

Post-Quantum PKE and More from a Noisy Unstructured Linear Algebraic Assumption: Beyond LWE and LPN

Riddhi Ghosal ^{*}
Nokia Bell Labs

Paul Lou [†]
Bocconi University

Amit Sahai [‡]
UCLA

Abstract

Noisy linear algebraic assumptions (NLAs), such as LWE and Alekhnovich’s LPN, have long served as the most reliable sources of post-quantum hardness. However, a series of recent classical attacks on assumptions once believed to be quantum-secure, coupled with the rapid progress in quantum technology, underscores the urgent need for assumptions that are plausibly *hard* even if both LWE and LPN turn out to be broken.

A fascinating recent work of Ghosal et al. (Eurocrypt 2025) takes an important first step in this direction, constructing a PKE scheme by combining two NLAs that remain plausibly secure even if LWE and Alekhnovich’s parameter regime for LPN were to be polynomially broken. A critical weakness of their work, however, is that it is still LPN-broken: namely, an oracle that breaks LPN with $n^{-\delta}$ noise rate can break their NLA, for some $\delta > 0$. We ask: *can this barrier be overcome?*

We answer this question positively, by introducing an NLA that is plausibly hard given an oracle that breaks even *constant* noise rate LPN, while still being plausibly hard given an oracle that breaks LWE. Along the way, we improve upon the work of Ghosal et al. significantly in several dimensions.

Specifically, we introduce a single new natural NLA and provide constructions of (i) public-key encryption and (ii) two-round maliciously secure oblivious transfer protocol from this single assumption. We also construct a one-time circular-secure secret-key encryption from a binary variant of this assumption.

Our assumption takes the form $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$, where $\mathbf{A}$ is a random unstructured *square* matrix, the secret $\mathbf{s}$ has *short and sparse entries*, and the error $\mathbf{e}$ is a mixture of two components: a *small-but-dense error* combined with a *large-but-sparse error*.

We provide a comprehensive initial cryptanalysis, ruling out (1) efficient combinatorial attacks that exploit sparsity such as information set decoding, (2) dimension reduction attacks, and (3) dual lattice attacks. We also prove, via a reduction, that our assumption, and its binary variant, is at-least as hard as LWE. Additionally, we show that for a broad range of parameters including those used by our constructions, our assumption does not appear to fall within the class of “lattice” assumptions, that is, there seems to be no reduction to approximate shortest vector problem solvers, and, crucially, it is plausibly hard even given an oracle that breaks LPN for *any noise density* parameter.

1 Introduction

In recent years, quantum technology has seen remarkable advances, with significant innovation in both algorithmic development [Sho94, Pre12, CHSH69, BCM⁺18, BKVV20, KMCVY22, KLVY23, JSW⁺24, Hal07, BBS09, YZ22, BDG⁺25, BBW25] and hardware capabilities [IBM25, Gam23, Nev24, Bol25]. This has made the construction of post-quantum secure public-key cryptography (with classical secrets and ciphertexts) a matter of utmost importance. Yet, despite years of effort, the number of plausible assumptions capable of yielding post-quantum security remains very limited. These include isogenies [ABLJK18, CLM⁺18, Cou06, RS06], multivariate polynomial based assumptions [Pat96, DS05a,

^{*}riddhi.ghosal@nokia-bell-labs.com

[†]paul.lou@unibocconi.it

[‡]sahai@cs.ucla.edu

TDTD13, KPG99, DS05b, PGC98, CFMR⁺17, ZT15], minrank based [CMV26, DAGNR25], graphs with hidden structure based assumptions [ABW10, BKR23, GHJS25], and noisy linear algebraic assumptions (NLAs) that includes the very popular Learning with Errors (LWE) [Reg09] and Learning Parity with Noise (LPN) [BFKL94, Ale03], among others [GJL⁺25].

The state of affairs has become more worrisome after recent discoveries of classical attacks on some isogeny [Rob23] and multivariate quadratic polynomial [Beu22] based assumptions that were initially believed to be quantum safe. This should raise alarms because the short history of modern cryptography has taught us that it is crucial to remain prepared for cryptanalytic breakthroughs. Indeed, it was not too long ago when Shor’s algorithm [Sho94] single-handedly quantumly broke the two most centrally used cryptographic assumptions at the time, the hardness of factoring and the hardness of discrete logarithm (including over elliptic curve groups). Certainly, if a similar algorithm that broke the (quantum) security of LWE or LPN were to appear, it could prove catastrophic to the world of post-quantum cryptography. In recent times, there certainly have been some serious (albeit failed) attempts to break LWE quantumly [ES16, Che24] and quantum speed-ups against certain LPN type assumptions [Has19, SOKB24]. Therefore, as things stand, there are too few well-understood hardness assumptions that can plausibly be quantum-safe, and our community’s reliance on LPN and especially LWE is extremely high.

At the same time, NLAs remain the most resilient, versatile and most well-studied family of assumptions. NLAs are assumptions of the form $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$ being computationally indistinguishable from $(\mathbf{A}, \mathbf{u})$ for polynomial-time adversaries, where $\mathbf{A} \in \mathbb{F}_q^{m \times n}$ is a typically an expanding matrix ($m > n$), $\mathbf{s}$ is typically uniform over $\mathbb{F}_q^n$, $\mathbf{u}$ is typically uniform over $\mathbb{F}_q^m$, and $\mathbf{e} \in \mathbb{F}_q^m$ is “noise” from an appropriate distribution with some underlying structure. Two key assumptions in this category are (1) Learning with Errors (LWE) [Reg09] where the error vector $\mathbf{e}$ has small integer entries from a discrete gaussian distribution centered at zero, and (2) Learning Parity with Noise [BFKL94, Ale03, IPS09] where $\mathbf{e}$ is a sparse vector with roughly $mn^{-\delta}$ many non-zero entries that are uniform over $\mathbb{F}_q$. In general, LPN is said to be δ -dense if the probability of non-zero entries is $n^{-\delta}$, and Alekhnovich’s LPN is the special case of $\delta = \frac{1}{2}$. Due to their simple mathematical structure, LWE and LPN can be used to build a plethora of primitives (in different parameter regimes) starting from public key encryption (PKE) [Reg09, Ale03, YZ16], non-interactive proof systems [PS19, BKM20, DJJ24], to advanced primitives like round optimal oblivious transfer [DGH⁺20], fully homomorphic encryption [Gen10, BV11, BGV12, GSW13], functional encryption [ABV⁺12, AFV11], indistinguishability obfuscation [JLS21], and more [GPV08, Lyu12, WZ17].

New Noisy Linear Assumptions. With this in mind, a natural direction to improve the state-of-affairs in terms of post-quantum secure hardness assumptions is to explore new noisy linear algebraic assumptions that can be plausibly hard even if LWE and LPN turn out to be broken. In a fascinating recent work, [GJL⁺25] took a significant step in this direction. They proposed two new NLAs, namely, (1) Learning with Two Errors (LW2E) where the error $\mathbf{e} = \mathbf{e}_1 + \mathbf{e}_2$, with $\mathbf{e}_1$ being short LWE error and $\mathbf{e}_2$ being sparse LPN error, and (2) Learning with Short and Sparse Errors (LWSSE) where $\mathbf{e}$ is a sparse error with low ℓ_2 norm, but is denser than Alekhnovich’s LPN, i.e., the probability of error is $n^{-\gamma}$ for $\gamma < \frac{1}{2}$. They showed that both assumptions together imply PKE in a parameter regime in which *both* assumptions are potentially quantum-secure even if LWE is broken and LPN is broken, albeit only in the Alekhnovich parameter regime of $\delta = \frac{1}{2}$. In more detail, they give evidence that the two new assumptions are not subject to complexity-theoretic reductions to lattice assumptions nor to typical cryptanalytical strategies applicable to lattice assumptions. Notably, however, they do concede that LWSSE *is broken* by an LPN solver for some constant $\delta > 0$.

We ask: *Is this limitation surmountable?* In other words, is it possible to develop an NLA that allows us to construct a PKE, yet remains secure in the presence of *any* LWE and LPN breaking oracle?

Our Contribution. In our work, we answer this question in the affirmative. In particular, we introduce a new noisy linear algebraic assumption and argue that for a wide range of parameters, it appears to be quantum-secure even if *both* LPN (with any parameters) and LWE are (quantum)-broken. Indeed,

we also show that in the same parameter regime, our assumption can be used to construct public key encryption and more!

To support our claim, we provide comprehensive analysis that it is unlikely that our assumption reduces to LPN with sparsity $n^{-\delta}$ for *any* $\delta > 0$. In fact, based on our analysis, we make a stronger conjecture that our assumption is likely to remain secure even if LPN with constant sparsity (probability of errors is $\frac{1}{2} - \varepsilon$) is broken. We provide a reduction from LWE to our assumption. Additionally, for our assumption, we give evidence that it is not subject to complexity-theoretic reductions to lattice assumptions or to typical cryptanalytical strategies applicable to lattice assumptions.

Our work also resolves another shortcoming of [GJL⁺25]. In order to construct the PKE, [GJL⁺25] needed to use *two* different new NLAs. Developing comprehensive cryptanalysis for *two new assumptions* requires significant effort. Indeed, it is essential to carefully understand the security guaranteed by both their underlying assumptions, and two separate assumptions also mean two potential points of failure.

In such a situation, our result improves the state of affairs by providing a single cryptanalytic target.

We also show that our assumption can be used to construct advanced primitives beyond public key encryption. Specifically, we also build two-round secure malicious oblivious transfer [CCM98, NP01, AIR01, PVW08, DGH⁺20].

Our assumption. Similar to LWE, LPN, LW2E and LWSSE, our assumption is also defined with respect to polynomial-time adversaries and holds for random *unstructured* matrices $\mathbf{A}$. We call such assumptions Noisy Unstructured Linear Assumptions (NULAs). This is in contrast to structured and/or subexponentially strong variants of LWE and LPN like [DIJL23, DJ24, CDMT24, HPS98, LPR10, DP12, LPR13, PRS17, McE78, YZ16]. Our PKE (and other constructions) also deviate from that of [GJL⁺25], in the sense that we are able to achieve the constructions by symmetric use of our assumption twice which we will describe later. We now define our hardness assumption.

The Square Noisy Unstructured Linear Algebraic Assumption (sqNULA). In this assumption, the adversary is given $\mathbf{A}$ and $\mathbf{A}\mathbf{s} + \mathbf{e}$, where $\mathbf{e} = \mathbf{e}_1 + \mathbf{e}_2$. Here $\mathbf{A}$ is a random *square* matrix over $\mathbb{F}_q^{n \times n}$, $\mathbf{s}$ is a sparse secret over $\mathbb{F}_q^n$ with small ℓ_2 norm, $\mathbf{e}_1$ is short LWE error, and $\mathbf{e}_2$ is sparse LPN error. The assumption states that no polynomial-time adversary can recover the short and sparse secret $\mathbf{s}$ except with negligible probability. The crucial observation is that we have picked $\mathbf{A}$ to be a square matrix, i.e., $m = n$. This choice plays a critical role in achieving our desired security in the face of LWE and LPN breakers. The number of rows in $\mathbf{A}$ determines the number of NLA samples given out corresponding to a secret $\mathbf{s}$. Intuitively, the fewer the number of samples, the harder the assumption. Indeed, reducing the number of samples beyond a certain point will make the assumption statistically close to random, and render it useless for cryptographic constructions. One of our innovations is the careful choice of the secret distribution that simultaneously allows us to use it for constructions while remaining plausibly hard even if both LWE and LPN are broken. Moreover, if the sparsity of the LPN error $\mathbf{e}_2$ is $n^{-\delta}$ and that of $\mathbf{s}$ is $n^{-\gamma}$, the only restriction on these parameters is that we need $\gamma + \delta \geq 1$ to have correctness of our cryptographic primitives.

As a part of the cryptanalysis, we consider and rule out several families of attacks such as combinatorial attacks that exploit sparsity like information set decoding, dimension reduction attacks, attacks for solving approximate worst-case lattice problems, and dual lattice attacks. A more detailed overview of our evidence for the plausible hardness of sqNULA is presented in Section 1.2. We now describe how to construct the PKE from our assumption.

1.1 PKE from sqNULA

We first give a brief background on the most common approach to construct PKE from NLAs and point out the difficulty to use it directly in our construction.

Starting point: Using the generic approach. Let us recall the usual recipe to construct a PKE from NLAs that has been used in numerous constructions including those from LWE due to Regev [Reg09], from LPN due to Alekhnovich [Ale03] and from LW2E and LWSSE [GJL⁺25]. These constructions crucially rely on the special distribution that the error vectors are sampled from to allow decryption. In general, the construction idea is typically as follows:

- **Key Generation:** Sample a random $\mathbf{A} \in \mathbb{F}_q^{m \times n}$, a random $\mathbf{s} \in \mathbb{F}_q^n$ and an error vector $\mathbf{e} \in \mathbb{F}_q^m$ as per the distribution defined by a noisy linear assumption. Set $(\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e})$ as the public key pk and set $\mathbf{s}, \mathbf{e}$ as the secret key sk.
- **Encryption:** To encrypt a 1, sample a uniform random $\mathbf{u}_1 \in \mathbb{F}_q^n$ and $u_2 \in \mathbb{F}_q$ and output $(\mathbf{u}_1, u_2)$. To encrypt a 0, sample a random vector $\mathbf{r} \in \mathbb{F}_q^m$ from either *the error distribution of a noisy linear assumption* as done in [Reg09, Ale03] or from *the error distribution of another appropriately chosen noisy linear assumption* as done in [GJL⁺25], and output $(\text{ct}_1^\top = \mathbf{r}^\top \mathbf{A}, \text{ct}_2 = \mathbf{r}^\top \mathbf{b})$.
- **Decryption:** Compute $\text{ct}_2 - \text{ct}_1^\top \mathbf{s}$. If the result is below some pre-determined threshold, then output 0; otherwise output 1.

Correctness of this construction intuitively works because $\text{ct}_2 - \text{ct}_1^\top \mathbf{s}$ is going to be uniform random over $\mathbb{F}_q$ if 1 was encrypted. On the other hand, if 0 was encrypted, then $\text{ct}_2 - \text{ct}_1^\top \mathbf{s} = \mathbf{r}^\top \mathbf{e}$. Now,

- For LWE, both $\mathbf{r}$ and $\mathbf{e}$ have small entries, and therefore their inner product is also small if we pick a prime modulus that is large enough.
- For LPN both $\mathbf{r}$ and $\mathbf{e}$ are sparse vectors, and their inner product will be 0 with a high probability if the sparsity parameters are picked appropriately.
- For the construction from LW2E and LWSSE, $\mathbf{e} = \mathbf{e}_1 + \mathbf{e}_2$ where $\mathbf{e}_1$ is small and $\mathbf{e}_2$ is sparse, and $\mathbf{r}$ is both sparse and is small at the non-zero positions. So, the inner product is $\mathbf{r}^\top \mathbf{e}_1 + \mathbf{r}^\top \mathbf{e}_2$, where the first term remains small as in the LWE case above, and the second term is likely zero as in the LPN case above.

The security proof typically has two phases: First, one replaces $\mathbf{b}$ in the public key with a uniform random $\mathbf{u} \in \mathbb{F}_q^m$. At this point, the view of the adversary is of the form $(\mathbf{A}, \mathbf{u}, \mathbf{r}^\top \mathbf{A}, \mathbf{r}^\top \mathbf{u})$. At a high level, we then need to argue that we can replace $(\mathbf{A} \parallel \mathbf{u}, \mathbf{r}^\top (\mathbf{A} \parallel \mathbf{u}))$ with $(\mathbf{A} \parallel \mathbf{u}, \mathbf{y}^\top)$, where $\mathbf{y}$ is a uniform random vector over $\mathbb{F}_q^n$. For LWE this boils down to invoking the leftover hash lemma; for LPN this involves using a dual form of the LPN assumption; and in the final case, this involves invoking the LWSSE assumption.

How would this recipe work with sqNULA? For a moment, let us ignore the error distribution for our assumption and focus on the dimension and distribution of $\mathbf{A}$. If we extend the above blueprint with sqNULA, then observe that $m = n$, i.e., $\mathbf{A} \in \mathbb{F}_q^{n \times n}$ and $\mathbf{A}$ is invertible with high probability. This means that upon seeing $\mathbf{A}$ and $\text{ct}_1^\top = \mathbf{r}^\top \mathbf{A}$, it is trivial to recover $\mathbf{r}^\top$ by computing $\text{ct}_1^\top \mathbf{A}^{-1}$. Therefore, there is simply no hope of getting any security. Let's go back to the drawing board.

Making use of $\mathbf{A}$'s squareness and using symmetry. Since we are anyway using NLAs, a natural way to make $(\mathbf{A}, \text{ct}_1^\top = \mathbf{r}^\top \mathbf{A})$ immune to the inversion attack mentioned above is to add some noise to ct_1 , i.e., sample $\mathbf{e}'^\top$ from *some distribution* and output $\text{ct}_1' = \text{ct}_1 + \mathbf{e}'^\top$. The hope is that if we pick $\mathbf{e}'^\top$ from a “good” distribution then we can appeal to the security of an NLA to argue pseudorandomness of $(\mathbf{A}, \mathbf{r}^\top \mathbf{A} + \mathbf{e}'^\top)$. Implementing that hope, we obtain the following construction, which does work:

- **Key Generation:** Sample a random invertible $\mathbf{A} \in \mathbb{F}_q^{n \times n}$, a random short and sparse $\mathbf{s} \in \mathbb{F}_q^n$ and an error vector is $\mathbf{e}_1 + \mathbf{e}_2 \in \mathbb{F}_q^n$ where $\mathbf{e}_1$ is short and $\mathbf{e}_2$ is sparse. Set $(\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}_1 + \mathbf{e}_2)$ as the public key pk and set $\mathbf{s}$ as the secret key sk.

- **Encryption:** To encrypt a 1, sample a uniform random $\mathbf{u}_1 \in \mathbb{F}_q^n$ and $u_2 \in \mathbb{F}_q$ and output $(\mathbf{u}_1, u_2)$. To encrypt a 0, sample a random short and sparse vector $\mathbf{s}' \in \mathbb{F}_q^n$ and an error vector $\mathbf{e}'_1 + \mathbf{e}'_2 \in \mathbb{F}_q^n$ where $\mathbf{e}'_1$ is short and $\mathbf{e}'_2$ is sparse, and output $(\text{ct}_1^\top = \mathbf{s}'^\top \mathbf{A} + \mathbf{e}'_1{}^\top + \mathbf{e}'_2{}^\top, \text{ct}_2 = \mathbf{s}'^\top \mathbf{b})$.
- **Decryption:** Compute $\text{ct}_2 - \text{ct}_1^\top \mathbf{s}$. If the result is below some pre-determined threshold, then output 0; otherwise output 1.

Does decryption work? When the ciphertext encrypts a 0, the decryption process outputs

$$\begin{aligned} (\text{ct}_2 - \text{ct}_1^\top \mathbf{s}) \bmod q &= (\mathbf{s}'^\top (\mathbf{A}\mathbf{s} + \mathbf{e}_1 + \mathbf{e}_2) - (\mathbf{s}'^\top \mathbf{A} + \mathbf{e}'_1{}^\top + \mathbf{e}'_2{}^\top) \mathbf{s}) \bmod q \\ &= (\mathbf{e}'_1{}^\top \mathbf{s}' + \mathbf{e}'_2{}^\top \mathbf{s}' - \mathbf{s}^\top \mathbf{e}'_1 - \mathbf{s}^\top \mathbf{e}'_2) \bmod q. \end{aligned}$$

We first consider $\mathbf{e}'_2{}^\top \mathbf{s}'$ and $\mathbf{s}^\top \mathbf{e}'_2$. Say that the density parameter of the LPN error components $\mathbf{e}_2, \mathbf{e}'_2$ is δ and the density parameter of the sqNULA secret component $\mathbf{s}, \mathbf{s}'$ is γ . As first argued by Alekhnovich [Ale03], if $\gamma + \delta > 1$, then $\mathbf{e}'_2{}^\top \mathbf{s}' = \mathbf{s}^\top \mathbf{e}'_2 = 0$ with high probability. Conditioning on the event that this happens, we will have

$$(\mathbf{e}'_1{}^\top \mathbf{s}' + \mathbf{e}'_2{}^\top \mathbf{s}' - \mathbf{s}^\top \mathbf{e}'_1 - \mathbf{s}^\top \mathbf{e}'_2) \bmod q = (\mathbf{e}'_1{}^\top \mathbf{s}' - \mathbf{s}^\top \mathbf{e}'_1) \bmod q$$

which will be a small value by inspection as long as q is large enough. Hence, decryption succeeds!

Arguing security via the hardcore lemma [GL89, GRS00]. To argue security first observe that we can use the hardness of decisional sqNULA hardness assumption to claim that $(\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}_1 + \mathbf{e}_2, \mathbf{s}'^\top \mathbf{A} + \mathbf{e}'_1{}^\top + \mathbf{e}'_2{}^\top, \mathbf{s}'^\top \mathbf{b})$ is computationally indistinguishable from $(\mathbf{A}, \mathbf{u}, \mathbf{s}'^\top \mathbf{A} + \mathbf{e}'_1{}^\top + \mathbf{e}'_2{}^\top, \mathbf{s}'^\top \mathbf{y})$ where $\mathbf{u}$ is picked at random from $\mathbb{F}_q^n$.

At this stage, observe that the decisional sqNULA also implies the hardness of search sqNULA that says given $(\mathbf{A}, \mathbf{s}'^\top \mathbf{A} + \mathbf{e}'_1{}^\top + \mathbf{e}'_2{}^\top)$, it is hard to recover $\mathbf{s}'$. This is essentially a one-way function, and therefore $\mathbf{s}'^\top \mathbf{u}$ is an hardcore “bit” with respect to the random predicate $\mathbf{u}$, and therefore indistinguishable from random. Thus, using the classic hardcore lemma [GL89, GRS00], we get $(\mathbf{A}, \mathbf{u}, \mathbf{s}'^\top \mathbf{A} + \mathbf{e}'_1{}^\top + \mathbf{e}'_2{}^\top, \mathbf{s}'^\top \mathbf{u})$ is computationally indistinguishable from $(\mathbf{A}, \mathbf{u}, \mathbf{s}'^\top \mathbf{A} + \mathbf{e}'_1{}^\top + \mathbf{e}'_2{}^\top, r)$ where r is a random $\mathbb{F}_q$ element. Finally, we can again use the decisional sqNULA assumption to replace $(\mathbf{A}, \mathbf{u}, \mathbf{s}'^\top \mathbf{A} + \mathbf{e}'_1{}^\top + \mathbf{e}'_2{}^\top, r)$ with $(\mathbf{A}, \mathbf{u}, \mathbf{u}'^\top, r)$ for a random $\mathbf{u}' \in \mathbb{F}_q^n$, completing the security argument. We refer the reader to Section 4 for formal details.

Other constructions. Along with a PKE, we also construct a two-round maliciously secure oblivious transfer protocol. The core idea behind both constructions is similar to our PKE construction. For OT, we show how to construct a two-round protocol with elementary sender security (See Section 2 for formal definition). This is a weak form of security that can then be bootstrapped in a black-box manner to full fledged secure (even UC secure) malicious OT using results from prior work [DGH⁺20]. The novel contribution in this construction is the use of an encoding mechanism with sufficient redundancies to ensure accurate message recovery by the receiver. We refer to Section 5 for formal details.

As an additional contribution, we construct a one-time circular-secure secret-key encryption from a variant of sqNULA where the secret is over $\{0, 1\}^n$. We show that this variant is also at least as hard as standard LWE.

1.2 Hardness of our assumption

We also perform an initial systematic study of the hardness of sqNULA and provide evidence that under a very wide range of parameter settings, it is potentially hard even in the presence of an oracle that can break worst-case approximate¹ lattice problems, and an oracle that can break LPN with any reasonable parameters. Here is the formal description of the assumption:

¹For any polynomial approximation factor n^ϵ , for $\epsilon > 0$, not counting potential NP-hardness reductions for approximation factors that are NP-hard.

Assumption 1.1 (The Square Unstructured Noisy Linear Algebraic Assumption (sqNULA)). *For all $n \in \mathbb{N}$, prime $q \in \mathbb{N}$, $\sigma = \text{poly}(n)$, $\xi = \text{poly}(n)$, $\delta, \gamma \in (0, 1)$, the decisional square noisy unstructured linear algebraic assumption (decisional-sqNULA), formally parameterized by $\text{sqNULA}_{n,q,\sigma,\xi,\delta,\gamma}$, states that the following distributions are computationally indistinguishable:*

1. $\left(\mathbf{A} \xleftarrow{\$} \text{GL}_q(n), \mathbf{A}\mathbf{s} + \mathbf{e}^{(1)} + \mathbf{e}^{(2)} \pmod{q} \right)$ where $\mathbf{s} \xleftarrow{\$} \mathcal{E}_{n,\xi,\gamma}^n$, $\mathbf{e}^{(1)} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}^n$, and $\mathbf{e}^{(2)} \leftarrow \mathcal{S}_{n,q,\delta}^n$.
2. $\left(\mathbf{A} \xleftarrow{\$} \text{GL}_q(n), \mathbf{b} \xleftarrow{\$} \mathbb{F}_q^n \right)$.

Here,

- $\text{GL}_q(n)$ denotes the group of $n \times n$ invertible matrices over $\mathbb{F}_q$.
- $\mathcal{D}_{\mathbb{Z},\sigma}$: For $\sigma \in \mathbb{R}^+$, we let $\mathcal{D}_{\mathbb{Z},\sigma}$ denote the discrete Gaussian distribution over the integer lattice $\mathbb{Z}$ with mean 0 and scale parameter $\sigma > 0$.
- $\mathcal{S}_{n,q,\delta}$: For $n, q \in \mathbb{N}$, $\delta \in (0, 1)$, we define $\mathcal{S}_{n,q,\delta}$ to be the distribution samples 0 with probability $1 - n^{-\delta}$ or a uniformly random element from $\mathbb{F}_q$ with probability $n^{-\delta}$.
- $\mathcal{E}_{n,\xi,\gamma}$: For $n, \xi \in \mathbb{N}$, $\gamma \in (0, 1)$, we define $\mathcal{E}_{n,\xi,\gamma}$ to be the distribution over $\mathbb{Z}$ which samples 0 with probability $1 - n^{-\gamma}$ or a random element from the distribution $\mathcal{D}_{\mathbb{Z},\xi}$ with probability $n^{-\gamma}$.

