

Cryptography with Auxiliary Input and Trapdoor from Constant-Noise LPN

Yu Yu* Jiang Zhang†

May 25, 2016

Abstract

Dodis, Kalai and Lovett (STOC 2009) initiated the study of the Learning Parity with Noise (LPN) problem with (static) exponentially hard-to-invert auxiliary input. In particular, they showed that under a new assumption (called Learning Subspace with Noise) the above is quasi-polynomially hard in the high (polynomially close to uniform) noise regime.

Inspired by the “sampling from subspace” technique by Yu (eprint 2009 / 467) and Goldwasser et al. (ITCS 2010), we show that standard LPN can work in a mode (reducible to itself) where the constant-noise LPN (by sampling its matrix from a random subspace) is robust against sub-exponentially hard-to-invert auxiliary input with comparable security to the underlying LPN. Plugging this into the framework of [DKL09], we obtain the same applications as considered in [DKL09] (i.e., CPA/CCA secure symmetric encryption schemes, average-case obfuscators, reusable and robust extractors) with resilience to a more general class of leakages, improved efficiency and better security under standard assumptions.

As a main contribution, under constant-noise LPN with certain sub-exponential hardness (i.e., $2^{\omega(n^{1/2})}$ for secret size n) we obtain a variant of the LPN with security on polylogarithmic entropy sources, which in turn implies CPA/CCA secure public-key encryption (PKE) schemes and oblivious transfer (OT) protocols. Prior to this, basing PKE and OT on constant-noise LPN had been an open problem since Alekhnovich’s work (FOCS 2003).

Keywords: Foundations, Cryptography with Auxiliary Input, Learning Parity with Noise, Post-quantum Cryptography, Public-Key Encryption.

1 Introduction

LEARNING PARITY WITH NOISE. The computational version of learning parity with noise (LPN) assumption with parameters $n \in \mathbb{N}$ (length of secret) and $0 < \mu < 1/2$ (noise rate) postulates that for any $q = \text{poly}(n)$ (number of queries) it is computationally infeasible for any probabilistic polynomial-time (PPT) algorithm to recover the random secret $\mathbf{x} \stackrel{\$}{\leftarrow} \{0, 1\}^n$ given $(\mathbf{A}, \mathbf{A} \cdot \mathbf{x} + \mathbf{e})$, where \mathbf{a} is a random $q \times n$ Boolean matrix, \mathbf{e} follows $\mathcal{B}_\mu^q = (\mathcal{B}_\mu)^q$, \mathcal{B}_μ denotes the Bernoulli distribution with parameter μ (i.e., $\Pr[\mathcal{B}_\mu = 1] = \mu$ and $\Pr[\mathcal{B}_\mu = 0] = 1 - \mu$), ‘ \cdot ’ denotes matrix vector multiplication over $\text{GF}(2)$ and ‘ $+$ ’ denotes bitwise addition over $\text{GF}(2)$. The decisional version of LPN simply assumes that $(\mathbf{A}, \mathbf{A} \cdot \mathbf{x} + \mathbf{e})$ is pseudorandom. The two versions are polynomially equivalent [8, 34, 4].

HARDNESS OF LPN. The computational LPN problem represents a well-known NP-complete problem “decoding random linear codes” [6] whose worst-case hardness is well-investigated. LPN was also extensively studied in learning theory, and it was shown in [21] that an efficient

*Shanghai Jiao Tong University. Email: yuyuathk@gmail.com.

†State Key Laboratory of Cryptology, P.O. Box 5159, Beijing 100878, China. Email: jiangzhang09@gmail.com.

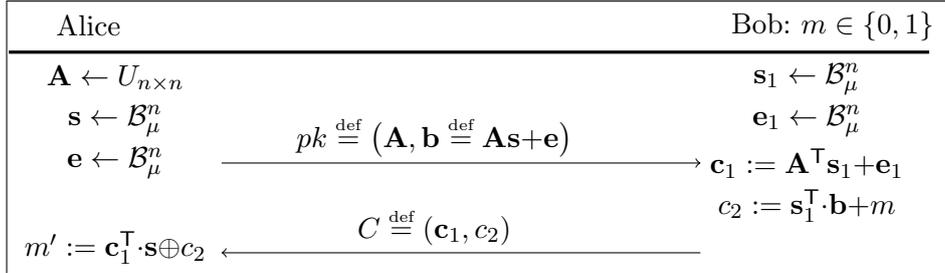


Figure 1: A two-pass protocol by which Bob transmits a message bit m to Alice with passive security and noticeable correctness (for proper choice of μ), where Bob receives $m' = m + (\mathbf{s}_1^\top \cdot \mathbf{e}) + (\mathbf{e}_1^\top \cdot \mathbf{s})$.

algorithm for LPN would allow to learn several important function classes such as 2-DNF formulas, juntas, and any function with a sparse Fourier spectrum. Under a constant noise rate, the best known LPN solvers [9, 39] require time and query complexity both $2^{O(n/\log n)}$. The time complexity goes up to $2^{O(n/\log \log n)}$ when restricted to $q = \text{poly}(n)$ queries [40], or even $2^{O(n)}$ given only $q = O(n)$ queries [42]. Under low noise rate $\mu = n^{-c}$ (for constant $0 < c < 1$), the best attacks [51, 12, 7, 38, 5] solve LPN with time complexity $2^{O(n^{1-c})}$ and query complexity $q = O(n)$ or more¹. The low-noise LPN is mostly believed a stronger assumption than constant-noise LPN. In noise regime $\mu = O(1/\sqrt{n})$, LPN can be used to build public-key encryption (PKE) schemes and oblivious transfer (OT) protocols (more discussions below). Quantum algorithms are not known to have any advantages over classic ones in solving LPN, which makes LPN a promising candidate for “post-quantum cryptography”. Furthermore, LPN enjoys simplicity and is more suited for weak-power devices (e.g., RFID tags) than other quantum-secure candidates such as LWE [48].

CRYPTOGRAPHY IN minicrypt. LPN was used as a basis for building lightweight authentication schemes against passive [29] and even active adversaries (e.g. [32, 34], see [1] for a more complete literature). Kiltz et al. [37] and Dodis et al. [18] constructed randomized MACs from LPN, which implies a two-round authentication scheme with man-in-the-middle security. Lyubashevsky and Masny [41] gave an more efficient three-round authentication scheme from LPN (without going through the MAC transformation) and recently Cash, Kiltz, and Tessaro [13] reduced the round complexity to 2 rounds. Applebaum et al. [3] used LPN to construct efficient symmetric encryption schemes with certain key-dependent message (KDM) security. Jain et al. [30] constructed an efficient perfectly binding string commitment scheme from LPN. We refer to a recent survey [47] on the current state-of-the-art about LPN.

CRYPTOGRAPHY BEYOND minicrypt. Alekhnovich [2] constructed the first (CPA secure) public-key encryption scheme from LPN with noise rate² $\mu = 1/\sqrt{n}$. By plugging the correlated products approach of [49] into Alekhnovich’s CPA secure PKE scheme, Döttling et al. [20] constructed the first CCA secure PKE scheme from low-noise LPN. After observing that the complexity of the scheme in [20] was hundreds of times worse than Alekhnovich’s original scheme, Kiltz et al. [36] proposed a neat and more efficient CCA secure construction by adapting the techniques from LWE-based encryption in [44] to the case of LPN. More recently, Döttling [19] constructed a PKE with KDM security. All the above schemes are based on LPN of noise rate $O(1/\sqrt{n})$. To see that noise rate $1/\sqrt{n}$ is inherently essential for PKE, we illustrate the (weakly correct) single-bit PKE protocol by Döttling et al. [20] in Figure 1, which is inspired

¹We are not aware of any non-trivial time-query tradeoff results to break low-noise LPN in time $2^{o(n^{1-c})}$ even with super-polynomial number of queries.

