdot (L_{\mathsf{add}} \cdot k^2 n)^{L_{\mathsf{mult}}} \cdot (kn)^{3/2})$ suffices.

Theorem 4. *Let $L_{\mathsf{add}} = \mathsf{poly}(\lambda), L_{\mathsf{mult}} = \widetilde{O}(1)$. The scheme is homomorphically correct with probability $1 - 2^{-\Omega(\lambda)}$ for depth-L_{mult} arithmetic circuits (with L_{add} additions at each level) according to Definition 7 if $q = \Omega(k \cdot (L_{\mathsf{add}} \cdot M)^{L_{\mathsf{mult}}} \cdot (\alpha q) \cdot \sigma \cdot m)$.* (The proof is available in Supplementary Material D.2.)

We analyze robustness properties in Supplementary Material E.2 and now focus on proving anamorphism and full asymmetry.

Theorem 5. *If $\sigma = 2\alpha q \cdot s$ with $s \geq \sqrt{1+m}$, the scheme is anamorphic under the $\mathsf{LWE}_{n,m,q,\chi}, \mathsf{HNF\text{-}LWE}_{n,m,q,\chi}$ and $\mathsf{HNF\text{-}LWE}_{n_0-1,n,q,\chi}$ assumptions with $\chi = D_{\mathbb{Z},\alpha q}$.*

Proof. The proof considers a sequence of games where we call W_i the event that the adversary outputs 1 in Game_i. The first game is the experiment RealG_{Π} of Definition 3 while the last game is $\mathsf{AnamorphicG}_{\Sigma}$.

Game_0**:** This is the real experiment RealG_{Π}. At each query $\mathcal{O}_e(\mathsf{pk}, \mu, \hat{\mu})$, the challenger samples $\mathbf{S} \hookleftarrow U(\mathbb{Z}_q^{n \times M})$, $\mathbf{E} \hookleftarrow (D_{\mathbb{Z}^{m+1},\sigma})^M$, and returns a ciphertext

$$\mathbf{C} = \begin{bmatrix} \mathbf{A}^\top \\ \mathbf{s}^\top \mathbf{A}^\top \end{bmatrix} \cdot \mathbf{S} + \mu \cdot \mathbf{G} + \mathbf{E}, \tag{13}$$

where $\mathbf{s} \sim \mathcal{P}^m$ is the secret key.

Game_1**:** We modify the encryption oracle of RealG_{Π}. At each query $\mathcal{O}_e(\mathsf{pk}, \mu, \hat{\mu})$, the challenger now samples $\mathbf{S} \hookleftarrow U(\mathbb{Z}_q^{n \times M})$, $\mathbf{E}_0 \hookleftarrow \chi^{m \times M}$, to compute

$$\mathbf{C}_0 = \big[\mathbf{C}_0[\cdot, 1] \mid \ldots \mid \mathbf{C}_0[\cdot, M]\big] = \mathbf{A}^\top \cdot \mathbf{S} + \mathbf{E}_0 \qquad \in \mathbb{Z}_q^{m \times M}. \tag{14}$$

Then, it uses the ReRand algorithm of Lemma 4 to compute

$$\mathbf{C}'_0 = \big[\mathbf{C}'_0[\cdot, 1] \mid \ldots \mid \mathbf{C}'_0[\cdot, M]\big] = \mathsf{ReRand}([\mathbf{I}_m \mid \mathbf{s}], \mathbf{C}_0, \alpha q, s) \in \mathbb{Z}_q^{(m+1) \times M}, \tag{15}$$

with $s \geq \sqrt{1+m}$ and where the right-hand-side member of (15) means that

$$\mathbf{C}'_0[\cdot, j] = \mathsf{ReRand}([\mathbf{I}_m \mid \mathbf{s}], \mathbf{C}_0[\cdot, j], \alpha q, s) \in \mathbb{Z}_q^{m+1} \qquad \forall j \in [M]. \tag{16}$$

Then, the ciphertext is obtained as

$$\mathbf{C} = \mathbf{C}'_0 + \mu \cdot \mathbf{G}, \tag{17}$$

and returned to $\mathcal{A}$. Since $\|\mathbf{s}\| \leq \sqrt{m}$, we can apply Lemma 4 to each column of $\mathbf{C}_0$ with $\mathbf{V} = [\mathbf{I}_m \mid \mathbf{s}] \in \mathbb{Z}^{m\times(m+1)}$, $\mathbf{b} = \mathbf{A}^\top \cdot \mathbf{S}[\cdot, j] \in \mathbb{Z}_q^{m+1}$, and $r = \alpha q$ to ensure that the distribution of $\mathbf{C}'_0[\cdot, j]$ is within statistical distance $2^{-\Omega(\lambda)}$ from the distribution obtained by computing

$$\mathbf{C}'_0[\cdot, j] = \begin{bmatrix} \mathbf{A}^\top \\ \mathbf{s}^\top \mathbf{A}^\top \end{bmatrix} \cdot \mathbf{S}[\cdot, j] + \mathbf{E}[\cdot, j],$$

where $\mathbf{E}[\cdot, j] \sim D_{\mathbb{Z}^{m+1}, 2\alpha q s}$. This implies $|\Pr[W_2] - \Pr[W_1]| \leq M \cdot Q \cdot 2^{-\Omega(\lambda)}$, where Q is the number of queries to the oracle $\mathcal{O}_e$.

Game_2**:** In this game, we change the output distribution of the encryption oracle. At each query $\mathcal{O}_e(\mathsf{pk}, \mu, \hat{\mu})$, instead of computing a ciphertext as in (14)-(17), the challenger replaces the pseudorandom matrix $\mathbf{C}_0 = \mathbf{A}^\top \cdot \mathbf{S} + \mathbf{E}_0 \in \mathbb{Z}_q^{m\times M}$ of (14) by a truly random one $\mathbf{C}_0 \hookleftarrow U(\mathbb{Z}_q^{m\times M})$. Then, it computes the ciphertext $\mathbf{C} \in \mathbb{Z}_q^{(m+1)\times M}$ as per (15)-(17). By a standard hybrid argument over the columns of $\mathbf{C}_0$ and the queries to the encryption oracle $\mathcal{O}_e$, we obtain $|\Pr[W_2] - \Pr[W_1]| \leq M \cdot Q \cdot \mathbf{Adv}^{\mathsf{LWE}}_{n,m,q,\chi}(\lambda)$.

Game_3**:** We change again the encryption oracle. At each query $\mathcal{O}_e(\mathsf{pk}, \mu, \hat{\mu})$, instead of computing the ciphertext as per (17), the challenger computes

$$\mathbf{C} = \mathsf{ReRand}([\mathbf{I}_m \mid \mathbf{s}], \mathbf{C}_0 + (\hat{\mu} - \mu) \cdot \bar{\mathbf{G}}_I, \alpha q, s) + \mu \cdot \mathbf{G}, \tag{18}$$

where $\bar{\mathbf{G}}_I \in \mathbb{Z}_q^{m\times M}$ denotes the matrix comprised of the first m rows of $\mathbf{G}_I$. We note that this change amounts to replacing $\mathbf{C}_0$ by $\mathbf{C}_0 + (\hat{\mu} - \mu) \cdot \bar{\mathbf{G}}_I$ in (15). However, this change does not affect the distribution of $\mathbf{C}$ since $\mathbf{C}_0 \sim U(\mathbb{Z}_q^{m\times M})$ and thus $\mathbf{C}_0 + (\hat{\mu} - \mu) \cdot \bar{\mathbf{G}}_I \sim U(\mathbb{Z}_q^{m\times M})$. This implies $\Pr[W_3] = \Pr[W_2]$. Moreover, in the last row of (18), we have $\mathbf{s}^\top \bar{\mathbf{G}}_I = \mathbf{0}^{1\times M}$, so that $\mathbf{s}^\top \cdot \mathbf{C}_0 = \mathbf{s}^\top \cdot (\mathbf{C}_0 + (\hat{\mu} - \mu) \cdot \bar{\mathbf{G}}_I)$. Hence, the ciphertext of (18) can equivalently be computed as

$$\mathbf{C} = \mathsf{ReRand}([\mathbf{I}_m \mid \mathbf{s}], \mathbf{C}_0, \alpha q, s) + \mu \cdot \mathbf{G}_{[m+1]\setminus I} + \hat{\mu} \cdot \mathbf{G}_I. \tag{19}$$

Game_4**:** We change again the output distribution of $\mathcal{O}_e$. At each query $\mathcal{O}_e(\mathsf{pk}, \mu, \hat{\mu})$, the challenger now computes $\mathbf{C}_0$ as

$$\mathbf{C}_0 = \big[\mathbf{C}_0[\cdot, 1] \mid \ldots \mid \mathbf{C}_0[\cdot, M]\big] = \mathbf{A}^\top \cdot \mathbf{S} + \mathbf{E}_0 \quad \in \mathbb{Z}_q^{m\times M} \tag{20}$$

where $\mathbf{S} \hookleftarrow \chi^{n\times M}$, $\mathbf{E}_0 \hookleftarrow \chi^{m\times M}$, before computing the ciphertext as per (19). At each query, it thus replaces the truly uniform $\mathbf{C}_0 \in \mathbb{Z}_q^{m\times M}$ in (19) by a matrix whose columns are HNF-LWE samples. Under the $\mathsf{HNF\text{-}LWE}_{n,m,q,\chi}$ assumption, this change does not affect $\mathcal{A}$'s view and we obtain the inequality $|\Pr[W_4] - \Pr[W_3]| \leq M \cdot Q \cdot \mathbf{Adv}^{\mathsf{HNF\text{-}LWE}}_{n,m,q,\chi}(\lambda)$ via a standard hybrid argument over the columns of $\mathbf{C}_0$ and the queries to $\mathcal{O}_e$.

Game_5**:** Now, the challenger answers all encryption queries $\mathcal{O}_e(\mathsf{pk}, \mu, \hat{\mu})$ by sampling $\mathbf{S} \hookleftarrow \chi^{n\times M}$, $\mathbf{E} \hookleftarrow (D_{\mathbb{Z}^{m+1},\sigma})^M$ and computing

$$\mathbf{C} = \begin{bmatrix} \mathbf{A}^\top \\ \mathbf{s}^\top \cdot \mathbf{A}^\top \end{bmatrix} \cdot \mathbf{S} + \mu \cdot \mathbf{G}_{[m+1]\setminus I} + \hat{\mu} \cdot \mathbf{G}_I + \mathbf{E} \tag{21}$$

By applying Lemma 4 in the same way as in the transition from Game_1 to Game_2 (but in the converse direction), we find that the distribution of ciphertexts is statistically close to that of Game_4. If Q is the number of queries to the oracle $\mathcal{O}_e$, we have $|\Pr[W_5] - \Pr[W_4]| \leq M \cdot Q \cdot 2^{-\Omega(\lambda)}$.

Game_6**:** This game is like Game_5 except that, at step 2 of KGen, the challenger aborts if $\mathbf{s} \hookleftarrow \mathcal{P}^m$ contains less than n_0 zeroes. We claim that this only happens with negligible probability. Since the entries of $\mathbf{s}$ are sampled independently from $\mathcal{P}$, we have $\mathbb{E}[|J|] = m/2$ at step 2. By Hoeffding's inequality, we then have

$$\Pr_{\mathbf{s}\hookleftarrow\mathcal{P}^m}\left[\ \left||J| - \tfrac{m}{2}\right| > \left\lceil\sqrt{\lambda m/2}\right\rceil\ \right] \leq 2 \cdot \exp(-\lambda),$$

which means that $|J| \geq \frac{m}{2} - \left\lceil\sqrt{\lambda m/2}\right\rceil \geq n_0$ with probability at least $1 - 2\exp(-\lambda)$. Therefore, $|\Pr[W_6] - \Pr[W_5]| \leq 2 \cdot \exp(-\lambda) \leq 2^{-\Omega(\lambda)}$

Game_7**:** In this game, the challenger modifies the key generation phase. Instead of sampling $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ uniformly, the challenger first samples $\mathbf{s} \hookleftarrow \mathcal{P}^m$ and then computes $\mathbf{t} \in \mathbb{Z}^m$ and $\mathbf{A}$ as in aGen, by first choosing a matrix $\bar{\mathbf{A}} \in \mathbb{Z}_q^{n_0\times m}$ of the form (12). Lemma 11 shows that Game_7 is indistinguishable from Game_6 under the $\mathsf{HNF\text{-}LWE}_{n_0-1,n,q,\chi}$ assumption.

In Game_7, the challenger is interacting with $\mathcal{A}$ as in the $\mathsf{AnamorphicG}_\Sigma$ experiment. We then obtain the following inequality which proves the result:

$$\begin{aligned}&\mathbf{Adv}^{\mathsf{Anamorphism}}_{\mathcal{D},\Pi,\Sigma}(\lambda) = |\Pr[W_0] - \Pr[W_7]| \\ &\leq M \cdot Q \cdot \left(\mathbf{Adv}^{\mathsf{LWE}}_{n,m,q,\chi}(\lambda) + \mathbf{Adv}^{\mathsf{HNF\text{-}LWE}}_{n,m,q,\chi}(\lambda) + 2^{-\Omega(\lambda)}\right) + \mathbf{Adv}^{\mathsf{HNF\text{-}LWE}}_{n_0-1,n,q,\chi}(\lambda).\end{aligned}$$

□

Lemma 11. *We have* $|\Pr[W_7] - \Pr[W_6]| \leq \mathbf{Adv}^{\mathsf{HNF\text{-}LWE}}_{n_0-1,n,q,\chi}(\lambda)$*, where* $\chi = D_{\mathbb{Z},\alpha q}$*.*

Proof. We will use the same notation as in the description of aGen to build an HNF-LWE distinguisher $\mathcal{D}$ from an adversary $\mathcal{A}$ that can tell apart Game_6 and Game_7 with non-negligible advantage.

Algorithm $\mathcal{D}$ first samples $\mathbf{s} \hookleftarrow \mathcal{P}^m$ as in aGen to obtain a set of indices $I \subseteq \{i \in [n-1] : s_i = 0\}$ of size $|I| = n_0$. Let $I = \{i_1, \ldots, i_{n_0}\}$. Then, $\mathcal{D}$ obtains an LWE challenge parsed as $\bar{\mathbf{A}} = \begin{bmatrix} \mathbf{A}_{\mathsf{LWE}} \\ \mathbf{b}^\top \end{bmatrix} \in \mathbb{Z}_q^{n_0\times n}$. It then samples $\mathbf{A}$ in such a way that, for each $j \in [n_0]$, the i_j-th row of $\mathbf{A}$ is the j-th row of $\bar{\mathbf{A}}$ (as in aGen). It then runs $\mathcal{A}$ on input of a public key containing $\mathbf{B}^\top = [\mathbf{A} \mid \mathbf{A} \cdot \mathbf{s}] \in \mathbb{Z}_q^{n\times(m+1)}$ and plays the role of $\mathcal{A}$'s challenger in Game_6. Clearly, if $\mathbf{b}$ is uniform over

$\mathbb{Z}_q^n$, $\mathcal{D}$ perfectly simulates the challenger of Game_6 since $\mathbf{A} \sim U(\mathbb{Z}_q^{n\times m})$. In contrast, if $\mathbf{b} = \mathbf{A}_{\mathsf{LWE}}^\top \cdot \mathbf{t}' + \mathbf{e}_I$ is an LWE sample for some $\mathbf{t}' \sim \chi^{n_0-1}$ and $\mathbf{e}_I \sim \chi^n$, then $\mathcal{D}$ perfectly simulates the challenger of Game_7 since $\bar{\mathbf{A}}$ has the same distribution as in aGen. Therefore, $\mathcal{D}$ can output whatever $\mathcal{A}$ outputs and succeed as a distinguisher if $\mathcal{A}$'s output distribution in Game_7 significantly differs from that of Game_6. This yields the stated inequality. □

Theorem 6. *The dual GSW scheme is fully asymmetric with respect to the anamorphic triplet above under the* $\mathsf{HNF\text{-}LWE}_{n_0-1,n,q,\chi}$ *and* $\mathsf{HNF\text{-}LWE}_{n,m+1,q,\chi}$ *assumptions with* $\chi = D_{\mathbb{Z},\alpha q}$ *if* $\sigma = 2\alpha q \cdot s$ *and* $\alpha q \cdot (2s-1) > \sqrt{M}$.