Our parameter settings for the constructions. We can use a very wide range of parameters. The following is simply one suggestion.

- Secret dimension of sqNULA: $n = \lambda$.
- Sparsity parameter for sqNULA noise $\mathbf{e}_2$: $\frac{1}{2} < \delta < 1$.
- sqNULA secret sparsity parameter: $\gamma < \frac{1}{2} : \gamma + \delta > 1$.
- Smallness parameter for sqNULA noise $\mathbf{e}_1$: $\sigma = n$.
- Prime modulus $q \geq n^4$.
- sqNULA secret smallness parameter: $\xi = n$.

The upper bound on γ is a very conservative choice, and we believe that our assumption remains secure in the presence of LPN and lattice problem breaking oracle for larger values of γ as well. Effectively, the only constraints we need are that $\gamma + \delta > 1$ and q is reasonably large.

Analyzing the hardness of the assumption. The sqNULA assumption has a natural dual equivalent, where we can swap the error and secret distribution. This is evident from the observation that given $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}_1 + \mathbf{e}_2)$, one can use $\mathbf{A}^{-1}$ and produce $(\mathbf{A}^{-1}, \mathbf{s} + \mathbf{A}^{-1}(\mathbf{e}_1 + \mathbf{e}_2))$. We call the latter d-sqNULA.

On Reductions to other problems. We can provably show that sqNULA, as well as its binary secret variant are at least as hard as LWE with appropriate parameters. We refer readers to Section 7.3 for details. We also show in Section 7.4 that sqNULA with binary secrets reduces to sqNULA with small secrets via a natural randomization technique. In Section 7.5, we show a reduction from sqNULA to a variant of LW2E with number of samples equal to the secret dimension, and with short secrets. At a high level, we flood the short and sparse secret of sqNULA with the short secret of LW2E. To prove that the resultant distribution of the secret is identical to that of LW2E, we use the Gaussian Smudging Lemma 7.11. In order for this lemma to apply, we require sqNULA to have a sub-exponential “modulus to secret” ratio, i.e., $\frac{q}{\xi} = n^{\omega(1)}$. We point out that this reduction does not apply to the parameter regimes that are used for our applications, where we use $q = \text{poly}(n)$.

Another natural question to explore is a search-to-decision (solve the search variant using a solver for the decision variant) reduction for sqNULA. A decision-to-search reduction is trivial. We prove the reverse direction in Section 7.2 via a predictor for the hardcore predicate over large fields. In particular, we convert an adversary for the decision problem to a predictor for the hardcore predicate with noticeable advantage. Then, we invoke the hardcore lemma from [GRS00] to convert this predictor to an inverter for the one-way function.

In Section 8 we present a detailed initial evaluation of the hardness of sqNULA even in the presence of various oracles, by trying natural approaches to formulate reductions to other standard NLAs and understanding why they fail to reveal attacks.

Discussion on the approaches.

- The fact the our matrix $\mathbf{A}$ is square seems to take it out of the realm of any lattice- or code-based problems. In particular, the only reason why our assumption is a computational assumption is because of the short and sparse nature of the secret. When $\mathbf{A}$ is square and the secret $\mathbf{s}$ is uniformly random, there is too much entropy in the secret which makes *all unstructured* NLAs like LWE, LPN, LW2E and LWSSE statistically indistinguishable from random, so a reduction to them is meaningless. Thus, any reduction must somehow transform our square $\mathbf{A}$ into a rectangular expanding $\mathbf{A}'$.

In Section 7.6, we show one approach to interpret sqNULA equivalently as a non-square problem of the form $\mathbf{A}\mathbf{s} + \mathbf{e}$, where $\mathbf{A} \in \mathbb{F}_q^{2n \times n}$, $\mathbf{s} \in \mathbb{F}_q^n$ are uniform random, and $\mathbf{e} = (\mathbf{s}' \mid \mathbf{e}^{(1)} + \mathbf{e}^{(2)})$ where $\mathbf{s}'$ is from the sqNULA secret distribution and $\mathbf{e}^{(1)} + \mathbf{e}^{(2)}$ are appropriate sqNULA errors. While this is a non-square regime, the “small plus sparse” component of the error take it out of scope for any LPN or LWE based attack as explored in [GJL⁺25]. With this formulation, intuitively, sqNULA appears to be LW2E with lower entropy in the error vector. We leave it open to come up with be formal reduction from our assumption to LW2E.

- Another such idea is to start with d-sqNULA. Then, take sparse (say, log-sparse, for instance) linear combinations of rows of $\mathbf{A}$ and the corresponding entries of $\mathbf{b}$ to generate fresh samples. The sparse linear combination will ensure that non-zero error entries are unlikely to be picked up, i.e., for the i^{th} new sample, we pick a log-sparse $\mathbf{v}_i \in \mathbb{F}_q^n$, then $\mathbf{v}_i^\top \mathbf{b} = \mathbf{v}_i^\top (\mathbf{A}(\mathbf{e}_1 + \mathbf{e}_2) + \mathbf{s}) = \mathbf{a}_i'^\top (\mathbf{e}_1 + \mathbf{e}_2) + \mathbf{v}_i^\top \mathbf{s}$. With high probability, $\mathbf{v}_i^\top \mathbf{s} = 0$. Therefore, $b'_i = \mathbf{a}_i'^\top (\mathbf{e}_1 + \mathbf{e}_2)$. This way one can generate an arbitrary number of samples and try to use an LPN solving oracle to break it. (The error distribution is not quite LPN, but it is still a sparse error, so one can be optimistic.) The problem here is that for *every* $\mathbf{b}$ (random or not), there exists some $\mathbf{x}$ such that $\mathbf{A}\mathbf{x} = \mathbf{b}$ and this is a valid noiseless LPN pre-image, and will continue being a pre-image even if we generate new samples. The LPN solving oracle can always return such a preimage because it is oblivious to the secret distribution. So, this approach also fails.
- What about lattice attacks? What we can show is that, when the prime modulus q is extremely (pathologically) small, in particular $q = O(n^\epsilon)$, for $\epsilon \in (0, 1)$, there is a plausibly heuristically successful way to use an approximate SVP oracle to distinguish a sqNULA sample from random. However, this parameter setting is far beyond what we recommend using for the constructions. For larger values of q , all our attempts to embed an sqNULA instance into a lattice to mount an attack failed.

Furthermore, even in the case where $q = o(n)$, we will need an approximate SVP oracle with approximation factor $O(n^\epsilon)$, for $0 < \epsilon < \frac{1}{2}$ to attack sqNULA, which is a very strong oracle indeed. We conjecture that it is unlikely that an SVP oracle with a larger approximation factor will be useful even for these parameters. For comparison, LWE that implies PKE can be broken with an approximation factor $O(\sqrt{n})$ SVP oracle [BV14, SD16]. Please refer to Section 8.1 for further details and discussion.

- Can we use combinatorial attacks like Information Set Decoding (ISD) that exploits the sparse structure of the error or secret? In Section 8.6, we explore attack paradigms that involve guessing the non-zero locations of error $\mathbf{e}_2$ or secret $\mathbf{s}$. We show that every attack that we explored requires $2^{\text{poly}(n)}$ time. The intuition for this is that guessing the non-zero locations already takes $\binom{n}{n^\epsilon} \approx 2^{n^\epsilon \log n}$, for some constant $\epsilon \in (0, 1)$.

Discussion on lack of separation results: Note that if we were able to actually prove that no efficient algorithm given an approximate SVP oracle (or LPN breaking oracle) can break sqNULA, this would imply that no efficient algorithm that does not use approximate SVP oracle (or LPN breaking oracle) can break sqNULA, implying $P \neq NP$. Also, note that an oracle separation is meaningless, because neither LWE or LPN nor sqNULA are defined with respect to oracles. What we are able to argue is that natural approaches to use oracles that can solve LWE or LPN do not break sqNULA and we find that natural cryptanalysis approaches do not work unless the sparse error components are extremely sparse (constant sparsity) and the field size q is also limited with respect to the dimension n .

A perspective on the distinction between sqNULA and other NLA assumptions. At first glance, sqNULA might appear very similar to other NLAs like LWE, LPN, LW2E or LWSSE. Indeed, they are all structurally similar, with a different error distribution. In fact, this “qualitative similarity” is strongly desirable because it allows one to leverage the extensive cryptanalysis of both LWE and LPN for the new assumptions.

On the other hand, the question of what should count as qualitative similarity” is very deep and subtle. For instance, LPN and LWE have essentially the same structural template, differing only in the error distribution—sparse random error for LPN and small random error for LWE. Should we believe that these are qualitatively similar”? They also share some cryptanalytic tools, e.g., the BKW algorithm. Yet, they are believed to be qualitatively very distinct in hardness, e.g., LWE with parameters that yield PKE is broken given a $\sqrt{n}$ -GapCVP oracle, while LPN is not known to be. Indeed, for implications to cryptography, these assumptions are totally different. For instance, LWE implies FHE, ABE, NIZKs, and a myriad of other cryptographic primitives. On the other hand, from LPN with inverse-polynomial noise, we struggle to build any interesting advanced functionality beyond PKE. Complexity-theoretically, LWE and its worst-case generalizations sit in a relatively structured complexity class ($NP \cap \text{coNP}$). On the other hand, LPN does not have known natural proofs of unsatisfiability.

Furthermore, the multivariate quadratic (MQ) and minrank problems can also be interpreted as NLAs with tensor error and low rank error respectively. We have no evidence to believe that these two assumptions are similar to those discussed above. Essentially, it is the choice of the error distribution that creates the distinction between these assumptions. LWSSE is another addition to this family that is likely harder than both LWE and LPN.

2 Preliminaries

General Linear Group: $\text{GL}_q(n)$ denotes the group of $n \times n$ invertible matrices over $\mathbb{F}_q$.

Lemma 2.1 (A result about the distribution of a quadratic form). *Let $\mathbb{F}_q$ be the finite field of size q , and let $n \geq 1$. Fix nonzero vectors $\mathbf{x}, \mathbf{v} \in \mathbb{F}_q^n$. If $\mathbf{A}$ is drawn uniformly at random from $\text{GL}_q(n)$, then*

$$\Pr[\mathbf{Ax} = \mathbf{v}] = \frac{1}{q^n - 1} = \Theta(q^{-n}).$$

In particular, conditional on $\mathbf{x} \neq 0$, the random vector $\mathbf{Ax}$ is uniform over $\mathbb{F}_q^n \setminus \{0\}$.

Proof. Choose $\mathbf{R} \in \text{GL}_q(n)$ with $\mathbf{Rx} = \mathbf{e}_1$ (extend $\mathbf{x}$ to a basis and send it to the standard basis). The map $\mathbf{A} \mapsto \mathbf{AR}^{-1}$ is a bijection of $\text{GL}_q(n)$, so it preserves uniformity. Moreover,

$$\mathbf{Ax} = \mathbf{v} \iff (\mathbf{AR}^{-1})\mathbf{e}_1 = \mathbf{v}.$$

Writing $\mathbf{B} := \mathbf{A}\mathbf{R}^{-1}$, the event $\{\mathbf{A}\mathbf{x} = \mathbf{v}\}$ is equivalent to the event that “the first column of $\mathbf{B}$ equals $\mathbf{v}$,” where $\mathbf{B}$ is still distributed uniformly on $\text{GL}_q(n)$.

Now let

$$S_{\mathbf{v}} := \{ \mathbf{B} \in \text{GL}_q(n) : \text{the first column of } \mathbf{B} \text{ is } \mathbf{v} \}.$$

A straightforward counting argument gives

$$\begin{aligned} |S_{\mathbf{v}}| &= (q^n - q)(q^n - q^2) \cdots (q^n - q^{n-1}). \\ |\text{GL}_q(n)| &= (q^n - 1)(q^n - q)(q^n - q^2) \cdots (q^n - q^{n-1}). \end{aligned}$$

Therefore,

$$\Pr [\text{first column of } \mathbf{B} \text{ equals } \mathbf{v}] = \frac{|S_{\mathbf{v}}|}{|\text{GL}_q(n)|} = \frac{(q^n - q) \cdots (q^n - q^{n-1})}{(q^n - 1)(q^n - q) \cdots (q^n - q^{n-1})} = \frac{1}{q^n - 1}.$$

Since this probability does not depend on $\mathbf{v} \neq \mathbf{0}$, it follows that $\mathbf{B}\mathbf{e}_1$ (hence $\mathbf{A}\mathbf{x}$) is uniform over $\mathbb{F}_q^n \setminus \{\mathbf{0}\}$, as claimed. $\square$

2.1 Lattices

Notation. We define some notation for distributions:

- $\mathcal{D}_{\mathbb{Z},\sigma}$: For $\sigma \in \mathbb{R}^+$, we let $\mathcal{D}_{\mathbb{Z},\sigma}$ denote the discrete Gaussian distribution over the integer lattice $\mathbb{Z}$ with mean 0 and scale parameter $\sigma > 0$.
- $\mathcal{S}_{n,q,\delta}$: For $n, q \in \mathbb{N}, \delta \in (0, 1)$, we define $\mathcal{S}_{n,q,\delta}$ to be the distribution samples 0 with probability $1 - n^{-\delta}$ or a uniformly random element from $\mathbb{F}_q$ with probability $n^{-\delta}$.
- $\mathcal{E}_{n,\xi,\gamma}$: For $n, \xi \in \mathbb{N}, \gamma \in (0, 1)$, we define $\mathcal{E}_{n,\xi,\gamma}$ to be the distribution over $\mathbb{Z}$ which samples 0 with probability $1 - n^{-\gamma}$ or a random element from the distribution $\mathcal{D}_{\mathbb{Z},\xi}$ with probability $n^{-\gamma}$.

Lemma 2.2 ([Lyu12]. Lemma 4.4). *For any $k > 1$, $\Pr [\|\mathbf{z}\|_2 > k\sigma \cdot \sqrt{m} : \mathbf{z} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}^m] \leq k^m e^{m(1-k^2)/2}$.*

2.2 Hardness Assumptions

Definition 2.3 (Decisional Learning With Errors Assumption [Reg09]). *The decisional learning with errors assumption (decisional-LWE), formally parameterized as $\text{LWE}_{n,m,q,\mathcal{D}_{\mathbb{Z},\sigma}}$, states that the following distributions are computationally indistinguishable:*

1. $(\mathbf{A} \leftarrow \mathbb{F}_q^{m \times n}, \mathbf{A}\mathbf{s} + \mathbf{e} \pmod{q})$ where $\mathbf{s} \leftarrow \mathbb{F}_q^n$ and $\mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}^m$.
2. $(\mathbf{A} \leftarrow \mathbb{F}_q^{m \times n}, \mathbf{b} \leftarrow \mathbb{F}_q^m)$.

Definition 2.4 (Decisional Learning Parity with Noise Assumption over Large Fields [Ale03, JLS21]). *The decisional learning parity with noise assumption over large fields (decisional-LPN), formally parameterized as $\text{LPN}_{n,m,q,\delta}$, states that the following distributions are computationally indistinguishable:*

1. $(\mathbf{A} \leftarrow \mathbb{F}_q^{m \times n}, \mathbf{A}\mathbf{s} + \mathbf{e} \pmod{q})$ where $\mathbf{s} \leftarrow \mathbb{F}_q^n$ and $\mathbf{e} \leftarrow \mathcal{S}_{n,q,\delta}^m$.
2. $(\mathbf{A} \leftarrow \mathbb{F}_q^{m \times n}, \mathbf{b} \leftarrow \mathbb{F}_q^m)$.

Definition 2.5 (Learning With Two Errors Assumption [GJL⁺25]). *For all $n \in \mathbb{N}, m = \text{poly}(n)$, prime $q \in \mathbb{N}, \sigma = \text{poly}(n), \delta \in (0, 1)$, the decisional learning with two errors assumption (decisional-LW2E), formally parameterized as $\text{LW2E}_{n,m,q,\mathcal{D}_{\mathbb{Z},\sigma},\delta}$, states that the following distributions are computationally indistinguishable:*

1. $\left(\mathbf{A} \xleftarrow{\$} \mathbb{F}_q^{m \times n}, \mathbf{A}\mathbf{s} + \mathbf{e}^{(1)} + \mathbf{e}^{(2)} \pmod{q} \right)$ where $\mathbf{s} \xleftarrow{\$} \mathbb{F}_q^n$, $\mathbf{e}^{(1)} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^m$, and $\mathbf{e}^{(2)} \leftarrow \mathcal{S}_{n, q, \delta}^m$.
2. $\left(\mathbf{A} \xleftarrow{\$} \mathbb{F}_q^{m \times n}, \mathbf{b} \xleftarrow{\$} \mathbb{F}_q^m \right)$.

Similarly, the search-LW2E assumption, formally parameterized by $\text{Search-LW2E}_{n, m, q, \mathcal{D}_{\mathbb{Z}, \sigma}, \delta}$, states that every polynomial time adversary, given Item 1 as input, outputs $\mathbf{s}$ with at most negligible probability.

Definition 2.6 (Decisional Learning With Short and Sparse Errors Assumption [GJL⁺25]). *For all $n \in \mathbb{N}$, $m = \text{poly}(n) > n$, prime $q \in \mathbb{N}$, $\xi = \text{poly}(n)$, $\gamma \in (0, 1)$, the decisional learning with short and sparse errors assumption (decisional-LWSSE), formally parameterized by $\text{LWSSE}_{n, m, q, \xi, \gamma}$, states that the following distributions are computationally indistinguishable:*

1. $\left(\mathbf{A}^\perp \xleftarrow{\$} \mathbb{F}_q^{(m-n) \times m}, \mathbf{s}^\top \mathbf{A}^\perp + \mathbf{e}^\top \pmod{q} \right)$ where $\mathbf{s} \xleftarrow{\$} \mathbb{F}_q^{m-n}$, $\mathbf{e} \leftarrow \mathcal{E}_{m, n, \xi, \gamma}^m$.
2. $\left(\mathbf{A}^\perp \xleftarrow{\$} \mathbb{F}_q^{(m-n) \times m}, \mathbf{b}^\top \xleftarrow{\$} \mathbb{F}_q^m \right)$.

Similarly, the search-LWSSE assumption, formally parameterized by $\text{Search-LWSSE}_{n, m, q, \xi, \gamma}$, states that every polynomial time adversary, given Item 1 as input, outputs $\mathbf{s}^\top$ with at most negligible probability.

2.3 Two Round Oblivious Transfer

We borrow the following definitions verbatim from [DGH⁺20]. A two-round oblivious transfer (OT) protocol is given by algorithms (Setup, Receiver₁, Sender, Receiver₂), where the setup algorithm Setup generates a CRS value $\text{crs} \xleftarrow{\$} \text{Setup}(1^\lambda)$. The receiver runs the algorithm Receiver₁ which takes crs and a choice bit $c \in \{0, 1\}$ as input and outputs (otr, st) . The receiver then sends otr to the sender, who obtains ots by evaluating $\text{Sender}(1^\lambda, \text{otr}, m_0, m_1)$, where $m_0 \in \{0, 1\}^\lambda$ and $m_1 \in \{0, 1\}^\lambda$ are its inputs. The sender then sends ots to the receiver who obtains m_c by evaluating $\text{Receiver}_2(1^\lambda, \text{st}, \text{ots})$.

Correctness We say that a two-round OT scheme is δ -correct, if with probability $1 - \text{negl}(\lambda)$ over the choice of $\text{crs} \xleftarrow{\$} \text{Setup}(1^\lambda)$ the following holds: for every choice bit $c \in \{0, 1\}$ of the receiver and input messages m_0 and m_1 of the sender, and for any $(\text{otr}, \text{st}) \in \text{Receiver}_1(\text{crs}, c)$ and $\text{ots} \in \text{Sender}(\text{crs}, \text{otr}, m_0, m_1)$, we have

$$\Pr[\text{Receiver}_2(\text{st}, \text{ots}) = m_c] \geq 1 - \delta.$$

Receiver's indistinguishability security. For every non-uniform polynomial-time adversary $\mathcal{A}$:

$$|\Pr[\mathcal{A}(\text{crs}, \text{Receiver}_1(\text{crs}, 0)) = 1] - \Pr[\mathcal{A}(\text{crs}, \text{Receiver}_1(\text{crs}, 1)) = 1]| = \text{negl}(\lambda),$$

where $\text{crs} \xleftarrow{\$} \text{Setup}(1^\lambda)$.

2.4 Sender's Security Notions

We consider a weak but sufficiently interesting notion of sender's security, in which we slightly modify the syntax of Sender as was done in [DGH⁺20]. Namely, in the following experiments, the algorithm Sender* does not take as input two messages and instead outputs two masks y_0, y_1 , only one of which is given to the receiver.

Sender's Elementary Security. The elementary sender security corresponds to a weak security notion. This requires that the receiver compute both the strings y_0 and y_1 used by the sender. Let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be an adversary. Consider the following experiment $\text{Exp}_{\text{eOT}}^\lambda(\mathcal{A})$:

1. Run $\text{crs} \xleftarrow{\$} \text{Setup}(1^\lambda)$.
2. Run $(\text{otr}, \text{st}) \xleftarrow{\$} \mathcal{A}_1(1^\lambda, \text{crs})$.
3. Compute $(\text{ots}, y_0, y_1) \xleftarrow{\$} \text{Sender}^*(\text{crs}, \text{otr})$.
4. Compute $(y_0^*, y_1^*) \xleftarrow{\$} \mathcal{A}_2(\text{st}, \text{ots})$ and output 1 iff $(y_0^*, y_1^*) = (y_0, y_1)$.

We say that a scheme satisfies eOT security if

$$\Pr[\text{Exp}_{\text{eOT}}^\lambda(\mathcal{A}) = 1] = \text{negl}(\lambda).$$

2.5 UC Secure OT

We copy verbatim the definitions of receiver's UC-security and sender's UC-security from [DGH⁺20]. Note that our paper considers security in the CRS model since two-round simulation-secure malicious OT is impossible in the plain model.

We work in Canetti's UC framework with static corruptions [Can01]. We assume familiarity with this model. We use $\mathcal{Z}$ for denoting the underlying environment. For a real protocol Π and an adversary $\mathcal{A}$, we use $\text{EXEC}_{\Pi, \mathcal{A}, \mathcal{Z}}$ to denote the real-world ensemble. Also, for an ideal functionality $\mathcal{F}$ and an adversary $\mathcal{S}$ we denote $\text{IDEAL}_{\mathcal{F}, \mathcal{S}, \mathcal{Z}}$ to denote the ideal-world ensemble.

Definition 2.7 (Receiver's UC-security). *We say that an OT protocol OT is receiver-UC secure if for any adversary $\mathcal{A}$ corrupting the sender, there exists a simulator $\mathcal{S}$ such that for all environments $\mathcal{Z}$:*

$$\text{IDEAL}_{\mathcal{F}_{\text{OT}}, \mathcal{S}, \mathcal{Z}} \equiv \text{EXEC}_{\text{OT}, \mathcal{A}, \mathcal{Z}},$$

where the ideal functionality $\mathcal{F}_{\text{OT}}$ is defined in Figure 1.

$\mathcal{F}_{\text{OT}}$ interacts with an ideal sender **S** and an ideal receiver **R**.

1. On input $(\text{sid}, \text{sender}, m_0, m_1)$ from the sender, store (m_0, m_1) .
2. On input $(\text{sid}, \text{receiver}, b)$, check if a pair of inputs (m_0, m_1) has been already recorded for session sid ; if so, send m_b to **R** and send sid to the adversary and halt; else, send nothing.

Figure 1: Ideal Functionality $\mathcal{F}_{\text{OT}}$

Since our OT protocols are in the CRS model, we also give the $\mathcal{F}_{\text{CRS}}$ ideal functionality below in Figure 2.

Definition 2.8 (Sender's UC-security). *We say that an OT protocol OT is sender-UC secure if for any adversary $\mathcal{A}$ corrupting the receiver, there exists a simulator $\mathcal{S}$ such that for all environments $\mathcal{Z}$:*

$$\text{IDEAL}_{\mathcal{F}_{\text{OT}}, \mathcal{S}, \mathcal{Z}} \equiv \text{EXEC}_{\text{OT}, \mathcal{A}, \mathcal{Z}},$$

where the ideal functionality $\mathcal{F}_{\text{OT}}$ is defined in Figure 1.

Lemma 2.9 (See Section 10.2.1 in [DGH⁺20]). *If there exists a two-round OT protocol with receiver indistinguishability security and elementary sender security and inverse polynomial correctness error, then there exists a two-round UC OT protocol.*

$\mathcal{F}_{\text{CRS}}^{\mathcal{D}}$: parameterized over a distribution $\mathcal{D}$, run by parties $P_1, \dots, P_n$, and an adversary $\mathcal{S}$:

- Whenever receiving a message (sid, P_i, P_j) from party P_i , sample $\text{crs} \leftarrow \mathcal{D}$ and send (sid, crs) to P_i and send $(sid, \text{crs}, P_i, P_j)$ to $\mathcal{S}$. Whenever receiving the message (sid, P_i, P_j) from P_j , send (sid, crs) to P_j and $\mathcal{S}$.

Figure 2: Ideal Functionality $\mathcal{F}_{\text{CRS}}^{\mathcal{D}}$ [CR03]

2.6 Circular Secure Secret Key Encryption

We borrow these definitions from [KMT19, KM20].

Definition 2.10 (Secret-Key Encryption). A secret-key encryption (SKE) scheme SKE consists of the three PPT algorithms (K, E, D) :

- K is the key generation algorithm that takes 1^λ as input, and outputs a secret key sk . We assume that the security parameter λ determines the secret key space $\mathcal{K}$ and the message space $\mathcal{M}$.
- E is the encryption algorithm that takes a secret key sk and a plaintext m as input, and outputs a ciphertext ct .
- D is the (deterministic) decryption algorithm that takes a secret key sk and a ciphertext ct as input, and outputs a plaintext m or the invalid symbol $\perp \notin \mathcal{M}$.