²More precisely, Alekhnovich’s PKE is based on a variant called the Exact LPN whose noise vector is sampled from $\chi_{\mu q}^q$ for $\mu = \frac{1}{\sqrt{n}}$ (i.e., uniform random distribution over q -bit strings of Hamming weight μq), which is implied by LPN with noise rate μ .

by the counterparts based on LWE [48, 23]. First, the decisional $\text{LPN}_{\mu,n}$ assumption implies that $(\mathbf{A}, \mathbf{A}\mathbf{x} + \mathbf{e})$ is pseudorandom even when \mathbf{x} is drawn from $X \sim \mathcal{B}_{\mu}^n$ (instead of $X \sim U_n$), which can be shown by a simple reduction [20]. Second, the passive security of the protocol is straightforward as (pk, \mathbf{c}_1) is pseudorandom even when concatenated with the Goldreich-Levin³ hardcore bit $\mathbf{s}_1^{\top} \cdot \mathbf{b}$ (replacing \mathbf{b} with U_n by a hybrid argument). The final and most challenging part is correctness, i.e., m' needs to correlate with m at least noticeably. It is not hard to see for $n\mu^2 = O(1)$ and $\mathbf{e}, \mathbf{s} \leftarrow \mathcal{B}_{\mu}^n$ we have $\Pr[\langle \mathbf{e}, \mathbf{s} \rangle = 0] \geq 1/2 + \Omega(1)$, and thus noise rate $\mu = O(1/\sqrt{n})$ seems an inherent barrier⁴ for the PKE to be correct. The scheme is “weak” in the sense that correctness is only $1/2 + \Omega(1)$ and it can be transformed into a standard CPA scheme (that encrypts multiple-bit messages with overwhelming correctness) using standard techniques (e.g., [20, 15]). Notice a correct PKE scheme (with certain properties) yields also a (weak form of) 2-round oblivious transfer protocol against honest-but-curious receiver. Suppose that Alice has a choice $i \in \{0, 1\}$, and she samples pk_i with trapdoor \mathbf{s} (as described in the protocol) and a uniformly random pk_{1-i} without trapdoor. Upon receiving pk_0 and pk_1 , Bob uses the scheme to encrypt two bits σ_0 and σ_1 under pk_0 and pk_1 respectively, and sends them to Alice. Alice can then recover σ_i and but knows nothing about σ_{1-i} . David et al. [16] constructed a universally composable OT under LPN with noise rate $1/\sqrt{n}$. Therefore, basing PKE (and OT) on LPN with noise rate $\mu = n^{-1/2+\epsilon}$ (and ideally a constant $0 < \mu < 1/2$) remains an open problem for the past decade.

LPN WITH AUXILIARY INPUT. Despite being only sub-exponentially secure, LPN is known to be robust against any constant-fraction of static linear leakages, i.e., for any constant $0 < \alpha < 1$ and any $f(\mathbf{x}; \mathbf{Z}) = (\mathbf{Z}, \mathbf{Z}\mathbf{x})$ it holds that

$$(f(\mathbf{x}), \mathbf{A}, \mathbf{A}\mathbf{x} + \mathbf{e}) \stackrel{\epsilon}{\sim} (f(\mathbf{x}), \mathbf{A}, U_q), \quad (1)$$

where \mathbf{Z} is any $(1 - \alpha)n \times n$ matrix (that can be sampled in polynomial time and independent of \mathbf{A}). The above can be seen by a change of basis so that the security is reducible from the LPN assumption with the same noise rate on a uniform secret of size αn . Motivated by this, Dodis, Kalai and Lovett [17] further conjectured that LPN is secure against any polynomial-time computable f such that 1) \mathbf{x} given $f(\mathbf{x})$ has average min-entropy αn ; or even 2) any f that is $2^{-\alpha n}$ -hard-to-invert for PPT algorithms (see Definition 2.2 for a formal definition). Note the distinction between the two types of leakages: the former f is a lossy function and the latter can be even injective (the leakage $f(\mathbf{x})$ may already determine \mathbf{x} in an information theoretical sense). However, they didn’t manage to prove the above claim (i.e., LPN with auxiliary input) under standard LPN. Instead, they introduced a new assumption called Learning Subspace with Noise (LSN) as below, where the secret to be learned is the random subspace \mathbf{V} .

Assumption 1 (The LSN assumption [17]) *For any constant $\beta > 0$, there exists a polynomial $p = p_{\beta}(n)$ such that for any polynomial $q = \text{poly}(n)$ the following two distributions are computationally indistinguishable:*

$$\left((\mathbf{a}_1, \mathbf{V}\mathbf{a}_1 + U_n^{(1)}E_1), \dots, (\mathbf{a}_q, \mathbf{V}\mathbf{a}_q + U_n^{(q)}E_q) \right) \stackrel{\epsilon}{\sim} \left((\mathbf{a}_1, U_n^{(1)}), \dots, (\mathbf{a}_q, U_n^{(q)}) \right),$$

where $\mathbf{V} \sim U_{n \times \beta n}$ is a random $n \times \beta n$ matrix, $\mathbf{a}_1, \dots, \mathbf{a}_q$ are vectors i.i.d. to $U_{\beta n}$, and E_1, \dots, E_q are Boolean variables (determining whether the respective noise is uniform randomness or nothing) i.i.d. to $\mathcal{B}_{1-\frac{1}{p}}$.

³ Typically (in the context of one-way functions), the Goldreich-Levin Theorem [25] assumes a uniformly random secret \mathbf{s} , which is however not necessary. A Markov argument suggests that \mathbf{s} can follow any polynomial-time sampleable distribution, as long as f on \mathbf{s} is hard to invert.

⁴In fact, $\mu = O(\sqrt{\log n/n})$ is sufficient to have a noticeable correctness, i.e., $1/2 + 1/\text{poly}(n)$, but known PKE constructions avoid the strong noise by assuming noise rate $n^{-1/2}$ [36] or even lower rate $n^{-1/2-\epsilon}$ [2, 20].

Then, the authors of [17] showed that LSN with parameters β and $p_\beta = p_\beta(n)$ implies the decisional LPN (as in (1)) under noise rate $\mu = (\frac{1}{2} - \frac{1}{4p_\beta})$ holds with $2^{-\alpha n}$ -hard-to-invert auxiliary input (for any constant $\alpha > \beta$). Further, this yields many interesting applications such as CPA/CCA secure symmetric encryption schemes, average-case obfuscators for the class of point functions, reusable and robust extractors, all remain secure with exponentially hard-to-invert auxiliary input (see [17] for more details). We note that [17] mainly established the feasibility about cryptography with auxiliary input, and there remain issues to be addressed or improved. First, to counteract $2^{-\alpha n}$ -hard-to-invert auxiliary input one needs to decide in advance the noise rate $1/2 - 1/4p_\beta$ (recall the constraint $\beta < \alpha$). Second, Raz showed that for any constant β , $p_\beta = n^{\Omega(1)}$ is necessary (otherwise LSN can be broken in polynomial-time) and even with $p_\beta = n^{\Theta(1)}$ there exist quasi-polynomial attacks (see the full version of [17] for more discussions about Raz’s attacks). Therefore, the security reduction in [17] is quite loose. As the main end result of [17], one needs a high-noise LPN for $\mu = 1/2 - 1/\text{poly}(n)$ (and thus low efficiency due to the redundancy needed to make a correct scheme) only to achieve quasi-polynomial security (due to Raz’s attacks) against $2^{-\alpha n}$ -hard-to-invert leakage for some constant α (i.e., not any exponentially hard-to-invert leakage). Third, LSN is a new (and less well-studied) assumption and it was left as an open problem in [17] whether the aforementioned cryptographic applications can be based on the hardness of standard LPN, ideally admitting more general class of leakages, such as sub-exponentially or even quasi-polynomially hard-to-invert auxiliary input.

THE MAIN OBSERVATION. Yu [52] introduced the “sampling from subspace” technique to prove the above “LPN with auxiliary input” conjecture under standard LPN but the end result of [52] was invalid due to a flawed intermediate step. A similar idea was also used by Goldwasser et al. [26] in the setting of LWE, where the public matrix was drawn from a (noisy) random subspace. Informally, the observation (in our setting) is that, the decisional LPN with constant noise rate $0 < \mu < 1/2$ implies that for any constant $0 < \alpha < 1$, any 2^{-2n^α} -hard-to-invert f and any $q' = \text{poly}(n)$ it holds that

$$(f(\mathbf{x}), \mathbf{A}', \mathbf{A}' \cdot \mathbf{x} + \mathbf{e}) \stackrel{c}{\sim} (f(\mathbf{x}), \mathbf{A}', U_{q'}), \quad (2)$$

where $\mathbf{x} \sim U_n$ ⁵, $\mathbf{e} \sim \mathcal{B}_\mu^{q'}$, and \mathbf{A}' is a $q' \times n$ matrix with rows sampled from a random subspace of dimension $\lambda = n^\alpha$. Further, if the underlying LPN is $2^{\omega(n^{\frac{1}{1+\beta}})}$ -hard⁶ for any constant $\beta > 0$, then by setting $\lambda = \log^{1+\beta} n$, (2) holds for any $q' = \text{poly}(n)$ and any $2^{-2 \log^{1+\beta} n}$ -hard-to-invert f . The rationale is that distribution \mathbf{A}' can be considered as the multiplication of two random matrices $\mathbf{A} \stackrel{\$}{\leftarrow} \{0,1\}^{q' \times \lambda}$ and $\mathbf{V} \stackrel{\$}{\leftarrow} \{0,1\}^{\lambda \times n}$, i.e., $\mathbf{A}' \sim (\mathbf{A} \cdot \mathbf{V})$, where \mathbf{V} constitutes the basis of the λ -dimensional random subspace and \mathbf{A} is the random coin for sampling from \mathbf{V} . Unlike the LSN assumption whose subspace \mathbf{V} is secret, the \mathbf{V} and \mathbf{A} in (2) are public coins (implied by \mathbf{A}' , see Remark 3.1). We have by the associative law $\mathbf{A}' \cdot \mathbf{x} = \mathbf{A}(\mathbf{V} \cdot \mathbf{x})$ and by the Goldreich-Levin theorem $\mathbf{V} \cdot \mathbf{x}$ is a pseudorandom secret (even conditioned on \mathbf{V} and $f(\mathbf{x})$), and thus (2) is reducible from the standard decisional LPN on noise rate μ , secret size λ and query complexity q' . Concretely, assume that the LPN problem is $2^{\omega(n^{3/4})}$ -hard then by setting $\lambda = n^{2/3}$ (resp., $\lambda = \log^{4/3} n$) we have that (2) is $2^{\Omega(n^{1/2})}$ -secure (resp., $n^{\omega(1)}$ -secure) with any auxiliary input that is $2^{-2n^{2/3}}$ -hard (resp., $2^{-2 \log^{4/3} n}$ -hard) to invert. Plugging (2) into the framework of [17] we obtain the same applications (CPA/CCA secure symmetric encryption schemes, average-case obfuscators for point functions, reusable and robust extractors) under standard (constant-noise) LPN with improved efficiency (as the noise is constant rather than