Proof. We use a sequence of games that starts with the real $\mathsf{Fasym}^b_{\mathcal{A},\Pi,\Sigma}(\lambda)$ experiment. We gradually modify this experiment to reach a game where the challenge ciphertext does not carry any information on the bit b. For each i, we call W_i^b the event that the adversary $\mathcal{A}$ outputs 1 in Game_i^b. We only analyze cases conditioning on aGen succeeding as, if aGen aborts, $\mathcal{A}$ has no advantage.

Game_0^b**:** This is the $\mathsf{Fasym}^b_{\mathcal{A},\Pi,\Sigma}(\lambda)$ game where, in the challenge phase, $\mathcal{A}$ chooses pairs $(\mu_0, \hat{\mu}_0)$, $(\mu_1, \hat{\mu}_1)$ and obtains an anamorphic encryption of $(\mu_b, \hat{\mu}_b)$.

Game_1^b**:** This is like Game_0^b except that, in the execution of aGen, the challenger replaces the pseudorandom $\bar{\mathbf{A}}^\top$ of (12) by a uniform matrix $\bar{\mathbf{A}}^\top \hookleftarrow U(\mathbb{Z}_q^{n_0\times n})$. Under the $\mathsf{HNF\text{-}LWE}_{n_0-1,n,q,\chi}$ assumption (since the secret $\mathbf{t}' \in \mathbb{Z}^{n_0-1}$ is sampled from the distribution χ), Game_1^b is indistinguishable from Game_0^b and we have $|\Pr[W_1^b] - \Pr[W_0^b]| \leq \mathbf{Adv}^{\mathsf{LWE}}_{n_0-1,n,q,\chi}(\lambda)$. In Game_0^b, we note that aGen generates $\mathbf{A}$ by re-arranging the columns of $\bar{\mathbf{A}}$ and adding random columns. This implies that, in Game_1^b, $\mathbf{A}$ is uniform over $\mathbb{Z}_q^{n\times m}$.

Game_2^b**:** This game is like Game_1^b except that, during aGen, the challenger replaces $\mathbf{B}$ by a uniform matrix in $\mathbb{Z}_q^{(n+1)\times m}$ instead of choosing the last row as $\mathbf{A}\cdot\mathbf{s} \in \mathbb{Z}_q^n$ with $\mathbf{s} \hookleftarrow \mathcal{P}^m$. Since $\mathbf{A}$ is uniformly distributed in Game_1^b, we can use Lemma 1 (with the trivial function f) to argue that $\mathbf{A}\cdot\mathbf{s}$ is statistically uniform. Namely, since $\mathbf{A}\cdot\mathbf{s} = \sum_{j\in I}\mathbf{A}[\cdot,j]\cdot\mathbf{s}[j] + \sum_{j\in[m]\setminus I}\mathbf{A}[\cdot,j]\cdot\mathbf{s}[j] \in \mathbb{Z}_q^n$, we can rely on the min-entropy of $\{\mathbf{s}[j]\}_{j\in[m]\setminus I}$ (which is not affected by the revealed $\{\mathbf{s}[j]\}_{j\in I}$ since the entries of $\mathbf{s}$ are sampled independently from $\mathcal{P}$) to claim that the distribution of $\sum_{j\in[m]\setminus I}\mathbf{A}[\cdot,j]\cdot\mathbf{s}[j] \bmod q$ is statistically close to $U(\mathbb{Z}_q^n)$ due to our choice of $|[m]\setminus I| = m - n_0 > n\log q + 2\lambda$. The joint distribution $(\mathbf{A} \mid \mathbf{A}\cdot\mathbf{s} \bmod q)$ is thus itself statistically close to $U(\mathbb{Z}_q^{n\times m}\times\mathbb{Z}_q^n)$ and we have $|\Pr[W_2^b] - \Pr[W_1^b]| \leq 2^{-\lambda}$.

Game_3^b**:** In this game, we now compute the challenge ciphertext as

$$\mathbf{C} = \underbrace{\left(\mathbf{B}\cdot\mathbf{S} + \mathbf{E}'\right)}_{\triangleq\, \mathbf{C}'} + \mu_b \cdot \mathbf{G}_{[m+1]\setminus I} + \hat{\mu}_b \cdot \mathbf{G}_I + \mathbf{E}''. \tag{22}$$

where $\mathbf{S} \hookleftarrow \chi^{n\times M}$, $\mathbf{E}' \hookleftarrow \chi^{(m+1)\times M}$, and $\mathbf{E}'' \hookleftarrow \left(D_{\mathbb{Z}^{m+1},\sqrt{\sigma^2-(\alpha q)^2}}\right)^M$. When applying Lemma 5 with $\sigma_1 = \alpha q$ and $\sigma_2 = \sqrt{\sigma^2 - (\alpha q)^2}$, we obtain $\tau =$

$\sigma_1\sigma_2/\sqrt{\sigma_1^2+\sigma_2^2} = \alpha q\sqrt{4s^2-1} \geq \alpha q \cdot (2s-1) > \sqrt{M} > \eta_{2^{-M}}(\mathbb{Z}^M)$, so that the distribution of $\mathbf{C}$ in (22) is statistically close to the distribution of

$$\mathbf{C} = \mathbf{B}\cdot\mathbf{S} + \mu_b \cdot \mathbf{G}_{[m+1]\setminus I} + \hat{\mu}_b \cdot \mathbf{G}_I + \mathbf{E},$$

where $\mathbf{S} \hookleftarrow \chi^{n\times M}$, $\mathbf{E} \hookleftarrow (D_{\mathbb{Z}^{m+1},\sigma})^M$, which is the ciphertext distribution of Game_2^b. This implies $|\Pr[W_3^b] - \Pr[W_2^b]| \leq 2^{-\Omega(\lambda)}$.

Game_4^b: This game is identical to Game_3^b except that, in the challenge ciphertext, the challenger replaces the pseudorandom matrix $\mathbf{C}' = \mathbf{B}\cdot\mathbf{S}+\mathbf{E}'$ by a uniform $\mathbf{C}' \hookleftarrow U\left(\mathbb{Z}_q^{(m+1)\times M}\right)$ in (22). Under the $\mathsf{HNF\text{-}LWE}_{n,m+1,q,\chi}$ assumption, Game_4^b is indistinguishable from Game_3^b. By a standard hybrid argument over the columns of $\mathbf{C}'$ we have $|\Pr[W_4^b] - \Pr[W_3^b]| \leq M \cdot \mathbf{Adv}^{\mathsf{HNF\text{-}LWE}}_{n,m+1,q,\chi}(\lambda)$.

In Game_4^b, the pair $(\mu_b, \hat{\mu}_b)$ is perfectly hidden by the uniform matrix $\mathbf{C}'$. Therefore, we have $\Pr[W_4^0] = \Pr[W_4^1]$ and the triangle inequality then implies

$$\mathbf{Adv}^{\mathsf{Fasym}}_{\mathcal{A},\Pi,\Sigma}(\lambda) \leq 2\cdot\left(2^{-\Omega(\lambda)} + M\cdot\mathbf{Adv}^{\mathsf{HNF\text{-}LWE}}_{n,m+1,q,\chi}(\lambda) + \mathbf{Adv}^{\mathsf{HNF\text{-}LWE}}_{n_0-1,n,q,\chi}(\lambda)\right).$$

□

References

1. Agrawal, S., Boneh, D., Boyen, X.: Efficient lattice (H)IBE in the standard model. In: Eurocrypt (2010). https://doi.org/10.1007/978-3-642-13190-5_28
2. Agrawal, S., Gentry, G., Halevi, S., Sahai, A.: Discrete gaussian leftover hash lemma over infinite domains. In: Asiacrypt (2013). https://doi.org/10.1007/978-3-642-42033-7_6
3. Agrawal, S., Libert, B., Stehlé, D.: Fully secure functional encryption for inner products from standard assumptions. In: Crypto (2016). https://doi.org/10.1007/978-3-662-53015-3_12
4. Alwen, J., Krenn, S., Pietrzak, K., Wichs, D.: Learning with rounding, revisited - new reduction, properties and applications. In: Crypto (2013). https://doi.org/10.1007/978-3-642-40041-4_4
5. Ananth, P., Deshpande, A., Tauman Kalai, Y., Lysyanskaya, A.: Fully homomorphic NIZK and NIWI proofs. In: TCC (2019). https://doi.org/10.1007/978-3-030-36033-7_14
6. Applebaum, B., Cash, D., Peikert, C., Sahai, A.: Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In: Crypto (2009). https://doi.org/10.1007/978-3-642-03356-8_35
7. Banaszczyk, W.: New bounds in some transference theorems in the geometry of number. Mathematische Annalen **296**, 625–635 (1993)
8. Banfi, F., Gegier, K., Hirt, M., Maurer, U., Rito, G.: Anamorphic encryption, revisited. In: Eurocrypt (2024). https://doi.org/10.1007/978-3-031-58723-8_1
9. Barak, B., Goldreich, O., Impagliazzo, R., Rudich, S., Sahai, A., Vadhan, S., Yang, K.: On the (im)possibility of obfuscating programs. In: Crypto (2001). https://doi.org/10.1007/3-540-44647-8_1

10. Bellare, M., Rogaway, P.: Optimal asymmetric encryption. In: Eurocrypt (1994). https://doi.org/10.1007/BFB0053428
11. Blum, M., Feldman, P., Micali, S.: Non-interactive zero-knowledge and its applications. In: STOC (1988). https://doi.org/10.1145/62212.62222
12. Boneh, D., Canetti, R., Halevi, S., Katz, J.: Chosen-ciphertext security from identity-based encryption. SIAM Journal on Computing **36**(5) (2007). https://doi.org/10.1137/S009753970544713X
13. Boneh, D., Freeman, D.: Linearly homomorphic signatures over binary fields and new tools for lattice-based signatures. In: PKC (2011). https://doi.org/10.1007/978-3-642-19379-8_1
14. Brakerski, Z., Halevi, S., Polychroniadou, A.: Four round secure computation without setup. In: TCC (2017). https://doi.org/10.1007/978-3-319-70500-2_22
15. Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors. In: STOC (2013). https://doi.org/10.1145/2488608.2488680
16. Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) LWE. In: FOCS (2011). https://doi.org/10.1109/FOCS.2011.12
17. Brakerski, Z., Vaikuntanathan, V.: Constrained key-homomorphic PRFs from LWE (or) how to secretly embed a circuit in your PRF. In: TCC (2015). https://doi.org/10.1007/978-3-662-46497-7_1
18. Catalano, D., Giunta, E., Migliaro, F.: Anamorphic encryption: New constructions and homomorphic realizations. In: Eurocrypt (2024). https://doi.org/10.1007/978-3-031-58723-8_2
19. Catalano, D., Giunta, E., Migliaro, F.: Limits of black-box anamorphic encryption. In: Crypto (2024). https://doi.org/10.1007/978-3-031-68379-4_11
20. Catalano, D., Giunta, E., Migliaro, F.: Generic anamorphic encryption, revisited: New limitations and constructions. In: Eurocrypt (2025), https://eprint.iacr.org/2024/1119
21. Cramer, R., Shoup, V.: A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In: Crypto (1998). https://doi.org/10.1007/BFB0055717
22. Dodis, Y., Goldwasser, S., Kalai, Y., Peikert, C., Vaikuntanathan, V.: Public key encryption schemes with auxiliary inputs. In: TCC (2010). https://doi.org/10.1007/978-3-642-11799-2_22
23. Gentry, C., Halevi, S., Vaikuntanathan, V.: i-Hop Homomorphic Encryption and Rerandomizable Yao Circuits. In: Crypto (2010). https://doi.org/10.1007/978-3-642-14623-7_9
24. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: STOC (2008). https://doi.org/10.1145/1374376.1374407
25. Gentry, C., Sahai, A., Waters, B.: Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In: Crypto (2013). https://doi.org/10.1007/978-3-642-40041-4_5
26. Goldwasser, S., Kalai, Y., Peikert, C., Vaikuntanathan, V.: Robustness of the Learning with Errors assumption. In: ICS (2010)
27. Goldwasser, S., Micali, S.: Probabilistic encryption and how to play mental poker keeping secret all partial information. In: STOC (1982). https://doi.org/10.1145/800070.802212
28. Katsumata, S., Yamada, S.: Partitioning via non-linear polynomial functions: More compact IBEs from ideal lattices and bilinear maps. In: Asiacrypt (2016). https://doi.org/10.1007/978-3-662-53890-6_23

29. Katsumata, S., Yamada, S., Yamakawa, T.: Tighter security proofs for GPV-IBE in the quantum random oracle model. In: Asiacrypt (2018). https://doi.org/10.1007/978-3-030-03329-3_9
30. Koppula, V., Waters, B.: Realizing chosen ciphertext security generically in attribute-based encryption and predicate encryption. In: Crypto (2019). https://doi.org/10.1007/978-3-030-26951-7_23
31. Kutylowski, M., Persiano, G., Phan, D.H., Yung, M., Zawada, M.: The self-anti-censorship nature of encryption: On the prevalence of anamorphic cryptography. In: PoPETS. No. 4 (2023). https://doi.org/10.56553/POPETS-2023-0104
32. Libert, B., Sakzad, A., Stehlé, D., Steinfeld, R.: All-but-many lossy trapdoor functions and selective opening chosen-ciphertext security from LWE. In: Crypto (2017). https://doi.org/10.1007/978-3-319-63697-9_12
33. Litvak, A., Pajor, A., Rudelson, M., Tomczak-Jaegermann, N.: Smallest singular value of random matrices and geometry of random polytopes. Advances in Mathematics **195**(2) (2005)
34. Lyubashevsky, V.: Lattice signatures without trapdoors. In: Eurocrypt (2012). https://doi.org/10.1007/978-3-642-29011-4_34
35. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: Eurocrypt (2010). https://doi.org/10.1007/978-3-642-13190-5_1
36. Micciancio, D.: Fully composable homomorphic encryption. Cryptology ePrint Archive Report 2024/1545
37. Micciancio, D., Peikert, C.: Trapdoors for lattices: Simpler, tighter, faster, smaller. In: Eurocrypt (2012). https://doi.org/10.1007/978-3-642-29011-4_41
38. Micciancio, D., Regev, O.: Worst-case to average-case reductions based on Gaussian measures. SIAM J. Comput **37**(1) (2007). https://doi.org/10.1137/S0097539705447360
39. Naor, M., Yung, M.: Public-key cryptosystems provably secure against chosen ciphertext attacks. In: STOC (1990). https://doi.org/10.1145/100216.100273
40. Paillier, P.: Public-key cryptosystems based on composite degree residuosity classes. In: Eurocrypt (1999). https://doi.org/10.1007/3-540-48910-X_16
41. Peikert, C., Shiehian, S.: Non-interactive Zero Knowledge for NP from (Plain) Learning With Errors. In: Crypto (2019). https://doi.org/10.1007/978-3-030-26948-7_4
42. Peikert, C., Vaikuntanathan, V., Waters, B.: A framework for efficient and composable oblivious transfer. In: Crypto (2008). https://doi.org/10.1007/978-3-540-85174-5_31
43. Persiano, G., Phan, D.H., Yung, M.: Anamorphic encryption: Private communication against a dictator. In: Eurocrypt (2022). https://doi.org/10.1007/978-3-031-07085-3_2
44. Persiano, G., Phan, D.H., Yung, M.: Public-Key Anamorphism in (CCA-Secure) Public-Key Encryption. In: Crypto (2024). https://doi.org/10.1007/978-3-031-68379-4_13
45. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: STOC (2005). https://doi.org/10.1145/1060590.1060603
46. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. J. ACM **56**(6) (2009). https://doi.org/10.1145/1568318.1568324
47. Schnorr, C.: A hierarchy of polynomial time lattice basis reduction algorithms. Theor. Comput. Sci. **53** (1987). https://doi.org/10.1016/0304-3975(87)90064-8

48. Wang, Y., Chen, R., Huang, X., Yung, M.: Sender-anamorphic encryption reformulated: Achieving robust and generic constructions. In: Asiacrypt (2023). https://doi.org/10.1007/978-981-99-8736-8_5
49. Waters, B.: A new approach for non-interactive zero-knowledge from learning with errors. In: STOC (2024). https://doi.org/10.1145/3618260.3649683
50. Yao, A.: How to generate and exchange secrets. In: FOCS (1986). https://doi.org/10.1109/SFCS.1986.25
51. Zhang, J., Yu, Y., Fan, S., Zhang, Z.: Improved lattice-based CCA2-secure PKE in the standard model. Sci. China Inf. Sci. **63**(8) (2020). https://doi.org/10.1007/S11432-019-9861-3, https://eprint.iacr.org/2019/149, Cryptology ePrint Archive Report 2019/149.