An SKE scheme $\text{SKE} = (K, E, D)$ is said to be correct if for all $\text{sk} \in \mathcal{K}$ and all $m \in \mathcal{M}$, it holds that

$$D(\text{sk}, E(\text{sk}, m)) = m.$$

Definition 2.11 (KDM Security). Let $\text{SKE} = (K, E, D)$ be an SKE scheme with a secret key space $\mathcal{K}$ and a plaintext space $\mathcal{M}$. For a family of functions $\mathcal{F}$ with domain $\mathcal{K}$ and range $\mathcal{M}$ and an adversary $\mathcal{A}$, consider the experiment $\text{Expt}_{\text{SKE}, \mathcal{F}, \mathcal{A}}^{\text{kdm}}(\lambda)$ defined as in Figure below (left), where the KDM-encryption oracle $\mathcal{O}_{\text{kdm}}$ is described in Figure below (right).

We say that SKE is $\mathcal{F}$ -KDM secure if for all PPT adversaries $\mathcal{A}$, we have

$$\text{Adv}_{\text{SKE}, \mathcal{F}, \mathcal{A}}^{\text{kdm}}(\lambda) := 2 \cdot \left| \Pr[\text{Expt}_{\text{SKE}, \mathcal{F}, \mathcal{A}}^{\text{kdm}}(\lambda) = 1] - \frac{1}{2} \right| = \text{negl}(\lambda).$$

Furthermore, we say that SKE is one-time $\mathcal{F}$ -KDM secure if we have $\text{Adv}_{\text{SKE}, \mathcal{F}, \mathcal{A}}^{\text{kdm}}(\lambda) = \text{negl}(\lambda)$ for all PPT adversaries $\mathcal{A}$ that make a single KDM-encryption query.

$\begin{aligned} \text{Expt}_{\text{SKE}, \mathcal{F}, \mathcal{A}}^{\text{kdm}}(\lambda) : \\ & \text{sk} \leftarrow K(1^\lambda) \\ & b \xleftarrow{r} \{0, 1\} \\ & b' \leftarrow \mathcal{A}^{\mathcal{O}_{\text{kdm}}(\cdot, b)}(1^\lambda) \\ & \text{return } (b' \stackrel{?}{=} b) \end{aligned}$	$\begin{aligned} \mathcal{O}_{\text{kdm}}(f \in \mathcal{F}, b) : \\ & \text{ct}_0 \leftarrow E(\text{sk}, f(\text{sk})) \\ & \text{ct}_1 \leftarrow E(\text{sk}, 0) \\ & \text{return } \text{ct}_b \end{aligned}$
--	---

Figure 3: The KDM security experiment for an SKE (left) scheme and the KDM-encryption oracle used in the KDM security experiment (right).

KDM Security for Projection Function. A function is said to be a projection function if each of its output bits depends on at most a single bit of its input.

Circular Security for Bit-SKE Circular security is a special case of KDM security for projection functions. Here, every function takes as input the secret key and an index i , and outputs the i^{th} bit of the secret key. (We point out that in the KDM literature the term “circular security” can also refer a different notion of just encrypting key cycles.)

3 The New Hardness Assumption

Assumption 3.1 (The Square-Noisy Unstructured Linear Algebraic Assumption (sqNULA)). *For all $n \in \mathbb{N}$, prime $q \in \mathbb{N}$, $\sigma = \text{poly}(n)$, $\xi = \text{poly}(n)$, $\delta, \gamma \in (0, 1)$, the decisional square-noisy unstructured linear algebraic assumption (decisional-sqNULA), formally parameterized by $\text{sqNULA}_{n,q,\sigma,\xi,\delta,\gamma}$, states that the following distributions are computationally indistinguishable:*

1. $\left(\mathbf{A} \xleftarrow{\$} \text{GL}_q(n), \mathbf{A}\mathbf{s} + \mathbf{e}^{(1)} + \mathbf{e}^{(2)} \pmod{q} \right)$ where $\mathbf{s} \xleftarrow{\$} \mathcal{E}_{n,\xi,\gamma}^n$, $\mathbf{e}^{(1)} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}^n$, and $\mathbf{e}^{(2)} \leftarrow \mathcal{S}_{n,q,\delta}^n$.
2. $\left(\mathbf{A} \xleftarrow{\$} \text{GL}_q(n), \mathbf{b} \xleftarrow{\$} \mathbb{F}_q^n \right)$.

Here,

- $\mathcal{D}_{\mathbb{Z},\sigma}$: For $\sigma \in \mathbb{R}^+$, we let $\mathcal{D}_{\mathbb{Z},\sigma}$ denote the discrete Gaussian distribution over the integer lattice $\mathbb{Z}$ with mean 0 and scale parameter $\sigma > 0$.
- $\mathcal{S}_{n,q,\delta}$: For $n, q \in \mathbb{N}$, $\delta \in (0, 1)$, we define $\mathcal{S}_{n,q,\delta}$ to be the distribution samples 0 with probability $1 - n^{-\delta}$ or a uniformly random element from $\mathbb{F}_q$ with probability $n^{-\delta}$.
- $\mathcal{E}_{n,\xi,\gamma}$: For $n, \xi \in \mathbb{N}$, $\gamma \in (0, 1)$, we define $\mathcal{E}_{n,\xi,\gamma}$ to be the distribution over $\mathbb{Z}$ which samples 0 with probability $1 - n^{-\gamma}$ or a random element from the distribution $\mathcal{D}_{\mathbb{Z},\xi}$ with probability $n^{-\gamma}$.

Similarly, the search-sqNULA assumption, formally parameterized by $\text{Search-sqNULA}_{n,q,\sigma,\xi,\delta,\gamma}$, states that every polynomial time adversary, given Item 1 as input, outputs $\mathbf{s}$ with at most negligible probability.

Remark 3.2. Observe that the error term $\mathbf{e}^{(1)}$ is sampled from an LWE error distribution, and the error term $\mathbf{e}^{(2)}$ is sampled from an LPN over large fields error distribution. Additionally, $\mathbf{s}$ is sampled from a distribution that has small entries in sparse locations, and is zero otherwise.

4 CPA secure Public Key Encryption

Suggested Parameters: We denote the security parameter as λ and suggest the following choices of parameters:

- Decryption threshold parameter: $\zeta \geq 4 \in \mathbb{N}$.
- Secret dimension of sqNULA: $n = \lambda$.
- Sparsity parameter for sqNULA noise $\mathbf{e}_2$: $\frac{1}{2} < \delta < 1$.
- sqNULA secret sparsity parameter: $\gamma < \frac{1}{2} : \gamma + \delta > 1$.
- Smallness parameter for sqNULA noise $\mathbf{e}_1$: $\sigma = n$.
- Number of sqNULA samples: n .
- Prime modulus $q \geq n^4$.
- sqNULA secret smallness parameter: $\xi = n$.

Parameters: $n, q, \delta, \gamma, \xi, \sigma, \zeta$.

Error Distributions:

1. **Sparse Error for sqNULA.** Let $S_{n,q,\delta}$ be a distribution over $\mathbb{F}_q$ such that sampling results in 0 with probability $1 - n^{-\delta}$ and a uniform random element from $\mathbb{F}_q$ with probability $n^{-\delta}$.
2. **Gaussian Error for sqNULA.** Let $\mathcal{D}_{\mathbb{Z},\sigma}$ be a discrete Gaussian with width σ .
3. **Small and Sparse secret for sqNULA.** Let $\mathcal{E}_{n,\xi,\gamma}$ be a distribution over $\mathbb{F}_q$ such that sampling results in 0 with probability $1 - n^{-\gamma}$ and in a random value from $\mathcal{D}_{\mathbb{Z},\xi}$ with probability $n^{-\gamma}$.

Algorithms:

★ $(pk, sk) \leftarrow \text{Gen}(1^\lambda)$:

1. Sample a uniform random matrix $\mathbf{A} \in \text{GL}_q(n)$.
2. Sample $\mathbf{s} \sim \mathcal{E}_{n,\xi,\gamma}^n$, sample $\mathbf{e}_1 \sim \mathcal{D}_\sigma^n$, sample $\mathbf{e}_2 \sim S_{n,q,\delta}^n$.
3. Set $pk \leftarrow (\mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e}_1 + \mathbf{e}_2)$. Set $sk \leftarrow \mathbf{s}$.
4. Output (pk, sk) .

★ $ct \leftarrow \text{Enc}(1^\lambda, pk, b \in \{0, 1\})$:

1. Parse pk as $(\mathbf{A}, \mathbf{y})$.
2. Sample $\mathbf{s}' \sim \mathcal{E}_{n,\xi,\gamma}^n$, $\mathbf{e}'_1 \sim \mathcal{D}_\sigma^n$, and $\mathbf{e}'_2 \sim S_{n,q,\delta}^n$.
3. If $b = 0$, output $(\mathbf{A}^\top \mathbf{s}' + \mathbf{e}'_1 + \mathbf{e}'_2, \mathbf{y}^\top \mathbf{s}')$.
4. If $b = 1$, output $(\mathbf{u}_1, u_2)$ for uniform randomly sampled $\mathbf{u}_1^\top \in \mathbb{F}_q^n$, and $u_2 \in \mathbb{F}_q$.

★ $b' \leftarrow \text{Dec}(1^\lambda, sk, ct)$:

1. Parse sk as $\mathbf{s}$, and parse ct as $(ct_1 \in \mathbb{F}_q^n, ct_2 \in \mathbb{F}_q)$
2. $x \leftarrow (ct_2 - \mathbf{s}^\top ct_1) \bmod q$
3. If $x \leq \lfloor \frac{q}{\zeta} \rfloor$ or $x \geq q - \lfloor \frac{q}{\zeta} \rfloor$, then output $b' = 0$, otherwise output $b' = 1$.

Correctness.

Lemma 4.1. *Let $\lambda \in \mathbb{N}$ be the security parameter. Then there exists a constant $c < 1/2$, such that for all $n = n(\lambda) = \text{poly}(\lambda) \in \mathbb{N}$, $\delta \in (0, 1)$, $\gamma \geq 1 - \delta$, $b \in \{0, 1\}$, $\zeta \geq 4 \in \mathbb{N}$, $\mu, \nu > 0$, $\sigma = \sigma(\lambda) > 0$, $\xi = \xi(\lambda) > 0$, prime $q = q(\lambda) > 2\zeta n \xi^{1+\nu} \sigma^{1+\mu}$,*

$$\Pr[(pk, sk) \leftarrow \text{Gen}(1^\lambda) \wedge \text{Dec}(sk, \text{Enc}(pk, b)) = b] \geq 1 - c.$$

Proof. We proceed by analyzing the decryption error for the cases of $b = 1$ and $b = 0$ separately.

Case 1: If $b = 1$, in Step 2 of the decryption procedure,

$$x = ct_2 - \mathbf{s}^\top ct_1 = u_2 - \sum_{i \in [m]} s_i u_{1,i} \bmod q.$$

Since, each $u_{1,i}$ and s_i are independent and uniformly sampled from $\mathbb{F}_q$, y is also distributed uniformly over $\mathbb{F}_q$, thus

$$\Pr \left[q - \lfloor \frac{q}{\zeta} \rfloor \geq x \geq \lfloor \frac{q}{\zeta} \rfloor \right] \geq 1 - \frac{2}{\zeta}.$$

Thus,

$$\Pr[(pk, sk) \leftarrow \text{Gen}(1^\lambda) \wedge \text{Dec}(sk, \text{Enc}(pk, 1)) = 1] = 1 - \frac{2}{\zeta}.$$

Case 2: If $b = 0$, then in Step 2 of the decryption procedure,

$$x = (ct_2 - s^\top ct_1) \bmod q = (e_1^\top s' + e_2^\top s' - s^\top e'_1 - s^\top e'_2) \bmod q.$$

We first compute the probability that the event $e_2^\top s' - s^\top e'_2 = 0$ occurs.

To do so, observe that the expected number of non-zero entries in $s \sim \mathcal{E}_{n, \xi, \gamma}^n$ is $n^{1-\gamma}$. By a standard Chernoff bound, with an all but negligible probability in n over the choice of s , the number of non-zero entries in s is at most $2n^{1-\gamma}$. Clearly $s^\top e'_2 = 0$ if s has $2n^{1-\gamma}$ non-zero elements in the worst case, and all the corresponding entries in e'_2 at these indices are 0. Therefore, the probability that $s^\top e'_2 = 0$ is upper bounded by the probability that e'_2 has 0 at the $2n^{1-\gamma}$ many non-zero entries of s . The probability of this is

$$(1 - n^{-\delta})^{2n^{1-\gamma}} = O(1)$$

for any γ, δ such that $\gamma + \delta \geq 1$.

Similarly, with the same probability, $e_2^\top s' = 0$.

Conditioning on the event that $e_2^\top s' - s^\top e'_2 = 0$, we have that $x = e_1^\top s' - s^\top e'_1 \bmod q$. For correctness to hold, we need to show that this value of x will be sufficiently small with high probability. First, we note that the entries of e_1 and e'_1 will all be upper bounded by $\sigma^{1+\mu}$ with overwhelming probability. Let $e_{1,i}$ denote the i^{th} coordinate of the vector e_1 . From Lemma 2.2, we know that for a fixed $e_{1,i}$ and small constant $\mu > 0$,

$$\Pr[e_{1,i} > \sigma^{1+\mu}] \leq 2e^{-\pi\sigma^{2\mu}}.$$

Taking a union bound,

$$\Pr[\exists i \in [n] : e_{1,i} > \sigma^{1+\mu}] \leq 2n \cdot e^{-\pi\sigma^{2\mu}}.$$

This bound also applies to e'_1 .

Similarly applying the same lemma on s gives us

$$\Pr[\exists i \in [n] : s_i > \xi^{1+\nu}] \leq 2n \cdot e^{-\pi\sigma^{2\nu}}.$$

Therefore it must be that $|s^\top e'_1 - s'^\top e_1| \leq 2n\xi^{1+\nu}\sigma^{1+\mu}$ with probability at least $1 - 4n \cdot e^{-\pi\sigma^{2\mu}} - 4n \cdot e^{-\pi\sigma^{2\nu}}$. Since we have chosen $q > 2\zeta n\xi^{1+\nu}\sigma^{1+\mu}$, x must be less than $\lfloor \frac{q}{\zeta} \rfloor$ or larger than $q - \lfloor \frac{q}{\zeta} \rfloor$.

Then,

$$\Pr[(pk, sk) \leftarrow \text{Gen}(1^\lambda) \wedge \text{Dec}(sk, \text{Enc}(pk, 0)) = 0] \geq 1 - c.$$

Combining Case 1 and Case 2 gives us our desired correctness statement. $\square$

Remark 4.2. While we achieve only constant decryption error in the above lemma, correctness can be amplified to achieve negligible decryption error by considering an encryption process that outputs several fresh encryptions for a single bit and a decryption process that outputs the majority value of their decryptions.

Lemma 4.3 (Semantic Security). *Let λ be the security parameter. Then for every $n = \text{poly}(\lambda)$, prime $q(\lambda)$, $\xi, \sigma = \text{poly}(\lambda)$, $\gamma, \delta \in (0, 1)$, and for all polynomial-sized adversaries $\mathcal{A}$, assuming the hardness of the Decisional-sqNULA $_{n,q,\sigma,\xi,\delta,\gamma}$ problem (Assumption 1.1), there exists a negligible function $\mu(\cdot) : \mathbb{N} \rightarrow [0, 1]$ such that*

$$\left| \Pr \left[(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda); \mathcal{A}(\text{pk}, \text{Enc}(\text{pk}, 0)) = 1, \right] \right. \\ \left. - \Pr \left[(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda); \mathcal{A}(\text{pk}, \text{Enc}(\text{pk}, 1)) = 1 \right] \right| \leq \mu(\lambda),$$

where the probability statement is over the coins of Gen, the coins of Enc, and the coins of $\mathcal{A}$.

Proof. We proceed by arguing the indistinguishability of several hybrids.

Hybrid $\mathcal{H}_0(1^\lambda)$:

1. $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$ where $\text{pk} = (\mathbf{A}, \mathbf{y} := \mathbf{A} \cdot \mathbf{s} + \mathbf{e}_1 + \mathbf{e}_2)$ and $\mathbf{A} \xleftarrow{\$} \mathbb{F}_q^{n \times n}$, and $\mathbf{s}, \mathbf{e}_1, \mathbf{e}_2$ are sampled from the appropriate distributions as in $\text{Gen}(1^\lambda)$.
2. $\text{ct} \leftarrow \text{Enc}(\text{pk}, 0)$ where
$$\text{ct} \leftarrow (\mathbf{A}^\top \mathbf{s}' + \mathbf{e}'_1 + \mathbf{e}'_2, \mathbf{y}^\top \mathbf{s}')$$
where $\mathbf{s}' \sim \mathcal{E}_{n,\xi,\gamma}^n$, $\mathbf{e}'_1 \sim \mathcal{D}_\sigma^n$, and $\mathbf{e}'_2 \sim S_{n,q,\delta}^n$.

Hybrid $\mathcal{H}_{0,\$}(1^\lambda)$:

1. $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$ where $\text{pk} = (\mathbf{A}, \mathbf{y})$ and $\mathbf{A} \xleftarrow{\$} \mathbb{F}_q^{n \times n}$, $\mathbf{y} \xleftarrow{\$} \mathbb{F}_q^n$.
2. $\text{ct} \leftarrow \text{Enc}(\text{pk}, 0)$ where
$$\text{ct} \leftarrow (\mathbf{A}^\top \mathbf{s}' + \mathbf{e}'_1 + \mathbf{e}'_2, \mathbf{y}^\top \mathbf{s}')$$
where $\mathbf{s}' \sim \mathcal{E}_{n,\xi,\gamma}^n$, $\mathbf{e}'_1 \sim \mathcal{D}_\sigma^n$, and $\mathbf{e}'_2 \sim S_{n,q,\delta}^n$.
3. Output (pk, ct) .

Hybrid $\mathcal{H}_{0.5,\$}(1^\lambda)$:

1. $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$ where $\text{pk} = (\mathbf{A}, \mathbf{y})$ and $\mathbf{A} \xleftarrow{\$} \mathbb{F}_q^{n \times n}$, $\mathbf{y} \xleftarrow{\$} \mathbb{F}_q^n$.
2. $\text{ct} \leftarrow \text{Enc}(\text{pk}, 1)$ so that
$$\text{ct} \leftarrow (\mathbf{A}^\top \mathbf{s}' + \mathbf{e}'_1 + \mathbf{e}'_2, u_2),$$
where, $\mathbf{s}' \sim \mathcal{E}_{n,\xi,\gamma}^n$, $\mathbf{e}'_1 \sim \mathcal{D}_\sigma^n$, and $\mathbf{e}'_2 \sim S_{n,q,\delta}^n$, $u_2 \xleftarrow{\$} \mathbb{F}_q$.
3. Output (pk, ct) .

Hybrid $\mathcal{H}_{1,\$}(1^\lambda)$:

1. $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$ where $\text{pk} = (\mathbf{A}, \mathbf{y})$ and $\mathbf{A} \xleftarrow{\$} \mathbb{F}_q^{n \times n}$, $\mathbf{y} \xleftarrow{\$} \mathbb{F}_q^n$.
2. $\text{ct} \leftarrow \text{Enc}(\text{pk}, 1)$ so that
$$\text{ct} \leftarrow (\mathbf{u}_1, u_2),$$

where, $\mathbf{u}_1 \xleftarrow{\$} \mathbb{F}_q^n, u_2 \xleftarrow{\$} \mathbb{F}_q$.

3. Output (pk, ct) .

Hybrid $\mathcal{H}_1(1^\lambda)$:

1. $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$ where $\text{pk} = (\mathbf{A}, \mathbf{y} := \mathbf{A} \cdot \mathbf{s} + \mathbf{e}_1 + \mathbf{e}_2)$ and $\mathbf{A} \xleftarrow{\$} \mathbb{F}_q^{n \times n}$, and $\mathbf{s}, \mathbf{e}_1, \mathbf{e}_2$ are sampled from the appropriate distributions as in $\text{Gen}(1^\lambda)$.

2. $\text{ct} \leftarrow \text{Enc}(\text{pk}, 1)$ so that

$$\text{ct} \leftarrow (\mathbf{u}_1, u_2),$$

where, $\mathbf{u}_1 \xleftarrow{\$} \mathbb{F}_q^n, u_2 \xleftarrow{\$} \mathbb{F}_q$.

3. Output (pk, ct) .

Lemma 4.4 ($\mathcal{H}_0 \approx_{\xi_0} \mathcal{H}_{0,\$}$). For every $n = \text{poly}(\lambda)$, $\sigma(\lambda), \xi(\lambda) > 0$, $\delta, \gamma < 1$, prime $q(\lambda)$, and every polynomial sized adversary $\mathcal{A}$, if the hardness of the decisional-sqNULA $_{n,q,\sigma,\xi,\delta,\gamma}$ problem holds (Assumption 1.1), then there exists a negligible function $\xi_0 : \mathbb{N} \rightarrow [0, 1]$ such that for all sufficiently large $\lambda \in \mathbb{N}$,

$$\left| \Pr [\mathcal{A}(1^\lambda, \mathcal{H}_0(1^\lambda)) = 1] - \Pr [\mathcal{A}(1^\lambda, \mathcal{H}_{0,\$}(1^\lambda)) = 1] \right| \leq \xi_0(\lambda).$$

Proof. Suppose there exists a polynomial-sized $\mathcal{A}$ that distinguishes $\mathcal{H}_0(1^\lambda)$ and $\mathcal{H}_{0,\$}(1^\lambda)$ with probability $\frac{1}{2} + \text{adv}(\lambda)$ for non-negligible $\text{adv} : \mathbb{N} \rightarrow [0, 1]$. We now construct a distinguisher $\mathcal{D}$ that breaks the decisional-sqNULA problem (Assumption 1.1). $\mathcal{D}$ takes as input $\mathbf{A} \xleftarrow{\$} \mathbb{F}_q^{n \times n}$, and $\mathbf{y}$ is sampled either as $\mathbf{A} \cdot \mathbf{s} + \mathbf{e}_1 + \mathbf{e}_2 \bmod q$ or uniformly at random from $\mathbb{F}_q^n$.

1. $\mathcal{D}$ sets $\text{pk} \leftarrow (\mathbf{A}, \mathbf{y})$.
2. Sample random $\mathbf{s}' \sim \mathcal{E}_{n,\xi,\gamma}^n$, $\mathbf{e}'_1 \sim \mathcal{D}_\sigma^n$, and $\mathbf{e}'_2 \sim \mathcal{S}_{n,q,\delta}^n$.
3. $\mathcal{D}$ sets $\text{ct} \leftarrow (\mathbf{A}^\top \mathbf{s}' + \mathbf{e}'_1 + \mathbf{e}'_2, \mathbf{y}^\top \mathbf{s}')$.
4. Output $\mathcal{A}(\text{pk}, \text{ct})$.

Observe that if $\mathbf{y}$ is uniform randomly sampled, then (pk, ct) is exactly from the distribution $\mathcal{H}_{0,\$}(1^\lambda)$ and if $\mathbf{y}$ is of the form $\mathbf{A} \cdot \mathbf{s} + \mathbf{e}_1 + \mathbf{e}_2 \bmod q$, then (pk, ct) is exactly from the distribution $\mathcal{H}_0(1^\lambda)$. Therefore, $\mathcal{D}$ distinguishes between the two settings with non-negligible advantage $\text{adv}(\lambda)$, contradicting Assumption 1.1. Therefore, no such $\mathcal{A}$ can exist. $\square$

Lemma 4.5 ($\mathcal{H}_{0,\$} \approx_{\xi_1} \mathcal{H}_{0.5,\$}$). Assuming hardness of Search-sqNULA $_{n,q,\sigma,\xi,\delta,\gamma}$, for all polynomial-sized adversaries $\mathcal{A}$, there exists a negligible function $\xi_1(\cdot) : \mathbb{N} \rightarrow [0, 1]$ such that for all sufficiently large $\lambda \in \mathbb{N}$,

$$\left| \Pr [\mathcal{A}(1^\lambda, \mathcal{H}_{0,\$}(1^\lambda)) = 1] - \Pr [\mathcal{A}(1^\lambda, \mathcal{H}_{1,\$}(1^\lambda)) = 1] \right| \leq \xi_1(\lambda).$$

Proof. The indistinguishability of these two hybrids follows from the pseudorandomness of the hardcore bit that is implied by the hardness of the Search-sqNULA $_{n,q,\sigma,\xi,\delta,\gamma}$ assumption. Observe that the search-sqNULA $_{n,q,\sigma,\xi,\delta,\gamma}$ assumption also implies that given $\mathbf{A}^\top, \mathbf{b} = \mathbf{A}^\top \mathbf{s}' + \mathbf{e}'_1 + \mathbf{e}'_2, \mathbf{y} \xleftarrow{\$} \mathbb{F}_q^n$, it is hard to computationally distinguish $\mathbf{y}^\top \mathbf{s}'$ and $u_2 \xleftarrow{\$} \mathbb{F}_q$. This is a direct consequence of the generalization of the Goldreich-Levin hardcore bit lemma [GL89] over large fields [GRS00].

Suppose there exists a polynomial-sized $\mathcal{A}$ that distinguishes $\mathcal{H}_{0.5,\$}(1^\lambda)$ and $\mathcal{H}_{1,\$}(1^\lambda)$ with probability $\frac{1}{2} + \text{adv}(\lambda)$ for non-negligible $\text{adv} : \mathbb{N} \rightarrow [0, 1]$, we construct a distinguisher $\mathcal{D}$ that breaks the pseudorandomness of the hardcore bit for the Search-sqNULA $_{n,q,\sigma,\xi,\delta,\gamma}$ assumption as follows.

$\mathcal{D}$ takes as input $((\mathbf{A}, \mathbf{y}), (\mathbf{b} = \mathbf{A}^\top \mathbf{s}' + \mathbf{e}'_1 + \mathbf{e}'_2, r))$ for which it must distinguish having received a sample from one of two settings. In the first setting, $r \leftarrow \mathbf{y}^\top \mathbf{s}'$, In the second setting, $r = u_2 \xleftarrow{\$} \mathbb{F}_q$. $\mathcal{D}$ proceeds as follows.

1. Set $\text{pk} \leftarrow (\mathbf{A}, \mathbf{y})$, $\text{ct} \leftarrow (\mathbf{b}, r)$.
2. Output $\mathcal{A}(\text{pk}, \text{ct})$.