⁵We assume $\mathbf{x} \sim U_n$ to be in line with [17], but actually our results hold for any efficiently sampleable \mathbf{x} as long as \mathbf{x} given $f(\mathbf{x})$ is $2^{-2\lambda}$ -hard-to-invert.

⁶Informally, we say that a cryptographic scheme/problem Π is T -secure/hard if every probabilistic adversary of time (and query, if applicable) complexity T achieve advantage no more than $1/T$ in breaking/solving Π .

polynomially close to uniform) and tighter security against sub-exponentially (or even quasi-polynomially) hard-to-invert auxiliary input.

PKE FROM CONSTANT-NOISE LPN. More surprisingly, we show a connection from “LPN with auxiliary input” to “basing PKE on (constant-noise) LPN”. The feasibility can be understood by the single-bit weak PKE in [Figure 1](#) with some modifications: assume that LPN is $2^{\omega(n^{\frac{1}{2}})}$ -hard (i.e., $\beta = 1$), then for $\lambda = \log^2 n/4$ we have that [\(2\)](#) holds on any $\mathbf{x} \sim X$ with min-entropy $\mathbf{H}_\infty(X) \geq \log^2 n/2$. Therefore, by replacing the uniform matrix \mathbf{A} with $\mathbf{A}' \sim (U_{n \times \lambda} \cdot U_{\lambda \times n})$, and sampling $\mathbf{s}, \mathbf{s}_1 \leftarrow X$ and $\mathbf{e}, \mathbf{e}_1 \leftarrow \mathcal{B}_\mu^n$ for constant μ and $X \sim \chi_{\log n}^n$ ⁷, we get that $\mathbf{s}_1^\top \mathbf{e}$ and $\mathbf{e}_1^\top \mathbf{s}$ are both $(1/2 + 1/\text{poly}(n))$ -biased to 0 independently, and thus the PKE scheme has noticeable correctness. We then transform the weak PKE into a full-fledged CPA secure scheme, where the extension is not trivial (more than a straightforward parallel repetition plus error-correction codes). In particular, neither $X \sim \chi_{\log n}^n$ or $X \sim \mathcal{B}_{\log n/n}^n$ can guarantee security and correctness simultaneously and thus additional ideas are needed (more details deferred to [Section 4.3](#)).

PKE WITH CCA SECURITY. Once we have a CPA scheme based on constant-noise LPN, we can easily extend it to a CCA one by using the techniques in [\[20\]](#), and thus suffer from the same performance slowdown as that in [\[20\]](#). A natural question is whether we can construct a simpler and more efficient CCA scheme as that in [\[36\]](#). Unfortunately, the techniques in [\[36\]](#) do not immediately apply to the case of constant-noise LPN. The reason is that in order to employ the ideas from the LWE-based encryption scheme [\[44\]](#), the scheme in [\[36\]](#) has to use a variant of LPN (called knapsack LPN), and the corresponding description key is exactly the secret of some knapsack LPN instances. Even though there is a polynomial time reduction [\[43\]](#) from the LPN problem to the knapsack LPN problem, such a reduction will map the noise distribution of the LPN problem into the secret distribution of the knapsack LPN problem. If we directly apply the techniques in [\[36\]](#), the resulting scheme will not have any guarantee of correctness because the corresponding decryption key follows the Bernoulli distribution with constant parameter μ . Recall that for the correctness of our CPA secure PKE scheme, the decryption key cannot simply be chosen from either $\chi_{\log n}^n$ or $\mathcal{B}_{\log n/n}^n$. Fortunately, based on several new observations and some new technical lemmas, we manage to adapt the idea of [\[44, 36\]](#) to construct a simpler and efficient CCA secure PKE scheme from constant-noise LPN.

OT FROM CONSTANT-NOISE LPN. PKE and OT are incomparable in general [\[24\]](#). But if the considered PKE scheme has some additional properties, then we can build OT protocol from it in a black-box way [\[24\]](#). More concretely, Gertner et al. [\[24\]](#) showed that if the public key of some CPA secure PKE scheme can be sampled (without knowing the corresponding secret key) from a distribution that is indistinguishable from that produced by the key generation algorithm, then we can use it to construct an OT protocol with honest parties (and thus can be transformed into a standard OT protocol by using zero-knowledge proof). It is easy to check that our CPA secure PKE scheme satisfies this property under the LPN assumption. Besides, none of the techniques used in transforming Alekhnovich’s CPA secure PKE scheme into a universally composable OT protocol [\[16\]](#) prevent us from obtaining a universally composable OT protocol from our CPA secure PKE scheme. In summary, our results imply that there exists (universally composable) OT protocol under constant-noise LPN assumption. We omit the details, and refer to [\[24, 16\]](#) for more information.

⁷Recall that for $m \ll n$ we have by Stirling’s approximation that $\binom{n}{m} \approx n^m/m!$ and thus $\chi_{\log n}^n$ (uniform distribution over n -bit strings of Hamming weight $\log n$) is of min-entropy roughly $\log^2 n - \log n \log \log n \geq \log^2 n/2$.

2 Preliminaries

NOTATIONS AND DEFINITIONS. We use capital letters (e.g., X, Y) for random variables and distributions, standard letters (e.g., x, y) for values, and calligraphic letters (e.g. \mathcal{X}, \mathcal{E}) for sets and events. Vectors are used in the column form and denoted by bold lower-case letters (e.g., \mathbf{a}). We treat matrices as the sets of its column vectors and denoted by bold capital letters (e.g., \mathbf{A}). The support of a random variable X , denoted by $\text{Supp}(X)$, refers to the set of values on which X takes with non-zero probability, i.e., $\{x : \Pr[X = x] > 0\}$. For set \S and binary string s , $|\S|$ denotes the cardinality of \S and $|s|$ refers to the Hamming weight of s . We use \mathcal{B}_μ to denote the Bernoulli distribution with parameter μ , i.e., $\Pr[\mathcal{B}_\mu = 1] = \mu$, $\Pr[\mathcal{B}_\mu = 0] = 1 - \mu$, while \mathcal{B}_μ^q denotes the concatenation of q independent copies of \mathcal{B}_μ . We use χ_i^n to denote a uniform distribution over $\{\mathbf{e} \in \{0, 1\}^n : |\mathbf{e}| = i\}$. We denote by $\mathcal{D}_\lambda^{n_1 \times n} \stackrel{\text{def}}{=} (U_{n_1 \times \lambda} \cdot U_{\lambda \times n})$ to be a matrix distribution induced by multiplying two random matrices. For $n, q \in \mathbb{N}$, U_n (resp., $U_{q \times n}$) denotes the uniform distribution over $\{0, 1\}^n$ (resp., $\{0, 1\}^{q \times n}$) and independent of any other random variables in consideration, and $f(U_n)$ (resp., $f(U_{q \times n})$) denotes the distribution induced by applying function f to U_n (resp., $U_{q \times n}$). $X \sim D$ denotes that random variable X follows distribution D . We use $s \leftarrow S$ to denote sampling an element s according to distribution S , and let $s \stackrel{\S}{\leftarrow} \S$ denote sampling s uniformly from set \S .