Supplementary Material

A Deferred Proofs for the Public-Key Anamorphic Dual Regev Scheme 3

A.1 Norm Bound for a Random Ternary Matrix

Lemma 12. *Let a matrix* $\mathbf{R}' \in \{-1,0,1\}^{m\times n}$ *sampled from the distribution* $\mathcal{P}^{m\times n}$. *There is a universal constant* $C > 0$ *(the same constant as in Lemma 8) such that*

$$\Pr_{\mathbf{R}'}[\|\mathbf{R}'\| > C\sqrt{m+n}] \le 2\cdot\exp(-(m+n)).$$

Proof. The distribution $\mathcal{P}$ can be obtained as the halved difference between two independent uniform samples from $U(\{-1,1\})$. In other words, we may write $\mathbf{R}' = (\mathbf{S}-\mathbf{T})/2$ where $\mathbf{S},\mathbf{T} \sim U(\{-1,1\})^{m\times n}$. By the triangle inequality, $\|\mathbf{R}'\| \le (\|\mathbf{S}\| + \|\mathbf{T}\|)/2$. Setting $C_{m,n} = C\sqrt{m+n}$ we have

$$\begin{aligned}\Pr_{\mathbf{R}'}[\|\mathbf{R}'\| > C_{m,n}] &\le \Pr_{\mathbf{S},\mathbf{T}}[\|\mathbf{S}\| + \|\mathbf{T}\| > 2\cdot C_{m,n}]\\ &\le \Pr_{\mathbf{S},\mathbf{T}}[\|\mathbf{S}'\| > C_{m,n} \vee \|\mathbf{T}\| > C_{m,n}]\\ &\le \Pr_{\mathbf{S}}[\|\mathbf{S}\| > C_{m,n}] + \Pr_{\mathbf{T}}[\|\mathbf{T}\| > C_{m,n}]\end{aligned}$$

by a union bound. Lemma 8 then implies the result. □

A.2 Proof of Lemma 9

Proof. We first show the correctness of aDec with overwhelming probability over the randomness of aGen and aEnc. Lemma 3 implies that $\|\mathbf{s}\|_\infty \le \alpha q\sqrt{2\lambda}$ with probability $\ge 1 - 2n\cdot\exp(-\lambda)$ over the randomness of aEnc, in which case the decode_p algorithm of Lemma 7 recovers $\hat{\boldsymbol{\mu}} = \mathsf{decode}_p(\mathbf{s} + \mathsf{encode}_p(\hat{\boldsymbol{\mu}}))$.

We now show that Invert correctly recovers $(\mathbf{s},\mathbf{e}_0)$ from $\mathbf{c}_0 = \mathbf{A}^\top\cdot\mathbf{s} + \mathbf{e}_0$ with overwhelming probability over the randomness of aGen and aEnc. Recall that Hoeffding's inequality says that, if $(X_i)_{i\in[\ell]}$ are random independent real

variables satisfying $\Pr[a_i \le X_i \le b_i] = 1$ for real numbers $(a_i)_{i\in[\ell]}$, $(b_i)_{i\in[\ell]}$ such that $a_i < b_i$ for all $i \in [\ell]$, then $S_\ell \triangleq \sum_{i=1}^{\ell} X_i$ satisfies

$$\Pr\left[|S_\ell - \mathbb{E}[S_\ell]| \ge t\right] \le 2 \cdot \exp\left(-\frac{2t^2}{\sum_{i=1}^{\ell}(a_i - b_i)^2}\right)$$

By the same argument as in [1, Lemma 16], we can view each row of $\mathbf{R}^\top \cdot \mathbf{e}_{0,1}$ as a sum of random variables $X_i \in \{-\mathbf{e}_{0,1}[i], 0, \mathbf{e}_{0,1}[i]\}$ such that $\mathbb{E}[X_i] = 0$ for each $i \in [\bar{m}]$. When applying Hoeffding with $t = \sqrt{2\lambda} \cdot \|\mathbf{e}_{0,1}\|$, we obtain

$$\Pr_{\mathbf{r}_i \hookleftarrow \mathcal{P}^\ell}\left[\left|\langle \mathbf{r}_i, \mathbf{e}_{0,1}\rangle\right| > \sqrt{2\lambda} \cdot \|\mathbf{e}_{0,1}\|\right] \le 2 \cdot \exp(-\lambda).$$

where $\mathbf{r}_i$ denotes the i-th row of $\mathbf{R}^\top$. By a union bound over the independent rows of $\mathbf{R}^\top$, the probability that $\|\mathbf{R}^\top \cdot \mathbf{e}_{0,1}\|_\infty > \sqrt{2\lambda} \cdot \|\mathbf{e}_{0,1}\|$ is at most $2nk \cdot \exp(-\lambda)$. Lemma 3 also implies $\|\mathbf{e}_{0,1}\| \le \sigma\sqrt{2\bar{m}}$ except with probability $\le (2/\exp(1))^{\bar{m}/2}$. Therefore we have $\|\mathbf{R}^\top \cdot \mathbf{e}_{0,1}\|_2 \le 2\sigma \cdot \sqrt{\lambda \bar{m} nk}$ except with negligible probability $2nk \cdot \exp(-\lambda) + (2/\exp(1))^{\bar{m}/2}$ over the randomness of aGen and aEnc. Since $\|\mathbf{e}_{0,2}\| \le \sigma\sqrt{2nk}$ with probability $\ge 1 - (2/\exp(1))^{nk/2}$ by Lemma 3, the inequality (3) is satisfied with overwhelming probability over the randomness of aGen and aEnc.

We now show the correctness of Dec for the suggested parameters. The normal decryption algorithm computes $\mathbf{c}_1 - \mathbf{E}^\top \cdot \mathbf{c}_0 = \Delta \cdot \boldsymbol{\mu} + (\mathbf{e}_1 - \mathbf{E}^\top \cdot \mathbf{e}_0)$ and correctly recovers $\boldsymbol{\mu}$ if $\|\mathbf{e}_1 - \mathbf{E}^\top \cdot \mathbf{e}_0\|_\infty \le q/(2p)$. By Hoeffding and an argument identical to above, we know that $\|\mathbf{E}^\top \cdot \mathbf{e}_0\|_\infty \le \sqrt{2\lambda} \cdot \|\mathbf{e}_0\|$ except with probability $\le 2n \cdot \exp(-\lambda)$ over the random choice of $\mathbf{E}$. Lemma 3 implies $\|\mathbf{e}_0\| \le \sigma\sqrt{2m}$ with probability $\ge 1 - (2/\exp(1))^{m/2}$ and $\|\mathbf{e}_1\|_\infty \le \sigma\sqrt{2\lambda}$ with probability $\ge 1 - 2n \cdot \exp(-\lambda)$. This implies $\|\mathbf{e}_1 - \mathbf{E}^\top \cdot \mathbf{e}_0\|_\infty < 3\sigma\sqrt{\lambda m} = O(m \cdot n^{1/2})$, which is smaller than $q/(2p) = \Theta(m^2)$ as required. □

B Deferred Proofs for the Primal-Regev-Based Construction

B.1 Proof of Lemma 10

Proof. In the anamorphic mode, the matrix $\mathbf{A}$ produced by aGen is of the form $\mathbf{A}^\top = \mathbf{B}^\top\mathbf{C} + \mathbf{F}$ with $\mathbf{C} \sim U(\mathbb{Z}_q^{\ell\times n})$, $\mathbf{B} \sim D_{\mathbb{Z}^m,\alpha q}^{\ell}$, $\mathbf{F} \sim D_{\mathbb{Z}^m,\alpha q}^{n}$. The anamorphic encryption algorithm generates ciphertexts of the form

$$\begin{aligned}
\mathbf{c}_0 &= \mathbf{A} \cdot \mathbf{r} + \mathbf{C}^\top \cdot \mathsf{encode}_p(\hat{\mu}) \\
&= (\mathbf{C}^\top\mathbf{B} + \mathbf{F}^\top) \cdot \mathbf{r} + \mathbf{C}^\top \cdot \mathsf{encode}_p(\hat{\mu}) \\
&= \mathbf{C}^\top \underbrace{(\mathbf{B} \cdot \mathbf{r} + \mathsf{encode}_p(\hat{\mu}))}_{\triangleq\ \hat{\mathbf{s}}'} + \underbrace{(\mathbf{F}^\top \cdot \mathbf{r})}_{\triangleq\ \mathbf{e}_0}
\end{aligned} \tag{23}$$

By Lemma 3 and the Cauchy-Schwarz inequality, we have $\|\mathbf{B} \cdot \mathbf{r}\|_\infty \le \sqrt{2}\alpha q m$ and $\|\mathbf{F}^\top \cdot \mathbf{r}\|_\infty \le \sqrt{2}\alpha q m$ with overwhelming probability $\ge 1 - (2/\exp(1))^{m/2}$

over the random choice of $\mathbf{B}$ and $\mathbf{F}$. In (23), if we parse $\mathbf{e}_0 = \mathbf{F}^\top \mathbf{r} \in \mathbb{Z}^n$ as $\mathbf{e}_0 = (\mathbf{e}_{0,1}^\top \mid \mathbf{e}_{0,2}^\top)^\top \in \mathbb{Z}^{\ell k + 2\lambda} \times \mathbb{Z}^{\ell k}$, we then have $\|\mathbf{e}_{0,1}\| \leq \alpha q m \sqrt{2(\ell k + 2\lambda)}$ and $\|\mathbf{e}_{0,2}\| \leq \alpha q m \sqrt{2\ell k}$. By the same argument as in Lemma 9, we get $\|\mathbf{R}_C^\top \cdot \mathbf{e}_{0,1}\|_\infty \leq \sqrt{2\lambda} \cdot \|\mathbf{e}_{0,1}\|$ and thus $\|\mathbf{R}_C^\top \cdot \mathbf{e}_{0,1}\| \leq \sqrt{2\lambda \ell k} \cdot \|\mathbf{e}_{0,1}\| < 2\alpha q m (\ell k + 2\lambda)\sqrt{\lambda}$. With overwhelming probability over the randomness of aGen and aEnc, we then have $\|\mathbf{e}_{0,2} - \mathbf{R}_C^\top \cdot \mathbf{e}_{0,1}\| < 3\alpha q m (\ell k + 2\lambda)\sqrt{\lambda}$. In order for Invert to recover $\hat{\mathbf{s}}'$ from $\mathbf{c}_0$ by applying Lemma 6, we can satisfy the condition $\|\mathbf{e}_{0,2} - \mathbf{R}_C^\top \cdot \mathbf{e}_{0,1}\| < q/(2\sqrt{k})$ if we choose q so that $q > 6\alpha q m (\ell \lambda^{1/2} k^{3/2} + 2\lambda^{3/2} k^{1/2})$, which is implied by $q > 18 \cdot \alpha q m \cdot \ell^{3/2} k^{3/2}$ if $\ell \geq \lambda$.

In order to guarantee that decode_p recovers $\hat{\boldsymbol{\mu}}$ from $\hat{\mathbf{s}}' = \mathbf{B} \cdot \mathbf{r} + \mathsf{encode}_p(\hat{\boldsymbol{\mu}})$, we need to have $\|\mathbf{B} \cdot \mathbf{r}\|_\infty \leq \frac{q-(p-1)p}{p}$ by Lemma 7. Since $\|\mathbf{B} \cdot \mathbf{r}\|_\infty \leq \sqrt{2}\alpha q m$, the latter condition is implied by $q > 18 \cdot \alpha q m \cdot \ell^{3/2} k^{3/2}$ if we set $p = O(1)$.

In a normal decryption, Dec computes $\mathbf{c}_1 - \mathbf{S}^\top \mathbf{c}_0 = \mathbf{E}^\top \mathbf{r} + \Delta \cdot \boldsymbol{\mu} \bmod q$ and outputs the correct plaintext $\boldsymbol{\mu} \in \mathbb{Z}_p^\ell$ as long as $\|\mathbf{E}^\top \mathbf{r}\|_\infty \leq q/(2p)$. Since $\|\mathbf{E}^\top \cdot \mathbf{r}\|_\infty \leq \sqrt{2}\alpha q m$ with overwhelming probability over the randomness of KGen, the latter condition is fulfilled if we set $p = O(1)$ and $q > 18 \cdot \alpha q m \cdot \ell^{3/2} k^{3/2}$. The plaintext modulus p can be increased to $p = \Theta(\ell)$ while keeping other parameters asymptotically unchanged. □

C Trapdoor-less Anamorphic Dual Regev Encryption

We now show that, in its non-packed version, the dual Regev system also admits an anamorphic mode that does not rely on lattice trapdoors. Although this modified anamorphic mode is no longer public-key anamorphic in the sense of [44], it remains fully asymmetric in the model of [18].

This variant has the property that its bandwidth rate [43] (i.e., the ratio between the length of anamorphic plaintexts and that of regular plaintexts) is larger than 1. For the packed schemes, the bandwidth rate was also 1 but only because we restricted regular plaintexts to be shorter than they could be. Indeed, these schemes could encrypt longer regular messages by increasing the length of secret keys as in packed LWE-based encryption schemes (in contrast, the length of anamorphic messages is limited by the LWE dimension).

In the construction below, anamorphic messages are naturally longer than regular ones. In [31], Kutylowski *et al.* showed that the Goldwasser-Micali cryptosystem [27] can similarly achieve a bandwidth rate larger than 1. However, it is not known to admit a fully asymmetric anamorphic mode. So far, the only known fully asymmetric anamorphic scheme offering a bandwidth rate larger than 1 was the construction of [44], which builds on the Koppula-Waters cryptosystem [30] and the technique used in its proof of CCA security. Our dual-Regev-based construction is very different and does not rely on a CCA-secure encryption scheme.

$\mathsf{KGen}(1^\lambda)$: Given a security parameter 1^λ,

1. Sample moduli $p, q = \mathsf{poly}(\lambda)$, dimensions $m, n = \mathsf{poly}(\lambda)$ and k such that $k \leq \frac{m}{2} - \lceil \sqrt{\lambda m/2} \rceil$, and $m - k > n \log q + 2\lambda$. Also, sample an error

rate $\alpha \in (0,1)$ such that $\alpha q = \Theta(\sqrt{n})$ and a standard deviation $\sigma > \alpha q$. Set $\chi = D_{\mathbb{Z},\alpha q}$ and define parameters $\mathsf{par} = (q, p, m, n, \alpha, \chi, \sigma)$.