Observe that given an input from the first setting, (pk, ct) is exactly from the distribution $\mathcal{H}_{0.5, \$}(1^\lambda)$. Given an input from the second setting, then (pk, ct) is exactly from the distribution $\mathcal{H}_{1, \$}(1^\lambda)$. Therefore $\mathcal{D}$ distinguishes with advantage $\text{adv}(\lambda)$.

Therefore no such $\mathcal{A}$ exists, implying that for all polynomial-sized adversaries $\mathcal{A}$ there exists some negligible function ξ_1 such that $\mathcal{A}$ has at most a $\xi_1(1^\lambda)$ distinguishing advantage for all $\lambda \in \mathbb{N}$. $\square$

Lemma 4.6 ($\mathcal{H}_{1, \$}(1^\lambda) \approx_{\xi_2} \mathcal{H}_{0.5, \$}(1^\lambda)$). *For every $n = \text{poly}(\lambda)$, $\sigma(\lambda), \xi(\lambda) > 0$, $\delta, \gamma < 1$, prime $q(\lambda)$, and every polynomial sized adversary $\mathcal{A}$, if the hardness of the decisional-sqNULA $_{n, q, \sigma, \xi, \delta, \gamma}$ problem holds (Assumption 1.1), then there exists a negligible function $\xi_2 : \mathbb{N} \rightarrow [0, 1]$ such that for all sufficiently large $\lambda \in \mathbb{N}$,*

$$\left| \Pr [\mathcal{A}(1^\lambda, \mathcal{H}_{1, \$}(1^\lambda)) = 1] - \Pr [\mathcal{A}(1^\lambda, \mathcal{H}_1(1^\lambda)) = 1] \right| \leq \xi_2(\lambda).$$

Proof. The indistinguishability of these two hybrids is proven in similarly as Lemma 4.4 (by symmetry). $\square$

Lemma 4.7 ($\mathcal{H}_{1, \$}(1^\lambda) \approx_{\xi_3} \mathcal{H}_1(1^\lambda)$). *For every $n = \text{poly}(\lambda)$, $\sigma(\lambda), \xi(\lambda) > 0$, $\delta, \gamma < 1$, prime $q(\lambda)$, and every polynomial sized adversary $\mathcal{A}$, if the hardness of the decisional-sqNULA $_{n, q, \sigma, \xi, \delta, \gamma}$ problem holds (Assumption 1.1), then there exists a negligible function $\xi_3 : \mathbb{N} \rightarrow [0, 1]$ such that for all sufficiently large $\lambda \in \mathbb{N}$,*

$$\left| \Pr [\mathcal{A}(1^\lambda, \mathcal{H}_{1, \$}(1^\lambda)) = 1] - \Pr [\mathcal{A}(1^\lambda, \mathcal{H}_1(1^\lambda)) = 1] \right| \leq \xi_3(\lambda).$$

Proof. The indistinguishability of these two hybrids is proven in exactly the same manner as Lemma 4.4. $\square$

Combining Lemmas 4.4, 4.5, 4.6, 4.7 gives the desired statement. $\square$

5 Malicious Oblivious Transfer in Two Rounds

We present a construction of two round elementary OT. Our parameter choice suggestions remain identical to that of the PKE.

Parameters: $n, q, \delta, \gamma, \xi, \sigma, \zeta, \ell$.

Rounding Function: For every $x \in \mathbb{F}_q$, $\text{Round}(x)$ is defined as follows:

- Let $\ell \in \mathbb{N}$ and let $q' < q$ such that $\ell \mid q'$.
- Partition $\mathbb{F}_q$ into $\frac{q'}{\ell} + 1$ blocks such that all but the last one has length ℓ , i.e., the blocks form a set of intervals as follows: $\{[0, \ell], [\ell, 2\ell], \dots, [\frac{q'}{\ell}\ell, q - 1]\}$
- Find k^* such that $k^*\ell \leq x \leq (k^* + 1)\ell$.
- $\text{Round}(x) = k^*\ell + \lfloor \frac{\ell}{2} \rfloor \pmod{q}$ if $x \leq q'$, else $\text{Round}(x) = q - 1$.

Algorithms:

- ★ $\text{crs} \leftarrow \text{Gen}(1^\lambda)$:
 1. Sample a uniform random matrix $\mathbf{A} \in \text{GL}_q(n)$.
 2. Sample a uniform random vector $\mathbf{v} \in \mathbb{F}_q^n$.
 3. Output $\text{crs} := (\mathbf{A}, \mathbf{v})$
- ★ $(\text{otr}, \text{st}) \leftarrow \text{Receiver}_1(\text{crs}, c \in \{0, 1\})$:
 1. Sample $\mathbf{s} \sim \mathcal{E}_{n, \xi, \gamma}^n$, $\mathbf{e}_1 \sim \mathcal{D}_\sigma^n$, and $\mathbf{e}_2 \sim S_{n, q, \delta}^n$.
 2. $\text{otr} \leftarrow \mathbf{A}\mathbf{s} + \mathbf{e}_1 + \mathbf{e}_2 + c \cdot \lfloor \frac{q}{2} \rfloor \cdot \mathbf{v}$.
 3. $\text{st} = (\mathbf{s}, \mathbf{e}_1, \mathbf{e}_2)$.
 4. Output otr, st .
- ★ $\text{ots} \leftarrow \text{Sender}(\text{crs}, \text{otr})$:
 1. $x_0 = \text{otr}, x_1 = \text{otr} - \lfloor \frac{q}{2} \rfloor \cdot \mathbf{v}$.
 2. Sample $\mathbf{s}' \sim \mathcal{E}_{n, \xi, \gamma}^n$, $\mathbf{e}'_1 \sim \mathcal{D}_\sigma^n$, and $\mathbf{e}'_2 \sim S_{n, q, \delta}^n$.
 3. $y_0 \leftarrow \mathbf{s}'^\top x_0, y_1 \leftarrow \mathbf{s}'^\top x_1$.
 4. $m_0 \leftarrow \text{Round}(y_0), m_1 \leftarrow \text{Round}(y_1)$.
 5. $\text{ots} \leftarrow \mathbf{A}^\top \mathbf{s}' + \mathbf{e}'_1 + \mathbf{e}'_2$.
 6. Output ots .
- ★ $m' \leftarrow \text{Receiver}_2(\text{crs}, \text{otr}, \text{ots}, \text{st})$:
 1. $x \leftarrow \text{ots}^\top \mathbf{s}$.
 2. $m' \leftarrow \text{Round}(x)$.

Correctness

Lemma 5.1. *Let $\lambda \in \mathbb{N}$ be the security parameter. Then, there exists a positive constant $\alpha \leq \frac{1}{2}$ such that for all $n = \text{poly}(\lambda)$, $\delta \in (0, 1)$, $\gamma \geq 1 - \delta$, $c \in \{0, 1\}$, $\eta = \text{poly}(\lambda)$, $\mu, \nu > 0$, $\sigma = \sigma(\lambda) > 0$, $\xi = \xi(\lambda) > 0$, $\ell = \ell(\lambda) \geq \eta n \sigma^{1+\mu} \xi^{1+\nu}$, prime $q = q(\lambda) = \text{poly}(\ell)$,*

$$\Pr[m_c = \text{Receiver}_2(\text{crs}, \text{otr}, \text{ots}, \text{st}) \mid \text{crs} \leftarrow \text{Gen}(1^\lambda), (\text{otr}, \text{st}) \leftarrow \text{Receiver}_1(\text{crs}, c), \text{ots} \leftarrow \text{Sender}(\text{crs}, \text{otr})] \geq 1 - \alpha.$$

Here (m_0, m_1) is the message pair set by Sender in Step 4 of its execution.

Proof. First we compute the the two purported messages m_0 and m_1 .

Case I. $c = 0$:

In this case, $\text{otr} = \mathbf{A}\mathbf{s} + \mathbf{e}_1 + \mathbf{e}_2$. Thus,

$$y_0 = \mathbf{s}'^\top (\mathbf{A}\mathbf{s} + \mathbf{e}_1 + \mathbf{e}_2).$$

Since $\gamma + \delta \geq 1$, $\mathbf{s}'^\top \mathbf{e}_2 = 0$ with probability $\Theta(1)$. This follows from the the same arguments as made in the proof of Lemma 4.1.

Therefore, with probability $1 - \Theta(1)$,

$$y_0 = \mathbf{s}'^\top \mathbf{A}\mathbf{s} + \mathbf{s}'^\top \mathbf{e}_1.$$

Similarly, with probability $1 - \Theta(1)$,

$$y_1 = \mathbf{s}'^T \mathbf{A} \mathbf{s} + \mathbf{s}'^T \mathbf{e}_1 - \lfloor \frac{q}{2} \rfloor \cdot \mathbf{s}'^T \mathbf{v}.$$

Observe that with all but negligible probability, $\mathbf{s}'^T \mathbf{e}_1 \leq n\sigma^{1+\mu}\xi^{1+\nu}$, for some small constant $\mu, \nu \in (0, 1)$.

Moreover, for any fixed $t \in \mathbb{F}_q$, the probability that, $\mathbf{s}'^T \mathbf{A} \mathbf{s} = t$ is statistically close to $\frac{1}{q}$.

To see why this is the case, first observe that $\mathbf{s}$ and $\mathbf{s}'$ are non-zero with probability exponential close to 1. Now, for uniform random invertible $\mathbf{A}$ and any fixed non-zero $\mathbf{s}$, $\mathbf{A} \mathbf{s}$ is uniformly distributed over $\mathbb{F}_q^n \setminus \{\mathbf{0}\}$. This result is stated as Lemma 2.1. Then, $\mathbf{s}'^T \mathbf{A} \mathbf{s}$ is identical to the inner product of non-zero $\mathbf{s}'$ with a random non-zero vector over $\mathbb{F}_q^n \setminus \{\mathbf{0}\}$ which is uniformly distributed over $\mathbb{F}_q$.

Therefore, the probability that there exists an $k < \frac{q'}{\ell}$ such that $\mathbf{s}'^T \mathbf{A} \mathbf{s} - k\ell \geq n\sigma^{1+\mu}\xi^{1+\nu}$ and $(k+1)\ell - \mathbf{s}'^T \mathbf{A} \mathbf{s} \geq n\sigma^{1+\mu}\xi^{1+\nu}$ is at least $1 - \frac{2n\sigma^{1+\mu}\xi^{1+\nu} \cdot \frac{q'}{\ell} + \ell}{q}$. This is lower bounded by $1 - \frac{2n\sigma^{1+\mu}\xi^{1+\nu}}{\ell} - \frac{\ell}{q}$.

We have chosen $\ell \geq \eta n\sigma^{1+\mu}\xi^{1+\nu}$, where $\eta = \text{poly}(\lambda)$ and $q = \text{poly}(\ell)$. Thus, the above probability becomes $1 - \frac{1}{\eta} - \frac{1}{\text{poly}(\ell)}$.

Let $k^* < \frac{q'}{\ell}$ be such that $\mathbf{s}'^T \mathbf{A} \mathbf{s} - k^*\ell \geq n\sigma^{1+\mu}\xi^{1+\nu}$ and $(k^*+1)\ell - \mathbf{s}'^T \mathbf{A} \mathbf{s} \geq n\sigma^{1+\mu}\xi^{1+\nu}$, we can conclude that $\text{Round}(y_0) = k^*\ell + \lfloor \frac{\ell}{2} \rfloor$.

Also, in that case the recovered x by the receiver is

$$x = \mathbf{s}'^T \mathbf{A} \mathbf{s} + \mathbf{e}_1^T \mathbf{s} + \mathbf{e}_2^T \mathbf{s}.$$

Again, with probability $(1 - \Theta(1))$

$$x = \mathbf{s}'^T \mathbf{A} \mathbf{s} + \mathbf{e}_1^T \mathbf{s}.$$

Conditioned on the event that $\mathbf{e}_1^T \mathbf{s} \leq n\sigma^{1+\mu}\xi^{1+\nu}$ and $\mathbf{s}'^T \mathbf{A} \mathbf{s} - k^*\ell \geq n\sigma^{1+\mu}\xi^{1+\nu}$ and $(k^*+1)\ell - \mathbf{s}'^T \mathbf{A} \mathbf{s} \geq n\sigma^{1+\mu}\xi^{1+\nu}$, it must be that $m' = \text{Round}(x) = m_0$.

Case II. $c = 1$:

In this case, $\text{otr} = \mathbf{A} \mathbf{s} + \mathbf{e}_1 + \mathbf{e}_2 + \lfloor \frac{q}{2} \rfloor \cdot \mathbf{v}$. Thus,

$$y_0 = \mathbf{s}'^T (\mathbf{A} \mathbf{s} + \mathbf{e}_1 + \mathbf{e}_2 + \lfloor \frac{q}{2} \rfloor \cdot \mathbf{v}).$$

Therefore, with probability $1 - \Theta(1)$,

$$y_0 = \mathbf{s}'^T \mathbf{A} \mathbf{s} + \mathbf{s}'^T \mathbf{e}_1 + \lfloor \frac{q}{2} \rfloor \cdot \mathbf{s}'^T \mathbf{v}.$$

Similarly,

$$y_1 = \mathbf{s}'^T \mathbf{A} \mathbf{s} + \mathbf{s}'^T \mathbf{e}_1.$$

From here, the observation that $m' = m_1$ follows identically to the above case.

□

Remark 5.2. We can again amplify to achieve negligible correctness error by repeating several times taking a majority value of the outputs of calls to Receiver2.

Lemma 5.3 (Receiver Security against Malicious Sender). *Assuming the hardness of decision-sqNULA_{n,q,σ,δ,γ}, Construction 5 satisfies receiver's indistinguishability security.*

Proof. The receiver's bit c is masked by an sqNULA sample $\mathbf{A}\mathbf{s} + \mathbf{e}_1 + \mathbf{e}_2$, so pseudorandomness of this sample directly implies receiver's security. $\square$

Lemma 5.4 (Elementary Sender Security against Malicious Receiver). *Assuming the hardness of search-sqNULA_{n,q,σ,ξ,δ,γ}, Construction 5 satisfies sender's elementary OT security.*

Proof. First, for all $x, y \in \mathbb{F}_q$, define $d(x, y) = \min((x - y) \bmod q, q - (x - y) \bmod q)$.

If there exists a receiver that can break elementary sender security, i.e., output both m_0 and m_1 , we can use it to build an adversary $\mathcal{A}$ that can break the search-sqNULA assumption. Specifically, we will show that if there exists an adversary $(\mathcal{B}_1, \mathcal{B}_2)$ that can break elementary sender security, then we can use it to build another adversary $\mathcal{A}$ that breaks pseudorandomness of the hardcore bit with respect to the search sqNULA assumption.

Say $\mathcal{A}$ receives as input $(\mathbf{A}^\top, \mathbf{v}, \mathbf{b}, z)$ where $\mathbf{b} = \mathbf{A}\mathbf{s}' + \mathbf{e}'_1 + \mathbf{e}'_2$, and z is either uniform over $\mathbb{F}_q$ or $\mathbf{v}^T \mathbf{s}'$. $\mathcal{A}$ forwards $(\mathbf{A}, \mathbf{v})$ as crs to $\mathcal{B}_1$. Let $\mathcal{B}_1$ output X as its first message. The sender imagines $y_0 = \mathbf{s}'^T \mathbf{X}$ and $y_1 = \mathbf{s}'^T (\mathbf{X} + \lfloor \frac{q}{2} \rfloor \cdot \mathbf{v})$, and imagine $m_0 = \text{Round}(y_0)$, $m_1 = \text{Round}(y_1)$, and then forwards $\mathbf{b}$ to $\mathcal{B}_2$ as its message. Let the response from $\mathcal{B}_2$ be (m'_0, m'_1) . Output 1 if $|d(0, \lfloor \frac{q}{2} \rfloor \cdot z) - d(m'_1, m'_0)| \leq 2\ell$, else output 0.

Observe that $d(y_1, y_0) = d(0, \lfloor \frac{q}{2} \rfloor \cdot \mathbf{s}'^T \mathbf{v})$ and therefore

$$\begin{aligned} d(\text{Round}(y_1), \text{Round}(y_0)) \\ \leq d(y_1, y_0) + d(\text{Round}(y_1), y_1) + d(\text{Round}(y_0), y_0) \leq d\left(0, \left\lfloor \frac{q}{2} \right\rfloor \langle \mathbf{s}', \mathbf{v} \rangle\right) + 2\ell. \end{aligned}$$

This last inequality is because for all $i \in \{0, 1\}$, $|y_i - \text{Round}(y_i)| \leq \ell$.

Therefore, if $m'_0 = \text{Round}(y_0)$ and $m'_1 = \text{Round}(y_1)$ and $z = \mathbf{s}'^T \mathbf{v}$, then $|d(0, \lfloor \frac{q}{2} \rfloor \cdot z) - d(m'_1, m'_0)| = |d(0, \lfloor \frac{q}{2} \rfloor \cdot \mathbf{s}'^T \mathbf{v}) - d(m'_1, m'_0)| \leq 2\ell$ using triangle inequality. On the other hand, if z is a random value then $z = \mathbf{s}'^T \mathbf{v}$, then $|d(0, \lfloor \frac{q}{2} \rfloor \cdot z) - d(m'_1, m'_0)| = |\$ - d(m'_1, m'_0)|$ which will be randomly distributed over $\mathbb{F}_q$. $\square$

Corollary 5.5. *Assuming the hardness of search-sqNULA, there exists two round UC-secure OT.*

Proof. This corollary follows directly through a series of transformations into two round UC-secure OT (see Figure 1 and Section 10.2.1 in [DGH⁺20] for an overview of these transformations²). Namely, an OT satisfying receiver's indistinguishability security (Lemma 5.3) along with sender's elementary security (Lemma 5.4) is enough to imply the existence of UC-secure OT. $\square$

6 One-Time Circular secure Bit-Secret Key Encryption

Our construction of circular secure SKE relies on a mild variant of the sqNULA assumption. Specifically, we change the distribution of the short and sparse secret. Instead of sampling from $\mathcal{E}_{n,\xi,\gamma}$, it is now sampled as a random vector over $\mathbb{F}_2^n$ such that the probability of 1 is $n^{-\gamma}$. We denote this distribution as $\mathcal{H}_{n,\gamma}$. We call this assumption binary secret-sqNULA

Parameters: $n, \delta, \gamma, \sigma, \zeta, q > 2n$.

Algorithms:

²A reader might notice that we need perfectly correct CPA-secure PKE in one of the steps of the transformation. This is obtained first by constructing perfectly correct OT via a generic PRG trick described in Section 10.2.1 of [DGH⁺20] rather than directly using our PKE construction.

★ $\text{sk} \leftarrow \text{Gen}(1^\lambda)$: 1. Sample $\text{s} \sim \mathcal{H}_{n,\gamma}$. 2. Output $\text{sk} := \text{s}$. ★ $\text{ct} \leftarrow \text{Enc}(1^\lambda, \text{sk}, b \in \{0, 1\})$: 1. Sample a uniform random matrix $\mathbf{A} \in \mathbb{F}_q^{n \times n}$. 2. Sample $\text{s}' \sim \mathcal{H}_{n,\gamma}$, $\mathbf{e}_1 \sim \mathcal{D}_\sigma^n$, and $\mathbf{e}_2 \sim S_{n,q,\delta}^n$, and $\mathbf{e}'_1 \sim \mathcal{D}_\sigma^n$ and $\mathbf{e}'_2 \sim S_{n,q,\delta}^n$. 3. Output $(\mathbf{A}, \mathbf{A}^\top \text{s}' + \mathbf{e}'_1 + \mathbf{e}'_2, \text{s}'^\top (\mathbf{A}\text{s} + \mathbf{e}_1 + \mathbf{e}_2) + \lfloor \frac{q}{2} \rfloor b)$. ★ $b' \leftarrow \text{Dec}(1^\lambda, \text{sk}, \text{ct})$: 1. Parse sk as s, and parse ct as $(\mathbf{A}, \text{ct}_1 \in \mathbb{F}_q^n, \text{ct}_2 \in \mathbb{F}_q)$ 2. $x \leftarrow (\text{ct}_2 - \text{s}^\top \text{ct}_1) \bmod q$ 3. If $x \leq \lfloor \frac{q}{\zeta} \rfloor$, then output $b' = 0$, otherwise output $b' = 1$.

Decryption Correctness

Lemma 6.1. *Let $\lambda \in \mathbb{N}$ be the security parameter. Then there exists a constant $c < 1/2$, such that for all $n = n(\lambda) = \text{poly}(\lambda) \in \mathbb{N}$, $\delta \in (0, 1)$, $\gamma = 1 - \delta$, $b \in \{0, 1\}$, $\zeta \in \mathbb{N}$, $\zeta \geq 2$, $\mu, \nu > 0$, $\sigma = \sigma(\lambda) > 0$, $\xi = \xi(\lambda) > 0$ prime $q = q(\lambda) > 2\zeta n \xi^{1+\nu} \sigma^{1+\mu}$,*

$$\Pr[\text{sk} \leftarrow \text{Gen}(1^\lambda) \wedge \text{Dec}(\text{sk}, \text{Enc}(\text{sk}, b)) = b] \geq 1 - c.$$

Proof. (Sketch) At a high level, when $b = 1$, the decrypted value will be close to $\lfloor \frac{q}{2} \rfloor$. On the other hand, when $b = 0$, the decrypted value will be much smaller for appropriately chosen parameters.

Formally, the proof will follow a similar blueprint as that of Lemma 4.1. $\square$

Semantic Security

Lemma 6.2. *Let λ be the security parameter. Then for every $n = \text{poly}(\lambda)$, prime $q(\lambda)$, $\gamma \in (0, 1)$, and for all polynomial-sized adversaries $\mathcal{A}$, assuming the hardness of the Decisional binary secret-sqNULA $_{n,q,\sigma,\delta,\gamma}$ problem (Assumption 1.1), there exists a negligible function $\mu(\cdot) : \mathbb{N} \rightarrow [0, 1]$ such that*

$$\left| \Pr[\text{sk} \leftarrow \text{Gen}(1^\lambda); \mathcal{A}(\text{Enc}(\text{sk}, 0)) = 1] - \Pr[\text{sk} \leftarrow \text{Gen}(1^\lambda); \mathcal{A}(\text{Enc}(\text{sk}, 1)) = 1] \right| \leq \mu(\lambda),$$

where the probability statement is over the coins of Gen , the coins of Enc , and the coins of $\mathcal{A}$.

Proof. (Sketch) The proof is identical to that of Lemma 4.3. Essentially, the encryptions of 0 and 1 are indistinguishable due to decisional binary secret-sqNULA $_{n,q,\sigma,\delta,\gamma}$ and the hardcore bit property corresponding to search binary secret-sqNULA $_{n,q,\sigma,\delta,\gamma}$ $\square$

One-Time circular Security Note that since the secret key s is over $\mathbb{F}_2^n$, we can assume without loss of generality that any projection function f that runs on input s will output $\text{s}_i \in \mathbb{F}_2$, for some $i \in [n]$.

Lemma 6.3. *For all $q > n$, assuming the hardness of decisional binary secret-sqNULA $_{n,q,\sigma,\delta,\gamma}$, Construction 6 is One-Time Circular secure.*

Proof. We will use the following claim in the proof. We proceed with the main proof assuming this claim that we shall prove later.

Claim 1. Let $\mathbf{r} \in \mathbb{F}_q^n$ be nonzero, and let $\mathbf{v} \in \mathbb{F}_q^n$ be arbitrary. Let $\mathbf{A} \leftarrow \text{GL}_q(n)$ be uniformly random. Sample $\mathbf{R} \in \mathbb{F}_q^{n \times n}$ uniformly subject to the affine constraint $\mathbf{r}^\top \mathbf{R} = \mathbf{v}^\top$. Set $\mathbf{C} := \mathbf{A} + \mathbf{R}$. Then $\mathbf{C}$ is uniform over the set

$$\mathcal{C}_{\mathbf{r}, \mathbf{v}} = \left\{ \mathbf{C} \in \mathbb{F}_q^{n \times n} : \mathbf{r}^\top \mathbf{C} \neq \mathbf{v}^\top \right\}.$$

In other words,

$$\Delta(\mathbf{C}, \text{Unif}(\mathbb{F}_q^{n \times n})) = q^{-n}.$$

We now prove one-time circular security. It suffices to show that an encryption of a projection function of the secret key is computationally indistinguishable from random.

We will show via a sequence of hybrids that the encryption of s_i is computationally indistinguishable from the encryption of 0.

Hybrid H_1 . This is the real encryption of s_i . Namely, sample

$$\mathbf{C} \leftarrow \mathbb{F}_q^{n \times n}, \quad \mathbf{s} \leftarrow H_{n, \gamma}, \quad \mathbf{s}' \leftarrow H_{n, \gamma},$$

sample the error vectors

$$\mathbf{e}^{(1)}, \mathbf{e}'^{(1)}, \mathbf{e}^{(2)}, \mathbf{e}'^{(2)}$$

as in Construction 6, and output

$$\left(\mathbf{C}, \mathbf{C}^\top \mathbf{s}' + \mathbf{e}^{(1)} + \mathbf{e}'^{(2)}, \mathbf{s}'^\top (\mathbf{C}\mathbf{s} + \mathbf{e}^{(1)} + \mathbf{e}^{(2)}) + ts_i \right).$$

Hybrid H_2 . This hybrid is sampled as follows.

First sample

$$\mathbf{A} \leftarrow \text{GL}_q(n).$$

Then sample $\mathbf{s} \leftarrow H_{n, \gamma}$, $\mathbf{s}' \leftarrow H_{n, \gamma}$ conditioned on $\mathbf{s}' \neq 0$. Sample $\mathbf{e}^{(1)}, \mathbf{e}^{(2)}$ as in the binary-secret square noisy unstructured linear algebraic assumption, and define $\mathbf{z} := \mathbf{A}\mathbf{s} + \mathbf{e}^{(1)} + \mathbf{e}^{(2)}$. Let $t = \lfloor \frac{q}{2} \rfloor$. Now sample $\mathbf{R} \in \mathbb{F}_q^{n \times n}$ uniformly subject to $\mathbf{s}'^\top \mathbf{R} = -t\mathbf{u}_i^\top$. Here $\mathbf{u}_i$ denotes the elementary vector that has 1 at the i^{th} position.