ENTROPY NOTIONS. For $0 < \mu < 1/2$, the binary entropy function is defined as $\mathbf{H}(\mu) \stackrel{\text{def}}{=} \mu \log(1/\mu) + (1 - \mu) \log(1/(1 - \mu))$. We define the Shannon entropy and min-entropy of a random variable X respectively, i.e.,

$$\mathbf{H}_1(X) \stackrel{\text{def}}{=} \sum_{x \in \text{Supp}(X)} \Pr[X = x] \log \frac{1}{\Pr[X = x]}, \quad \mathbf{H}_\infty(X) \stackrel{\text{def}}{=} \min_{x \in \text{Supp}(X)} \log(1/\Pr[X = x]) .$$

Note that $\mathbf{H}_1(\mathcal{B}_\mu) = \mathbf{H}(\mu)$. The average min-entropy of a random variable X conditioned on another random variable Z is defined as

$$\mathbf{H}_\infty(X|Z) \stackrel{\text{def}}{=} -\log \left(\mathbb{E}_{z \leftarrow Z} \left[2^{-\mathbf{H}_\infty(X|Z=z)} \right] \right) .$$

INDISTINGUISHABILITY AND STATISTICAL DISTANCE. We define the (t, ε) - *computational distance* between random variables X and Y , denoted by $X \underset{(t, \varepsilon)}{\sim} Y$, if for every probabilistic distinguisher \mathcal{D} of running time t it holds that

$$| \Pr[\mathcal{D}(X) = 1] - \Pr[\mathcal{D}(Y) = 1] | \leq \varepsilon .$$

The *statistical distance* between X and Y , denoted by $\text{SD}(X, Y)$, is defined by

$$\text{SD}(X, Y) \stackrel{\text{def}}{=} \frac{1}{2} \sum_x |\Pr[X = x] - \Pr[Y = x]| .$$

Computational/statistical indistinguishability is defined with respect to distribution ensembles (indexed by a security parameter). For example, $X \stackrel{\text{def}}{=} \{X_n\}_{n \in \mathbb{N}}$ and $Y \stackrel{\text{def}}{=} \{Y_n\}_{n \in \mathbb{N}}$ are computationally indistinguishable, denoted by $X \stackrel{c}{\sim} Y$, if for every $t = \text{poly}(n)$ there exists $\varepsilon = \text{negl}(n)$ such that $X \underset{(t, \varepsilon)}{\sim} Y$. X and Y are statistically indistinguishable, denoted by $X \stackrel{s}{\sim} Y$, if $\text{SD}(X, Y) = \text{negl}(n)$.

SIMPLIFYING NOTATIONS. To simplify the presentation, we use the following simplified notations. Throughout, n is the security parameter and most other parameters are functions of n , and we often omit n when clear from the context. For example, $q = q(n) \in \mathbb{N}$, $t = t(n) > 0$, $\varepsilon = \varepsilon(n) \in (0, 1)$, and $m = m(n) = \text{poly}(n)$, where poly refers to some polynomial.

Definition 2.1 (Learning Parity with Noise) The *decisional* $\text{LPN}_{\mu,n}$ problem (with secret length n and noise rate $0 < \mu < 1/2$) is hard if for every $q = \text{poly}(n)$ we have

$$(\mathbf{A}, \mathbf{A} \cdot \mathbf{x} + \mathbf{e}) \stackrel{c}{\sim} (\mathbf{A}, U_q) , \quad (3)$$

where $q \times n$ matrix $\mathbf{A} \sim U_{q \times n}$, $\mathbf{x} \sim U_n$ and $\mathbf{e} \sim \mathcal{B}_\mu^q$. The *computational* $\text{LPN}_{\mu,n}$ problem is hard if for every $q = \text{poly}(n)$ and every PPT algorithm \mathcal{D} we have

$$\Pr[\mathcal{D}(\mathbf{A}, \mathbf{A} \cdot \mathbf{x} + \mathbf{e}) = \mathbf{x}] = \text{negl}(n) , \quad (4)$$

where $\mathbf{A} \sim U_{q \times n}$, $\mathbf{x} \sim U_n$ and $\mathbf{e} \sim \mathcal{B}_\mu^q$.

LPN WITH SPECIFIC HARDNESS. We say that the decisional (resp., computational) $\text{LPN}_{\mu,n}$ is T -hard if for every $q \leq T$ and every probabilistic adversary of running time T the distinguishing (resp., inverting) advantage in (3) (resp., (4)) is upper bounded by $1/T$.

Definition 2.2 (Hard-to-invert function) Let n be the security parameter and let $\kappa = \omega(\log n)$. A polynomial-time computable function $f : \{0, 1\}^n \rightarrow \{0, 1\}^l$ is $2^{-\kappa}$ -hard-to-invert if for every PPT adversary \mathcal{A}

$$\Pr_{\mathbf{x} \sim U_n} [\mathcal{A}(f(\mathbf{x})) = \mathbf{x}] \leq 2^{-\kappa} .$$

Lemma 2.1 (Chernoff bound) For any $n \in \mathbb{N}$, let X_1, \dots, X_n be independent random variables and let $\bar{X} = \sum_{i=1}^n X_i$, where $\Pr[0 \leq X_i \leq 1] = 1$ holds for every $1 \leq i \leq n$. Then, for any $\Delta_1 > 0$ and $0 < \Delta_2 < 1$,

$$\begin{aligned} \Pr[\bar{X} > (1 + \Delta_1) \cdot \mathbb{E}[\bar{X}]] &< \exp^{-\frac{\min(\Delta_1, \Delta_1^2)}{3} \mathbb{E}[\bar{X}]} , \\ \Pr[\bar{X} < (1 - \Delta_2) \cdot \mathbb{E}[\bar{X}]] &< \exp^{-\frac{\Delta_2^2}{2} \mathbb{E}[\bar{X}]} . \end{aligned}$$

Lemma 2.2 (Union bound) Let $\mathcal{E}_1, \dots, \mathcal{E}_l$ be any (not necessarily independent) events such that $\Pr[\mathcal{E}_i] \geq (1 - \epsilon_i)$ for every $1 \leq i \leq l$, then we have

$$\Pr[\mathcal{E}_1 \wedge \dots \wedge \mathcal{E}_l] \geq 1 - (\epsilon_1 + \dots + \epsilon_l) .$$

We will use the following (essentially the Hoeffding's) bound on the Hamming weight of a high-noise Bernoulli vector.

Lemma 2.3 For any $0 < p < 1/2$ and $\delta \leq (\frac{1}{2} - p)$, we have

$$\Pr[|\mathcal{B}_\delta^q| > (\frac{1}{2} - \frac{p}{2})q] < \exp^{-\frac{p^2 q}{8}} .$$

3 Learning Parity with Noise with Auxiliary Input

3.1 Leaky Sources and (Pseudo)randomness Extraction

We define below two types of leaky sources and recall two technical lemmas for (pseudo)randomness extraction from the respective sources, where \mathbf{x} for TYPE-II source is assumed to be uniform only for alignment with [17] (see Footnote 3).

Definition 3.1 (Leaky sources) Let \mathbf{x} be any random variable over $\{0, 1\}^n$ and let $f : \{0, 1\}^n \rightarrow \{0, 1\}^l$ be any polynomial-time computable function. $(\mathbf{x}, f(\mathbf{x}))$ is called an (n, κ) TYPE-I (resp., TYPE-II) leaky source if it satisfies condition 1 (resp., condition 2) below:

1. *Min-entropy leaky sources.* $\mathbf{H}_\infty(\mathbf{x}|f(\mathbf{x})) \geq \kappa$ and $f(\mathbf{x})$ is polynomial-time sampleable.

2. **Hard-to-invert leaky sources.** $\mathbf{x} \sim U_n$ and f is $2^{-\kappa}$ -hard-to-invert.

Lemma 3.1 (Goldreich-Levin Theorem [25]) *Let n be a security parameter, let $\kappa = \omega(\log n)$ be polynomial-time computable from n , and let $f : \{0, 1\}^n \rightarrow \{0, 1\}^l$ be any polynomial-time computable function that is $2^{-\kappa}$ -hard-to-invert. Then, for any constant $0 < \beta < 1$ and $\lambda = \lceil \beta \kappa \rceil$, it holds that*

$$(f(\mathbf{x}), \mathbf{V}, \mathbf{V} \cdot \mathbf{x}) \stackrel{c}{\sim} (f(\mathbf{x}), \mathbf{V}, U_\lambda) ,$$

where $\mathbf{x} \sim U_n$ and $\mathbf{V} \sim U_{\lambda \times n}$ is a random $\lambda \times n$ Boolean matrix.

Lemma 3.2 (Leftover hash lemma [28]) *Let $(X, Z) \in \mathcal{X} \times \mathcal{Z}$ be any joint random variable with $\mathbf{H}_\infty(X|Z) \geq k$, and let $\mathcal{H} = \{h_{\mathbf{V}} : \mathcal{X} \rightarrow \{0, 1\}^l, \mathbf{V} \in \{0, 1\}^s\}$ be a family of universal hash functions, i.e., for any $x_1 \neq x_2 \in \mathcal{X}$, $\Pr_{\mathbf{V} \leftarrow \{0, 1\}^s} [h_{\mathbf{V}}(x_1) = h_{\mathbf{V}}(x_2)] \leq 2^{-l}$. Then, it holds that*

$$\text{SD} \left((Z, \mathbf{V}, h_{\mathbf{V}}(X)) , (Z, \mathbf{V}, U_l) \right) \leq 2^{l-k} ,$$

where $\mathbf{V} \sim U_s$.