2. Sample $\mathbf{A} \hookleftarrow U(\mathbb{Z}_q^{n \times m}), \mathbf{s} \hookleftarrow \mathcal{P}^m$ and set

$$\mathbf{B} = \begin{bmatrix} \mathbf{A}^\top \\ \mathbf{s}^\top \cdot \mathbf{A}^\top \bmod q \end{bmatrix} \in \mathbb{Z}_q^{(m+1) \times n}.$$

Return $\mathsf{pk} = (\mathsf{par}, \mathbf{B})$ and $\mathsf{sk} = \mathbf{s}$.

aGen(1^λ)**:** Given a security parameter 1^λ,

1. Run step 1 of KGen to generate parameters $\mathsf{par} = (q, p, m, n, \alpha, \chi, \sigma)$ and k such that $k \leq \frac{m}{2} - \lceil \sqrt{\lambda m / 2} \rceil$, and $m - k > n \log q + 2\lambda$.
2. Sample $\mathbf{s} \hookleftarrow \mathcal{P}^m$ and define $J \triangleq \{j \in [m] : s_j = 0\}$. If $|J| < k$, then abort. Otherwise, let $I \subseteq J$ be a uniformly chosen subset consisting of k elements of J and continue to the following steps.
3. Sample $\bar{\mathbf{A}} \hookleftarrow U(\mathbb{Z}_q^{(m-k) \times n}), \mathbf{T} \hookleftarrow \mathcal{P}^{(m-k) \times k}$, and set $\mathbf{A}'$ such that

$$(\mathbf{A}')^\top = \begin{bmatrix} \bar{\mathbf{A}} \\ \mathbf{T}^\top \cdot \bar{\mathbf{A}} \end{bmatrix} \in \mathbb{Z}_q^{m \times n}. \tag{24}$$

4. Write $I = \{i_1, \ldots, i_k\}$ and $[m] \setminus I = \{i_{k+1}, \ldots, i_{|J|}, \ldots, i_m\}$. Define $\mathbf{A}$ to be the matrix where

$$\begin{cases} \mathbf{A}^\top[i_{k+j}, \cdot] = (\mathbf{A}')^\top[j, \cdot] & \text{for } \ j \in [m-k] \\ \mathbf{A}^\top[i_j, \cdot] = (\mathbf{A}')^\top[m-k+j, \cdot] & \text{for } \ j \in [k] \end{cases}$$

and set

$$\mathbf{B} = \begin{bmatrix} \mathbf{A}^\top \\ \mathbf{s}^\top \cdot \mathbf{A}^\top \end{bmatrix} \in \mathbb{Z}_q^{(m+1) \times n}.$$

5. Set $\mathbf{T}' := [\ -\mathbf{T}^\top \mid \mathbf{I}_k\ ] \in \mathbb{Z}_q^{k \times m}$ and then define $\widetilde{\mathbf{T}} \in \mathbb{Z}^{k \times m}$ by reordering the columns of $\mathbf{T}'$ according to the following rules:
   - $\widetilde{\mathbf{T}}[\cdot, i_{k+j}] = \mathbf{T}'[\cdot, j] = \mathbf{T}^\top[\cdot, j]$ for $j \in [m-k]$
   - $\widetilde{\mathbf{T}}[\cdot, i_j] = \mathbf{T}'[\cdot, m-k+j]$ for $j \in [k]$ (i.e., $\widetilde{\mathbf{T}}[\cdot, i_j]$ is the j-th unit vector of dimension k).

   Note that by construction, $[\widetilde{\mathbf{T}} \mid \mathbf{0}] \cdot \mathbf{B} = \widetilde{\mathbf{T}} \cdot \mathbf{A}^\top = \mathbf{T}' \cdot \mathbf{A'}^\top = \mathbf{0}^{k \times n} \bmod q$.

Output $\mathsf{apk} = (\mathsf{par}, \mathbf{B})$, $\mathsf{ask} = \mathbf{s}$, $\mathsf{dk} = (k, I)$ and $\mathsf{tk} = \widetilde{\mathbf{T}}$.

Enc(pk, μ)**:** To encrypt $\mu \in \mathbb{Z}_p$ under $\mathsf{pk} = (\mathsf{par}, \mathbf{B})$, sample $\mathbf{r} \hookleftarrow U(\mathbb{Z}_q^n), \mathbf{e}_0 \hookleftarrow D_{\mathbb{Z}^m,\sigma}$, $e_1 \hookleftarrow D_{\mathbb{Z},\sigma}$ and set

$$\mathbf{c} = \begin{bmatrix} \mathbf{c}_0 \\ c_1 \end{bmatrix} = \mathbf{B} \cdot \mathbf{r} + \begin{bmatrix} \mathbf{e}_0 \\ e_1 \end{bmatrix} + \Delta \cdot \begin{bmatrix} \mathbf{0}^m \\ \mu \end{bmatrix},$$

where $\Delta = \lfloor q/p \rceil$.

aEnc$(\mathsf{apk}, \mathsf{dk}, \mu, \hat{\boldsymbol{\mu}})$**:** Given $\mathsf{apk} = (\mathsf{par}, \mathbf{B})$, $\mathsf{dk} = (k, I)$ and messages $\mu \in \mathbb{Z}_p$, $\hat{\boldsymbol{\mu}} = (\hat{\mu}_1, \ldots, \hat{\mu}_k)^\top \in \mathbb{Z}_p^k$, let $I = \{i_1, \ldots, i_k\}$ and $[m] \setminus I = \{i_{k+1}, \ldots, i_m\}$. Define $\mathbf{f}^{\hat{\boldsymbol{\mu}}} \in \mathbb{Z}_p^m$ such that $\mathbf{f}^{\hat{\boldsymbol{\mu}}}[i_j] = \hat{\mu}_j$ for all $j \in [k]$ and $\mathbf{f}^{\hat{\boldsymbol{\mu}}}[i_j] = 0$ for all $j \in [k+1, m]$. Sample $\mathbf{r} \hookleftarrow U(\mathbb{Z}_q^n), \mathbf{e}_0 \hookleftarrow D_{\mathbb{Z}^m,\sigma}$, $e_1 \hookleftarrow D_{\mathbb{Z},\sigma}$ and set

$$\mathbf{c} = \begin{bmatrix} \mathbf{c}_0 \\ c_1 \end{bmatrix} = \mathbf{B} \cdot \mathbf{r} + \begin{bmatrix} \mathbf{e}_0 \\ e_1 \end{bmatrix} + \Delta \cdot \begin{bmatrix} \mathbf{f}^{\hat{\boldsymbol{\mu}}} \\ \mu \end{bmatrix},$$

where $\Delta = \lfloor q/p \rceil$.

Dec$(\mathsf{sk}, \mathsf{ct})$**:** Given a ciphertext $\mathsf{ct} = \mathbf{c} \in \mathbb{Z}_q^{m+1}$ and $\mathsf{sk} = \mathbf{s} \in \{-1, 0, 1\}^m$, compute $\nu = [-\mathbf{s}^\top \mid 1\,] \cdot \mathbf{c} \bmod q$ and output $\mu = \lfloor \nu/\Delta \rceil \in \mathbb{Z}_p^n$.

aDec$(\mathsf{dk}, \mathsf{tk}, \mathsf{ask}, \mathsf{act})$**:** Given $\mathsf{act} = \mathbf{c} \in \mathbb{Z}_q^{m+1}$, parse tk as $\mathsf{tk} = \widetilde{\mathbf{T}} \in \mathbb{Z}^{k \times m}$ and $\mathsf{dk} = (k, I)$ with $I = \{i_1, \ldots, i_k\}$. Then, compute

$$\hat{\boldsymbol{\nu}} = [\,\widetilde{\mathbf{T}} \mid \mathbf{0}\,] \cdot \mathbf{c} \bmod q$$

and output $\hat{\boldsymbol{\mu}} = \lfloor \hat{\boldsymbol{\nu}}/\Delta \rceil \in \mathbb{Z}_p^k$.

The above description assumes secret keys sampled from the distribution $\mathcal{P}$, but uniform binary secret keys $\mathbf{s}$ may also be used without any significant change to the security proofs.

Correctness. We first note that aGen only fails with negligible probability, as shown in the proof of Theorem 7. For correct decryption in the normal mode, we require

$$\|[-\mathbf{s}^\top \mid 1] \cdot [\mathbf{e}_0^\top \mid e_1]^\top\| < q/2p.$$

Note that $\|[-\mathbf{s}^\top \mid 1]\| \leq \sqrt{m-k+1}$ and $\|[\mathbf{e}_0^\top \mid e_1]\| \leq \sigma \cdot \sqrt{2(m+1)}$ with overwhelming probability $1 - (2/\exp(1))^{m/2}$ by Lemma 3. Using these bounds, we can ensure correct decryption with overwhelming probability by setting

$$\tfrac{q}{2p} > \sigma \cdot \sqrt{2(m-k+1) \cdot (m+1)}.$$

Moreover, to analyze correctness for aDec, we note that

$$[\,\widetilde{\mathbf{T}} \mid \mathbf{0}\,] \cdot \mathbf{c} = [\,\widetilde{\mathbf{T}} \mid \mathbf{0}\,] \cdot \begin{bmatrix} \mathbf{e}_0 \\ e_1 \end{bmatrix} + \widetilde{\mathbf{T}} \cdot \mathbf{f}^{\hat{\boldsymbol{\mu}}} = \widetilde{\mathbf{T}} \cdot \mathbf{e}_0 + \Delta \cdot \hat{\boldsymbol{\mu}}$$

where the second equality holds since, for each $i \in [k]$, the i-th row of $\widetilde{\mathbf{T}} \cdot \mathbf{f}^{\hat{\boldsymbol{\mu}}}$ is

$$\sum_{j=1}^{m} \widetilde{\mathbf{T}}[i, i_j] \cdot \mathbf{f}^{\hat{\boldsymbol{\mu}}}[i_j] = \sum_{j=1}^{m} \widetilde{\mathbf{T}}[i, i_j] \cdot \hat{\mu}_j$$

and $\widetilde{\mathbf{T}}[\cdot, i_j]$ is the j-th unit vector of dimension k for each $j \in [k]$. Since each row of $\widetilde{\mathbf{T}}$ also has at most $m-k+1$ non-zero entries that belong to $\{-1, 1\}$, the same condition as above also implies correct anamorphic decryption.

PARAMETERS. Setting $n = \Theta(\lambda)$ and $k = \frac{m}{2} - \lceil\sqrt{\lambda m/2}\rceil = \Theta(m)$ for the Hoeffding bound, we may take $m = \widetilde{\Theta}(\lambda)$ (ignoring logarithmic factors) whilst satisfying the bound required for the Leftover Hash Lemma. Next, recall that we have $\sigma = 2\alpha qs$ with the requirement that $\alpha q = \omega(\sqrt{\log m})$ so as to satisfy the conditions of Lemma 4. Taking $s = \sqrt{1+m}$, we aim to take $\alpha q > 2\sqrt{n}$ for $q > 2\sqrt{n}$. This allows for reductions from standard lattice problems with approximation factor $\gamma = \widetilde{O}(\sqrt{n}q)$. The correct decryption condition can then be written as $q/p = \Theta(\alpha qs \cdot \sqrt{m\lambda}) = \widetilde{\Theta}(m\sqrt{n\lambda}) = \widetilde{\Theta}(\lambda^2)$. By taking $p = O(1)$, we get a modulus $q = \widetilde{\Theta}(\lambda^2)$.

COMPARISON WITH THE TRAPDOOR-BASED CONSTRUCTION. The length of anamorphic messages can be larger than in Section 3 for a given LWE dimension n since we can use up to $k = \frac{m}{2} - \sqrt{\lambda m/2} = \Theta(m)$ message slots instead of $k = n$. By not relying on lattice trapdoors, the scheme can use somewhat smaller matrices and a smaller modulus (namely, $q = \widetilde{\Theta}(\lambda^2)$ instead of $q = \widetilde{\Theta}(\lambda^{5/2})$) for the same length of anamorphic plaintexts. If we set $k = n$, we just need $m = n(\log q + 1) + 2\lambda$ instead of $m > 2n\log q + 2\lambda$ in Section 3.

On the other hand, the scheme does not extend to a packed version with $\ell = \mathsf{poly}(\lambda)$ regular plaintext slots. The reason is that such a packed scheme [24] requires secret keys comprised of ℓ vectors and anamorphic messages should be encoded in the positions of $\mathbf{f}^{\hat{\boldsymbol{\mu}}}$ where the ℓ secret key vectors all contain zeroes.

Theorem 7. *The scheme is anamorphic in the sense of Definition 3 under the* $\mathsf{LWE}_{n,m,q,\chi}$ *assumption with* $\chi = D_{\mathbb{Z},\alpha q}$ *if* $\sigma = 2\alpha q \cdot s$ *for some* $s \geq \sqrt{1+m}$.

Proof. The proof considers again a sequence of games where we call W_i the event that the adversary outputs 1 in Game_i. The first game is the experiment RealG_{Π} of Definition 3 while the last game is $\mathsf{AnamorphicG}_{\Sigma}$.

Game_0**:** This is the real experiment RealG_{Π}. At each query $\mathcal{O}_e(\mathsf{pk}, \mu, \hat{\boldsymbol{\mu}})$, the challenger returns a ciphertext $\mathbf{c} = (\mathbf{c}_0, c_1) \leftarrow \mathsf{Enc}(\mathsf{pk}, \mu)$. We call W_0 the event that $\mathcal{A}$ outputs 1 when it halts.

Game_0'**:** This game is like Game_0 except that, at step 2 of KGen, the challenger aborts if it samples a secret key $\mathbf{s} \hookleftarrow \mathcal{P}^m$ that contains strictly less than k zeroes. We claim that this only happens with negligible probability. Since each entry of $\mathbf{s}$ is sampled independently from $\mathcal{P}$, we have $\mathbb{E}[|J|] = m/2$ at step 2. Then, Hoeffding's inequality implies

$$\Pr_{\mathbf{s}\hookleftarrow\mathcal{P}^m}\left[\ \left||J| - \tfrac{m}{2}\right| > \left\lceil\sqrt{\lambda m/2}\right\rceil\ \right] \leq 2\cdot\exp(-\lambda).$$

With probability $\geq 1 - 2\exp(-\lambda)$, we thus have $|J| \geq \frac{m}{2} - \lceil\sqrt{\lambda m/2}\rceil \geq k$, meaning that $|\Pr[W_0'] - \Pr[W_0]| \leq 2\cdot\exp(-\lambda) \leq 2^{-\Omega(\lambda)}$

Game_1**:** We modify the encryption oracle of RealG_{Π}. At each query $\mathcal{O}_e(\mathsf{pk}, \mu, \hat{\boldsymbol{\mu}})$, instead of running the real Enc algorithm, the challenger samples $\mathbf{r} \hookleftarrow U(\mathbb{Z}_q^n)$,

$\mathbf{x} \hookleftarrow D_{\mathbb{Z}^m,\alpha q}$ and uses the ReRand algorithm of Lemma 4 to compute

$$\begin{aligned} \mathbf{c}'_0 &= \mathbf{A}^\top \cdot \mathbf{r} + \mathbf{x} \in \mathbb{Z}_q^m \\ \begin{bmatrix} \mathbf{c}_0 \\ c_1 \end{bmatrix} &= \mathsf{ReRand}([\mathbf{I}_m \mid \mathbf{s}], \mathbf{c}'_0, \alpha q, s) + \begin{bmatrix} \mathbf{0}^m \\ \Delta \cdot \mu \end{bmatrix} \in \mathbb{Z}_q^{m+1}, \end{aligned} \tag{25}$$

where $s \geq \sqrt{1+m} \geq \|[\mathbf{I}_m \mid \mathbf{s}]\|$ since $\|\mathbf{s}\| \leq \sqrt{m}$. We can thus apply Lemma 4 with $\mathbf{V} = [\mathbf{I}_m \mid \mathbf{s}] \in \mathbb{Z}^{m \times (m+1)}$, $\mathbf{b} = \mathbf{A}^\top \cdot \mathbf{r} \in \mathbb{Z}_q^m$, and $r = \alpha q$ to obtain that the distribution of $(\mathbf{c}_0, c_1)$ is statistically close to the distribution obtained by computing

$$\begin{bmatrix} \mathbf{c}_0 \\ c_1 \end{bmatrix} = \begin{bmatrix} \mathbf{A}^\top \\ \mathbf{s}^\top \cdot \mathbf{A}^\top \end{bmatrix} \cdot \mathbf{r} + \begin{bmatrix} \mathbf{e}_0 \\ e_1 \end{bmatrix} + \begin{bmatrix} \mathbf{0}^m \\ \Delta \cdot \mu \end{bmatrix},$$

where $\mathbf{e}_0 \sim D_{\mathbb{Z}^m,2\alpha qs}$, $e_1 \sim D_{\mathbb{Z},2\alpha qs}$. This implies $|\Pr[W_1] - \Pr[W'_0]| \leq Q \cdot 2^{-\Omega(\lambda)}$, where Q is the number of queries to the oracle $\mathcal{O}_e$.