This affine constraint has a solution because $\mathbf{s}' \neq 0$, as shown in Claim 1.

Set $\mathbf{C} := \mathbf{A} + \mathbf{R}$. Finally, sample fresh errors $\mathbf{e}'^{(1)}, \mathbf{e}'^{(2)}$ as in Construction 6, and output

$$\left(\mathbf{C}, \mathbf{C}^\top \mathbf{s}' + \mathbf{e}^{(1)} + \mathbf{e}'^{(2)}, \mathbf{s}'^\top \mathbf{z} \right).$$

Observe that

$$\begin{aligned} \mathbf{s}'^\top \mathbf{z} &= \mathbf{s}'^\top (\mathbf{A}\mathbf{s} + \mathbf{e}^{(1)} + \mathbf{e}^{(2)}) \\ &= \mathbf{s}'^\top ((\mathbf{C} - \mathbf{R})\mathbf{s} + \mathbf{e}^{(1)} + \mathbf{e}^{(2)}) \\ &= \mathbf{s}'^\top (\mathbf{C}\mathbf{s} + \mathbf{e}^{(1)} + \mathbf{e}^{(2)}) - \mathbf{s}'^\top \mathbf{R}\mathbf{s}. \end{aligned}$$

By the definition of $\mathbf{R}$,

$$\mathbf{s}'^\top \mathbf{R} = -t\mathbf{u}_i^\top.$$

Hence

$$\mathbf{s}'^\top \mathbf{R}\mathbf{s} = -t\mathbf{u}_i^\top \mathbf{s} = -ts_i.$$

Substituting this into the previous equality gives

$$\mathbf{s}'^\top \mathbf{z} = \mathbf{s}'^\top (\mathbf{C}\mathbf{s} + \mathbf{e}^{(1)} + \mathbf{e}^{(2)}) + ts_i.$$

Therefore, the third component in H_2 is exactly the third component of an encryption of s_i . The only remaining difference between H_1 and H_2 is the distribution of $\mathbf{C}$. In H_1 , $\mathbf{C}$ is uniform over $\mathbb{F}_q^{n \times n}$. In

H_2 , by Claim 1 with $\mathbf{r} = \mathbf{s}'$ and $\mathbf{v} = -t\mathbf{u}_i$, the matrix $\mathbf{C} = \mathbf{A} + \mathbf{R}$ is q^{-n} -close to uniform over $\mathbb{F}_q^{n \times n}$. Therefore

$$\Delta(H_1, H_2) \leq \text{negl}(n).$$

Hybrid H_3 . This hybrid is identical to H_2 , except that $\mathbf{z}$ is sampled uniformly from $\mathbb{F}_q^n$ instead of being set to $\mathbf{A}\mathbf{s} + \mathbf{e}^{(1)} + \mathbf{e}^{(2)}$. Thus H_3 outputs

$$(\mathbf{C}, \mathbf{C}^\top \mathbf{s}' + \mathbf{e}'^{(1)} + \mathbf{e}'^{(2)}, \mathbf{s}'^\top \mathbf{z}),$$

where

$$\mathbf{z} \leftarrow \mathbb{F}_q^n, \mathbf{A} \leftarrow \text{GL}_q(n), \mathbf{s}' \leftarrow H_{n,\gamma} \mathbf{s}'^\top \mathbf{R} = -t\mathbf{u}_i^\top, \mathbf{C} = \mathbf{A} + \mathbf{R}.$$

We show that H_2 and H_3 are computationally indistinguishable under decisional binary-secret square noisy unstructured linear algebraic assumption.

Construct a distinguisher $\mathcal{D}_1$ as follows. It receives a challenge $(\mathbf{A}, \mathbf{z})$, where $\mathbf{A} \leftarrow \text{GL}_q(n)$, and either $\mathbf{z} = \mathbf{A}\mathbf{s} + \mathbf{e}^{(1)} + \mathbf{e}^{(2)}$ for $\mathbf{s} \leftarrow H_{n,\gamma}$, or $\mathbf{z} \leftarrow \mathbb{F}_q^n$.

The distinguisher $\mathcal{D}_1$ runs the hybrid distinguisher $\mathcal{B}$. When $\mathcal{B}$ makes its unique KDM query s_i , $\mathcal{D}_1$ samples $\mathbf{s}' \leftarrow H_{n,\gamma}$, samples $\mathbf{R}$ uniformly subject to $\mathbf{s}'^\top \mathbf{R} = -t\mathbf{u}_i^\top$, sets $\mathbf{C} := \mathbf{A} + \mathbf{R}$, samples fresh errors $\mathbf{e}'^{(1)}, \mathbf{e}'^{(2)}$, and answers the KDM query with

$$(\mathbf{C}, \mathbf{C}^\top \mathbf{s}' + \mathbf{e}'^{(1)} + \mathbf{e}'^{(2)}, \mathbf{s}'^\top \mathbf{z}).$$

If $\mathbf{z}$ is a real binary-secret sqNULA vector, then the answer is distributed exactly as H_2 . If $\mathbf{z}$ is uniform, then the answer is distributed exactly as H_3 . Thus the two hybrids are indistinguishable by decisional sqNULA.

Hybrid H_4 . This hybrid is the same as H_3 , except that we change the affine constraint on $\mathbf{R}$. Instead of sampling $\mathbf{R}$ subject to $\mathbf{s}'^\top \mathbf{R} = -t\mathbf{u}_i^\top$, we sample $\mathbf{R}$ subject to $\mathbf{s}'^\top \mathbf{R} = 0$. and set $\mathbf{C} := \mathbf{A} + \mathbf{R}$, and output $(\mathbf{C}, \mathbf{C}^\top \mathbf{s}' + \mathbf{e}'^{(1)} + \mathbf{e}'^{(2)}, \mathbf{s}'^\top \mathbf{z})$, where $\mathbf{z} \leftarrow \mathbb{F}_q^n$.

The only difference between H_3 and H_4 is the distribution of $\mathbf{C}$.

In H_3 , by Claim 1, $\mathbf{C}$ is uniform over $\{\mathbf{C} \in \mathbb{F}_q^{n \times n} : \mathbf{s}'^\top \mathbf{C} \neq -t\mathbf{u}_i^\top\}$. Therefore its distance from uniform over all of $\mathbb{F}_q^{n \times n}$ is q^{-n} .

In H_4 , by Claim 1 with $\mathbf{v} = 0$, $\mathbf{C}$ is uniform over $\{\mathbf{C} \in \mathbb{F}_q^{n \times n} : \mathbf{s}'^\top \mathbf{C} \neq 0\}$. Therefore its distance from uniform over all of $\mathbb{F}_q^{n \times n}$ is also q^{-n} .

Thus, by the triangle inequality we have $\Delta(H_3, H_4) \leq 2q^{-n}$.

Hybrid H_5 . This hybrid is identical to H_4 , except that $\mathbf{z}$ is changed back from uniform to a real binary-secret sqNULA vector. Hybrids H_4 and H_5 are computationally indistinguishable under decisional binary-secret sqNULA.

Observe that Hybrid H_5 is identical to the true encryption of 0 except the distribution of $\mathbf{C}$. Hence, if H_6 denotes the real encryption of 0, then using Claim 1 again, we get $\Delta(H_5, H_6) \leq q^{-n}$.

Combining all hybrid transitions gives us the desired proof. $\square$

Proof of Claim 1. First, note that matrices $\mathbf{R}$ satisfying $\mathbf{r}^\top \mathbf{R} = \mathbf{v}^\top$ exist and are easy to find. Indeed, this is a system of n linear equations in n^2 variables.

Additionally, since $\mathbf{r} \neq 0$, the linear map $\mathbf{x} \mapsto \mathbf{r}^\top \mathbf{x}$ from $\mathbb{F}_q^n$ to $\mathbb{F}_q$ is nonzero and hence surjective. Therefore, for every $v_j \in \mathbb{F}_q$, the equation $\mathbf{r}^\top \mathbf{R}_j = v_j$ has exactly q^{n-1} solutions in $\mathbf{R}_j$. Hence the set of matrices $\mathbf{R}$ satisfying $\mathbf{r}^\top \mathbf{R} = \mathbf{v}^\top$ is a nonempty affine subspace of $\mathbb{F}_q^{n \times n}$ of size $(q^{n-1})^n = q^{n(n-1)}$. We will use this result in the latter part of this proof.

Now fix any $\mathbf{A} \in \text{GL}_q(n)$. Since $\mathbf{R}$ is uniform over the affine subspace $\{\mathbf{R} : \mathbf{r}^\top \mathbf{R} = \mathbf{v}^\top\}$, the matrix $\mathbf{C} = \mathbf{A} + \mathbf{R}$ is uniform over the affine set $\{\mathbf{C} \in \mathbb{F}_q^{n \times n} : \mathbf{r}^\top \mathbf{C} = \mathbf{r}^\top \mathbf{A} + \mathbf{v}^\top\}$.

Now, since $\mathbf{A} \leftarrow \text{GL}_q(n)$ uniformly and $\mathbf{r} \neq 0$, the row vector $\mathbf{r}^\top \mathbf{A}$ is uniform over $\mathbb{F}_q^n \setminus \{0\}$. This implies $\mathbf{r}^\top \mathbf{A} + \mathbf{v}^\top$ is uniform over $\mathbb{F}_q^n \setminus \{\mathbf{v}^\top\}$. Therefore $\mathbf{r}^\top \mathbf{C}$ is uniform over $\mathbb{F}_q^n \setminus \{\mathbf{v}^\top\}$.

Thus $\mathbf{C}$ is uniform over $\{\mathbf{C} \in \mathbb{F}_q^{n \times n} : \mathbf{r}^\top \mathbf{C} \neq \mathbf{v}^\top\}$.

It remains to compute the statistical distance from uniform over all matrices. As seen above, the set $\{\mathbf{C} \in \mathbb{F}_q^{n \times n} : \mathbf{r}^\top \mathbf{C} = \mathbf{v}^\top\}$ has size $q^{n(n-1)}$. Thus the excluded set has relative density $\frac{q^{n(n-1)}}{q^{n^2}} = q^{-n}$. Therefore $\Delta(\mathbf{C}, \text{Unif}(\mathbb{F}_q^{n \times n})) = q^{-n}$. This proves Claim 1.

7 Relating sqNULA to other assumptions via reductions

7.1 Equivalent Dual Assumption

The sqNULA assumption has an equivalently hard dual form. We formally state this relation below:

Assumption 7.1 (Dual Square-Noisy Unstructured Linear Algebraic Assumption(d-sqNULA)). *For all $n \in \mathbb{N}$, prime $q \in \mathbb{N}$, $\sigma = \text{poly}(n)$, $\xi = \text{poly}(n)$, $\delta, \gamma \in (0, 1)$, the decisional dual square-noisy unstructured linear algebraic(decisional-dual sqNULA), formally parameterized by $\text{d-sqNULA}_{n,q,\sigma,\xi,\delta,\gamma}$, states that the following distributions are computationally indistinguishable:*

1. $\left(\mathbf{A} \xleftarrow{\$} \text{GL}_q(n), \mathbf{A}(\mathbf{e}^{(1)} + \mathbf{e}^{(2)}) + \mathbf{s} \pmod{q} \right)$ where $\mathbf{s} \xleftarrow{\$} \mathcal{E}_{n,\xi,\gamma}^n$, $\mathbf{e}^{(1)} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}^n$, and $\mathbf{e}^{(2)} \leftarrow \mathcal{S}_{n,q,\delta}^n$.
2. $\left(\mathbf{A} \xleftarrow{\$} \text{GL}_q(n), \mathbf{b} \xleftarrow{\$} \mathbb{F}_q^n \right)$.

Here,

- $\mathcal{D}_{\mathbb{Z},\sigma}$: For $\sigma \in \mathbb{R}^+$, we let $\mathcal{D}_{\mathbb{Z},\sigma}$ denote the discrete Gaussian distribution over the integer lattice $\mathbb{Z}$ with mean 0 and scale parameter $\sigma > 0$.
- $\mathcal{S}_{n,q,\delta}$: For $n, q \in \mathbb{N}$, $\delta \in (0, 1)$, we define $\mathcal{S}_{n,q,\delta}$ to be the distribution samples 0 with probability $1 - n^{-\delta}$ or a uniformly random element from $\mathbb{F}_q$ with probability $n^{-\delta}$.
- $\mathcal{E}_{n,\xi,\gamma}$: For $n, \xi \in \mathbb{N}$, $\gamma \in (0, 1)$, we define $\mathcal{E}_{n,\xi,\gamma}$ to be the distribution over $\mathbb{Z}$ which samples 0 with probability $1 - n^{-\gamma}$ or a random element from the distribution $\mathcal{D}_{\mathbb{Z},\xi}$ with probability $n^{-\gamma}$.

Lemma 7.2 (Primal and Dual Equivalence). *There are polynomial time reductions in both directions between $\text{sqNULA}_{n,q,\sigma,\delta,\gamma}$ and $\text{d-sqNULA}_{n,q,\sigma,\delta,\gamma}$.*

Proof. We first reduce $\text{sqNULA}_{n,q,\sigma,\delta,\gamma}$ to $\text{d-sqNULA}_{n,q,\sigma,\delta,\gamma}$. On given a sqNULA instance $(\mathbf{A} \in \mathbb{F}_q^{n \times n}, \mathbf{b} \in \mathbb{F}_q^n)$, by standard linear algebra operations, we can compute $\mathbf{A}^{-1} \in \mathbb{F}_q^{n \times n}$ such that $\mathbf{A}^{-1}\mathbf{A} = \mathbf{I}_n$. We then output $(\mathbf{A}^{-1}, \mathbf{A}^{-1}\mathbf{b})$ as the d-sqNULA instance. Since $\mathbf{A}$ is sampled from $\text{GL}_q(n)$, it is invertible, and hence the above operation is always valid.

To see correctness of this reduction (assuming it does not abort), note that because $\mathbf{A}$ is full rank, if $\mathbf{b}$ is uniformly random, then so is $(\mathbf{A}^{-1}, \mathbf{A}^{-1}\mathbf{b})$, showing correctness in the “null” case. In the “planted” case, we have

$$\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}^{(1)} + \mathbf{e}^{(2)} \pmod{q},$$

where $\mathbf{s} \xleftarrow{\$} \mathcal{E}_{n,\xi,\gamma}^n$, $\mathbf{e}^{(1)} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}^n$, and $\mathbf{e}^{(2)} \leftarrow \mathcal{S}_{n,q,\delta}^n$. Multiplying by $\mathbf{A}^{-1}$ on the left gives

$$\mathbf{A}^{-1}\mathbf{b} = \mathbf{s} + \mathbf{A}^{-1}(\mathbf{e}^{(1)} + \mathbf{e}^{(2)}),$$

which matches the distribution of the “planted” case in d-sqNULA. The reduction in the other direction also works identically. □

7.2 Search to decision reduction

We give a reduction from search-sqNULA to decision-sqNULA for polynomial prime q . The reduction preserves the number of samples, the secret dimension, and the secret and error distributions.

Theorem 7.3. *Let q be prime with $q = \text{poly}(n)$, and let $\sigma, \xi = \text{poly}(n)$ and $\delta, \gamma \in (0, 1)$. Suppose a polynomial-time adversary $\mathcal{A}$ distinguishes a sqNULA sample from random with advantage at least $\varepsilon \geq 1/\text{poly}(n)$. Then there is an algorithm $\mathcal{B}^{\mathcal{A}}$ outputs $\mathbf{s} \pmod{q}$ with non-negligible probability.*

Intuition. We will use the Goldreich-Levin hardcore lemma over large field [GRS00] for this reduction. In detail, we first convert $\mathcal{A}$ into an adversary that upon input a planted sample $(\mathbf{A}, \mathbf{b})$, and a uniform $\mathbf{u} \in \mathbb{F}_q^n$, can predict $\mathbf{u}^\top \mathbf{s}$ with probability noticeably larger than $1/q$, and then convert this predictor to an inverter of the search problem via the Goldreich-Levin over large fields proof technique. The latter part is a black-box invocation of the result from [GRS00].

Proof. Assume that $\mathcal{A}$ outputs 1 in the planted case. We will now build a predictor $\mathcal{P}$ for the hardcore bit. On input a planted sample $(\mathbf{A}, \mathbf{b})$, and $\mathbf{u}$, it does as follows.

1. Sample independent $\mathbf{r} \leftarrow \mathbb{F}_q^n$ and $\beta \leftarrow \mathbb{F}_q$.
2. Set $\mathbf{A}' = \mathbf{A} + \mathbf{r}\mathbf{u}^\top$ and $\mathbf{b}' = \mathbf{b} + \beta\mathbf{r}$.
3. If $\mathbf{A}'$ is singular, set $a = 0$. Otherwise set $a \leftarrow \mathcal{A}(\mathbf{A}', \mathbf{b}')$.
4. If $a = 1$, output β . If $a = 0$, output a uniformly random element of $\mathbb{F}_q \setminus \{\beta\}$.

If $\mathcal{P}$ is correct with probability noticeably more than $\frac{1}{q}$, then we can directly apply the hardcore lemma over large fields [GRS00], which states for a one-way function f , given $\mathbf{u}$ and $f(\mathbf{s})$, $\mathbf{u}^\top \mathbf{s}$ must be unpredictable. Thus, the contrapositive implies that a predictor for $\mathbf{u}^\top \mathbf{s}$ implies the existence of an inverter $\mathcal{B}$ for $f(\mathbf{s})$. Moreover, the proof is constructive so we get an explicit breaker for the search problem.

Thus, it suffices to prove the correctness of the predictor.

We begin by showing that the transformed samples $(\mathbf{A}', \mathbf{b}')$ are indeed valid sqNULA samples. In particular, we prove the following lemma:

Lemma 7.4. *Sample $\mathbf{A} \leftarrow \text{GL}_q(n)$ and independent uniform $\mathbf{r}, \mathbf{u} \leftarrow \mathbb{F}_q^n$. Then,*

$$\Pr[\mathbf{A} + \mathbf{r}\mathbf{u}^\top \in \text{GL}_q(n)] = 1 - \frac{1}{q} + \frac{1}{q^{n+1}}.$$

Conditioned on invertibility, $\mathbf{A} + \mathbf{r}\mathbf{u}^\top$ is uniform over $\text{GL}_q(n)$.

Proof. Set $\mathbf{v} = \mathbf{A}^{-1}\mathbf{r}$. Since $\mathbf{r}$ is uniform, $\mathbf{v}$ is uniform and independent of $\mathbf{A}$ and $\mathbf{u}$. Using matrix determinant properties, we know that

$$\det(\mathbf{A} + \mathbf{r}\mathbf{u}^\top) = \det(\mathbf{A})(1 + \mathbf{u}^\top \mathbf{v}).$$

If $\mathbf{v} = \mathbf{0}$, it is easy to see that the resultant matrix is invertible. Otherwise, $\mathbf{u}^\top \mathbf{v}$ is uniform over $\mathbb{F}_q$, so the matrix is invertible with probability $1 - 1/q$, i.e., when $\mathbf{u}^\top \mathbf{v} \neq 1$. Therefore,

$$\Pr[\mathbf{A} + \mathbf{r}\mathbf{u}^\top \in \text{GL}_q(n)] = \frac{1}{q^n} + \left(1 - \frac{1}{q^n}\right) \left(1 - \frac{1}{q}\right) = 1 - \frac{1}{q} + \frac{1}{q^{n+1}}.$$

For the proof of uniformity, note that

$$\mathbf{A} + \mathbf{r}\mathbf{u}^\top = \mathbf{A}(\mathbf{I}_n + \mathbf{v}\mathbf{u}^\top).$$

Observe that conditioned on $\mathbf{A} + \mathbf{r}\mathbf{u}^\top$ being invertible, it must be that $\mathbf{I}_n + \mathbf{v}\mathbf{u}^\top$ is invertible.

Now fix any $(\mathbf{v}, \mathbf{u})$ for which $\mathbf{I}_n + \mathbf{v}\mathbf{u}^\top$ is invertible. The matrix $\mathbf{A}$ remains uniform. For any desired output $\mathbf{B} \in \text{GL}_q(n)$, exactly one value of $\mathbf{A}$ produces it: $\mathbf{A} = \mathbf{B}(\mathbf{I}_n + \mathbf{v}\mathbf{u}^\top)^{-1}$. Thus, every invertible output $\mathbf{B}$ has probability $1/|\text{GL}_q(n)|$.

This holds for every fixed pair $(\mathbf{v}, \mathbf{u})$ satisfying the invertibility condition. Averaging over those pairs therefore gives the same uniform distribution.

The invertibility event depends only on $(\mathbf{v}, \mathbf{u})$, which is independent of $\mathbf{A}$. For every fixed pair satisfying this event, $\mathbf{I}_n + \mathbf{v}\mathbf{u}^\top$ is invertible. Right multiplication by this fixed matrix preserves the uniform distribution on $\text{GL}_q(n)$. Averaging over such pairs proves the claimed conditional uniformity. $\square$

The case $\beta = \mathbf{u}^\top \mathbf{s}$. Since β is independent and uniform, the event $\beta = \mathbf{u}^\top \mathbf{s}$ has probability $1/q$

In this case,

$$\begin{aligned} \mathbf{b}' &= \mathbf{A}'\mathbf{s} + \mathbf{e}^{(1)} + \mathbf{e}^{(2)} + (\beta - \mathbf{u}^\top \mathbf{s})\mathbf{r} \\ &= \mathbf{A}'\mathbf{s} + \mathbf{e}^{(1)} + \mathbf{e}^{(2)}. \end{aligned}$$

Lemma 7.4 shows that, conditioned on invertibility, $\mathbf{A}'$ is uniform over $\text{GL}_q(n)$ and independent of the secret and errors. Since $a = 0$ on singular matrices,

$$\begin{aligned} \Pr[a = 1 \mid \beta = \mathbf{u}^\top \mathbf{s}] \\ = \Pr[\mathbf{A}' \in \text{GL}_q(n)] \Pr[\mathcal{A}(\mathbf{A}', \mathbf{A}'\mathbf{s} + \mathbf{e}^{(1)} + \mathbf{e}^{(2)}) = 1]. \end{aligned} \tag{1}$$

The case $\beta \neq \mathbf{u}^\top \mathbf{s}$. Let

$$c = \beta - \mathbf{u}^\top \mathbf{s}, \mathbf{v} = c\mathbf{r}, \mathbf{u}' = c^{-1}\mathbf{u}.$$

Since β is independent and uniform, conditioned on this case, c is uniform over $\mathbb{F}_q \setminus \{0\}$ and independent of the original matrix, secret, errors, and vectors $\mathbf{r}, \mathbf{u}$. For each fixed nonzero c , multiplication by c and by c^{-1} preserves uniformity. Thus, $\mathbf{v}, \mathbf{u}'$ are independent uniform vectors, independent of the original matrix, secret, and errors. Moreover observe that,

$$\mathbf{A}' = \mathbf{A} + \mathbf{v}\mathbf{u}'^\top, \mathbf{b}' = \mathbf{A}'\mathbf{s} + \mathbf{e}^{(1)} + \mathbf{e}^{(2)} + \mathbf{v}.$$

We show that in this case, will all but negligible probability, $a = 1$ when $\mathcal{A}$ outputs a 1 on a random sample and $\mathbf{A}'$ is invertible.

We first condition on $\mathbf{v} \neq \mathbf{0}$ and on $\mathbf{A}'$ being invertible. For any fixed invertible $\mathbf{B}$ and nonzero $\mathbf{v}$, setting $\mathbf{A}' = \mathbf{B}$ determines $\mathbf{A} = \mathbf{B} - \mathbf{v}\mathbf{u}'^\top$. By the matrix determinant properties, this matrix is invertible exactly when $\mathbf{u}'^\top \mathbf{B}^{-1} \mathbf{v} \neq 1$. Since $\mathbf{B}^{-1} \mathbf{v} \neq \mathbf{0}$, the number of choices of $\mathbf{u}'$ satisfying this condition is the same for every such pair $(\mathbf{B}, \mathbf{v})$. Thus, under the conditioning, $\mathbf{A}'$ is uniform over $\text{GL}_q(n)$ and $\mathbf{v}$ is independently uniform over $\mathbb{F}_q^n \setminus \{\mathbf{0}\}$. They are also independent of the secret and errors. A uniform nonzero vector has statistical distance q^{-n} from a uniform vector in $\mathbb{F}_q^n$. Therefore, under this conditioning, $(\mathbf{A}', \mathbf{b}')$ has statistical distance at most q^{-n} from the uniform experiment, in which the matrix is uniform over $\text{GL}_q(n)$ and the response is an independent uniform vector in $\mathbb{F}_q^n$.

The excluded event $\mathbf{v} = \mathbf{0}$ has probability q^{-n} and always gives an invertible $\mathbf{A}' = \mathbf{A}$. Thus, the above conditioning yields a loss of at most q^{-n} . Moreover, $a = 0$ whenever $\mathbf{A}'$ is singular, thus it follows that

$$\begin{aligned} &\left| \Pr[a = 1 \mid \beta \neq \mathbf{u}^\top \mathbf{s}] \right. \\ &\quad \left. - \Pr[\mathbf{A}' \in \text{GL}_q(n)] \Pr[\mathcal{A}(\mathbf{A}, \mathbf{b}) = 1] \right| \leq cq^{-n}, \end{aligned} \tag{2}$$

for some constant c .

Probability of a correct prediction. Assume that $\mathcal{A}$ accepts the planted distribution with probability at least ε greater than its acceptance probability in the uniform experiment; otherwise, complement its output.