3.2 The Main Technical Lemma and Immediate Applications

Inspired by [52, 26], we state a technical lemma below where the main difference is that we sample from a random subspace of sublinear-sized dimension (rather than linear-sized one [52] or from a noisy subspace in the LWE setting [26]).

Theorem 3.1 (LPN with hard-to-invert auxiliary input) *Let n be a security parameter and let $0 < \mu < 1/2$ be any constant. Assume that the decisional $\text{LPN}_{\mu, n}$ problem is hard, then for every constant $0 < \alpha < 1$, $\lambda = n^\alpha$, $q' = \text{poly}(n)$, and every $(n, 2\lambda)$ TYPE-I or TYPE-II leaky source $(\mathbf{x}, f(\mathbf{x}))$, we have*

$$(f(\mathbf{x}), \mathbf{A}', \mathbf{A}' \cdot \mathbf{x} + \mathbf{e}) \stackrel{c}{\sim} (f(\mathbf{x}), \mathbf{A}', U_{q'}) , \quad (5)$$

where $\mathbf{e} \sim \mathcal{B}_\mu^{q'}$, and $\mathbf{A}' \sim \mathcal{D}_\lambda^{q' \times n}$ is a $q' \times n$ matrix, i.e., $\mathbf{A}' \sim (\mathbf{A} \cdot \mathbf{V})$ for random matrices $\mathbf{A} \leftarrow \{0, 1\}^{q' \times \lambda}$ and $\mathbf{V} \leftarrow \{0, 1\}^{\lambda \times n}$.

Furthermore, if the $\text{LPN}_{\mu, n}$ problem is $2^{\omega(n^{\frac{1}{1+\beta}})}$ -hard for any constant $\beta > 0$ and any superconstant hidden by $\omega(\cdot)$ then the above holds for any $\lambda = \Theta(\log^{1+\beta} n)$, any $q' = \text{poly}(n)$ and any $(n, 2\lambda)$ TYPE-I/TYPE-II leaky source.

Remark 3.1 (Closure under composition) *The random subspace \mathbf{V} and the random coin \mathbf{A} can be public as well, which is seen from the proof below but omitted from (5) to avoid redundancy (since they are implied by \mathbf{A}'). That is, there exists a PPT Simu such that $(\mathbf{A}', \text{Simu}(\mathbf{A}'))$ is $2^{-\Omega(n)}$ -close to $(\mathbf{A}', (\mathbf{A}, \mathbf{V}))$. Therefore, (5) can be written in an equivalent form that is closed under composition, i.e., for any $q' = \text{poly}(n)$ and $l = \text{poly}(n)$*

$$\left(f(\mathbf{x}), \mathbf{V}, (\mathbf{A}_i, (\mathbf{A}_i \cdot \mathbf{V}) \cdot \mathbf{x} + \mathbf{e}_i)_{i=1}^l \right) \stackrel{c}{\sim} \left(f(\mathbf{x}), \mathbf{V}, (\mathbf{A}_i, U_{q'}^{(i)})_{i=1}^l \right) ,$$

where $\mathbf{A}_1, \dots, \mathbf{A}_l \leftarrow \{0, 1\}^{q' \times \lambda}$, $\mathbf{e}_1, \dots, \mathbf{e}_l \sim \mathcal{B}_\mu^{q'}$ and $\mathbf{V} \leftarrow \{0, 1\}^{\lambda \times n}$. This will be a useful property for constructing symmetric encryption schemes w.r.t. sub-exponentially hard-to-invert auxiliary input (details deferred to [Appendix C](#)).

Proof of Theorem 3.1. We have by the assumption of $(\mathbf{x}, f(\mathbf{x}))$ and Lemma 3.1 or Lemma 3.2 that

$$\begin{aligned} & (f(\mathbf{x}), \mathbf{V}, \mathbf{V} \cdot \mathbf{x}) \stackrel{c}{\sim} (f(\mathbf{x}), \mathbf{V}, \mathbf{y}) \\ \Rightarrow & (f(\mathbf{x}), (\mathbf{A}, \mathbf{V}), (\mathbf{A} \cdot \mathbf{V}) \cdot \mathbf{x} + \mathbf{e}) \stackrel{c}{\sim} (f(\mathbf{x}), (\mathbf{A}, \mathbf{V}), \mathbf{A} \cdot \mathbf{y} + \mathbf{e}) . \end{aligned}$$

where $\mathbf{y} \sim U_\lambda$. Next, consider T -hard decisional $\text{LPN}_{\mu, \lambda}$ problem on uniform secret \mathbf{y} of length λ (instead of n), which postulates that for any $q' \leq T$

$$\begin{aligned} & (\mathbf{A}, \mathbf{A} \cdot \mathbf{y} + \mathbf{e}) \underset{T, 1/T}{\sim} (\mathbf{A}, U_{q'}) \\ \Rightarrow & (f(\mathbf{x}), (\mathbf{A}, \mathbf{V}), \mathbf{A} \cdot \mathbf{y} + \mathbf{e}) \underset{T - \text{poly}(n), 1/T}{\sim} (f(\mathbf{x}), (\mathbf{A}, \mathbf{V}), U_{q'}) . \end{aligned}$$

Under the LPN assumption with standard asymptotic hardness (i.e., $T = \lambda^{\omega(1)}$) and by setting parameter $\lambda = n^\alpha$ we have $T = n^{\omega(1)}$, which suffices for our purpose since for any $q' = \text{poly}(n)$, any PPT adversary wins the above distinguishing game with advantage no greater than $n^{-\omega(1)}$. In case that $\text{LPN}_{\mu, \lambda}$ is $2^{\omega(n^{\frac{1}{1+\beta}})}$ -hard, substitution of $\lambda = \Theta(\log^{1+\beta} n)$ into $T = 2^{\omega(\lambda^{\frac{1}{1+\beta}})}$ also yields $T = n^{\omega(1)}$. Therefore, in both cases the above two ensembles are computationally indistinguishable in security parameter n . The conclusion then follows by a triangle inequality. \square

A COMPARISON WITH [17]. The work of [17] proved results similar to Theorem 3.1. In particular, [17] showed that the LSN assumption with parameters β and $p = \text{poly}_\beta(n)$ implies LPN with $2^{-\alpha n}$ -hard auxiliary input (for constant $\alpha > \beta$), noise rate $\mu = 1/2 - 1/4p$ and quasi-polynomial security (in essentially the same form as (5) except for a uniform matrix \mathbf{A}'). In comparison, by sampling \mathbf{A}' from a random subspace of sublinear dimension $\lambda = n^\alpha$ (for $0 < \alpha < 1$), constant-noise LPN implies that (5) holds with $2^{-\Omega(n^\alpha)}$ -hard auxiliary input, constant noise and comparable security to the underlying LPN. Furthermore, assume constant-noise LPN with $2^{\omega(n^{\frac{1}{1+\beta}})}$ -hardness (for constant $\beta > 0$), then (2) holds for $2^{-\Omega(\log^{1+\beta})}$ -hard auxiliary input, constant noise and quasi-polynomial security.

IMMEDIATE APPLICATIONS. This yields the same applications as considered in [17], such as CPA/CCA secure symmetric encryption schemes, average-case obfuscators for point functions, reusable and robust extractors, all under standard (constant-noise) LPN with improved efficiency (by bringing down the noise rate) and tighter security against sub-exponentially (or even quasi-polynomially) hard-to-invert auxiliary input. The proofs simply follow the route of [17] and can be informally explained as: the technique (by sampling from random subspace) implicitly applies pseudorandomness extraction (i.e., $\mathbf{y} = \mathbf{V} \cdot \mathbf{x}$) so that the rest of the scheme is built upon the security of $(\mathbf{A}, \mathbf{A} \cdot \mathbf{y} + \mathbf{e})$ on secret \mathbf{y} (which is pseudorandom even conditioned on the leakage), and thus the task is essentially to obtain the aforementioned applications from standard LPN (without auxiliary input). In other words, our technique allows to transform any applications based on constant-noise LPN into the counterparts with auxiliary input under the same assumption. Therefore, we only sketch some applications in Appendix C and refer to [17] for the redundancy.

4 CPA Secure PKE from Constant-Noise LPN

We show a more interesting application, namely, to build public-key encryption schemes from constant-noise LPN, which has been an open problem since the work of [2]. We refer to Appendix A.2 for standard definitions of public-key encryption schemes, correctness and CPA/CCA security.

4.1 Technical Lemmas

We use the following technical tool to build PKE scheme from constant-noise LPN. It would have been an immediate corollary of [Theorem 3.1](#) for sub-exponential hard LPN on squared-logarithmic min-entropy sources (i.e., $\beta = 1$), except for the fact that the leakage is also correlated with noise. Notice that we lose the ‘‘closure under composition’’ property by allowing leakage to be correlated with noise, and thus our PKE scheme will avoid this property.