Game_2: In this game, we change again the encryption oracle. At each query $\mathcal{O}_e(\mathsf{pk}, \mu, \hat{\boldsymbol{\mu}})$, instead of computing a ciphertext as in (25), the challenger replaces the pseudorandom $\mathbf{c}'_0$ by a random vector and computes

$$\begin{aligned} \mathbf{c}'_0 &\hookleftarrow U(\mathbb{Z}_q^m) \\ \begin{bmatrix} \mathbf{c}_0 \\ c_1 \end{bmatrix} &= \mathsf{ReRand}([\mathbf{I}_m \mid \mathbf{s}], \mathbf{c}'_0, \alpha q, s) + \begin{bmatrix} \mathbf{0}^m \\ \Delta \cdot \mu \end{bmatrix} \in \mathbb{Z}_q^{m+1}, \end{aligned} \tag{26}$$

and returns $(\mathbf{c}_0, c_1)$. By a standard hybrid argument over all queries to the encryption oracle $\mathcal{O}_e$, we have $|\Pr[W_2] - \Pr[W_1]| \leq Q \cdot \mathbf{Adv}^{\mathsf{LWE}}_{n,m,q,\chi}(\lambda)$.

Game_3: We change the generation of ciphertexts. At each query $\mathcal{O}_e(\mathsf{pk}, \mu, \hat{\boldsymbol{\mu}})$, the challenger now samples $\mathbf{u} \hookleftarrow U(\mathbb{Z}_q^m)$ uniformly and computes

$$\begin{aligned} \mathbf{c}'_0 &= \mathbf{u} + \Delta \cdot \mathbf{f}^{\hat{\boldsymbol{\mu}}}, \\ \begin{bmatrix} \mathbf{c}_0 \\ c_1 \end{bmatrix} &= \mathsf{ReRand}([\mathbf{I}_m \mid \mathbf{s}], \mathbf{c}'_0, \alpha q, s) + \begin{bmatrix} \mathbf{0}^m \\ \Delta \cdot \mu \end{bmatrix} \in \mathbb{Z}_q^{m+1}, \end{aligned} \tag{27}$$

where $\mathbf{f}^{\hat{\boldsymbol{\mu}}} \in \mathbb{Z}_p^m$ is defined as in aEnc. This change does not modify the distribution of $(\mathbf{c}_0, c_1)$ since still we have $\mathbf{c}'_0 \sim U(\mathbb{Z}_q^m)$ as in (26). Hence, $\Pr[W_3] = \Pr[W_2]$.

Game_4: We change again the output distribution of $\mathcal{O}_e$. At each query $\mathcal{O}_e(\mathsf{pk}, \mu, \hat{\boldsymbol{\mu}})$, the challenger now samples $\mathbf{r} \hookleftarrow U(\mathbb{Z}_q^n)$, $\mathbf{x} \hookleftarrow D_{\mathbb{Z}^m,\alpha q}$ and computes

$$\begin{aligned} \mathbf{c}'_0 &= (\mathbf{A}^\top \cdot \mathbf{r} + \mathbf{x}) + \Delta \cdot \mathbf{f}^{\hat{\boldsymbol{\mu}}}, \\ \begin{bmatrix} \mathbf{c}_0 \\ c_1 \end{bmatrix} &= \mathsf{ReRand}([\mathbf{I}_m \mid \mathbf{s}], \mathbf{c}'_0, \alpha q, s) + \begin{bmatrix} \mathbf{0}^m \\ \Delta \cdot \mu \end{bmatrix} \in \mathbb{Z}_q^{m+1}, \end{aligned} \tag{28}$$

and thus replaces the uniformly random vector $\mathbf{u}$ of (27) by an LWE sample $\mathbf{u} = \mathbf{A}^\top \cdot \mathbf{r} + \mathbf{x}$ at each query. Under the LWE assumption, this change does not affect $\mathcal{A}$'s view and we have $|\Pr[W_4] - \Pr[W_3]| \leq Q \cdot \mathbf{Adv}^{\mathsf{LWE}}_{n,m,q,\chi}(\lambda)$ via a standard hybrid argument over all queries to $\mathcal{O}_e$.

Game_5**:** Now, the challenger answers all encryption queries $\mathcal{O}_e(\mathsf{pk}, \mu, \hat{\boldsymbol{\mu}})$ by sampling $\mathbf{r} \hookleftarrow U(\mathbb{Z}_q^n)$, $\mathbf{e}_0 \hookleftarrow D_{\mathbb{Z}^m,2\alpha qs}$, $e_1 \hookleftarrow D_{\mathbb{Z},2\alpha qs}$ and computing

$$\begin{aligned} \mathbf{c}_0 &= \mathbf{A}^\top \cdot \mathbf{r} + \mathbf{e}_0 + \Delta \cdot \mathbf{f}^{\hat{\boldsymbol{\mu}}}, \\ c_1 &= \mathbf{s}^\top \cdot (\mathbf{A}^\top \cdot \mathbf{r}) + e_1 + \Delta \cdot \mu, \end{aligned} \tag{29}$$

By applying Lemma 4 again with $\mathbf{V} = [\mathbf{I}_m \mid \mathbf{s}] \in \mathbb{Z}^{m\times(m+1)}$, $\mathbf{b} = \mathbf{A}^\top \cdot \mathbf{r} \in \mathbb{Z}_q^m$, and $r = \alpha q$, the distribution of $(\mathbf{c}_0, c_1)$ is statistically close to the one of Game_4. Note that we crucially use the fact that $\mathbf{s}^\top \cdot (\mathbf{A}^\top \cdot \mathbf{r} + \Delta \cdot \mathbf{f}^{\hat{\boldsymbol{\mu}}}) = \mathbf{s}^\top \cdot \mathbf{A}^\top \cdot \mathbf{r}$ here by virtue of the fact that $\mathbf{s}^\top \cdot \mathbf{f}^{\hat{\boldsymbol{\mu}}} = 0$. Therefore, $|\Pr[W_5] - \Pr[W_4]| \leq Q \cdot 2^{-\Omega(\lambda)}$, where Q is the number of queries to the oracle $\mathcal{O}_e$.

Game_6**:** In this game, the challenger modifies the key generation phase. Instead of choosing the matrix $\mathbf{A}$ uniformly over $\mathbb{Z}_q^{n\times m}$, the challenger first generates a matrix $\mathbf{A}' \in \mathbb{Z}_q^{n\times m}$ as in (24) and then generates $\mathbf{A}$ by permuting the columns of $\mathbf{A}'$ and computing $\mathbf{B}$ as in step 4 of aGen. Since parameters are chosen in such a way that $m - k > n \log q + 2\lambda$, Lemma 1 implies that the matrix $\mathbf{A}'$ is statistically close to uniform over $\mathbb{Z}_q^{n\times m}$ and so is $\mathbf{A}$. Consequently, $|\Pr[W_6] - \Pr[W_5]| \leq 2^{-\Omega(\lambda)}$.

In Game_6, the challenger is interacting with the adversary exactly as in the $\mathsf{AnamorphicG}_\Sigma$ experiment. By combining the above, we obtain the following inequality which proves the result:

$$\mathbf{Adv}^{\mathsf{Anamorphism}}_{\mathcal{D},\Pi,\Sigma}(\lambda) = |\Pr[W_0] - \Pr[W_6]| \leq 2Q\Big(\mathbf{Adv}^{\mathsf{LWE}}_{n,m,q,\chi}(\lambda) + 2^{-\Omega(\lambda)}\Big) + 2^{-\Omega(\lambda)}.$$

□

Theorem 8. *The above scheme is fully asymmetric with respect to the anamorphic triplet above under the* $\mathsf{LWE}_{n,m+1,q,\chi'}$*, where* $\chi' = D_{\mathbb{Z},\sigma}$ *assumption.*

Proof. We will use a sequence of games that first replaces the key material given to the adversary $\mathcal{A}$ in the $\mathsf{Fasym}^b_{\mathcal{A},\Pi,\Sigma}(\lambda)$, and then replaces the challenge ciphertext by a uniformly random vector. Throughout the sequence, W_i^b will denote the event that the adversary $\mathcal{A}$ outputs 1 in Game_i^b. We only analyze the case where aGen does not abort at step 2 because, if it does, the adversary has no advantage.

Game_0^b**:** This is precisely the $\mathsf{Fasym}^b_{\mathcal{A},\Pi,\Sigma}(\lambda)$ game.

Game_1^b**:** This is the same as Game_0^b except that, when running aGen, the challenger replaces the pseudorandom $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ by a uniform matrix sampled from $U(\mathbb{Z}_q^{n\times m})$. We note that, in Game_0^b, aGen generates $\mathbf{A}$ by permuting the columns of the matrix $\mathbf{A}'$ from (24). However, the latter is statistically indistinguishable from a uniform matrix by the Leftover Hash Lemma. Since $m - k > n \log q + 2\lambda$, we can apply Lemma 2 for the source $\mathbf{T} \sim \mathcal{P}^{(m-k)\times k}$ so as to obtain $|\Pr[W_1^b] - \Pr[W_0^b]| \leq k \cdot 2^{-\lambda}$.

Game_2^b**:** This is the same as Game_1^b except that, during aGen, the challenger replaces the matrix $\mathbf{B}$ by a uniform matrix in $\mathbb{Z}_q^{(n+1)\times m}$ instead of choosing the last row as $\mathbf{A}\cdot\mathbf{s}\in\mathbb{Z}_q^n$ with $\mathbf{s}\hookleftarrow\mathcal{P}^m$. Since $\mathbf{A}$ is uniformly distributed in both Game_1^b and Game_2^b, we can use Lemma 1 with the trivial function f. More precisely, since $\mathbf{A}\cdot\mathbf{s}=\sum_{j\in I}\mathbf{A}[\cdot,j]\cdot\mathbf{s}[j]+\sum_{j\in[m]\setminus I}\mathbf{A}[\cdot,j]\cdot\mathbf{s}[j]\in\mathbb{Z}_q^n$, we can rely on the min-entropy of $\{\mathbf{s}[j]\}_{j\in[m]\setminus I}$ (which is not affected by the revealed $\{\mathbf{s}[j]\}_{j\in I}$ since the entries of $\mathbf{s}$ are sampled independently from $\mathcal{P}$) to claim that the distribution of $\sum_{j\in[m]\setminus I}\mathbf{A}[\cdot,j]\cdot\mathbf{s}[j] \bmod q$ is statistically close to $U(\mathbb{Z}_q^n)$ due to our choice of $|[m]\setminus I|=m-k>n\log q+2\lambda$. The sum $\mathbf{A}\cdot\mathbf{s} \bmod q$ is thus itself statistically close to uniform over $\mathbb{Z}_q^n$ and we conclude that $|\Pr[W_2^b]-\Pr[W_1^b]|\leq 2^{-\lambda}$.

Game_3^b**:** This is the same as Game_2^b apart from that when answering the adversary's challenge, the challenger replaces the ciphertext $\mathbf{c}$ by $\mathbf{u}+\Delta\cdot\begin{bmatrix}\mathbf{f}^{\hat{\boldsymbol{\mu}}_b}\\ \mu_b\end{bmatrix}$ where $\mathbf{u}$ is a uniformly chosen element of $\mathbb{Z}_q^{m+1}$. Lemma 13 shows that $|\Pr[W_3^b]-\Pr[W_2^b]|\leq\mathbf{Adv}^{\mathsf{LWE}}_{n,m+1,q,\chi'}(\lambda)$, where $\chi'=D_{\mathbb{Z},\sigma}$.

Game_4^b**:** This game is like Game_3^b but the challenge ciphertext is replaced by uniform $\mathbf{u}'\hookleftarrow U(\mathbb{Z}_q^{m+1})$. Clearly, the ciphertext follows a uniform distribution in both Game_3^b and Game_4^b, meaning that $\Pr[W_4^b]=\Pr[W_3^b]$

We note that Game_4^b is totally independent of the challenge bit b so that $\Pr[W_4^0]=\Pr[W_4^1]$. By the triangle inequality, we obtain the following bound which concludes the proof:

$$\mathbf{Adv}^{\mathsf{Fasym}}_{\mathcal{A},\Pi,\Sigma}(\lambda)=|\Pr[W_0^0]-\Pr[W_0^1]|\leq 2\cdot\big((k+1)\cdot 2^{-\Omega(\lambda)}+\mathbf{Adv}^{\mathsf{LWE}}_{n,m+1,q,\chi'}(\lambda)\big).$$

□

Lemma 13. *We have* $|\Pr[W_3^b]-\Pr[W_2^b]|\leq\mathbf{Adv}^{\mathsf{LWE}}_{n,m+1,q,\chi'}(\lambda)$*, where* $\chi'=D_{\mathbb{Z},\sigma}$.

Proof. We note that the matrix $\mathbf{B}$ is uniformly distributed in both Game_2^b and Game_3^b. We can thus build a straightforward reduction $\mathcal{B}$ that takes as input an LWE instance $(\mathbf{B},\mathbf{u})\in\mathbb{Z}_q^{(m+1)\times n}\times\mathbb{Z}_q^{m+1}$ and has to decide if $\mathbf{u}\sim U(\mathbb{Z}_q^{m+1})$ or $\mathbf{u}=\mathbf{B}\cdot\mathbf{r}+\mathbf{e}$ for some $\mathbf{r}\sim U(\mathbb{Z}_q^n)$ and $\mathbf{e}\sim D_{\mathbb{Z}^{m+1},\sigma}$.

To do this, $\mathcal{B}$ uses its input matrix $\mathbf{B}$ to build the public key apk and constructs a challenge ciphertext by setting

$$\begin{bmatrix}\mathbf{c}_0\\ c_1\end{bmatrix}=\mathbf{u}+\Delta\cdot\begin{bmatrix}\mathbf{f}^{\hat{\boldsymbol{\mu}}_b}\\ \mu_b\end{bmatrix},$$

When the adversary $\mathcal{A}$ halts, $\mathcal{B}$ outputs whatever $\mathcal{A}$ outputs.

If $\mathbf{u}=\mathbf{B}\cdot\mathbf{r}+\mathbf{e}$, then the distribution of $(\mathbf{c}_0,c_1)$ is identical to that of Game_2^b. If $\mathbf{u}\sim U(\mathbb{Z}_q^{m+1})$, then $(\mathbf{c}_0,c_1)$ is distributed as in Game_3^b. □

On the Use of Gaussian Secret Keys. The scheme extends to the case of secret keys sampled from a discrete Gaussian distribution $D_{\mathbb{Z}^m,\sigma}$ with standard

deviation $\sigma = O(\sqrt{m})$ as in the dual Regev variant of [24]. In this case, we need to assume $k = O(1)$. If each entry of $\mathbf{s} = (s_1, \ldots, s_m)$ is sampled independently from a one-dimensional discrete Gaussian $D_{\mathbb{Z},\sigma}$ with standard deviation σ, the probability of one s_i to be zero is $1/\rho_\sigma(\mathbb{Z})$, which is at least $\frac{1}{1+\sigma}$, as shown by Lemma 14. Except with probability $\leq \left(1 - \frac{1}{1+\sigma}\right)^m$ over the random choice of $\mathbf{s}$, there thus exists $i \in [m]$ such that $s_i = 0$.