If $\beta = \mathbf{u}^\top \mathbf{s}$, the predictor is correct exactly when $a = 1$. Otherwise, it is correct with probability $1/(q-1)$ when $a = 0$. Hence

$$\begin{aligned} & \Pr[\text{predictor outputs } \mathbf{u}^\top \mathbf{s}] \\ &= \frac{1}{q} \Pr[a = 1 \mid \beta = \mathbf{u}^\top \mathbf{s}] + \frac{q-1}{q} \frac{1}{q-1} \Pr[a = 0 \mid \beta \neq \mathbf{u}^\top \mathbf{s}] \\ &= \frac{1}{q} + \frac{1}{q} \left(\Pr[a = 1 \mid \beta = \mathbf{u}^\top \mathbf{s}] - \Pr[a = 1 \mid \beta \neq \mathbf{u}^\top \mathbf{s}] \right) \\ &\geq \frac{1}{q} + \frac{\varepsilon}{q} \Pr[\mathbf{A}' \in \text{GL}_q(n)] - \frac{c}{q^{n+1}}. \end{aligned} \tag{3}$$

We know that

$$\Pr[\mathbf{A}' \in \text{GL}_q(n)] \geq 1 - \frac{1}{q} \geq \frac{1}{2}.$$

Thus

$$\Pr[\text{predictor outputs } \mathbf{u}^\top \mathbf{s}] \geq \frac{1}{q} + \frac{\varepsilon}{2q} - \frac{c}{q^{n+1}}.$$

In particular, when ε is inverse polynomial, for all sufficiently large n ,

$$\Pr[\text{predictor outputs } \mathbf{u}^\top \mathbf{s}] \geq \frac{1}{q} + \frac{\varepsilon}{4q}.$$

This completes the proof. □

Remark 7.5. The same reduction carries over to binary secret sqNULA as well.

7.3 Reduction from LWE

Next we show that sqNULA is at least as hard as LWE via a concrete reduction. For simplicity, we will first consider the sqNULA distribution conditioned on exactly k coordinates of the secret sampled from the Gaussian distribution, and the rest set to 0. Equivalently, choose a uniform k -element set $J \subseteq [n]$, sample the coordinates of $\mathbf{s}_J$ independently from $\mathcal{D}_{\mathbb{Z}, \xi}$, and set the other coordinates to zero.

Theorem 7.6. Suppose $1 \leq k \leq n$ and $\xi = \sigma$. A distinguisher of advantage ε for sqNULA conditioned on exactly k coordinates sampled from $\mathcal{D}_{\mathbb{Z}, \xi}$ can be used to distinguish $\text{LWE}_{k, n+k, q, \mathcal{D}_{\mathbb{Z}, \sigma}}$ with advantage

$$\varepsilon \left(1 - \frac{2}{q} \right)^2.$$

Proof. Let the LWE instance have $n+k$ samples and secret dimension k such that it can be partitioned into its first k rows and remaining n rows. Specifically, write these as $(\mathbf{A}_0, \mathbf{b}_0)$ and $(\mathbf{A}_1, \mathbf{b}_1)$. The reduction works as follows:

1. Abort if $\mathbf{A}_0$ is singular.
2. Otherwise define

$$\mathbf{B} = -\mathbf{A}_1 \mathbf{A}_0^{-1}, \mathbf{b}' = \mathbf{b}_1 - \mathbf{A}_1 \mathbf{A}_0^{-1} \mathbf{b}_0.$$

3. Abort if $\mathbf{B}$ is not full column rank.
4. Create an $n \times n$ matrix $\mathbf{A}$ in the following way: Choose a uniform k -element set $J \subseteq [n]$ and insert the columns of $\mathbf{B}$ at those positions in the set J . Sample the remaining $n-k$ columns of $\mathbf{A}$ uniformly until the resulting matrix belongs to $\text{GL}_q(n)$.
5. Finally draw an independent $\mathbf{e}^{(2)} \leftarrow \mathcal{S}_{n, q, \delta}^n$ and give $(\mathbf{A}, \mathbf{b}' + \mathbf{e}^{(2)})$ to the sqNULA distinguisher.

Correctness: For a planted LWE input, let $\mathbf{s}'$ be its uniform secret. Let its two Gaussian error blocks as $\mathbf{e}_0^{(1)}$ and $\mathbf{e}_1^{(1)}$. Then we know that

$$\mathbf{b}_0 = \mathbf{A}_0 \mathbf{s}' + \mathbf{e}_0^{(1)}, \mathbf{b}_1 = \mathbf{A}_1 \mathbf{s}' + \mathbf{e}_1^{(1)}, \mathbf{b}' = \mathbf{B} \mathbf{e}_0^{(1)} + \mathbf{e}_1^{(1)}.$$

Conditional on $\mathbf{A}_0$ being invertible, $\mathbf{B}$ is uniform in $\mathbb{F}_q^{n \times k}$ and independent of both error blocks. Conditional on $\mathbf{B}$ being full column rank, it is uniform among full-column-rank matrices. After this, it is easy to see that the completed $\mathbf{A}$ is uniform in $\text{GL}_q(n)$ and independent of $J, \mathbf{e}_0^{(1)}, \mathbf{e}_1^{(1)}$. Set $\mathbf{s}_J = \mathbf{e}_0^{(1)}$ and the other coordinates to zero. Thus, $\mathbf{b}' = \mathbf{A} \mathbf{s}_J + \mathbf{e}_1^{(1)} + \mathbf{e}^{(2)}$. This is exactly a valid sqNULA sample, because $\xi = \sigma$.

For a uniform input response, $\mathbf{b}'$ remains uniform conditional. Adding $\mathbf{e}^{(2)}$ preserves uniformity.

Thus, it suffices to ensure that the rank checks satisfy with non-negligible probability. We use the following rank results:

- A uniformly random $d \times d$ matrix over $\mathbb{F}_q$ is invertible with probability $\prod_{i=1}^d (1 - q^{-i}) \geq 1 - \frac{2}{q}$.
- A uniformly random $n \times k$ matrix, where $k \leq n$, has full column rank with probability $\prod_{i=0}^{k-1} (1 - q^{i-n}) \geq \prod_{i=0}^n (1 - q^{i-n}) = \prod_{i=0}^n (1 - q^{-i}) \geq 1 - \frac{2}{q}$.

This completes the proof. $\square$

Next, we extend the above reduction to the true sqNULA secret distribution.

Theorem 7.7. Fix $0 < \gamma < 1$, let $k = \lfloor n^{1-\gamma}/2 \rfloor$, and suppose $\xi = \sigma$. A distinguisher of advantage ε for $\text{sqNULA}_{n,q,\sigma,\xi,\delta,\gamma}$ gives a distinguisher for $\text{LWE}_{k,2n,q,\mathcal{D}_{\mathbb{Z},\sigma}}$ with advantage at least

$$\left(1 - \frac{2}{q}\right)^2 (\varepsilon - \text{negl}(n)). \quad (4)$$

Proof. Let the LWE instance be $(\mathbf{A}, \mathbf{b})$ where $\mathbf{A}$ has $2n$ rows and k columns. The reduction works as follows:

- Sample $j \xleftarrow{\$} \text{Bin}(n, n^{-\gamma})$. Abort if $j < k$.
- Draw a uniform $2n \times (j - k)$ matrix $\mathbf{B}$ and an independent uniform vector $\mathbf{s}' \in \mathbb{F}_q^{j-k}$.
- Set $\mathbf{A}^* = [\mathbf{A} \mid \mathbf{B}]$, $\mathbf{b}^* = \mathbf{b} + \mathbf{B} \mathbf{s}'$. Observe that $\mathbf{b} + \mathbf{B} \mathbf{s}' = [\mathbf{A} \mid \mathbf{B}][\mathbf{s} \mid \mathbf{s}'] + \mathbf{e}$ if $\mathbf{b}$ is an LWE sample of the form $\mathbf{A} \mathbf{s} + \mathbf{e}'$.

The above is exactly a j -dimensional LWE instance in the planted case and has a uniform response in the uniform case. Keep its first $n + j \leq 2n$ rows and now apply Theorem 7.6. Observe that j is no longer a fixed value, rather a random variable that induces the exact same sparsity distribution as an sqNULA secret.

The rest of the proof follows naturally. $\square$

Remark 7.8. We get a reduction from binary secret sqNULA to binary secret LWE using the same proof technique above. Following the fact that binary secret LWE is at least as hard as LWE [Mic18], this shows us that binary secret sqNULA is at least as hard as LWE with appropriate parameters.

7.4 Relating binary secret and small secret sqNULA

Lemma 7.9. *For every PPT distinguisher for the original sqNULA, there exists a PPT distinguisher for the binary-secret variant whose probability of one is*

$$n^{-\gamma} \left(1 - \Pr_{x \leftarrow \mathcal{D}_{\mathbb{Z}, \xi}} [x = 0 \bmod q] \right).$$

The error distributions and all other parameters are unchanged.

Proof. Sample a diagonal matrix $\mathbf{R}$ whose diagonal entries are independent samples from $\mathcal{D}_{\mathbb{Z}, \xi}$ conditioned on being nonzero. This can be done in polynomial time. Map $(\mathbf{A}, \mathbf{b})$ to $(\mathbf{A}\mathbf{R}^{-1}, \mathbf{b})$. If $\mathbf{s}'$ is the unknown binary secret, the transformed secret is $\mathbf{s} = \mathbf{R}\mathbf{s}'$. For each nonzero $x \in \mathbb{F}_q$, we have

$$\Pr[s_i = x] = n^{-\gamma} \Pr_{y \leftarrow \mathcal{D}_{\mathbb{Z}, \xi}} [y = x \bmod q].$$

This is exactly the coordinate with probability distribution $\mathcal{E}_{n, \xi, \gamma}$. Since $\mathbf{A}$ is uniform invertible, $\mathbf{A}\mathbf{R}^{-1}$ is uniform invertible and independent of $\mathbf{R}, \mathbf{s}'$. Thus, the distinguishing advantage of the binary secret adversary is same as that of the standard sqNULA adversary. $\square$

7.5 Reducing sqNULA with subexponential error to modulus ratio to a variant of LW2E

Reduction from sqNULA to square LW2E with short secrets via secret flooding. We give a concrete reduction from search-sqNULA to a *square* ($m = n$) search-LW2E variant with *short* secrets. The reduction follows the same noise-flooding paradigm as Lemma 7.11, except that we flood the *secret* (rather than the error). Consequently, we require a superpolynomial flooding width and therefore a superpolynomial modulus.

Short-secret square-LW2E. Fix parameters $(n, q, \sigma, \delta, \sigma^*)$ where q is prime and $m = n$. An instance of $\text{LW2E}_{n, q, \sigma, \delta, \sigma^*}^{\text{sh}}$ consists of

$$\mathbf{A} \xleftarrow{\$} \mathbb{F}_q^{n \times n}, \quad \mathbf{b} = \mathbf{A}\mathbf{u} + \mathbf{e}^{(1)} + \mathbf{e}^{(2)} \pmod{q},$$

where $\mathbf{u} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma^*}^n$ (then reduced modulo q), $\mathbf{e}^{(1)} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^n$, and $\mathbf{e}^{(2)} \leftarrow \mathcal{S}_{n, q, \delta}^n$. The goal is to recover $\mathbf{u} \in \mathbb{F}_q^n$ given $(\mathbf{A}, \mathbf{b})$.

Lemma 7.10 (Search-sqNULA reduces to short-secret square-LW2E). *Let λ be the security parameter and let $n = n(\lambda)$. Let $\sigma = \text{poly}(\lambda)$, $\xi = \text{poly}(\lambda)$, and constants $\delta, \gamma \in (0, 1)$. Let $q = q(\lambda)$ be prime and superpolynomial in λ , and let $\sigma^* = \sigma^*(\lambda)$ be superpolynomial in λ such that*

$$\sigma^* = 2^{\omega(\log \lambda)}, \quad q = \text{poly}(\sigma^*) \geq 2\sigma^*, \quad \frac{\sigma^*}{\xi} = 2^{\omega(\log \lambda)}.$$

If there is a PPT algorithm $\mathcal{A}$ that solves search (decision) $\text{LW2E}_{n, q, \sigma, \delta, \sigma^*}^{\text{sh}}$ with non-negligible probability, then there is a PPT algorithm $\mathcal{B}$ that solves search (decision)-sqNULA $_{n, q, \sigma, \xi, \gamma, \delta}$ with non-negligible probability.

Proof. Let $(\mathbf{A}, \mathbf{b})$ be a search-sqNULA instance, i.e.,

$$\mathbf{A} \xleftarrow{\$} \text{GL}_q(n), \quad \mathbf{s} \xleftarrow{\$} \mathcal{E}_{n, \xi, \gamma}^n, \quad \mathbf{e}^{(1)} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^n, \quad \mathbf{e}^{(2)} \leftarrow \mathcal{S}_{n, q, \delta}^n,$$

and

$$\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}^{(1)} + \mathbf{e}^{(2)} \pmod{q}.$$

We construct an input to $\mathcal{A}$ as follows.

Sample an independent flooding vector $\mathbf{s}' \leftarrow \mathbf{D}_{\mathbb{Z}, \sigma^*}^n$ and define

$$\mathbf{b}' := \mathbf{b} + \mathbf{A}\mathbf{s}' \pmod{q}.$$

Then

$$\mathbf{b}' = \mathbf{A}(\mathbf{s} + \mathbf{s}') + \mathbf{e}^{(1)} + \mathbf{e}^{(2)} \pmod{q}.$$

Let $\mathbf{u} := \mathbf{s} + \mathbf{s}' \in \mathbb{F}_q^n$. Feed $(\mathbf{A}, \mathbf{b}')$ to $\mathcal{A}$, obtaining $\hat{\mathbf{u}} \in \mathbb{F}_q^n$, and output

$$\hat{\mathbf{s}} := \hat{\mathbf{u}} - \mathbf{s}' \pmod{q}.$$

Correctness. We will use the following lemma in the proof:

Lemma 7.11 (Gaussian Smudging Lemma [LLW21]). *Let $n \in \mathbb{N}$, $\forall \sigma > \omega(\sqrt{\log n})$, for any $\mathbf{c} \in \mathbb{Z}^n$,*

$$\Delta(\mathcal{D}_{\mathbb{Z}^n, \sigma}, \mathcal{D}_{\mathbb{Z}^n, \sigma, \mathbf{c}}) \leq \frac{\|\mathbf{c}\|_2}{\sigma}.$$

The pair $(\mathbf{A}, \mathbf{b}')$ has exactly the LW2E two-error form with secret $\mathbf{u}$ and errors distributed as in $\text{LW2E}_{n, q, \sigma, \delta, \sigma^*}^{\text{sh}}$, except that the secret is $\mathbf{s} + \mathbf{s}'$ rather than $\mathbf{s}'$.

We now argue that $\mathbf{s} + \mathbf{s}'$ is statistically close to $\mathbf{s}'$. By the definition of $\mathcal{E}_{n, \xi, \gamma}^n$, each coordinate of $\mathbf{s}$ is supported on integers of magnitude at most $\xi \cdot \text{poly}(\lambda)$ except with negligible probability. Since $\sigma^*/\xi = 2^{\omega(\log \lambda)}$, the shift by $\mathbf{s}$ is exponentially smaller than the width σ^* . A standard discrete-Gaussian flooding/smoothing argument (Lemma 7.11) applied to the *secret* implies that the statistical distance between the distributions of $\mathbf{s}'$ and $\mathbf{s}' + \mathbf{s}$ is negligible in λ .

Therefore, the distribution of $(\mathbf{A}, \mathbf{b}')$ produced by $\mathcal{B}$ is statistically negligibly close to a genuine instance of $\text{LW2E}_{n, q, \sigma, \delta, \sigma^*}^{\text{sh}}$. It follows that $\mathcal{A}$ outputs $\hat{\mathbf{u}} = \mathbf{u}$ with non-negligible probability, and in that case $\hat{\mathbf{s}} = \hat{\mathbf{u}} - \mathbf{s}' = \mathbf{s} \pmod{q}$, which solves the search-sqNULA instance. $\square$

Remark 7.12. Note that all of our applications use sqNULA with parameters that have a polynomial modulus to noise ratio, and therefore do not admit the above reduction.

7.6 An equivalent formulation

Theorem 7.13. *The original sqNULA distribution is equivalent to the following noisy linear experiment. Choose independent $\mathbf{B}_0, \mathbf{B}_1 \xleftarrow{\$} \text{GL}_q(n)$ and a uniform $\mathbf{s}' \in \mathbb{F}_q^n$. Independently draw $\mathbf{s}, \mathbf{e}^{(1)}, \mathbf{e}^{(2)}$ from the appropriate distributions for sqNULA, and give*

$$\left(\begin{pmatrix} \mathbf{B}_0 \\ \mathbf{B}_1 \end{pmatrix}, \begin{pmatrix} \mathbf{B}_0 \\ \mathbf{B}_1 \end{pmatrix} \mathbf{s}' + \begin{pmatrix} \mathbf{s} \\ \mathbf{e}^{(1)} + \mathbf{e}^{(2)} \end{pmatrix} \right). \quad (5)$$

The goal is to distinguish this from uniform in $\mathbb{F}_q^{2n}$. There is no advantage loss in either direction.

Proof. Given $(\mathbf{A}, \mathbf{b})$, draw independent $\mathbf{U} \xleftarrow{\$} \text{GL}_q(n)$ and $\mathbf{r} \xleftarrow{\$} \mathbb{F}_q^n$, and output

$$\left(\begin{pmatrix} \mathbf{U} \\ -\mathbf{A}\mathbf{U} \end{pmatrix}, \begin{pmatrix} \mathbf{0} \\ \mathbf{b} \end{pmatrix} + \begin{pmatrix} \mathbf{U} \\ -\mathbf{A}\mathbf{U} \end{pmatrix} \mathbf{r} \right).$$

The two matrix blocks are independent uniform invertible matrices. In the planted case, its response is

$$\begin{pmatrix} \mathbf{U} \\ -\mathbf{A}\mathbf{U} \end{pmatrix} (\mathbf{r} - \mathbf{U}^{-1}\mathbf{s}) + \begin{pmatrix} \mathbf{s} \\ \mathbf{e}^{(1)} + \mathbf{e}^{(2)} \end{pmatrix}.$$

Conditional on all original secrets, errors, and matrices, $\mathbf{r} - \mathbf{U}^{-1}\mathbf{s}$ is uniform. It is therefore independent of the public matrix and both error blocks. In the uniform case, $(\mathbf{r}, \mathbf{b}) \mapsto (\mathbf{U}\mathbf{r}, \mathbf{b} - \mathbf{A}\mathbf{U}\mathbf{r})$ is a bijection, so the full response is uniform.

For the other direction, partition the response in (5) as $(\mathbf{b}_0, \mathbf{b}_1)$ and form

$$\mathbf{A} = -\mathbf{B}_1\mathbf{B}_0^{-1}, \quad \mathbf{b} = \mathbf{b}_1 - \mathbf{B}_1\mathbf{B}_0^{-1}\mathbf{b}_0.$$

Thus we get correct distribution mapping in both directions. It is easy to verify that the uniform case also maps appropriately for both sides. $\square$

8 Analyzing the hardness of sqNULA

8.1 Hardness of sqNULA in the presence of Approximate SVP breaking oracle

We now define some lattice notation and the approximate Shortest Vector Problem (SVP) problem. For a lattice $\mathcal{L}$, we use the notation

$$\lambda_1(\mathcal{L}) := \min_{\mathbf{v} \in \mathcal{L} \setminus \{\mathbf{0}\}} \|\mathbf{v}\|_2$$

to denote the minimum distance of the lattice $\mathcal{L}$.

Definition 8.1. *The SVP_β problem: Given an instance $\mathbf{A} \in \mathbb{Z}^{m \times m}$ where $\mathbf{A}$ is a basis for a lattice $\mathcal{L} \subseteq \mathbb{Z}^m$, find some $\mathbf{x} \in \mathcal{L} \setminus \{\mathbf{0}\}$ such that $\|\mathbf{x}\|_2 \leq \beta \cdot \lambda_1(\mathcal{L})$, for $\beta > 1$.*

We now present a detailed discussion of how the plausibility of a reduction from sqNULA to SVP_β relates to the various choices of parameters. We will also compare this to the parameters that we use in our constructions above. Let us recall notations for some useful parameters:

- Secret dimension of sqNULA: n .
- Number of sqNULA samples: n .
- sqNULA secret sparsity parameter: γ .
- sqNULA secret smallness parameter: ξ .
- Smallness parameter for sqNULA noise $\mathbf{e}_1$: σ .
- Sparsity parameter for sqNULA noise $\mathbf{e}_2$: δ .
- Prime modulo: q .

Here is a potential reduction idea:

Given an sqNULA sample $(\mathbf{A} \in \mathbb{F}_q^{n \times n}, \mathbf{b} \in \mathbb{F}_q^n)$, a distinguisher $\mathcal{D}$ does the following:

1. Set up an integer lattice with basis matrix

$$\mathbf{B} = \begin{pmatrix} qI_n & -\mathbf{A} & \mathbf{b} \\ 0_{n \times n} & I_n & 0_{n \times 1} \\ 0_{1 \times n} & 0_{1 \times n} & 1 \end{pmatrix} \in \mathbb{Z}^{(2n+1) \times (2n+1)}.$$

2. Send $\mathbf{B}$ to the SVP_β solver and let $\mathbf{v}$ be its response.

3. If $\|\mathbf{v}\| \leq \beta \cdot \sqrt{(8n^{1-\delta}q^2 + n\sigma^{2(1+\nu)} + 2n^{1-\gamma}\xi^{2(1+\nu)})}$, output 1, else output 0.

In order for this reduction to be correct, we want that

- when $\mathbf{b}$ is an sqNULA sample, then $\|\mathbf{v}\|^2 \leq \beta^2 \cdot (8n^{1-\delta}q^2 + n\sigma^{2(1+\nu)} + 2n^{1-\gamma}\xi^{2(1+\nu)})$, and
- when $\mathbf{b}$ is random, then $\|\mathbf{v}\|^2 > \beta^2 \cdot (8n^{1-\delta}q^2 + n\sigma^{2(1+\nu)} + 2n^{1-\gamma}\xi^{2(1+\nu)})$.

Observe that when $\mathbf{b}$ is an sqNULA sample, i.e., $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}_1 + \mathbf{e}_2$, then the lattice generated by $\mathbf{B}$ contains the vector $(\mathbf{e}_1 + \mathbf{e}_2 \| \mathbf{s} \| 1) \in \mathbb{F}_q^{2n+1}$. To see why this is the case, we know that $-\mathbf{A}\mathbf{s} + \mathbf{b} \equiv \mathbf{e}_1 + \mathbf{e}_2 \pmod{q}$, and therefore over the integers, $-\mathbf{A}\mathbf{s} + \mathbf{b} = \mathbf{e}_1 + \mathbf{e}_2 + q*$ for some value of $*$. Thus, $\mathbf{B}\mathbf{x} = (\mathbf{e}_1 + \mathbf{e}_2 \| \mathbf{s} \| 1)$ where $\mathbf{x} = (* \| \mathbf{s} \| 1)$. Now,

$$\left\| (\mathbf{e}_1 + \mathbf{e}_2 \| \mathbf{s} \| 1) \right\|_2^2 = \|\mathbf{e}_1 + \mathbf{e}_2\|_2^2 + \|\mathbf{s}\|_2^2 + 1$$

We know that with all but negligible probability every entry of $\mathbf{e}_1$, i.e., $e_{1,i} < \sigma^{1+\mu}$ for some positive constant μ and every entry of $\mathbf{s}$, i.e., $s_i < \xi^{1+\nu}$ for some positive constant ν . Moreover, with overwhelming probability $\mathbf{e}_2$ has at most $2n^{1-\delta}$ many non-zero entries and $\mathbf{s}$ has at most $2n^{1-\gamma}$ many non-zero entries. Therefore,

$$\begin{aligned}\|\mathbf{e}_1 + \mathbf{e}_2\|_2^2 + \|\mathbf{s}\|_2^2 + 1 &\leq 2n^{1-\delta}(4q^2) + n\sigma^{2(1+\nu)} + 2n^{1-\gamma}\xi^{2(1+\nu)} + 1 \\ &\leq 8n^{1-\delta}q^2 + n\sigma^{2(1+\nu)} + 2n^{1-\gamma}\xi^{2(1+\nu)}\end{aligned}$$

We have to compare this to the norm of the shortest vector in the lattice generated by $\mathbf{B}$ when $\mathbf{b}$ is random. A good and widely accepted estimate of the value for the shortest vector in this case is given by the Gaussian Heuristic [ADD⁺20] that states

$$\lambda_1(\mathcal{L}) \approx \sqrt{\frac{2n+1}{2\pi e}} \text{Vol}(\mathcal{L})^{\frac{1}{2n+1}}.$$

We can use the following lemma about the volume of a lattice.

Lemma 8.2. (See Claim 2.12 in [Pei13]) The volume of a lattice $\mathcal{L}$, given by $\text{Vol}(\mathcal{L})$ is $\det(\mathcal{L}) = |\det(\mathbf{B})|$, where $\mathbf{B}$ is an basis of $\mathcal{L}$.

The above expression then becomes

$$\lambda_1(\mathcal{L}) \approx \sqrt{\frac{2n+1}{2\pi e}} q^{\frac{n}{2n+1}}.$$

This value is approximately $\sqrt{2n+1}q^{\frac{1}{2}}$. In fact, we also know a well-known result about λ_1 for a lattice.

Lemma 8.3. (See Corollary 2.15 in [Pei13]) For any lattice $\mathcal{L} \subset \mathbb{Z}^d$, we have $\lambda_1(\mathcal{L}) \leq \sqrt{d} \cdot \det(\mathcal{L})^{1/d}$.

Thus, we know that λ_1 of the lattice spanned by $\mathbf{B}$ when $\mathbf{b}$ is random is at most $\sqrt{2n+1}q^{\frac{n}{2n+1}} \leq \sqrt{2n+1}q^{\frac{1}{2}}$. Via the above mentioned natural approach of using the SVP_β solving oracle,

- there is plausibly no attack when

$$(2n+1)q \leq 8n^{1-\delta}q^2 + n\sigma^{2(1+\mu)} + 2n^{1-\gamma}\xi^{2(1+\nu)}.$$

- there is a potential attack using an SVP_β breaking oracle when

$$(2n+1)q \geq \beta^2 \cdot 8n^{1-\delta}q^2 + n\sigma^{2(1+\mu)} + 2n^{1-\gamma}\xi^{2(1+\nu)}.$$

Now, let us work out the parameters that will satisfy the above constraints. Say $q = n^\tau$, for some $\tau > 0$.