Theorem 4.1 (LPN on squared-log entropy) *Let n be a security parameter and let $0 < \mu < 1/2$ be any constant. Assume that the computational $\text{LPN}_{\mu,n}$ problem is $2^{\omega(n^{\frac{1}{2}})}$ -hard (for any superconstant hidden by $\omega(\cdot)$), then for every $\lambda = \Theta(\log^2 n)$, $q' = \text{poly}(n)$, and every polynomial-time sampleable $\mathbf{x} \in \{0, 1\}^n$ with $\mathbf{H}_\infty(\mathbf{x}) \geq 2\lambda$ and every probabilistic polynomial-time computable function $f : \{0, 1\}^{n+q'} \times \mathcal{Z} \rightarrow \{0, 1\}^{O(\log n)}$ with public coin Z , we have*

$$(f(\mathbf{x}, \mathbf{e}; Z), Z, \mathbf{A}', \mathbf{A}' \cdot \mathbf{x} + \mathbf{e}) \stackrel{c}{\sim} (f(\mathbf{x}, \mathbf{e}; Z), Z, \mathbf{A}', U_{q'}),$$

where noise vector $\mathbf{e} \sim \mathcal{B}_\mu^{q'}$ and $q' \times n$ matrix $\mathbf{A}' \sim \mathcal{D}_\lambda^{q' \times n}$.

Proof sketch. It suffices to adapt the proof of [Theorem 3.1](#) as follows. First, observe that (by the chain rule of min-entropy)

$$\mathbf{H}_\infty(\mathbf{x} | f(\mathbf{x}, \mathbf{e}; Z), Z, \mathbf{e}) \geq \mathbf{H}_\infty(\mathbf{x} | Z, \mathbf{e}) - O(\log n) = \mathbf{H}_\infty(\mathbf{x}) - O(\log n) \geq 2\lambda - O(\log n).$$

For our convenience, write $\mathbf{A}' \sim (\mathbf{A} \cdot \mathbf{V})$ for $\mathbf{A} \sim U_{q' \times \lambda}$, $\mathbf{V} \sim U_{\lambda \times n}$, and let $\mathbf{y}, \mathbf{r} \sim U_\lambda$. Then, we have by [Lemma 3.2](#)

$$\begin{aligned} & (f(\mathbf{x}, \mathbf{e}; Z), Z, \mathbf{e}, \mathbf{V}, \mathbf{V} \cdot \mathbf{x}) \stackrel{s}{\sim} (f(\mathbf{x}, \mathbf{e}; Z), Z, \mathbf{e}, \mathbf{V}, \mathbf{y}) \\ \Rightarrow & (f(\mathbf{x}, \mathbf{e}; Z), Z, (\mathbf{A} \cdot \mathbf{V}), (\mathbf{A} \cdot \mathbf{V}) \cdot \mathbf{x} + \mathbf{e}) \stackrel{s}{\sim} (f(\mathbf{x}, \mathbf{e}; Z), Z, (\mathbf{A} \cdot \mathbf{V}), \mathbf{A} \cdot \mathbf{y} + \mathbf{e}). \end{aligned}$$

Next, $2^{\omega(\lambda^{\frac{1}{2}})}$ -hard computational $\text{LPN}_{\mu,\lambda}$ problem with secret size λ postulates that for any $q' \leq 2^{\omega(\lambda^{\frac{1}{2}})} = n^{\omega(1)}$ (recall $\lambda = \Theta(\log^2 n)$) and any probabilistic $\mathcal{D}, \mathcal{D}'$ of running time $n^{\omega(1)}$

$$\begin{aligned} & \Pr[\mathcal{D}'(\mathbf{A}, \mathbf{A} \cdot \mathbf{y} + \mathbf{e}) = \mathbf{y}] = n^{-\omega(1)} \\ \Rightarrow & \Pr[\mathcal{D}'(f(\mathbf{x}, \mathbf{e}; Z), Z, \mathbf{A}, \mathbf{A} \cdot \mathbf{y} + \mathbf{e}) = \mathbf{y}] = n^{-\omega(1)} \\ \Rightarrow & (f(\mathbf{x}, \mathbf{e}; Z), Z, \mathbf{A}, \mathbf{A} \cdot \mathbf{y} + \mathbf{e}, \mathbf{r}, \mathbf{r}^\top \cdot \mathbf{y}) \stackrel{c}{\sim} (f(\mathbf{x}, \mathbf{e}; Z), Z, \mathbf{A}, \mathbf{A} \cdot \mathbf{y} + \mathbf{e}, \mathbf{r}, U_1) \\ \Rightarrow & (f(\mathbf{x}, \mathbf{e}; Z), Z, \mathbf{A}, \mathbf{A} \cdot \mathbf{y} + \mathbf{e}) \stackrel{c}{\sim} (f(\mathbf{x}, \mathbf{e}; Z), Z, \mathbf{A}, U_{q'}) \\ \Rightarrow & (f(\mathbf{x}, \mathbf{e}; Z), Z, (\mathbf{A} \cdot \mathbf{V}), \mathbf{A} \cdot \mathbf{y} + \mathbf{e}) \stackrel{c}{\sim} (f(\mathbf{x}, \mathbf{e}; Z), Z, (\mathbf{A} \cdot \mathbf{V}), U_{q'}) \quad , \end{aligned}$$

where the first implication is trivial since Z is independent of $(\mathbf{A}, \mathbf{y}, \mathbf{e})$ and any $O(\log n)$ bits of leakage affects unpredictability by a fact of $\text{poly}(n)$, the second step is the Goldreich-Levin theorem [25] with $\mathbf{r} \sim U_\lambda$, and the third implication uses the sample-preserving reduction from [4] and is reproduced as [Lemma B.1](#). The conclusion follows by a triangle inequality. \square

We will use [Lemma 4.1](#) to estimate the noise rate of an inner product between Bernoulli-like vectors .

Lemma 4.1 *For any $0 < \mu \leq 1/8$ and $\ell \in \mathbb{N}$, let E_1, \dots, E_ℓ be Boolean random variables i.i.d. to \mathcal{B}_μ , then $\Pr[\bigoplus_{i=1}^\ell E_i = 0] > \frac{1}{2} + 2^{-(4\mu\ell+1)}$.*

Proof. We complete the proof by [Fact 1](#) and [Fact 2](#)

$$\Pr\left[\bigoplus_{i=1}^\ell E_i = 1\right] = \frac{1}{2}(1 - (1 - 2\mu)^\ell) < \frac{1}{2}(1 - 2^{-4\mu\ell}) = \frac{1}{2} - 2^{-(4\mu\ell+1)} .$$

\square

Fact 1 (Piling-up lemma) For $0 < \mu < 1/2$ and random variables E_1, E_2, \dots, E_ℓ that are i.i.d. to \mathcal{B}_μ we have $\bigoplus_{i=1}^\ell E_i \sim \mathcal{B}_\sigma$ with $\sigma = \frac{1}{2}(1 - (1 - 2\mu)^\ell)$.

Fact 2 (Mean value theorem) For any $0 < x \leq 1/4$ we have $1 - x > 2^{-2x}$.

We recall the following facts about the entropy of Bernoulli-like distributions. In general, there's no closed formula for binomial coefficient, but an asymptotic estimation like [Fact 3](#) already suffices for our purpose, where the binary entropy function can be further bounded by [Fact 4](#) (see also Footnote 7).

Fact 3 (Asymptotics for binomial coefficients (e.g. [27], p.492)) For any $0 < \mu < 1/2$, and any $n \in \mathbb{N}$

$$\binom{n}{\mu n} = 2^{n\mathbf{H}(\mu) - \frac{\log n}{2} + O(1)}.$$

Fact 4 For any $0 < \mu < 1/2$, we have $\mu \log(1/\mu) < \mathbf{H}(\mu) < \mu(\log(1/\mu) + \frac{3}{2})$.

4.2 Weakly Correct 1-bit PKE from Constant-Noise LPN

As stated in [Theorem 4.1](#), for any constant $0 < \mu < 1/2$, $2^{\omega(n^{\frac{1}{2}})}$ -hard $\text{LPN}_{\mu,n}$ implies that $(\mathbf{A}' \cdot \mathbf{x} + \mathbf{e})$ is pseudorandom conditioned on \mathbf{A}' for $\mathbf{x} \sim X$ with squared-log entropy, where the leakage due to f can be omitted for now as it is only needed for CCA security. If there exists X satisfying the following three conditions at the same time then the 1-bit PKE as in [Figure 1](#) instantiated with the square matrix $\mathbf{A}' \leftarrow \mathcal{D}_\lambda^{n \times n}$, $\mathbf{s}, \mathbf{s}_1 \leftarrow X$ and $\mathbf{e}, \mathbf{e}_1 \leftarrow \mathcal{B}_\mu^n$ will be secure and noticeably correct (since $\mathbf{s}_1^\top \mathbf{e}$ and $\mathbf{e}_1^\top \mathbf{s}$ are both $(1/2 + 1/\text{poly}(n))$ -biased to 0 independently).