Lemma 14. *Over the choice of $\mathbf{s} \hookleftarrow D_{\mathbb{Z}^m,\sigma}$, the probability that $s_i \neq 0$ for each $i \in [m]$ is at most $\left(1 - \frac{1}{1+\sigma}\right)^m$.*

Proof. If each s_i is sampled from a one-dimensional Gaussian with standard deviation σ, the probability that $s_i = 0$ is $1/\rho_\sigma(\mathbb{Z})$ by the definition of the discrete Gaussian distribution. We know from [38, Lemma 4.4] that, for any n-dimensional lattice Λ, any $\mathbf{x}', \mathbf{c} \in \mathbb{R}^n$ and any standard deviation σ satisfying $\sigma \geq \eta_{2^{-n}}(\Lambda)$, we have $\rho_{\sigma,\mathbf{c}}(\Lambda + \mathbf{x}') \in [1 - 2^{-n}, 1 + 2^{-n}] \cdot \sigma/\det(\Lambda)$. With $n = 1$, this implies $\rho_\sigma(\mathbb{Z}) \in [1/2, 3/2] \cdot \sigma$. In the case of $\Lambda = \mathbb{Z}$, we can prove a tighter bound $\sigma \leq \rho_\sigma(\mathbb{Z}) \leq \sigma + 1$.

First, by the Poisson summation formula, we have that for suitably nice function f with Fourier transform $\hat{f}$, $\sum_{\mathbf{x}\in\Lambda} f(\mathbf{x}) = (\sum_{\mathbf{y}\in\widehat{\Lambda}} \hat{f}(\mathbf{y}))/\mathsf{Vol}(\Lambda)$ where $\mathsf{Vol}(\Lambda)$ is the volume of the fundamental parallelepiped. Setting $f(\mathbf{x}) = \rho_s(\mathbf{x})$ and $\Lambda = \mathbb{Z}^n$ we have the following Lemma:

Lemma 15. *For any $\sigma > 0$, $\rho_\sigma(\mathbb{Z}^n) \geq \sigma^n$ and $\Pr_{\mathbf{e}\hookleftarrow D_{\mathbb{Z}^n,\sigma}}[\mathbf{e} = 0] \leq \sigma^{-n}$.*

Second, we also know that, if a function f is non-increasing between 0 and $+\infty$, we have the inequality $\sum_{x=0}^{\infty} f(x+1) \leq \int_0^\infty f(x)dx$ (which follows from viewing $\sum_{x=0}^{\infty} f(x+1)$ as the surface covered by rectangles of width 1 and height $f(x+1)$ below the Gaussian curve). In the specific case of $f(x) = \exp(-\pi x^2/\sigma^2)$, this implies

$$\sum_{x\in\mathbb{Z}} f(x) = f(0) + 2\sum_{x=0}^{\infty} f(x+1) \leq f(0) + 2\int_0^\infty f(x)dx$$
$$= \exp(0) + \int_{-\infty}^{\infty} f(x)dx = 1 + \sigma,$$

and thus $\rho_\sigma(\mathbb{Z}) \leq 1 + \sigma$. Since the components of $\mathbf{s} = (s_1, \ldots, s_m) \hookleftarrow D_{\mathbb{Z}^m,\sigma}$ are independent, the probability that $s_i \neq 0$ for each $i \in [m]$ is at most $(1 - \frac{1}{1+\sigma})^m$, which completes the proof of Lemma 14. □

If we choose $\sigma \approx \sqrt{m}$, the probability to have $s_i \neq 0$ for each $i \in [m]$ is bounded by $(1 - 1/\sqrt{m})^m < (1/\exp(1))^{\sqrt{m}}$, which is sub-exponentially small as a function of m and can be made exponentially small in λ by setting $m = \Theta(\lambda^2)$. Then, the proofs of Theorem 7 and Theorem 8 carry over for $k = 1$.

We note that, for concrete values such as $n \geq 1024$, $\log q \geq 10$ and $\lambda = 128$, we have $m \geq 10490$ and thus $1/(\exp(1))^{\sqrt{m}} < 2^{-102}$.

D Deferred Proofs for the Anamorphic Dual GSW Scheme

D.1 Correctness and Error Growth

Suppose we wish to evaluate circuits of multiplicative depth L (ignoring additions for now). Since we are taking the gadget matrix with decomposition base 2, we know that $\mathbf{G}^{-1}(\mathbf{C_2}) \in \{-1,0,1\}^{M\times M}$ where $M \triangleq k(m+1)$. Reusing the notation from the discussion above Theorem 4, after one multiplication the error term in $\mathbf{C}^{\times}$ was $\mathbf{E}_1 \cdot \mathbf{G}^{-1}(\mathbf{C_2}) + \mu_1(\mathbf{E}_2)_{[m+1]\setminus I} + \hat{\mu}_1(\mathbf{E}_2)_I$. Now, every column of $\mathbf{E}_1$ and $\mathbf{E}_2$ has Euclidean norm bounded by $\beta_0 \triangleq c_1\sigma\sqrt{m+1}$ w.h.p. by Lemma 3. Similarly, the rows are bounded by $\gamma_0 \triangleq c_2\sigma\sqrt{M}$. It follows by the Cauchy-Schwarz inequality that the error matrix of $\mathbf{C}^{\times}$ has columns bounded in Euclidean norm by $\beta_1 = \gamma_0 \cdot \sqrt{M(m+1)} + (p-1)\cdot\beta_0$ with overwhelming probability. Similarly, the rows of the error matrix have a Euclidean bound of $\gamma_1 = \gamma_0 \cdot M + (p-1)\cdot\gamma_0$. We can use this analysis to obtain the recurrence relation

$$\beta_i = \gamma_{i-1}\cdot\sqrt{M(m+1)} + (p-1)\cdot\beta_{i-1}, \qquad \gamma_i = \gamma_{i-1}\cdot M + (p-1)\cdot\gamma_{i-1},$$

which holds for $i \in [L]$. This is easily reduced to

$$\beta_i = (M+p-1)^i \cdot \gamma_0 \cdot \sqrt{M(m+1)} + (p-1)\cdot\beta_{i-1}.$$

Generally, we will say a ciphertext $\mathbf{C} = \mathbf{B}_{[m+1]\setminus I}\cdot(\star) + \mathbf{B}_I\cdot(\star) + \mathbf{J}_{\mu,\hat{\mu}}\otimes\mathbf{g}^\top + \mathbf{E_C}$ has error matrix $\mathbf{E_C}$ so that $[-\mathbf{s}^\top \mid 1]\cdot\mathbf{C} = \mu\cdot[-\mathbf{s}^\top \mid 1]\cdot\mathbf{G} + [-\mathbf{s}^\top \mid 1]\cdot\mathbf{E_C}$. The parameters β_i and γ_i describe the norm bounds on the columns and rows of an error matrix after i multiplications (still ignoring homomorphic additions). We can then ensure correct decryption for normal messages if

$$\left\|[-\mathbf{s}^\top \mid 1]\cdot\mathbf{E_C}\cdot\mathbf{G}^{-1}(\Delta\cdot\hat{\mathbf{e}}_{m+1})\right\|_\infty$$

is less than $q/(2p)$. Since $\|[-\mathbf{s}^\top \mid 1]\| \le \sqrt{m-n_0+1}$ and $\mathbf{G}^{-1}(\Delta\cdot\hat{\mathbf{e}}_{m+1})$ has at most k non-zero ternary entries, we can write the correctness requirement for normal messages after L multiplications as $q/(2p) > k\cdot\sqrt{m-n_0+1}\cdot\beta_L$. Note that this conclusion holds whether or not the ciphertexts are anamorphic.

Unfortunately, the error matrix considered above is not the only thing that contributes to anamorphic decryption errors. This is due to the fact that the public key $\mathbf{B}$ satisfies $[\mathbf{t}^\top \mid 0]\cdot\mathbf{B} = [\mathbf{t}^\top \mid 0]\cdot\mathbf{B}_I = \mathbf{e}_I^\top \neq \mathbf{0}$. To calculate the anamorphic error arising from $\mathbf{B}_I$, we can see that the term of $\mathbf{C}^{\times}$ prefixed by $\mathbf{B}_I$ is $\mathbf{B}_I\cdot(\mathbf{S}_1\cdot\mathbf{G}^{-1}(\mathbf{C}_2) + \hat{\mu}_1\cdot\mathbf{S}_2)$. Note that the columns and rows of $\mathbf{S}_1, \mathbf{S}_2$ can be bounded in Euclidean norm by $\delta_0 \triangleq \bar{c}_1\cdot\alpha q\sqrt{n}$ and $\zeta_0 \triangleq \bar{c}_2\cdot\alpha q\sqrt{M}$ respectively. We can then bound the Euclidean norm of the columns of $(\mathbf{S}_1\cdot\mathbf{G}^{-1}(\mathbf{C}_2)+\hat{\mu}_1\cdot\mathbf{S}_2)$ by $\delta_1 = \zeta_0\cdot\sqrt{Mn} + (p-1)\cdot\delta_0$. Furthermore, its rows are bounded by $\zeta_1 = \zeta_0\cdot M + (p-1)\cdot\zeta_0$. Now, noting that if $\overline{\mathbf{C}}_1 = \mathbf{B}_I\cdot\overline{\mathbf{S}}_1 + \cdots$ and $\overline{\mathbf{C}}_2 = \mathbf{B}_I\cdot\overline{\mathbf{S}}_2 + \cdots$, then $\overline{\mathbf{C}}_1\cdot\mathbf{G}^{-1}(\overline{\mathbf{C}}_2) = \mathbf{B}_I\cdot(\overline{\mathbf{S}}_1\cdot\mathbf{G}^{-1}(\overline{\mathbf{C}}_2) + \hat{\mu}_1\cdot\overline{\mathbf{S}}_2) + \cdots$, we can extend our analysis to i levels. Define the following recurrence relation:

$$\delta_i = \zeta_{i-1}\cdot\sqrt{Mn} + (p-1)\cdot\delta_{i-1}, \qquad \zeta_i = \zeta_{i-1}\cdot M + (p-1)\cdot\zeta_{i-1}.$$

A ciphertext resulting from a depth-i multiplication circuit then has the term $\mathbf{B}_I \cdot \mathbf{S}^{(i)}$ where δ_i and ζ_i are bounds on the Euclidean norms of the columns and rows of $\mathbf{S}^{(i)}$ respectively. Now, if $\mathbf{C}$ is a depth-L ciphertext encrypting covert message $\hat{\mu}$, we can say that

$$[\mathbf{t}^\top \mid 0] \cdot \mathbf{C} = [\mathbf{t}^\top \mid 0] \cdot \mathbf{B}_I \cdot \mathbf{S}^{(L)} + [\mathbf{t}^\top \mid 0] \cdot \mathbf{E}_{\mathbf{C}} + [\mathbf{t}^\top \mid 0] \cdot \hat{\mu} \cdot \mathbf{G}.$$

Noting that $[\mathbf{t}^\top \mid 0] \cdot \mathbf{B}_I = \mathbf{e}_I^\top$, we have correctness as long as

$$\left\| \left(\mathbf{e}_I^\top \cdot \mathbf{S}^{(L)} + [\mathbf{t}^\top \mid 0] \cdot \mathbf{E}_{\mathbf{C}} \right) \cdot \mathbf{G}^{-1}(\Delta \cdot \hat{\mathbf{e}}_{i_{n_0}}) \right\|_\infty < q/(2p).$$

Since $\|\mathbf{t}\| \leq c'\alpha q\sqrt{n_0 - 1}$ and $\|\mathbf{e}_I\| \leq c''\alpha q\sqrt{n}$ by Lemma 3, we can ensure correctness for covert message decryption (with overwhelming probability) if $q/(2p) > k \cdot \left(c''\alpha q\sqrt{n} \cdot \delta_L + c'\alpha q\sqrt{n_0 - 1} \cdot \beta_L \right)$. Asymptotically, since p is constant, we end up with $\beta_L = O(M^L \cdot \sigma\sqrt{m})$ and $\delta_L = O(M^L \cdot \alpha q\sqrt{n})$. Further, since $n_0 = \Theta(m)$ is larger than n and $\sigma > \alpha q$, we can then take $q/p = \Omega(k \cdot M^L \cdot (\alpha q) \cdot \sigma \cdot m)$ for correctness. To obtain L_{mult} multiplicative levels with at most L_{add} additions between multiplications, the modulus should be amended to $q/p = \Omega(k \cdot (L_{\mathsf{add}} \cdot M)^{L_{\mathsf{mult}}} \cdot (\alpha q) \cdot \sigma \cdot m)$ as the recurrence relations for $(\beta_i, \gamma_i, \delta_i, \zeta_i)$ require an additional factor of L_{add} at each of the L_{mult} levels.

D.2 Proof of Theorem 4

Proof. We simply reuse the correctness analysis from Section D.1 above. First, the probability of aborting in aGen due to I not existing is at most $2^{-\Omega(\lambda)}$ (see Game_6 in the proof of Theorem 5). Second, we note that (i) the bounds for the rows/columns of un-evaluated error matrices (i.e., the parameters $\beta_0 = c_1\sigma\sqrt{m+1}$, $\gamma_0 = c_2\sigma\sqrt{M}$) and (ii) the bounds on unevaluated secret matrices (i.e., the parameters $\delta_0 = \bar{c}_1\alpha q\sqrt{n}$, $\zeta_0 = \bar{c}_2\alpha q\sqrt{M}$) are each violated with probability at most $2^{-\Omega(\lambda)}$ by Lemma 3. Furthermore, the bounds on $\mathbf{t}$ and $\mathbf{e}_I$ (namely, $c'\alpha q\sqrt{n_0 - 1}$ and $c''\alpha q\sqrt{n}$) are also violated with probability at most $2^{-\Omega(\lambda)}$. We use a union bound over a possible $L_{\mathsf{add}} \cdot 2^{L_{\mathsf{mult}}}$ input ciphertexts (and the rows/columns of their secret/error matrices), and $(\mathbf{t}, \mathbf{e}_I)$. This shows that all of the aforementioned bounds are respected with probability at least $1 - (4 \cdot M \cdot L_{\mathsf{add}} \cdot 2^{L_{\mathsf{mult}}} + 2) \cdot 2^{-\Omega(\lambda)} = 1 - 2^{-\Omega(\lambda)}$ since $L_{\mathsf{add}} = \mathsf{poly}(\lambda)$ and $L_{\mathsf{mult}} = \widetilde{O}(1)$. Note the factor of $4 \cdot M$ due to the fact that each ciphertext secret/error matrix has at most $M = \mathsf{poly}(\lambda)$ rows/columns. Considering the abort probability in aGen and probability that all bounds are respected together proves the theorem. □

E Homomorphic Robustness

In [8], Banfi *et al.* introduced a notion of robustness for anamorphic PKE schemes. Its motivation is that the receiver should always output a reject symbol when it applies aDec to a ciphertext generated by the *normal* encryption

algorithm Enc. However, when considering (fully) homomorphic schemes, there are two classes of *normal* ciphertexts: (i) those generated directly from the normal encryption algorithm (we will call these normal, fresh ciphertexts) and (ii) those generated by performing some homomorphic computation on normal fresh ciphertexts. Therefore, in the context of homomorphic encryption, existing notions of robustness may not be as strong as one can hope for.

E.1 Definition

Recall that the existing form of robustness for anamorphic PKE (Definition 5) entails an adversary supplying a challenge oracle $\mathcal{O}_b$ with messages. On input μ from the adversary, the $\mathcal{O}_0$ oracle returns $\mathsf{aDec}(\mathsf{dk}, \mathsf{tk}, \mathsf{ask}, \mathsf{Enc}(\mathsf{apk}, \mu))$. On the other hand, the $\mathcal{O}_1$ oracle always returns an error symbol $\perp$. As discussed above, this definition may not suffice for applications involving homomorphically evaluated ciphertexts. We thus present a new homomorphic robustness definition.