- **Case 1.** $\tau < \delta$. We will assume that $\sigma^2 = n^{\epsilon_1}$ and $\xi^2 = n^{\epsilon_2}$. Therefore, for there to be an attack, we need

$$\begin{aligned}(2n+1)q &\geq \beta^2 \cdot 8n^{1-\delta}q^2 + n\sigma^{2(1+\mu)} + 2n^{1-\gamma}\xi^{2(1+\nu)} \\ \implies (2n+1)q &\geq \beta^2 \cdot 8n^{1-\delta}q^2 + n^{1+\epsilon_1(1+\mu)} + 2n^{1-\gamma+\epsilon_2(1+\nu)}\end{aligned}$$

The above holds when

$$\begin{aligned}q &\geq \beta^2 \cdot 8n^{-\delta}q^2 + n^{\epsilon_1(1+\mu)} + 2n^{-\gamma+\epsilon_2(1+\nu)} \\ \implies n^\tau &\geq \beta^2 \cdot 8n^{-\delta}n^{2\tau} + n^{\epsilon_1(1+\mu)} + 2n^{-\gamma+\epsilon_2(1+\nu)} \\ \implies n^{\delta+\tau} &\geq \beta^2 \cdot 8n^{2\tau} + n^{\epsilon_1(1+\mu)+\delta} + 2n^{\delta-\gamma+\epsilon_2(1+\nu)} \\ \implies n^\delta &\geq \beta^2 \cdot 8n^\tau + n^{\epsilon_1(1+\mu)+\delta-\tau} + 2n^{\delta-\tau-\gamma+\epsilon_2(1+\nu)}\end{aligned}$$

Let $\alpha := \max(\epsilon_1(1 + \mu), \epsilon_2(1 + \nu))$. Now if we additionally pick τ such that $\delta - \tau = \kappa$ and $\tau > \kappa + \alpha$, then it suffices to show that

$$n^\delta \geq \beta^2 \cdot 8n^\tau + n^{\alpha+\kappa} + 2n^{\kappa+\alpha},$$

and therefore, it suffices that

$$n^\delta \geq \beta^2 \cdot 8n^\tau \implies \beta \leq O(n^{\frac{\kappa}{2}}).$$

For the above constraint to be valid, we need $\delta - \kappa > \kappa + \alpha \implies 2\kappa \leq \delta - \alpha$ which is a feasible constraint.

To summarize, for there to be a potential attack using an approximate SVP oracle with approximation factor $\beta = O(n^{\frac{\kappa}{2}})$, we can pick:

- $\delta < 1$.
 - $\epsilon_1, \epsilon_2, \mu, \nu : \max(\epsilon_1(1 + \mu), \epsilon_2(1 + \nu)) < \delta$.
 - $\sigma^2 = n^{\epsilon_1}$ and $\xi^2 = n^{\epsilon_2}$.
 - $\kappa : 2\kappa \leq \delta - \max(\epsilon_1(1 + \mu), \epsilon_2(1 + \nu))$.
 - $q = n^{\delta-\kappa}$
- **Case 2** $\tau > \delta$. Then $8n^{1-\delta}q^2 = \omega(n^{1+\tau})$ and $(2n+1)q = O(n^{1+\tau})$ that implies

$$(2n+1)q \leq 8n^{1-\delta}q^2 + n\sigma^{2(1+\mu)} + 2n^{1-\gamma}\xi^{2(1+\nu)}.$$

Therefore, with these parameters, the above reduction idea does not work.

Attack using SVP_β oracle with parameters used in PKE For correctness of PKE we want $q = \omega(n^{1-\gamma}\sigma^{1+\mu}\xi^{1+\nu})$. This will be true if $q = \omega(n^{1-\gamma}n^{\epsilon_1+\epsilon_2})$. Since $q = n^{\delta-\kappa}$ for an attack, we need to pick $\delta - \kappa > 1 - \gamma + \epsilon_1 + \epsilon_2 \implies \gamma + \delta > 1 + \kappa + \epsilon_1 + \epsilon_2$. Therefore there are parameters where we can simultaneously have PKE³ and a potential attack using SVP_β oracle. Specifically, in addition to the constraints above, pick

- $\gamma : \gamma + \delta > 1 + \kappa + \epsilon_1 + \epsilon_2$.

Indeed, this attack can be avoided if we simply pick a prime modulus $q = \Omega(n^\delta)$ which is going to be true in most use cases.

8.2 Inapplicability of natural reductions from sqNULA to other Noisy Linear Algebraic (NLA) assumptions

Most natural NLA hardness assumptions like Learning with Errors (LWE), Learning Parity with Noise (LPN), Learning with Two Errors (LW2E) and Learning with Small and Sparse Errors (LWSSE) among others use secret $\mathbf{s}$ sampled uniformly at random from $\mathbb{F}_q^n$. So, a natural attempts to reduce sqNULA to any of these assumptions is to randomize the short and sparse secret to a uniform random one. For demonstration, let us try to use this approach to reduce sqNULA to LW2E.

Upon receiving $(\mathbf{A} \in \mathbb{F}_q^{n \times n}, \mathbf{b} \in \mathbb{F}_q^n)$, the adversary trying to break decisional sqNULA does the following: Sample a random $\mathbf{s}' \in \mathbb{F}_q^n$, and forward $(\mathbf{A}, \mathbf{b}' = \mathbf{b} + \mathbf{A}\mathbf{s}')$ to the LW2E breaking adversary. Observe that if $\mathbf{b}$ is an sqNULA sample of the form $\mathbf{A}\mathbf{s} + \mathbf{e}_1 + \mathbf{e}_2$, then $\mathbf{b} + \mathbf{A}\mathbf{s}' = \mathbf{A}(\mathbf{s} + \mathbf{s}') + \mathbf{e}_1 + \mathbf{e}_2$. Since, $\mathbf{s}'$ is uniform random, $\mathbf{s} + \mathbf{s}'$ is also distributed randomly over $\mathbb{F}_q^n$. Therefore, $\mathbf{b}'$ is correctly distributed as an LW2E sample.

³This can also be extended to parameters used in the construction of OT and KDM secure SKE. In case of OT, we can pick $\ell = n^{\delta-\kappa-\epsilon_3}$, for appropriate choice of ϵ_3 .

While, this reduction seems plausible, in the parameter regime where $\mathbf{A}$ is a square matrix, there is simply no distinguisher for the LW2E problem. This is because with these parameters the decisional LW2E problem is information theoretically secure. To see why this is the case, note that since $\mathbf{A}$ is invertible, for every $\mathbf{b} \in \mathbb{F}_q^n$, and for every error vector of the form $\mathbf{e}_1 + \mathbf{e}_2$, there will always exist some $\mathbf{s}$ such that $\mathbf{s} = \mathbf{A}^{-1}(\mathbf{b} - \mathbf{e}_1 - \mathbf{e}_2)$. Thus, the case where $\mathbf{b}$ is an LW2E sample vs. when $\mathbf{b}$ is uniform random are statistically indistinguishable.

Observe that the above argument is oblivious to the exact error distribution, therefore any hope to use this approach to reduce sqNULA to other NLA hardness assumptions seems unlikely. Crucially, a reduction strategy of this form converts an sqNULA instance to a target NLA sample with parameters where they are in the “total regime”.

Total Regime: For an NLA, totality is the regime where for all $\mathbf{b} \in \mathbb{F}_q^n$, there will exist a secret $\mathbf{s}$ an error $\mathbf{e}$ sampled from appropriate distributions respectively such that $\mathbf{A}\mathbf{s} + \mathbf{e} = \mathbf{b}$ with overwhelming probability over the choice of invertible $\mathbf{A} \in \mathbb{F}_q^{n \times n}$. Now, observe that if $\mathbf{s}$ is uniformly distributed over $\mathbb{F}_q^n$, then for every choice of $\mathbf{b}$ there will always exist a pair $\mathbf{s}, \mathbf{e}$ such that $\mathbf{A}\mathbf{s} + \mathbf{e} = \mathbf{b}$. In particular, for every choice of $\mathbf{b}, \mathbf{e}$, we can set $\mathbf{s} = \mathbf{A}^{-1}(\mathbf{b} - \mathbf{e})$. In this parameter regime, an NLA sample will be information theoretically indistinguishable from a random sample.

A reduction between the search version of these assumptions will also fail for the same reason. In the above example, there is no guarantee that the LW2E adversary will return the same secret that was used to generate the sqNULA sample, because there are exponentially many pairs of $(\mathbf{s}, \mathbf{e}_1 + \mathbf{e}_2)$ such that $\mathbf{A}\mathbf{s} + \mathbf{e}_1 + \mathbf{e}_2 = \mathbf{b}$.

8.3 Using a “dual” approach to break sqNULA

We start with a brief discussion on a standard way to distinguish an NLA sample from uniform random using a “dual” oracle.

In standard LWE with $m > n$ samples, say $(\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e})$, a dual attack searches for a short nonzero vector $\mathbf{w} \in \mathbb{Z}^m$ such that $\mathbf{w}^T \mathbf{A} \equiv \mathbf{0} \pmod{q}$, and then tests the distribution of $\mathbf{w}^T \mathbf{b} \equiv \mathbf{w}^T \mathbf{e} \pmod{q}$. In the planted case this depends only on the noise and is small, while in the random case this is uniform random over $\mathbb{F}_q$.

First, observe that in our setting $\mathbf{A} \in GL_q(n)$ is square and invertible, so its left kernel modulo q is trivial. So, the above approach seems inapplicable directly. However, a natural attempt is to generate new samples by *manufacturing* redundancy by taking linear combinations of some of the equations.

Derived instance. Fix $m' = n + d > n$ and choose a matrix $\mathbf{U} \in \{0, 1\}^{m' \times n}$ such that every row is sparse, say, row-weight k). Now, we define a derived instance

$$\mathbf{A}' := \mathbf{U}\mathbf{A} \pmod{q}, \quad \mathbf{b}' := \mathbf{U}\mathbf{b} \pmod{q},$$

If $\mathbf{b}$ is planted, i.e.

$$\mathbf{b} \equiv \mathbf{A}\mathbf{s} + \mathbf{e}^{(1)} + \mathbf{e}^{(2)} \pmod{q},$$

then

$$\mathbf{b}' \equiv \mathbf{A}'\mathbf{s} + \mathbf{e}^{(1)'} + \mathbf{e}^{(2)'} \pmod{q}$$

where $\mathbf{e}^{(1)'} = \mathbf{U}\mathbf{e}^{(1)}$ and $\mathbf{e}^{(2)'} = \mathbf{U}\mathbf{e}^{(2)}$.

Since $m' > n$, the matrix $\mathbf{A}'$ has a nontrivial left kernel modulo q . Consider the set

$$\mathcal{L}_{\text{rel}} := \{ \mathbf{w} \in \mathbb{Z}^{m'} : \mathbf{w}^T \mathbf{A}' \equiv \mathbf{0} \pmod{q} \wedge \mathbf{w} \text{ is short} \}.$$

Say we have access to such a “dual” oracle (eg., a solver for the SIS problem) that upon input $\mathbf{A}'$ can outputs $\mathbf{w} \in \mathcal{L}_{\text{rel}}$. Then we can compute

$$t := [\mathbf{w}^T \mathbf{b}']_q.$$

Here, we can be optimistic that since every row of $\mathbf{U}$ is sparse, we pick k such that $\mathbf{U}\mathbf{e}^{(2)} = 0$. Then, we know that, in the planted case, we expect

$$\mathbf{w}^T \mathbf{b}' \equiv \mathbf{w}^T (\mathbf{e}^{(1)'} + \mathbf{e}^{(2)'}) \equiv \mathbf{w}^T \mathbf{e}^{(1)'} \pmod{q},$$

so t should be “small” rather than uniform. This gives a potential reduction of sqNULA to SIS.

Why this fails in the square regime. Let us observe the case where $\mathbf{b} \leftarrow \mathbb{F}_q^n$ is distributed uniformly. Since $\mathbf{A}$ is invertible, for every such pair there exists

$$\mathbf{s}^* := \mathbf{A}^{-1} \mathbf{b} \pmod{q}$$

such that $\mathbf{b} \equiv \mathbf{A}\mathbf{s}^* \pmod{q}$ with *zero* error. Therefore, for the derived instance we have

$$\mathbf{b}' = \mathbf{U}\mathbf{b} \equiv \mathbf{U}\mathbf{A}\mathbf{s}^* = \mathbf{A}'\mathbf{s}^* \pmod{q},$$

so $\mathbf{b}'$ lies *exactly* in the column space of $\mathbf{A}'$ even in the uniform case. Consequently, for any $\mathbf{w} \in \mathcal{L}_{\text{rel}}$,

$$\mathbf{w}^T \mathbf{b}' \equiv \mathbf{w}^T \mathbf{A}' \mathbf{s}^* \equiv 0 \pmod{q}.$$

Hence $t = [\mathbf{w}^T \mathbf{b}']_q$ is *forced* to be small, and thus cannot distinguish planted from uniform.

This illustrates a natural dual-style approach that fails for a nontrivial reason: generating new samples preserves the property that b (and hence b') is in the column space of $\mathbf{A}$ (resp. $\mathbf{A}'$) whenever $\mathbf{A}$ is invertible. As a result, kernel relations annihilate both planted and uniform instances, collapsing the distinguisher.

8.4 Failed Reduction approaches from sqNULA to LPN

One can try to reduce d-sqNULA to LPN by making use of the fact that the error is sparse:

The d-sqNULA distinguisher $\mathcal{D}$ upon input $(\mathbf{A} \in \mathbb{F}_q^{n \times n}, \mathbf{b} \in \mathbb{F}_q^n)$ first samples polynomially many “sparse and small” vectors $\mathbf{v}_i \in \mathbb{F}_q^n$, for $i \in [\text{poly}(n)]$. For the sake of simplicity, assume that $\mathbf{v}$ is non-zero at a tiny constant number of positions. It then computes $\mathbf{a}'_i \leftarrow \mathbf{v}_i^\top \mathbf{A}$ and $b'_i \leftarrow \mathbf{v}_i^\top \mathbf{a}_i$ and sends

$$((\mathbf{A}^\top \|\mathbf{a}'_1\| \mathbf{a}'_2\| \cdots)^\top, (\mathbf{b}^\top \|\mathbf{b}'_1\| \mathbf{b}'_2\| \cdots)^\top)$$

to an LPN distinguisher with appropriate parameters.

The idea behind the reduction is that if $\mathbf{b}$ is a d-sqNULA sample of the form $\mathbf{A}(\mathbf{e}_1 + \mathbf{e}_2) + \mathbf{s}$, then $\mathbf{v}_i^\top \mathbf{b} = \mathbf{v}_i^\top (\mathbf{A}(\mathbf{e}_1 + \mathbf{e}_2) + \mathbf{s}) = \mathbf{a}'_i^\top (\mathbf{e}_1 + \mathbf{e}_2) + \mathbf{v}_i^\top \mathbf{s}$. If $\mathbf{v}_i$ is sparse enough, then with high probability, $\mathbf{v}_i^\top \mathbf{s} = 0$. Therefore, $b'_i = \mathbf{a}'_i^\top (\mathbf{e}_1 + \mathbf{e}_2)$ and $(\mathbf{A}^\top \|\mathbf{a}'_1\| \mathbf{a}'_2\| \cdots)^\top (\mathbf{e}_1 + \mathbf{e}_2) + (\mathbf{s}^\top \|0\|0\| \cdots)^\top = (\mathbf{b}^\top \|\mathbf{b}'_1\| \mathbf{b}'_2\| \cdots)^\top$. The error distribution is not exactly an LPN error distribution, but we can transform them into an LPN error using re-randomization tricks used by [GJL⁺25] (See Section 5.5 in their paper). Additionally, the samples are not independent, but let us be optimistic and assume that an LPN oracle can detect a sample in this format. Now in the case where $\mathbf{b}$ is uniform random, we know that there always exists some $\mathbf{x} \in \mathbb{F}_q^n$ such that $\mathbf{A}\mathbf{x} = \mathbf{b}$ because $\mathbf{A}$ is invertible. Indeed, $\mathbf{x}$ might not have the same distribution as $(\mathbf{e}_1 + \mathbf{e}_2)$. $(\mathbf{A}, \mathbf{b})$ is also a valid LPN sample in this case with error vector $\mathbf{0}$ and therefore observe that

$$(\mathbf{A}^\top \|\mathbf{a}'_1\| \mathbf{a}'_2\| \cdots)^\top \mathbf{x} + (0\|0\|0\| \cdots)^\top = (\mathbf{b}^\top \|\mathbf{b}'_1\| \mathbf{b}'_2\| \cdots)^\top.$$

While the distribution of $\mathbf{x}$ is different from the secret distribution of d-sqNULA, this does not matter because an LPN oracle does not take into account the special distribution of the secret and will return *any* $\mathbf{x}$ such that $\mathbf{A}\mathbf{x} + \mathbf{e} = \mathbf{b}$ and $\mathbf{e}$ is sparse (as per the parameters specified). Therefore, in our case, the LPN oracle will claim that $(\mathbf{A}, \mathbf{b})$ is an LPN sample even when $\mathbf{b}$ is sampled at random. Thus, the idea of extending the number of samples by taking linear combinations and using an LPN oracle does not seem to work.

8.5 Attempted reductions via Dimension reduction

To avoid running into this square problem, we can consider dimension reduction. For any learning problem with sparse secrets, such as sqNULA, we can slightly increase the relative sample count without taking linear combinations of existing rows while preserving the dimension, by guessing the zero coordinates of the short and sparse secret. This differs from the previous subsection in that previously we considered a learning problem with sparse errors, namely d-sqNULA. One can therefore view this dimension reduction as the dual approach of the prior technique that targeted the sparse error distribution in d-sqNULA to boost sample count. Guessing these zero coordinates of the sparse secret can be done in two equivalent manners:

1. Removing randomly selected columns of $\mathbf{A}$ while leaving $\mathbf{b}$ untouched.
2. Or concatenating additional random constraints of the form of elementary row vectors $\mathbf{e}_i^\top$ to the bottom of $\mathbf{A}$ for uniform random $i \in [n]$ while extending the column vector $\mathbf{b}$ with zero entries.

Through either process, the resulting instance is no longer square and has a non-trivial dual. As we will explain, the existence of such a dual does not seem to help one obtain a polynomial-time reduction to any known learning problem because the dimension reduction success probability is too low.

The Guessing Limit Concretely, the dimension reduction proceeds as follows. Suppose we have an instance

$$(\mathbf{A}, \mathbf{b} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e}_1 + \mathbf{e}_2)$$

and the secret $\mathbf{s} \in \mathbb{F}_q^n$ has sparsity parameter γ , so that $\mathbf{s}$ has $n^{1-\gamma}$ many nonzero positions in expectation. By a straightforward counting argument, an adversary can guess up to n^γ positions with constant probability⁴—for $\varepsilon \in (0, 1)$, the probability of selecting n^ε zero positions of the secret is

$$\frac{\binom{n-n^{1-\gamma}}{n^\varepsilon}}{\binom{n}{n^\varepsilon}} = 2^{-n^{\varepsilon-\gamma} + o(n^{\varepsilon-\gamma})}.$$

Therefore, for a constant probability of success, the largest dimension reduction possible is achieved with $\varepsilon = \gamma$. Conditioned on the success of this step, the resulting instance post-dimension reduction, depending on which of the above two methods used, results in:

1. An instance of the form

$$(\overline{\mathbf{A}}, \mathbf{b})$$

where $\overline{\mathbf{A}} \in \mathbb{F}_q^{n \times (n-n^\gamma)}$.

2. Or of the form

$$\left(\overline{\mathbf{A}}, \overline{\mathbf{b}} \triangleq \begin{bmatrix} \mathbf{b} \\ \mathbf{0}_{n^\gamma} \end{bmatrix} \right)$$

where $\overline{\mathbf{A}} \in \mathbb{F}_q^{(n+n^\gamma) \times n}$ and $\overline{\mathbf{b}} \in \mathbb{F}_q^{n+n^\gamma}$ is obtained by adding the necessary elementary row vectors to $\overline{\mathbf{A}}$ and the n^γ additional zero to the column vector $\mathbf{b}$ that enforce the constraint that certain positions of $\mathbf{s}$ are zero.

With this dimension reduction, one might hope to obtain a reduction from sqNULA to some known learning problem outside of the statistically close to uniform parameter setting by avoiding the barrier that arose in the square $\mathbf{A}$ setting above.

⁴For an inverse-polynomial guess probability, we can increase the number of guesses to $n^{\delta + (\log \log n) / \log n}$ but the additive term tends to 0 super-polynomially fast.

No Known Reduction to LPN Since we do not know how to perform dimension reduction on d-sqNULA, we consider instead sqNULA. An attempt at a reduction given a sqNULA instance

$$(\mathbf{A}, \mathbf{b} \triangleq \mathbf{A} \cdot \mathbf{s} + \mathbf{e}_1 + \mathbf{e}_2),$$

is that the reduction attempt performs the dimension reduction step in the first manner obtaining

$$(\overline{\mathbf{A}}, \mathbf{b}) \in \mathbb{F}_q^{n \times (n-n^\gamma)} \times \mathbb{F}_q^n$$

The problem here is that $\mathbf{e}_1$ is dense, so any attempt to deal with the smallness of $\mathbf{e}_1$ results in a large and dense error vector that we are currently unable to remove. For example, a concrete attempt to reduce to an LPN instance is to follow [GJL⁺25] and compute a random invertible diagonal matrix $\mathbf{D} = \text{diag}(d_1, \dots, d_{n+n^\gamma})$ and a random invertible matrix $\mathbf{R}$ and compute

$$(\mathbf{D} \cdot \overline{\mathbf{A}} \cdot \mathbf{R}, \mathbf{D} \cdot \mathbf{b}).$$

Observe that the error term above is

$$\mathbf{D} \cdot \mathbf{e}_1 + \mathbf{D} \cdot \mathbf{e}_2.$$

Since $\mathbf{e}_1$ is small but dense, the term $\mathbf{D} \cdot \mathbf{e}_1$ will not be sparse, and so, it is unclear how to handle this additional error term to obtain an LPN instance.

No Useful Dual While this instance now has a non-trivial dual instance, where by dual we refer to the dual code of that defined by the columns of $\overline{\mathbf{A}}$, the dual instance is statistically close to uniform in our parameter settings. Namely, the reduction can compute the pair

$$(\overline{\mathbf{A}}^\perp, \overline{\mathbf{A}}^\perp \cdot \mathbf{b}) \in \mathbb{F}_q^{n^\gamma \times n} \times \mathbb{F}_q^{n^\gamma}$$

where $\overline{\mathbf{A}}^\perp \cdot \mathbf{b} = \overline{\mathbf{A}}^\perp \cdot (\mathbf{e}_1 + \mathbf{e}_2)$. We observe now that the amount of entropy in $\mathbf{e}_1 + \mathbf{e}_2$ is approximately

$$n \log B + n^{1-\delta} \log q$$

for some B polynomially related to q , while the output size of the compressing random linear function $\overline{\mathbf{A}}^\perp \cdot \mathbf{D}$ is only n^γ . By a standard entropic argument, this RHS is statistically close to uniform if

$$n \log B + n^{1-\delta} \log q \geq n^\gamma \log q.$$

If $q = \text{poly}(n)$ and B is polynomially related to q , we note that the term $n \log B$ already asymptotically dominates $n^\gamma \log q$ for any choice of $\gamma < 1$. In the case of superpolynomial q , however, this reduction idea may bring us out of the information-theoretic regime, in which case we may be able to apply known reductions to dual LW2E assumption.

8.6 An Information Set Decoding(ISD) attack on sqNULA

Parameter regime. We analyze the regime used for the applications: $q = n^4$ is prime, $\delta = 0.6$, $\gamma = 0.4$, and $\sigma = \xi = n$. Let

$$t := \lceil n^{1-\delta} \rceil = \lceil n^{0.4} \rceil, \quad m := n - t.$$

Fix polynomial bounds $B_e, B_s < n^{1.1}$ such that in the planted case $\|\mathbf{e}^{(1)}\|_\infty \leq B_e$ and $\|\mathbf{s}\|_\infty \leq B_s$ with overwhelming probability (e.g., by discrete-Gaussian tails).

ISD template. Since $\mathbf{e}^{(2)}$ is sparse (each coordinate is nonzero with probability $n^{-\delta}$), an ISD-style attack guesses a set T of size t containing the support of $\mathbf{e}^{(2)}$. On the remaining coordinates $I := [n] \setminus T$ we expect $\mathbf{e}_I^{(2)} = 0$, so that

$$\mathbf{b}_I \equiv \mathbf{A}_I \mathbf{s} + \mathbf{e}_I^{(1)} \pmod{q}$$

contains only the “small” error $\mathbf{e}_I^{(1)}$.

Now, because $|I| = m < n$, we are now left with a system that has less samples than the secret dimension. To address that, we use that $\mathbf{s}$ is sparse: we guess some coordinates of $\mathbf{s}$ and hope that they are mostly 0.

Distinguisher $\mathcal{D}$. On input $(\mathbf{A}, \mathbf{b})$, the distinguisher $\mathcal{D}$:

1. Sample $T \subseteq [n]$ uniformly at random with $|T| = t$, and set $I := [n] \setminus T$.
2. Sample $J \subseteq [n]$ uniformly at random with $|J| = m$, and set $C := [n] \setminus J$ (so $|C| = t$). If $\mathbf{A}_{I,J} \in \mathbb{F}_q^{m \times m}$ is not invertible, resample J .
3. Enumerate all vectors $\mathbf{s}_C \in \mathbb{Z}^C$ with $\|\mathbf{s}_C\|_\infty \leq B_s$ and Hamming weight $\text{wt}(\mathbf{s}_C) \leq w_{\max}$, where $w_{\max}$ is a small constant. For each such $\mathbf{s}_C$, compute

$$\mathbf{y} := [\mathbf{b}_I - \mathbf{A}_{I,C} \mathbf{s}_C]_q \in \mathbb{Z}^m.$$

4. (Decoding step.) Guess $\mathbf{s}_J \in \mathbb{Z}^J$ with $\|\mathbf{s}_J\|_\infty \leq B_s$. Test whether

$$\|[\mathbf{y} - \mathbf{A}_{I,J} \mathbf{s}_J]_q\|_\infty \leq B_e.$$

If the test passes for some $\mathbf{s}_J$, define $\hat{\mathbf{s}} \in \mathbb{Z}^n$ by combining $(\mathbf{s}_J, \mathbf{s}_C)$, and compute the residual

$$\mathbf{r} := [\mathbf{b} - \mathbf{A} \hat{\mathbf{s}}]_q \in \mathbb{Z}^n.$$

If $\|\mathbf{r}_I\|_\infty \leq B_e$ and at least $t/2$ coordinates of $\mathbf{r}_T$ satisfy $|\mathbf{r}_i| > B_e$, then output “planted” and halt.