1. **(Efficiency)** $X \in \{0, 1\}^n$ can be sampled in polynomial time.
2. **(Security)** $\mathbf{H}_\infty(X) = \Omega(\log^2 n)$ as required by [Theorem 4.1](#).
3. **(Correctness)** $|X| = O(\log n)$ such that $\Pr[\langle X, \mathcal{B}_\mu^n \rangle = 0] \geq 1/2 + 1/\text{poly}(n)$.

Note that any distribution $X \in \{0, 1\}^n$ satisfying $|X| = O(\log n)$ implies that $\mathbf{H}_\infty(X) = O(\log^2 n)$ (as the set $\{\mathbf{x} \in \{0, 1\}^n : |\mathbf{x}| = O(\log n)\}$ is of size $2^{O(\log^2 n)}$), so the job is to maximize the entropy of X under constraint $|X| = O(\log n)$. The first candidate seems $X \sim \mathcal{B}_{\mu'}^n$ for $\mu' = \Theta(\frac{\log n}{n})$, but it does not meet the security condition because the noise rate μ' is so small that a Chernoff bound only ensures (see [Lemma B.2](#)) that $\mathcal{B}_{\mu'}^n$ is $(2^{-O(\mu'n)} = 1/\text{poly}(n))$ -close to having min-entropy $\Theta(n\mathbf{H}(\mu')) = \Theta(\log^2 n)$. In fact, we can avoid the lower-tail issue by letting $X \sim \chi_{\log n}^n$, namely, a uniform distribution of Hamming weight exact $\log n$, which is of min-entropy $\Theta(\log^2 n)$ by [Fact 3](#). Thus, $X \sim \chi_{\log n}^n$ is a valid option to obtain a single-bit PKE with noticeable correctness.

4.3 CPA Secure PKE from Constant-Noise LPN

Unlike [\[20\]](#) where the extension from the weak single-bit PKE to a fully correct scheme is almost immediate (by a parallel repetition and using error correcting codes), it is not trivial to amplify the noticeable correctness of the single-bit scheme to an overwhelming probability, in particular, the scheme instantiated with distribution $X \sim \chi_{\log n}^n$ would no longer work. To see the difficulty, we define below our CPA secure scheme $\Pi_X = (\text{KeyGen}, \text{Enc}, \text{Dec})$ that resembles the counterpart for low-noise LPN (e.g., [\[20, 15\]](#)), where distribution X is left undefined (apart from the entropy constraint).

Distribution X : X is a polynomial-time sampleable distribution with $\mathbf{H}_\infty(X) = \Omega(\log^2 n)$ and we set $\lambda = \Theta(\log^2 n)$ such that $2\lambda \leq \mathbf{H}_\infty(X)$.

KeyGen(1^n): Given a security parameter 1^n , it samples matrix $\mathbf{A} \sim \mathcal{D}_\lambda^{n \times n}$, column vectors $\mathbf{s} \sim X$, $\mathbf{e} \sim \mathcal{B}_\mu^n$, computes $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$ and sets $(pk, sk) := ((\mathbf{A}, \mathbf{b}), \mathbf{s})$.

Enc $_{pk}(\mathbf{m})$: Given the public key $pk = (\mathbf{A}, \mathbf{b})$ and a plaintext $\mathbf{m} \in \{0, 1\}^n$, **Enc** $_{pk}$ chooses

$$\mathbf{S}_1 \sim (X^{(1)}, \dots, X^{(q)}) \in \{0, 1\}^{n \times q}, \mathbf{E}_1 \sim \mathcal{B}_\mu^{n \times q}$$

where $X^{(1)}, \dots, X^{(q)}$ are i.i.d. to X . Then, it outputs $C = (\mathbf{C}_1, \mathbf{c}_2)$ as ciphertext, where

$$\begin{aligned} \mathbf{C}_1 &:= \mathbf{A}^\top \mathbf{S}_1 + \mathbf{E}_1 \in \{0, 1\}^{n \times q}, \\ \mathbf{c}_2 &:= \mathbf{S}_1^\top \mathbf{b} + \mathbf{G} \cdot \mathbf{m} \in \{0, 1\}^q, \end{aligned}$$

and $\mathbf{G} \in \{0, 1\}^{q \times n}$ is a generator matrix for an efficiently decodable code (with error correction capacity to be defined and analyzed in [Section 4.4](#)).

Dec $_{sk}(\mathbf{C}_1, \mathbf{c}_2)$: On secret key $sk = \mathbf{s}$, ciphertext $(\mathbf{C}_1, \mathbf{c}_2)$, it computes

$$\tilde{\mathbf{c}}_0 := \mathbf{c}_2 - \mathbf{C}_1^\top \mathbf{s} = \mathbf{G} \cdot \mathbf{m} + \mathbf{S}_1^\top \mathbf{e} - \mathbf{E}_1^\top \mathbf{s}$$

and reconstructs \mathbf{m} from the error $\mathbf{S}_1^\top \mathbf{e} - \mathbf{E}_1^\top \mathbf{s}$ using the error correction property of \mathbf{G} .

We can see that the CPA security of Π_X , for any X with $\mathbf{H}_\infty(X) = \Omega(\log^2 n)$, follows from applying [Theorem 4.1](#) twice (once for replacing the public key \mathbf{b} with uniform randomness, and again together with the Goldreich Levin Theorem for encrypting a single bit) and a hybrid argument (to encrypt many bits).

Theorem 4.2 (CPA Security) *Assume that the decisional LPN $_{\mu, n}$ problem is $2^{\omega(n^{\frac{1}{2}})}$ -hard for any constant $0 < \mu < 1/2$, then Π_X is IND-CPA secure.*

4.4 Which X Makes a Correct Scheme?

$X \sim \chi_{\log n}^n$ MAY NOT WORK. To make a correct scheme, we need to upper bound $|\mathbf{S}_1^\top \mathbf{e} - \mathbf{E}_1^\top \mathbf{s}|$ by $q(1/2 - 1/\text{poly}(n))$, but in fact we do not have any useful bound even for $|\mathbf{S}_1^\top \mathbf{e}|$. Recall that \mathbf{S}_1^\top is now a $q \times n$ matrix and parse $\mathbf{S}_1^\top \mathbf{e}$ as Boolean random variables W_1, \dots, W_q . First, although every W_i satisfies $\Pr[W_i = 0] \geq 1/2 + 1/\text{poly}(n)$, they are not independent (correlated through \mathbf{e}). Second, if we fix any $|\mathbf{e}| = \Theta(n)$, all W_1, \dots, W_q are now independent conditioned on \mathbf{e} , but then we could no longer guarantee that $\Pr[W_i = 0 | \mathbf{e}] \geq 1/2 + \text{poly}(n)$ as \mathbf{S}_1 follows $(\chi_{\log n}^n)^q$ rather than $(\mathcal{B}_{\log n/n}^n)^q$. Otherwise said, condition #3 (as in [Section 4.2](#)) is not sufficient for overwhelming correctness. We introduced a tailored version of Bernoulli distribution (with upper/lower tails chopped off).

Definition 4.1 (Distribution $\tilde{\mathcal{B}}_{\mu_1}^n$) *Define $\tilde{\mathcal{B}}_{\mu_1}^n$ to be distributed to $\mathcal{B}_{\mu_1}^n$ conditioned on $(1 - \frac{\sqrt{6}}{3})\mu_1 n \leq |\mathcal{B}_{\mu_1}^n| \leq 2\mu_1 n$. Further, we define an $n \times q$ matrix distribution, denoted by $(\tilde{\mathcal{B}}_{\mu_1}^n)^q$, where every column is i.i.d. to $\tilde{\mathcal{B}}_{\mu_1}^n$.*

$\tilde{\mathcal{B}}_{\mu_1}^n$ IS EFFICIENTLY SAMPLEABLE. $\tilde{\mathcal{B}}_{\mu_1}^n$ can be sampled in polynomial-time with exponentially small error, e.g., simply sample $\mathbf{e} \leftarrow \mathcal{B}_{\mu_1}^n$ and outputs \mathbf{e} if $(1 - \frac{\sqrt{6}}{3})\mu_1 n \leq |\mathbf{e}| \leq 2\mu_1 n$. Otherwise, repeat the above until such \mathbf{e} within the Hamming weight range is obtained or the experiment failed (then output \perp in this case) for a predefined number of times (e.g., n).

$\tilde{\mathcal{B}}_{\mu_1}^n$ IS OF MIN-ENTROPY $\Omega(\log^2 n)$. For $\mu_1 = \Omega(\log n/n)$, it is not hard to see that $\tilde{\mathcal{B}}_{\mu_1}^n$ is a convex combination of $\chi_{(1 - \frac{\sqrt{6}}{3})\mu_1 n}^n, \dots, \chi_{2\mu_1 n}^n$, and thus of min-entropy $\Omega(\log^2 n)$ by [Fact 3](#).