The new homomorphic robustness definition provides an adversary with three oracles: an encryption oracle, an evaluation oracle and a challenge oracle. These three oracles share a state that consists of *all ciphertexts* produced by the challenger throughout the game. The evaluation oracle is written in the context of single-output circuits, but a multi-output circuit can be performed via multiple evaluation oracle queries. Furthermore, the evaluation oracle only accepts computation on fresh ciphertexts. For an i-hop homomorphic encryption scheme [23] (which is a weaker property than full composability [36]), the evaluation of a valid circuit $\mathsf{C} \in \mathcal{C}$ is performed by sequentially running at most i evaluation substeps. Note that i depends on $\mathcal{C}$ and we implicitly assume that $\mathsf{C} \in \mathcal{C}$ describes how to decompose C into hops. Therefore, the fact that the evaluation oracle only accepts fresh ciphertexts does not restrict the adversary's power in the common case of deterministic evaluation since intermediate ciphertexts in a multi-hop computation can be obtained by further evaluation oracle queries.

Moreover, this formulation allows conveniently checking whether a circuit is valid in the context of homomorphic encryption when the circuit is restricted to belong to a specific family (e.g. additively homomorphic LWE encryptions with a restricted number of additions). For every ciphertext, the state stores a bit to indicate whether a ciphertext is the result of a fresh encryption or the result of homomorphic evaluation. We first recall the definition of an i-hop homomorphic encryption scheme [23] and then introduce the homomorphic robustness definition.

Definition 8. *Let $i = i(\lambda)$ be a function of the security parameter. A scheme $\mathsf{HE} = (\mathsf{HE.KGen}, \mathsf{HE.Enc}, \mathsf{HE.Dec}, \mathsf{HE.Eval})$ is an i-hop homomorphic encryption scheme if, for every compatible sequence $\boldsymbol{f} = \langle f_1, \ldots, f_t \rangle$ with $t \leq i$ functions and every input $\boldsymbol{\mu}$ to f_1,*

$$\Pr\left[\mathsf{HE.Dec}(\mathsf{sk}, \mathsf{HE.Eval}(\mathsf{pk}, \boldsymbol{f}, \boldsymbol{c})) \neq (f_t \circ \cdots \circ f_1)(\boldsymbol{\mu}) : \begin{smallmatrix}(\mathsf{pk},\mathsf{sk}) \hookleftarrow \mathsf{HE.KGen}(\lambda) \\ \boldsymbol{c} \hookleftarrow \mathsf{HE.Enc}(\mathsf{pk}, \boldsymbol{\mu})\end{smallmatrix}\right] \leq \eta(\lambda)$$

for some $\eta(\lambda) = \mathsf{negl}(\lambda)$.

$\mathcal{O}_{\mathsf{Enc}}(\mathsf{apk}, \mu; \mathsf{st}, \mathsf{ctr})$

$\mathsf{st} \leftarrow \mathsf{st} \parallel (1, \mathsf{HE.Enc}(\mathsf{apk}, \mu))$
$\mathsf{ctr} \leftarrow \mathsf{ctr} + 1$
return 1

$\mathcal{O}_0(\mathsf{ask}, \mathsf{dk}, \mathsf{tk}, i; \mathsf{st}, \mathsf{ctr})$

if $i \notin [\mathsf{ctr}]$ **return** $\perp$
else return $\mathsf{HE.aDec}(\mathsf{dk}, \mathsf{tk}, \mathsf{ask}, \mathsf{st}_i)$

$\mathcal{O}_1(\mathsf{ask}, \mathsf{dk}, \mathsf{tk}, i; \mathsf{st}, \mathsf{ctr})$

return $\perp$

$\mathcal{O}_{\mathsf{Eval}}(\mathsf{apk}, \mathsf{C}, \mathsf{ind}; \mathsf{st}, \mathsf{ctr})$

Parse $\mathsf{ind} = (i_1, \ldots, i_{p(\lambda)})$
if $\#\text{-inputs}(\mathsf{C}) \neq p(\lambda)$ **return** $\perp$
if $\mathsf{C} \notin \mathcal{C}$ **return** $\perp$
if $\exists j \in [p(\lambda)] : i_j \notin [\mathsf{ctr}]$ **return** $\perp$
for $j \in [p(\lambda)]$ **do** $(b_j, c_j) \leftarrow \mathsf{st}_{i_j}$
if $\wedge_{j \in p(\lambda)} b_j \neq 1$ **return** $\perp$
$c' = \mathsf{HE.Eval}(\mathsf{C}, (c_1, \ldots, c_{p(\lambda)}))$
$\mathsf{st} \leftarrow \mathsf{st} \| (0, c')$; $\mathsf{ctr} \leftarrow \mathsf{ctr} + 1$
return 1

Fig. 1. Oracles for the homomorphic robustness game in Definition 9

In the definition below, we denote the set of permissible circuits (or equivalently, i-hop "compatible sequences") as $\mathcal{C}$. Therefore, any $\mathsf{C} \in \mathcal{C}$ implicitly describes a compatible sequence for the i-hop homomorphic encryption scheme in question.

Definition 9. *Let the scheme* $\Pi = (\mathsf{HE.KGen}, \mathsf{HE.Enc}, \mathsf{HE.Dec}, \mathsf{HE.Eval})$ *be a* $\mathcal{C}$*-homomorphic, i-hop encryption scheme for sufficiently large* $i = i(\mathcal{C})$ *with deterministic* $\mathsf{HE.Eval}$*. An anamorphic triple* $\Sigma = (\mathsf{HE.aGen}, \mathsf{HE.aEnc}, \mathsf{HE.aDec})$ *for* Π *is* $\mathcal{C}$*-homomorphically robust if, for any PPT adversary* $\mathcal{A}$*, there exists a negligible function* $\eta(\lambda)$ *such that*

$$\mathsf{Adv}^{\mathsf{Hom\text{-}Rob}}_{\mathcal{A},\Pi,\Sigma}(\lambda) = \left|\Pr[\mathsf{Hom\text{-}Rob}^0_{\mathsf{Adv},\Pi,\Sigma}(\lambda) = 1] - \Pr[\mathsf{Hom\text{-}Rob}^1_{\mathsf{Adv},\Pi,\Sigma}(\lambda) = 1]\right| \leq \eta(\lambda)$$

where the experiments $\mathsf{Hom\text{-}Rob}^0_{\mathcal{A},\Pi,\Sigma}(\lambda)$ *and* $\mathsf{Hom\text{-}Rob}^1_{\mathcal{A},\Pi,\Sigma}(\lambda)$ *are defined as follows.*

$\mathsf{Hom\text{-}Rob}^b_{\mathcal{A},\Pi,\Sigma}(\lambda)$:

1. $((\mathsf{apk}, \mathsf{ask}), \mathsf{tk}, \mathsf{dk}) \leftarrow \mathsf{HE.aGen}(1^\lambda)$;
2. *Initialize* $\mathsf{ctr} = 0$,, *and state* $\mathsf{st} = ()$, *as an empty list;*
3. *Return* $\mathcal{A}^{\mathcal{O}_{\mathsf{Enc}}(\mathsf{apk},\cdot;\mathsf{st},\mathsf{ctr}),\mathcal{O}_{\mathsf{Eval}}(\mathsf{apk},\cdot,\cdot;\mathsf{st},\mathsf{ctr}),\mathcal{O}_b(\mathsf{ask},\mathsf{dk},\mathsf{tk},\cdot;\mathsf{st},\mathsf{ctr}),}(\mathsf{apk}, \mathsf{ask})$ *where the oracles* $\mathcal{O}_{\mathsf{Enc}}, \mathcal{O}_{\mathsf{Eval}}, \mathcal{O}_b$ *are described in Figure 1.*

If a scheme is $\mathcal{C}$-homomorphically robust, where $\mathcal{C}$ is a class of addition circuits, we simply say the scheme is additively homomorphically robust or, equivalently, that the scheme has additively homomorphic robustness. The presence of a bound on the number of additions will be implicitly understood in the case of dual Regev or GSW. There are also alternative scenarios where one may unknowingly perform homomorphic evaluation on a mixture of anamorphic and normal encryptions. The above definition of homomorphic robustness does not

capture this possibility as all ciphertexts are normal. Nonetheless, if all ciphertexts originate from a sender that does not attempt to send covert messages, the above notion of robustness ensures that, even after homomorphic evaluation, the result will not mistakenly be anamorphically decrypted to a covert message.

In order to consider evaluating a mixture of normal and anamorphic ciphertexts, we could consider the case where a computation involving at least one normal ciphertext destroys the covert message. This yields a stronger homomorphic robustness property that implies the weaker notion above. However, it is not clear what should happen when a circuit mixing normal and anamorphic ciphertexts does not depend on its normal ciphertext's plaintext message. For example, if we consider the Boolean circuit $\mathsf{C}(x_1, x_2) = (x_1 \vee \bar{x}_1) \wedge x_2$, its output clearly does not depend on x_1. So, if the ciphertext encrypting x_1 is normal, one may prefer keeping the covert message unaltered. Therefore, the exact robustness definition would depend on the notion of mixing that an application desires. If it is enough to destroy covert messages as soon as a normal ciphertext is non-trivially involved in a homomorphic computation, then the weaker definition above can be extended. However, if the application cares about whether the plaintexts underlying normal ciphertexts affect the output of the circuit, then things are more complicated. The main difficulty is that the ideal decryption oracle may not be able to efficiently decide whether a covert message should be destroyed because it cannot *efficiently* test whether a (possibly complex) circuit depends on particular inputs. Therefore, we leave formally defining stronger definitions of robustness to future work.

E.2 Analyzing Homomorphic Robustness in Dual Regev/GSW

The dual Regev case. We now show that the construction in Section 3 can be made homomorphically robust modulo a slight change in the aDec algorithm. To do this, we exploit the fact that Enc uses uniform randomness $\mathbf{s}$ whereas aEnc uses small Gaussian randomness. When it comes to applying aDec, the size of the recovered randomness candidate $\mathbf{s}$ is compared with a bound to decide whether the ciphertext was intended to have an anamorphic message or not. In what follows, we let L denote the maximum number of allowed additions to preserve correctness. As a result, the error rates in the scheme are assumed to be reduced by a factor of L compared to those in Section 3 to preserve correctness with respect to L additions. To achieve this, one can imagine taking L as an input to KGen and aGen and increasing the modulus q by a factor of L. This is the parametrisation referred to in Lemma 16. We also note that as homomorphic addition corresponds to adding ciphertexts, the dual Regev scheme immediately yields an L-hop homomorphic encryption scheme. The updated aDec algorithm goes as follows:

$\mathsf{aDec}(\mathsf{dk}, \mathsf{tk}, \mathsf{ask}, \mathsf{act})$: Given a ciphertext $\mathsf{act} = (\mathbf{c}_0, \mathbf{c}_1) \in \mathbb{Z}_q^{(m+n)}$ and the trapdoor key $\mathsf{tk} = \mathbf{R}_C \in \{-1, 0, 1\}^{(\ell k + 2\lambda) \times \ell k}$,

1. Compute $(\hat{\mathbf{s}}', \mathbf{e}_0) \leftarrow \mathsf{Invert}(\mathbf{R}_C, \mathbf{c}_0)$ using the Invert algorithm of Lemma 6. If $(\hat{\mathbf{s}}', \mathbf{e}_0) = (\perp, \perp)$ (meaning there is no $\|\mathbf{e}_0\| \leq B$ such that $\mathbf{c}_0 = \mathbf{A}^\top \hat{\mathbf{s}} + \mathbf{e}_0$ for some $\hat{\mathbf{s}} \in \mathbb{Z}_q^n$), return $\perp$.
2. Compute $\hat{\boldsymbol{\mu}} = \mathsf{decode}_p(\hat{\mathbf{s}}') \in \mathbb{Z}_p^n$ using the decoding algorithm of Section 2.3.
3. Then, compute $\mathbf{s} = \hat{\mathbf{s}} - \Delta \cdot \hat{\boldsymbol{\mu}} \bmod q$. If $\|\mathbf{s}\|_\infty > L \cdot \alpha q \cdot \sqrt{2\lambda}$, return $\mu = \perp$. Otherwise, return $\hat{\boldsymbol{\mu}} \in \mathbb{Z}_p^n$.

Correctness of aDec for decrypting anamorphic ciphertexts follows from the fact that $\|\mathbf{s}\|_\infty \leq \alpha q \cdot \sqrt{2\lambda}$ with overwhelming probability $1 - 2^{-\Omega(\lambda)}$ by Lemma 3.

Lemma 16. *The modified* aDec *algorithm provides additively homomorphic robustness for the anamorphic dual Regev triplet and parameters in Section 3 (with modulus scaled up by L) if q is prime, $n = \Omega(\lambda)$ and $2L\alpha\sqrt{2\lambda} + \frac{1}{q} \leq 1/2$.*

Proof. At the end of all of the adversary $\mathcal{A}$'s queries in the $\mathsf{Hom\text{-}Rob}^b$ game, we write the state st as two parts – one for ciphertexts computed as fresh ciphertexts using $\mathcal{O}_{\mathsf{Enc}}$ called $\mathsf{st}_{\mathsf{fresh}}$ and another part for the remaining ciphertexts $\mathsf{st}_{\mathsf{eval}}$. In other words, $\mathsf{st}_{\mathsf{fresh}}$ contains ciphertexts with indicator bit 1 and $\mathsf{st}_{\mathsf{eval}}$ contains ciphertexts with indicator bit 0. Assuming Q queries to $\mathcal{O}_{\mathsf{Enc}}$, we will write $\mathsf{st}_{\mathsf{fresh}} = \{(\mathbf{c}^{(i)} = (\mathbf{c}_0^{(i)}, \mathbf{c}_1^{(i)})) : i \in [Q]\}$. Any ciphertext in $\mathsf{st}_{\mathsf{eval}}$ is the result of adding at most L ciphertexts from $\mathsf{st}_{\mathsf{fresh}}$.

Suppose there are Q' queries to the decryption oracle. Then in the k-th such query, we can express the ciphertext that $\mathcal{A}$ wishes to be anamorphically decrypted via $\bar{\mathbf{c}}^{(k)} = \sum_{i=1}^{Q} x_i^{(k)} \cdot \mathbf{c}^{(i)}$ using a vector satisfying $\|\mathbf{x}^{(k)}\|_1 \leq L$ and $\mathbf{x}^{(k)} \neq \mathbf{0}$. We now bound the probability that $\mathbf{x}^{(k)}$ satisfies $\mathsf{aDec}\left(\mathsf{dk}, \mathsf{tk}, \mathsf{ask}, \bar{\mathbf{c}}^{(k)}\right) \neq \perp$ for any single value of $k \in [Q']$. We will denote the encryption randomness or "secret" vector that right multiplies the public key in $\mathbf{c}_0^{(i)}$ as $\mathbf{s}^{(i)}$. Similarly, we denote the error sampled for $\mathbf{c}_0^{(i)}$ as $\mathbf{e}_0^{(i)}$. It follows by homomorphism that $\bar{\mathbf{c}}^{(k)}$ has "secret" vector $\bar{\mathbf{s}}^{(k)} = \sum_{i=1}^{Q} x_i^{(k)} \cdot \mathbf{s}^{(i)}$ and error $\bar{\mathbf{e}}_0^{(k)} = \sum_{i=1}^{Q} x_i^{(k)} \cdot \mathbf{e}_0^{(i)}$. By correctness of Invert and the scheme's parameters, the anamorphic decryption algorithm applied to $\bar{\mathbf{c}}^{(k)}$ correctly recovers $\bar{\mathbf{e}}_0^{(i)}$ and $\bar{\mathbf{s}}^{(k)}$ with probability $1 - 2^{-\Omega(\lambda)}$. Furthermore, aDec does not output $\perp$ (and therefore the adversary may distinguish between $\mathcal{O}_0$ and $\mathcal{O}_1$) if and only if $\bar{\mathbf{s}}^{(k)} \in \mathsf{BAD} \triangleq \{\Delta \cdot \mathbf{v} : \mathbf{v} \in \mathbb{Z}_p^k\} + [-L\alpha q\sqrt{2\lambda}, L\alpha q\sqrt{2\lambda}]^n$. Since $\mathbf{x}^{(k)} \neq \mathbf{0}$, q is prime and $\mathbf{s}^{(i)}$ is uniform for each $i \in [Q]$, we have that $\bar{\mathbf{s}}^{(k)}$ is uniformly distributed. Therefore,

$$\Pr[\bar{\mathbf{s}}^{(k)} \in \mathsf{BAD}] = |\mathsf{BAD}| / q^n \leq p \cdot \left(2L\alpha\sqrt{2\lambda} + \tfrac{1}{q}\right)^n.$$

We then apply a union bound over all Q' queries to $\mathcal{O}_b$ to show that $\mathcal{A}$'s overall advantage in the homomorphic robustness game is $\leq Q' \cdot p \cdot (2L\alpha\sqrt{2\lambda} + \frac{1}{q})^n$ which is negligible if $2L\alpha\sqrt{2\lambda} + \frac{1}{q} \leq 1/2$. □

The dual GSW case. Using similar ideas as above, we now show that a tweak of the dual GSW scheme from Section 5 is homomorphically robust under addition. Again, if L denotes the maximum number of additions, the dual GSW scheme is L-hop since homomorphic addition corresponds to addition of ciphertexts. The only change to the anamorphic construction of dual GSW is in the aDec algorithm, which is presented next. All the remaining algorithms are as in Section 5.