If no trial outputs “planted”, output “uniform”.

Correctness (informal). We argue that the distinguisher accepts planted instances with noticeable probability (over its internal randomness) and accepts uniform instances only with negligible probability.

Planted case. Suppose $\mathbf{b} = \mathbf{A} \mathbf{s} + \mathbf{e}^{(1)} + \mathbf{e}^{(2)} \pmod{q}$ with $\mathbf{s} \sim \mathcal{E}_{n,\xi,\gamma}^n$, $\mathbf{e}^{(1)} \sim \mathcal{D}_{\mathbb{Z},\sigma}^n$, and $\mathbf{e}^{(2)} \sim \mathcal{S}_{n,q,\delta}^n$. Consider an outer trial in which the guessed set T contains the support of $\mathbf{e}^{(2)}$, i.e. $\mathbf{e}_I^{(2)} = 0$ for $I = [n] \setminus T$. This is correct with probability approximately $N^{-1} = \frac{1}{\binom{n}{t}}$. Then on the coordinates in I we have

$$\mathbf{b}_I \equiv \mathbf{A}_{I,J} \mathbf{s}_J + \mathbf{A}_{I,C} \mathbf{s}_C + \mathbf{e}_I^{(1)} \pmod{q}.$$

For any candidate $\mathbf{s}_C$ equal to the true restriction of $\mathbf{s}$ to C , the computed target $\mathbf{y} = [\mathbf{b}_I - \mathbf{A}_{I,C} \mathbf{s}_C]_q$ satisfies

$$\mathbf{y} \equiv \mathbf{A}_{I,J} \mathbf{s}_J + \mathbf{e}_I^{(1)} \pmod{q},$$

hence if $\mathbf{s}_J$ is guessed correctly, then $\|[\mathbf{y} - \mathbf{A}_{I,J} \mathbf{s}_J]_q\|_\infty = \|\mathbf{e}_I^{(1)}\|_\infty \leq B_e$ except with negligible probability (by the choice of B_e from discrete-Gaussian tails). Moreover, since $\mathbf{s}$ is sparse and $|C| = t$, there exists a true $\mathbf{s}_C$ typically has constant Hamming weight, so it is included in the enumeration $\text{wt}(\mathbf{s}_C) \leq w_{\max}$ with high probability (for an appropriate constant $w_{\max}$). Therefore, in this trial the inner search over $\mathbf{s}_J$ will find the true $\mathbf{s}_J$, yielding $\hat{\mathbf{s}} = \mathbf{s}$ and residual $\mathbf{r} = [\mathbf{b} - \mathbf{A} \hat{\mathbf{s}}]_q = [\mathbf{e}^{(1)} + \mathbf{e}^{(2)}]_q$. By construction, $\mathbf{r}_I = [\mathbf{e}_I^{(1)}]_q$ is bounded by B_e , while on T many coordinates satisfy $|\mathbf{r}_i| > B_e$ because

on $\text{supp}(\mathbf{e}^{(2)})$ the term $\mathbf{e}_i^{(2)}$ is uniform in $\mathbb{F}_q$ and so overwhelms the small Gaussian component. Hence the acceptance condition $\|\mathbf{r}_I\|_\infty \leq B_e$ and “many large coordinates on T ” holds with high probability.

Uniform case. If $\mathbf{b} \xleftarrow{\$} \mathbb{F}_q^n$ is uniform, then for any fixed choices of T, J and any fixed $\mathbf{s}_C, \mathbf{s}_J$ the value $[\mathbf{b} - \mathbf{A}\hat{\mathbf{s}}]_q$ is uniform in $\mathbb{F}_q^n$. In particular, the event $\|\mathbf{r}_I\|_\infty \leq B_e$ has probability at most

$$\left(\frac{2B_e + 1}{q}\right)^{|I|} = \left(\frac{2B_e + 1}{q}\right)^m,$$

and the additional requirement that many coordinates in T have $|r_i| > B_e$ only decreases this probability. Taking a union bound over all enumerated candidates $(\mathbf{s}_C, \mathbf{s}_J)$ and all $N = \binom{n}{t}$ trials still yields a negligible overall false-accept probability as long as q is sufficiently larger than B_e (as in our parameter regime $q = n^4$ and $B_e < n^{1.1}$).

Therefore the distinguisher is correct: it accepts planted instances with noticeable probability given enough trials, and accepts uniform instances only with negligible probability.

Running time is exponential. Fix any trial and any enumerated $\mathbf{s}_C$. The inner enumeration over $\mathbf{s}_J$ in Step (4) has size

$$(2B_s + 1)^{|J|} = (2B_s + 1)^m.$$

Since $B_s = \text{poly}(n)$ we have $(2B_s + 1) = n^{O(1)}$, and since $m = n - t = n - \lceil n^{0.4} \rceil = \Theta(n)$,

$$(2B_s + 1)^m = \exp(\Theta(n) \cdot \Theta(\log n)) = 2^{\Theta(n \log n)}.$$

Thus even one trial runs in time at least $2^{\Theta(n \log n)}$. Multiplying by the number of trials $N = \binom{n}{t}$ to make the support-guessing succeed with noticeable probability makes this worse. In particular,

$$\text{Time}(n) \geq (2B_s + 1)^m = 2^{\Theta(n \log n)}.$$

Remark 8.4. The above attack is a generalization and an improvement of the attack where one first guesses the non-zero entries of $\mathbf{e}^{(2)}$ and then guess the entire secret $\mathbf{s}$.

Remark 8.5. Another potential attack strategy could be to combine ISD with an approximate SVP oracle (resembling the BKZ attack paradigm [SE94]). First, we guess the non-zero positions of $\mathbf{e}^{(2)}$. Then, we can again guess $\mathbf{s}_C$. Then we can use an SVP oracle to solve the instance

$$\mathbf{y} = \mathbf{A}_{I,J} \mathbf{s}_J + \mathbf{e}_I^{(1)} \pmod{q}.$$

However, since the first guessing stage already takes $n^t \approx 2^{n^{0.4 \log n}}$ time, this attack also run in subexponential time.

8.7 Brute force attack on d-sqNULA:

Given a d-sqNULA sample $(\mathbf{A}, \mathbf{b} = \mathbf{A}(\mathbf{e}_1 + \mathbf{e}_2) + \mathbf{s})$, one can try to enumerate over all “small and sparse” error vectors $\mathbf{s}$ and then see if there is a “small plus sparse” secret corresponding to it. Specifically, for a choice of error $\mathbf{s}^*$, check if $\mathbf{A}^{-1}(\mathbf{b} - \mathbf{s}^*)$ is correctly distributed as a short plus sparse secret. This approach runs in time $O(n^{n^{1-\gamma}})$ in the worst case.

Acknowledgements and AI disclosure

The primary innovations of coming up with the new assumptions and the PKE construction (including proof of correctness and security) were conceived entirely by the human authors, without any AI assistance. GPT-5.6 Sol Ultra and GPT-6 Astra Ultra assisted with: (1) Verifying reduction approaches developed by the authors in Sections 7. (2) Simplifying and formalizing author-directed cryptanalysis

in Section 8. (3) Writing and editing throughout the draft. The human authors take sole responsibility for all claims and proofs.

This research was supported in part by a Laude Moonshot seed award, a Simons Investigator Award, a DARPA expMath award, NSF grant 2333935, BSF grant 2022370, a Xerox Faculty Research Award, a Google Faculty Research Award, an Okawa Foundation Research Grant, and the Symantec Chair of Computer Science. PL was supported by the European Research Council (ERC) under the EU's Horizon 2020 research and innovation programme (Grant agreement No. 101019547) and Cariplo CRYPTONOMEX grant.

9 References

- [ABLJK18] Reza Azarderakhsh, Elena Bakos Lang, David Jao, and Brian Koziel. Edsidh: supersingular isogeny diffie-hellman key exchange on edwards curves. In *International Conference on Security, Privacy, and Applied Cryptography Engineering*, pages 125–141. Springer, 2018. ¹
- [ABV⁺12] Shweta Agrawal, Xavier Boyen, Vinod Vaikuntanathan, Panagiotis Voulgaris, and Hoeteck Wee. Functional encryption for threshold functions (or fuzzy ibe) from lattices. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 280–297. Springer, Berlin, Heidelberg, May 2012. ²
- [ABW10] Benny Applebaum, Boaz Barak, and Avi Wigderson. Public-key cryptography from different assumptions. In Leonard J. Schulman, editor, *42nd ACM STOC*, pages 171–180. ACM Press, June 2010. ²
- [ADD⁺20] Martin R. Albrecht, Alex Davidson, Amit Deo, Benjamin R. Curtis, Damien Stehlé, Eamonn W. Postlethwaite, Elena Kirshanova, Fernando Virdia, Florian Göpfert, Gottfried Herold, John M. Schanck, Léo Ducas, Marc Stevens, Paul Kirchner, Pierre-Alain Fouque, Rachel Player, Sam Scott, Shi Bai, Thomas Wunderer, Vlad Gheorghiu, and Weiqiang Wen. Algorithms for lattice problems: In practice. Lecture notes, Simons Institute Lattice Boot Camp, January 2020. Based on joint work with many co-authors. ³³
- [AFV11] Shweta Agrawal, David Mandell Freeman, and Vinod Vaikuntanathan. Functional encryption for inner product predicates from learning with errors. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 21–40. Springer, Berlin, Heidelberg, December 2011. ²
- [AIR01] William Aiello, Yuval Ishai, and Omer Reingold. Priced oblivious transfer: How to sell digital goods. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 119–135. Springer, Berlin, Heidelberg, May 2001. ³
- [Ale03] Michael Alekhnovich. More on average case vs approximation complexity. In *44th FOCS*, pages 298–307. IEEE Computer Society Press, October 2003. ^{2, 4, 5, 9}
- [BBS09] László Babai, Robert Beals, and Ákos Seress. Polynomial-time theory of matrix groups. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 55–64. ACM Press, May / June 2009. ¹
- [BBW25] Marcello Benedetti, Harry Buhrman, and Jordi Weggemans. Complement sampling: Provable, verifiable and nisqable quantum advantage in sample complexity. *arXiv preprint arXiv:2502.08721*, 2025. ¹
- [BCM⁺18] Zvika Brakerski, Paul Christiano, Urmila Mahadev, Umesh V. Vazirani, and Thomas Vidick. A cryptographic test of quantumness and certifiable randomness from a single quantum device. In Mikkel Thorup, editor, *59th FOCS*, pages 320–331. IEEE Computer Society Press, October 2018. ¹
- [BDG⁺25] Pierre Briaud, Itai Dinur, Riddhi Ghosal, Aayush Jain, Paul Lou, and Amit Sahai. Quantum advantage via solving multivariate polynomials. *arXiv preprint arXiv:2509.07276*, 2025. ¹
- [Beu22] Ward Beullens. Breaking rainbow takes a weekend on a laptop. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 464–479. Springer, Cham, August 2022. ²

- [BFKL94] Avrim Blum, Merrick L. Furst, Michael J. Kearns, and Richard J. Lipton. Cryptographic primitives based on hard learning problems. In Douglas R. Stinson, editor, *CRYPTO'93*, volume 773 of *LNCS*, pages 278–291. Springer, Berlin, Heidelberg, August 1994. ²
- [BGV12] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In Shafi Goldwasser, editor, *ITCS 2012*, pages 309–325. ACM, January 2012. ²
- [BKM20] Zvika Brakerski, Venkata Koppula, and Tamer Mour. NIZK from LPN and trapdoor hash via correlation intractability for approximable relations. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part III*, volume 12172 of *LNCS*, pages 738–767. Springer, Cham, August 2020. ²
- [BKR23] Andrej Bogdanov, Pravesh K. Kothari, and Alon Rosen. Public-key encryption, local pseudorandom generators, and the low-degree method. In Guy N. Rothblum and Hoeteck Wee, editors, *TCC 2023, Part I*, volume 14369 of *LNCS*, pages 268–285. Springer, Cham, November / December 2023. ²
- [BKVV20] Zvika Brakerski, Venkata Koppula, Umesh Vazirani, and Thomas Vidick. Simpler proofs of quantumness. *arXiv preprint arXiv:2005.04826*, 2020. ¹
- [Bol25] Catherine Bolgar. Microsoft’s majorana 1 chip carves new path for quantum computing. <https://news.microsoft.com/source/features/ai/microsofts-majorana-1-chip-carves-new-path-for-quantum-computing/>, February 2025. Published February 19, 2025; Accessed: 2025-09-30. ¹
- [BV11] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In Rafail Ostrovsky, editor, *52nd FOCS*, pages 97–106. IEEE Computer Society Press, October 2011. ²
- [BV14] Zvika Brakerski and Vinod Vaikuntanathan. Lattice-based fhe as secure as pke. In *Proceedings of the 5th conference on Innovations in theoretical computer science*, pages 1–12, 2014. ⁷
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, October 2001. ¹¹
- [CCM98] Christian Cachin, Claude Crépeau, and Julien Marcil. Oblivious transfer with a memory-bounded receiver. In *39th FOCS*, pages 493–502. IEEE Computer Society Press, November 1998. ³
- [CDMT24] Kévin Carrier, Thomas Debris-Alazard, Charles Meyer-Hilfiger, and Jean-Pierre Tillich. Reduction from sparse LPN to LPN, dual attack 3.0. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part VII*, volume 14657 of *LNCS*, pages 286–315. Springer, Cham, May 2024. ³
- [CFMR⁺17] Antoine Casanova, Jean-Charles Faugere, Gilles Macario-Rat, Jacques Patarin, Ludovic Perret, and Jocelyn Ryckeghem. *GeMSS: a great multivariate short signature*. PhD thesis, UPMC-Paris 6 Sorbonne Universités; INRIA Paris Research Centre, MAMBA Team ..., 2017. ²
- [Che24] Yilei Chen. Quantum algorithms for lattice problems. *Cryptology ePrint Archive*, 2024. ²

- [CHSH69] John F Clauser, Michael A Horne, Abner Shimony, and Richard A Holt. Proposed experiment to test local hidden-variable theories. *Physical review letters*, 23(15):880, 1969. ¹
- [CLM⁺18] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. CSIDH: An efficient post-quantum commutative group action. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part III*, volume 11274 of *LNCS*, pages 395–427. Springer, Cham, December 2018. ¹
- [CMV26] Rohit Chatterjee, Changrui Mu, and Prashant Nalini Vasudevan. Public-key encryption from the minrank problem. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 449–480. Springer, 2026. ²
- [Cou06] Jean-Marc Couveignes. Hard homogeneous spaces. 2006. ¹
- [CR03] Ran Canetti and Tal Rabin. Universal composition with joint state. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 265–281. Springer, Berlin, Heidelberg, August 2003. ¹²
- [DAGNR25] Thomas Debris-Alazard, Philippe Gaborit, Romaric Neveu, and Olivier Ruatta. A minrank-based encryption scheme\à la alekhnovich-regev. *arXiv preprint arXiv:2510.07584*, 2025. ²
- [DGH⁺20] Nico Döttling, Sanjam Garg, Mohammad Hajiabadi, Daniel Masny, and Daniel Wichs. Two-round oblivious transfer from CDH or LPN. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 768–797. Springer, Cham, May 2020. ^{2, 3, 5, 10, 11, 21}
- [DIJL23] Quang Dao, Yuval Ishai, Aayush Jain, and Huijia Lin. Multi-party homomorphic secret sharing and sublinear MPC from sparse LPN. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part II*, volume 14082 of *LNCS*, pages 315–348. Springer, Cham, August 2023. ³
- [DJ24] Quang Dao and Aayush Jain. Lossy cryptography from code-based assumptions. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part III*, volume 14922 of *LNCS*, pages 34–75. Springer, Cham, August 2024. ³
- [DJJ24] Quang Dao, Aayush Jain, and Zhengzhong Jin. Non-interactive zero-knowledge from lpn and mq. In *Annual International Cryptology Conference*, pages 321–360. Springer, 2024. ²
- [DP12] Ivan Damgård and Sunoo Park. Is public-key encryption based on LPN practical? *IACR Cryptol. ePrint Arch.*, page 699, 2012. ³
- [DS05a] Jintai Ding and Dieter Schmidt. Cryptanalysis of HFEv and internal perturbation of HFE. In Serge Vaudenay, editor, *PKC 2005*, volume 3386 of *LNCS*, pages 288–301. Springer, Berlin, Heidelberg, January 2005. ²
- [DS05b] Jintai Ding and Dieter Schmidt. Rainbow, a new multivariable polynomial signature scheme. In *ACNS*, volume 5, pages 164–175. Springer, 2005. ²
- [ES16] Lior Eldar and Peter W Shor. An efficient quantum algorithm for a variant of the closest lattice-vector problem. *arXiv preprint arXiv:1611.06999*, 2016. ²
- [Gam23] Jay Gambetta. The hardware and software for the era of quantum utility is here. <https://www.ibm.com/quantum/blog/quantum-roadmap-2033>, December 2023. Accessed: 2025-09-30. ¹

- [Gen10] Craig Gentry. Toward basing fully homomorphic encryption on worst-case hardness. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 116–137. Springer, Berlin, Heidelberg, August 2010. ²
- [GHJS25] Riddhi Ghosal, Isaac M Hair, Aayush Jain, and Amit Sahai. Using the planted clique conjecture for cryptography: Public-key encryption from planted clique and noisy k-lin over expanders. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 1921–1932, 2025. ²
- [GJL⁺25] Riddhi Ghosal, Aayush Jain, Paul Lou, Amit Sahai, and Neekon Vafa. Post-quantum pke from unstructured noisy linear algebraic assumptions: Beyond lwe and alekhnovich’s lpn. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 64–93. Springer, 2025. ^{2, 3, 4, 7, 9, 10, 36, 38}
- [GL89] Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In *21st ACM STOC*, pages 25–32. ACM Press, May 1989. ^{5, 17}
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008. ²
- [GRS00] Oded Goldreich, Ronitt Rubinfeld, and Madhu Sudan. Learning polynomials with queries: The highly noisy case. *SIAM Journal on Discrete Mathematics*, 13(4):535–570, 2000. ^{5, 7, 17, 26}
- [GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 75–92. Springer, Berlin, Heidelberg, August 2013. ²
- [Hal07] Sean Hallgren. Polynomial-time quantum algorithms for pell’s equation and the principal ideal problem. *Journal of the ACM (JACM)*, 54(1):1–19, 2007. ¹
- [Has19] Matthew B Hastings. Classical and quantum bounded depth approximation algorithms. *arXiv preprint arXiv:1905.07047*, 2019. ²
- [HPS98] Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In Joe Buhler, editor, *Algorithmic Number Theory, Third International Symposium, ANTS-III, Portland, Oregon, USA, June 21-25, 1998, Proceedings*, volume 1423 of *Lecture Notes in Computer Science*, pages 267–288. Springer, 1998. ³
- [IBM25] IBM Quantum. Processor types. <https://quantum.cloud.ibm.com/docs/en/guides/processor-types>, 2025. Accessed: 2025-09-30. ¹
- [IPS09] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Secure arithmetic computation with no honest majority. In Omer Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 294–314. Springer, Berlin, Heidelberg, March 2009. ²
- [JLS21] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. In Samir Khuller and Virginia Vassilevska Williams, editors, *53rd ACM STOC*, pages 60–73. ACM Press, June 2021. ^{2, 9}
- [JSW⁺24] Stephen P Jordan, Noah Shutty, Mary Wootters, Adam Zalcman, Alexander Schmidhuber, Robbie King, Sergei V Isakov, and Ryan Babbush. Optimization by decoded quantum interferometry. *arXiv preprint arXiv:2408.08292*, 2024. ¹

- [KLVY23] Yael Kalai, Alex Lombardi, Vinod Vaikuntanathan, and Lisa Yang. Quantum advantage from any non-local game. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1617–1628, 2023. ¹
- [KM20] Fuyuki Kitagawa and Takahiro Matsuda. Circular security is complete for KDM security. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part I*, volume 12491 of *LNCS*, pages 253–285. Springer, Cham, December 2020. ¹²
- [KMCVY22] Gregory D Kahanamoku-Meyer, Soonwon Choi, Umesh V Vazirani, and Norman Y Yao. Classically verifiable quantum advantage from a computational bell test. *Nature Physics*, 18(8):918–924, 2022. ¹
- [KMT19] Fuyuki Kitagawa, Takahiro Matsuda, and Keisuke Tanaka. CCA security and trapdoor functions via key-dependent-message security. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 33–64. Springer, Cham, August 2019. ¹²
- [KPG99] Aviad Kipnis, Jacques Patarin, and Louis Goubin. Unbalanced Oil and Vinegar signature schemes. In Jacques Stern, editor, *EUROCRYPT’99*, volume 1592 of *LNCS*, pages 206–222. Springer, Berlin, Heidelberg, May 1999. ²
- [LLW21] Qiqi Lai, Feng-Hao Liu, and Zhedong Wang. New lattice two-stage sampling technique and its applications to functional encryption - stronger security and smaller ciphertexts. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part I*, volume 12696 of *LNCS*, pages 498–527. Springer, Cham, October 2021. ³¹
- [LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 1–23. Springer, Berlin, Heidelberg, May / June 2010. ³
- [LPR13] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for ring-LWE cryptography. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 35–54. Springer, Berlin, Heidelberg, May 2013. ³
- [Lyu12] Vadim Lyubashevsky. Lattice signatures without trapdoors. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 738–755. Springer, Berlin, Heidelberg, April 2012. ^{2,9}
- [McE78] Robert J. McEliece. A public-key cryptosystem based on algebraic coding theory. The deep space network progress report 42-44, Jet Propulsion Laboratory, California Institute of Technology, January/February 1978. https://ipnpr.jpl.nasa.gov/progress_report2/42-44/44N.PDF. ³
- [Mic18] Daniele Micciancio. On the hardness of learning with errors with binary secrets. *Theory of Computing*, 14(1):1–17, 2018. ²⁹
- [Nev24] Hartmut Neven. Meet willow, our state-of-the-art quantum chip. <https://blog.google/technology/research/google-willow-quantum-chip/>, December 2024. Accessed: 2025-09-30. ¹
- [NP01] Moni Naor and Benny Pinkas. Efficient oblivious transfer protocols. In S. Rao Kosaraju, editor, *12th SODA*, pages 448–457. ACM-SIAM, January 2001. ³
- [Pat96] Jacques Patarin. Hidden fields equations (HFE) and isomorphisms of polynomials (IP): Two new families of asymmetric algorithms. In Ueli M. Maurer, editor, *EUROCRYPT’96*, volume 1070 of *LNCS*, pages 33–48. Springer, Berlin, Heidelberg, May 1996. ²

- [Pei13] Chris Peikert. Lecture 1: Mathematical background. Lattices in Cryptography course notes, Georgia Institute of Technology, 2013. Scribe: Sara Krehbiel. ³³
- [PGC98] Jacques Patarin, Louis Goubin, and Nicolas Courtois. C_{-+}^* and HM: Variations around two schemes of T. Matsumoto and H. Imai. In Kazuo Ohta and Dingyi Pei, editors, *ASIACRYPT'98*, volume 1514 of *LNCS*, pages 35–49. Springer, Berlin, Heidelberg, October 1998. ²
- [Pre12] John Preskill. Quantum computing and the entanglement frontier. *arXiv preprint arXiv:1203.5813*, 2012. ¹
- [PRS17] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudorandomness of ring-LWE for any ring and modulus. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *49th ACM STOC*, pages 461–473. ACM Press, June 2017. ³
- [PS19] Chris Peikert and Sina Shiehian. Noninteractive zero knowledge for NP from (plain) learning with errors. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 89–114. Springer, Cham, August 2019. ²
- [PVW08] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 554–571. Springer, Berlin, Heidelberg, August 2008. ³
- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM (JACM)*, 56(6):1–40, 2009. ^{2, 4, 9}
- [Rob23] Damien Robert. Breaking SIDH in polynomial time. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part V*, volume 14008 of *LNCS*, pages 472–503. Springer, Cham, April 2023. ²
- [RS06] Alexander Rostovtsev and Anton Stolbunov. Public-key cryptosystem based on isogenies. *Cryptology ePrint Archive*, 2006. ¹
- [SD16] Noah Stephens-Davidowitz. Dimension-preserving reductions between lattice problems. <http://www.noahsd.com/latticeproblems.pdf>, 2016. ⁷
- [SE94] Claus-Peter Schnorr and Martin Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Mathematical programming*, 66(1):181–199, 1994. ⁴⁰
- [Sho94] Peter W. Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *35th FOCS*, pages 124–134. IEEE Computer Society Press, November 1994. ^{1, 2}
- [SOKB24] Alexander Schmidhuber, Ryan O'Donnell, Robin Kothari, and Ryan Babbush. Quartic quantum speedups for planted inference. *arXiv preprint arXiv:2406.19378*, 2024. ²
- [TDTD13] Chengdong Tao, Adama Diene, Shaohua Tang, and Jintai Ding. Simple matrix scheme for encryption. In *Post-Quantum Cryptography: 5th International Workshop, PQCrypto 2013, Limoges, France, June 4-7, 2013. Proceedings 5*, pages 231–242. Springer, 2013. ²
- [WZ17] Daniel Wichs and Giorgos Zirdelis. Obfuscating compute-and-compare programs under LWE. In Chris Umans, editor, *58th FOCS*, pages 600–611. IEEE Computer Society Press, October 2017. ²
- [YZ16] Yu Yu and Jiang Zhang. Cryptography with auxiliary input and trapdoor from constant-noise LPN. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 214–243. Springer, Berlin, Heidelberg, August 2016. ^{2, 3}

- [YZ22] Takashi Yamakawa and Mark Zhandry. Verifiable quantum advantage without structure. In *63rd FOCS*, pages 69–74. IEEE Computer Society Press, October / November 2022. ¹
- [ZT15] Wenbin Zhang and Chik How Tan. Mi-t-hfe, a new multivariate signature scheme. In *Cryptography and Coding: 15th IMA International Conference, IMACC 2015, Oxford, UK, December 15-17, 2015. Proceedings 15*, pages 43–56. Springer, 2015. ²