$\mathbf{aDec}\big(\mathsf{dk}, \mathsf{tk}, \mathsf{ask}, \mathsf{act}\big)$**:** Given a ciphertext $\mathsf{act} = \mathbf{C} \in \mathbb{Z}_q^{(m+1)\times M}$ and the trapdoor key $\mathsf{tk} = \mathbf{t} \in \mathbb{Z}^m$, define $\hat{\mathbf{e}}_{i_{n_0}} = (0, \ldots, 0, 1, 0, \ldots, 0)^\top \in \{0,1\}^{m+1}$ as the i_{n_0}-th unit vector and compute

$$\nu = [\ \mathbf{t}^\top \mid 0] \cdot \mathbf{C} \cdot \mathbf{G}^{-1}(\Delta \cdot \hat{\mathbf{e}}_{i_{n_0}}) \in \mathbb{Z}_q$$

where $\Delta = \lfloor q/p \rceil$. Then, set $\hat{\mu} = \lfloor \nu/\Delta \rceil \in \mathbb{Z}_p$ and compute

$$e = \big\| [\ \mathbf{t}^\top \mid 0] \cdot (\mathbf{C} - \hat{\mu} \cdot \mathbf{G}_I) \big\|_\infty .$$

If $e > q/2p$, then output $\perp$. Otherwise, output $\hat{\mu}$.

Correctness of aDec when decrypting anamorphic ciphertexts follows by virtue of the correctness analysis in Section D.1. In more detail, for an evaluated ciphertext with error $\mathbf{E}_\mathbf{C}$ (using the notation from the aforementioned section), the correctness analysis ensures that that $[\ \mathbf{t}^\top \mid 0] \cdot (\mathbf{C} - \hat{\mu} \cdot \mathbf{G}_I) = [\ \mathbf{t}^\top \mid 0] \cdot \mathbf{E}_\mathbf{C}$ has infinity norm smaller than $q/(2p)$. Lemma 17 addresses the robustness of ciphertexts resulting from homomorphic additions. We then discuss the difficulty or proving robustness under more involved homomorphic computations.

Lemma 17. *The modified* aDec *algorithm provides additively homomorphic robustness with the anamorphic dual GSW triplet and parameters in Section 5 if q is prime.*

Proof. At the end of the homomorphic robustness game, we denote $\mathsf{st}_\mathsf{fresh} = \{\mathbf{C} : (1, \mathbf{C}) \in \mathsf{st}\}$. In other words, $\mathsf{st}_\mathsf{fresh} = \{\mathbf{C}_1, \ldots, \mathbf{C}_Q\}$ contains all ciphertexts produced by the adversary $\mathcal{A}$'s queries to $\mathcal{O}_\mathsf{Enc}$. We assume parameters are correct for up to L_add homomorphic additions. When $b = 0$, any query to $\mathcal{O}_b$ runs $\mathsf{HE.aDec}\big(\mathsf{dk}, \mathsf{tk}, \mathsf{ask}, \mathsf{HE.Eval}(\mathsf{apk}, \mathsf{C}, (\bar{\mathbf{C}}_1, \ldots, \bar{\mathbf{C}}_\ell))\big)$ where $\mathsf{C} \in \mathcal{C}$ is some *addition* circuit with $\ell \le L_\mathsf{add}$ inputs, and $(\bar{\mathbf{C}}_1, \ldots, \bar{\mathbf{C}}_\ell) \in \mathsf{st}_\mathsf{fresh}^\ell$. Note that the upper bound L_add is for correctness only and does not play a meaningful role in the proof. We may then write the ciphertext to be anamorphically decrypted by a query to $\mathcal{O}_{b=0}$ as $\mathbf{C}_\mathsf{eval} = \mathsf{HE.Eval}(\mathsf{apk}, \mathsf{C}, \bar{\mathbf{C}}_1, \ldots, \bar{\mathbf{C}}_\ell) = \sum_{i\in[\ell]} \bar{\mathbf{C}}_i = \sum_{i\in[Q]} x_i \cdot \mathbf{C}_i$ for some $\mathbf{x} \neq \mathbf{0}$ with $\|\mathbf{x}\|_1 \le L_\mathsf{add}$. Recall that $\mathbf{C}_i = \mathbf{B}\cdot\mathbf{S}_i + \mathbf{E}_i + \mu_i \cdot \mathbf{G}$ where $\mathbf{S}_i \hookleftarrow U(\mathbb{Z}_q^{m\times M})$. We will refer to $\mathbf{S}_i$ and $\mathbf{E}_i$ as the ciphertext secret and error of $\mathbf{C}_i$ respectively. Now, the ciphertext secret and error for $\mathbf{C}_\mathsf{eval}$ is $\mathbf{S}_\mathsf{eval} = \sum_{i\in[Q]} x_i \cdot \mathbf{S}_i$ and $\mathbf{E}_\mathsf{eval} = \sum_{i\in[Q]} x_i \cdot \mathbf{E}_i$ respectively. Clearly, $\mathbf{S}_\mathsf{eval}$ follows the uniform distribution as q is prime, $\mathbf{x} \neq \mathbf{0}$ and the $\mathbf{S}_i$'s are uniform. During anamorphic decryption, the value

$$\nu = [\ \mathbf{t}^\top \mid 0]\cdot\mathbf{C}_\mathsf{eval}\cdot\mathbf{G}^{-1}(\Delta\cdot\hat{\mathbf{e}}_{i_{n_0}}) = \big(\mathbf{e}_I^\top \cdot \mathbf{S}_\mathsf{eval} + [\mathbf{t}^\top \mid 0] \cdot \mathbf{E}_\mathsf{eval}\big)\cdot\mathbf{G}^{-1}(\Delta\cdot\hat{\mathbf{e}}_{i_{n_0}}) + \Delta\cdot\mu$$

is computed, where we write $\mathbf{C}_{\mathsf{eval}} = \mathbf{B} \cdot \mathbf{S}_{\mathsf{eval}} + \mathbf{E}_{\mathsf{eval}} + \mu \cdot \mathbf{G}$. Next, decoding ν recovers the unique value $\hat{\mu}$ such that $\nu \in \Delta \cdot \hat{\mu} + (-\frac{\Delta}{2}, \frac{\Delta}{2}]$ and $\hat{\mu}$ is accepted as a valid decryption if

$$\begin{aligned}[\mathbf{t}^\top \mid 0] \cdot \big(\mathbf{B}_I \cdot \mathbf{S}_{\mathsf{eval}} + \mathbf{E}_{\mathsf{eval}} + (\mu - \hat{\mu}) \cdot \mathbf{G}_I\big) \\ = \mathbf{e}_I^\top \cdot \mathbf{S}_{\mathsf{eval}} + [\mathbf{t}^\top \mid 0] \cdot \big(\mathbf{E}_{\mathsf{eval}} + (\mu - \hat{\mu}) \cdot \mathbf{G}_I\big) \end{aligned} \tag{30}$$

has infinity norm less than $q/2p$. We now note that $\hat{\mu}$ depends only on columns $(i_{n_0} - 1) \cdot \lceil \log q \rceil + 1$ to $i_{n_0} \cdot \lceil \log q \rceil$ of $\mathbf{S}_{\mathsf{eval}}$ due to the definition of $\mathbf{G}^{-1}(\hat{\mathbf{e}}_{i_{n_0}})$ and is independent of the remaining columns. Lemma 18 then shows that all but the $((i_{n_0} - 1) \cdot \lceil \log q \rceil + 1)$-th to $(i_{n_0} \cdot \lceil \log q \rceil)$-th entries of $\mathbf{e}_I^\top \cdot \mathbf{S}_{\mathsf{eval}}$ are uniform and independent of $\hat{\mu}$ as long as $\mathbf{e}_I \neq \mathbf{0}$ (which happens with overwhelming probability $1 - 2^{\Omega(\lambda)}$ by Lemma 15). It then follows that these entries are also uniformly distributed in (30). Therefore, the probability that one of these $M - \lceil \log q \rceil$ uniformly distributed entries in (30) has absolute value at most $q/(2p)$ is at most $\left((2 \cdot \frac{q}{2p} + 1)/q\right) = \left(\frac{1}{p} + \frac{1}{q}\right)$. This allows us to conclude that $\hat{\mu}$ is accepted by aDec with probability at most $(1 - 2^{-\Omega(\lambda)}) \cdot \left(\frac{1}{p} + \frac{1}{q}\right)^{M - \lceil \log q \rceil} = 2^{-\Omega(\lambda)}$. In other words, for a single query to $\mathcal{O}_b$, one can distinguish between $b = 0$ and $b = 1$ with probability at most $2^{-\Omega(\lambda)}$. Applying a union bound over all queries to $\mathcal{O}_b$ completes the proof. □

Lemma 18. *Let q be a prime and take any non-zero vector $\mathbf{v} \in \mathbb{Z}_q^m$. If $\mathbf{X} \in \mathbb{Z}_q^{m \times M}$ is uniformly distributed, then so is $\mathbf{v}^\top \cdot \mathbf{X} \in \mathbb{Z}_q^M$.*

Proof. Let the j-th row of $\mathbf{X}$ be denoted by $\mathbf{x}^{(j)}$. Since $\mathbf{y} \triangleq \mathbf{v}^\top \cdot \mathbf{X} = \sum_{j \in [M]} v_j \cdot \mathbf{x}_j$, we have that $\mathbf{y} \in \mathbb{Z}_q^m$ is uniformly distributed since there is at least one non-zero v_j and q is prime. □

Arithmetic Circuits. Proving that the anamorphic dual GSW construction is homomorphically robust with respect to arithmetic circuits is less trivial. However, it appears possible to prove this more general case using a 1D-SIS assumption [17]. Reusing the notation from the proof for additive homomorphism, we can say that any arithmetic circuit query to $\mathcal{O}_b$ attempts to decrypt a ciphertext $\mathbf{C}_{\mathsf{eval}}$ with ciphertext secret $\mathbf{S}_{\mathsf{eval}}$ and short error matrix $\mathbf{E}_{\mathsf{eval}}$. Throughout this informal discussion, we take q to be a multiple of p ignoring complications that this may bring. The crucial observation is that $\mathbf{S}_{\mathsf{eval}} = [\mathbf{S}_1 \mid \cdots \mid \mathbf{S}_Q] \cdot \mathbf{H}$ for some short $\mathbf{H}$. Then, the anamorphic decryption is accepted if (30) has infinity norm at most $q/2p$. Right multiplying by $\mathbf{G}^{-1}(\Delta \cdot \hat{\mathbf{e}}_{i_{n_0}})$ implies that $\mathbf{e}_I^\top \cdot [\mathbf{S}_1 \mid \cdots \mid \mathbf{S}_Q] \cdot \big(\mathbf{H} \cdot \mathbf{G}^{-1}(\Delta \cdot \hat{\mathbf{e}}_{i_{n_0}})\big) \in \Delta \cdot \mathbb{Z} + [-t, t]$ for some "small" t. The task of finding a short $\mathbf{x}$ given a uniform $\mathbf{v}$ such that $\mathbf{v}^\top \cdot \mathbf{x} \in \Delta \cdot \mathbb{Z} + [-t, t]$ is known as the 1D-SIS-R problem [17]. Therefore, if one chooses uniform $[\mathbf{S}_1 \mid \cdots \mid \mathbf{S}_Q]$ such that $\mathbf{v}^\top = \mathbf{e}_I^\top \cdot [\mathbf{S}_1 \mid \cdots \mid \mathbf{S}_Q]$ is a 1D-SIS-R challenge, one may use a break of the homomorphic robustness game to obtain a solution to the 1D-SIS-R problem. By a reduction from 1D-SIS to 1D-SIS-R [17], homomorphic robustness

would follow. A major challenge of meaningfully applying these ideas are that 1D-SIS security requires an exponentially large modulus that factors into $\Theta(\lambda)$ co-prime integers. This manifests in an obvious deviation from a normal set of parameters for dual GSW. As a result, a 1D-SIS-based solution may raise a dictator's suspicion and hint that anamorphic encryption is being used. Therefore, we leave the task of obtaining a more satisfactory solution as an open question.

Trapdoor-less dual Regev. The anamorphic dual Regev scheme from Supplementary Material C can similarly achieve additively homomorphic robustness with only slight changes. Firstly, the matrix $(\mathbf{A}')$ in (24) should be amended to

$$(\mathbf{A}')^\top = \begin{bmatrix} \bar{\mathbf{A}} \\ \mathbf{T}^\top \cdot \bar{\mathbf{A}} + \mathbf{E}^\top \end{bmatrix} \tag{31}$$

where $\mathbf{E} \in \mathbb{Z}_q^{n \times k}$ is a small-norm matrix of $\mathbb{Z}_q$-rank $\min(n, k)$. For example, if $k \geq n$ (meaning that $\mathbf{E}^\top$ is tall), we will set $\mathbf{E} = [\mathbf{I}_n \mid \mathbf{0}^{n \times (k-n)}]$. Alternatively, if $k < n$ (i.e. $\mathbf{E}^\top$ is wide), we will set $\mathbf{E}^\top = [\mathbf{I}_k \mid \mathbf{0}^{k \times (n-k)}]$. This change is undetectable by the Leftover Hash Lemma since the distribution of $\mathbf{T}^\top \cdot \bar{\mathbf{A}}$ remains statistically close to $U(\mathbb{Z}_q^{k \times n})$. Then, the uniform vector $\mathbf{r} \in \mathbb{Z}_q^n$ in the aEnc algorithm should be replaced by a Gaussian vector sampled from the discrete Gaussian distribution $D_{\mathbb{Z}, \alpha q}$. Finally, the aDec algorithm can check whether the randomness $\mathbf{r}$ was large or not by inspecting the decoding error vector in the covert message positions. When applying aDec, the key term to be analyzed for robustness is the product $\mathbf{E}^\top \cdot \mathbf{r} \bmod q$. If $\mathbf{r} \hookleftarrow U(\mathbb{Z}_q^n)$ and $k \geq n$, then the first n entries of $\mathbf{E}^\top \cdot \mathbf{r}$ are uniform assuming the form of $\mathbf{E}^\top$ described above. If $k < n$, then all k entries of $\mathbf{E}^\top \cdot \mathbf{r}$ are uniform. Assuming $\min(n, k)$ is large enough (e.g. $\Omega(\lambda)$), one can argue robustness by considering the uniform entries of $\mathbf{E}^\top \cdot \mathbf{r} \bmod q$ and notice that they all land in a small interval with probability $\leq 2^{-\min(n,k)}$.

We also note that one can alternatively rely on the HNF LWE assumption in (31) and choose $\mathbf{T}$ and $\mathbf{E}$ to be Gaussian matrices. Assuming $n > k$, we could choose $n = \widetilde{\Omega}(k)$ and use [2, Lemma 9] to ensure $\mathbf{E}$ has rank k with overwhelming probability. The robustness argument would then be completed similarly to above, by considering each of the $k = \Omega(\lambda)$ uniform entries of $\mathbf{E}^\top \cdot \mathbf{r}$.