Therefore, Π_X when instantiated with $X \sim \tilde{\mathcal{B}}_{\mu_1}^n$ is CPA secure by [Theorem 4.3](#), and we proceed to the correctness of the scheme.

Lemma 4.2 For constants $\alpha > 0$, $0 < \mu \leq 1/10$ and $\mu_1 = \alpha \log n/n$, let $\mathbf{S}_1 \sim (\tilde{\mathcal{B}}_{\mu_1}^n)^q$, $\mathbf{e} \sim \mathcal{B}_\mu^n$, $\mathbf{E}_1 \sim \mathcal{B}_\mu^{n \times q}$ and $\mathbf{s} \sim \tilde{\mathcal{B}}_{\mu_1}^n$, we have

$$\Pr \left[|\mathbf{S}_1^\top \mathbf{e} - \mathbf{E}_1^\top \mathbf{s}| \leq \left(\frac{1}{2} - \frac{1}{2n^{3\alpha/2}} \right) q \right] \geq 1 - 2^{-\Omega(n^{-3\alpha}q)} .$$

Proof. It is more convenient to consider $|\mathbf{S}_1^\top \mathbf{e} - \mathbf{E}_1^\top \mathbf{s}|$ conditioned on $|\mathbf{e}| \leq 1.01\mu n$ (except for a $2^{-\Omega(n)}$ -fraction) and $|\mathbf{s}| \leq 2\mu n$. We have by [Lemma 4.3](#) and [Lemma 4.4](#) that $\mathbf{S}_1^\top \mathbf{e}$ and $\mathbf{E}_1^\top \mathbf{s}$ are identical distributed to $\mathcal{B}_{\delta_1}^q$ and $\mathcal{B}_{\delta_2}^q$ respectively, where $\delta_1 \leq 1/2 - n^{-\alpha/2}$ and $\delta_2 \leq 1/2 - n^{-\alpha}/2$. Thus, $(\mathbf{S}_1^\top \mathbf{e} - \mathbf{E}_1^\top \mathbf{s})$ follows \mathcal{B}_δ^q for $\delta \leq 1/2 - n^{-3\alpha/2}$ by the Piling-up lemma, and then we complete the proof with [Lemma 2.3](#). \square

CONCRETE PARAMETERS. Enc_{pk} simply uses a generator matrix $\mathbf{G} : \{0,1\}^{q \times n}$ that efficiently corrects up to a $(1/2 - n^{-3\alpha/2}/2)$ -fraction of bit flipping errors, which exists for $q = O(n^{3\alpha+1})$ (e.g., [\[22\]](#)). We can now conclude the correctness of the scheme since every encryption will be correctly decrypted with overwhelming probability and thus so is the event that polynomially many of them occur simultaneously (even when they are not independent, see [Lemma 2.2](#)).

Theorem 4.3 (Correctness) Let $0 < \mu \leq 1/10$ and $\alpha > 0$ be any constants, let $q = \Theta(n^{3\alpha+1})$ and $\mu_1 = \alpha \log n/n$, and let $X \sim \tilde{\mathcal{B}}_{\mu_1}^n$. Assume that the decisional $\text{LPN}_{\mu,n}$ problem is $2^{\omega(n^{1/2})}$ -hard, then Π_X is a correct scheme.

Lemma 4.3 For any $0 < \mu \leq 1/10$, $\mu_1 = O(\log n/n) \leq 1/8$ and any $\mathbf{e} \in \{0,1\}^n$ with $|\mathbf{e}| \leq 1.01\mu n$,

$$\Pr[\langle \tilde{\mathcal{B}}_{\mu_1}^n, \mathbf{e} \rangle = 0] \geq 1/2 + 2^{-\frac{\mu_1 n}{2}} .$$

Proof. Denote by \mathcal{E} the event $(1 - \frac{\sqrt{6}}{3})\mu_1 n \leq |\mathcal{B}_{\mu_1}^n| \leq 2\mu_1 n$ and thus $\Pr[\mathcal{E}] \geq (1 - 2 \exp^{-\mu_1 n/3})$ by the Chernoff bound. We have by [Lemma 4.1](#)

$$\begin{aligned} \frac{1}{2} + 2^{-(4.04\mu\mu_1 n+1)} &\leq \Pr[\langle \mathcal{B}_{\mu_1}^n, \mathbf{e} \rangle = 0] \\ &\leq \Pr[\mathcal{E}] \cdot \Pr[\langle \tilde{\mathcal{B}}_{\mu_1}^n, \mathbf{e} \rangle = 0] + \Pr[\neg\mathcal{E}] \cdot \Pr[\langle \mathcal{B}_{\mu_1}^n, \mathbf{e} \rangle = 0 | \neg\mathcal{E}] \\ &\leq \Pr[\langle \tilde{\mathcal{B}}_{\mu_1}^n, \mathbf{e} \rangle = 0] + \Pr[\neg\mathcal{E}] . \end{aligned}$$

For $0 < \mu \leq 1/10$, $\Pr[\langle \tilde{\mathcal{B}}_{\mu_1}^n, \mathbf{e} \rangle = 0] \geq 1/2 + 2^{-(4.04\mu\mu_1 n+1)} - 2 \exp^{-\mu_1 n/3} > 1/2 + 2^{-\mu_1 n/2}$. \square

Lemma 4.4 For any $0 < \mu \leq 1/8$, $\mu_1 = O(\log n/n)$, and any $\mathbf{s} \in \{0,1\}^n$ with $|\mathbf{s}| \leq 2\mu_1 n$, we have by [Lemma 4.1](#)

$$\Pr[\langle \mathcal{B}_\mu^n, \mathbf{s} \rangle = 0] \geq 1/2 + 2^{-(8\mu\mu_1 n+1)} \geq 1/2 + 2^{-(\mu_1 n+1)} .$$

5 CCA-Secure PKE from Constant-Noise LPN

In this section, we show how to construct CCA-secure PKE from constant-noise LPN. Our starting point is the construction of a tag-based PKE against selective tag and chosen ciphertext attacks from LPN, which can be transformed into a standard CCA-secure PKE by using known techniques [\[11, 35\]](#). We begin by first recalling the definitions of tag-based PKE.

5.1 Tag-Based Encryption

A tag-based encryption (TBE) scheme with tag-space \mathcal{T} and message-space \mathcal{M} consists of three PPT algorithms $\mathcal{TB}\mathcal{E} = (\text{KeyGen}, \text{Enc}, \text{Dec})$. The randomized key generation algorithm KeyGen takes the security parameter n as input, outputs a public key pk and a secret key sk , denoted as $(pk, sk) \leftarrow \text{KeyGen}(1^n)$. The randomized encryption algorithm Enc takes pk , a tag $\mathbf{t} \in \mathcal{T}$, and a plaintext $\mathbf{m} \in \mathcal{M}$ as inputs, outputs a ciphertext C , denoted as $C \leftarrow \text{Enc}(pk, \mathbf{t}, \mathbf{m})$. The deterministic algorithm Dec takes sk and C as inputs, outputs a plaintext \mathbf{m} , or a special symbol \perp , which is denoted as $\mathbf{m} \leftarrow \text{Dec}(sk, \mathbf{t}, C)$. For correctness, we require that for all $(pk, sk) \leftarrow \text{KeyGen}(1^n)$, any tag \mathbf{t} , any plaintext \mathbf{m} and any $C \leftarrow \text{Enc}(pk, \mathbf{t}, \mathbf{m})$, the equation $\text{Dec}(sk, \mathbf{t}, C) = \mathbf{m}$ holds with overwhelming probability.

We consider the following game between a challenger \mathcal{C} and an adversary \mathcal{A} given in [35].

Init. The adversary \mathcal{A} takes the security parameter n as inputs, and outputs a target tag \mathbf{t}^* to the challenger \mathcal{C} .

KeyGen. The challenger \mathcal{C} computes $(pk, sk) \leftarrow \text{KeyGen}(1^n)$, gives the public key pk to the adversary \mathcal{A} , and keeps the secret key sk to itself.

Phase 1. The adversary \mathcal{A} can make decryption queries for any pair (\mathbf{t}, C) for any polynomial time, with a restriction that $\mathbf{t} \neq \mathbf{t}^*$, and the challenger \mathcal{C} returns $\mathbf{m} \leftarrow \text{Dec}(sk, \mathbf{t}, C)$ to \mathcal{A} accordingly.

Challenge. The adversary \mathcal{A} outputs two equal length plaintexts $\mathbf{m}_0, \mathbf{m}_1 \in \mathcal{M}$. The challenger \mathcal{C} randomly chooses a bit $b^* \xleftarrow{\$} \{0, 1\}$, and returns the challenge ciphertext $C^* \leftarrow \text{Enc}(pk, \mathbf{t}^*, \mathbf{m}_{b^*})$ to the adversary \mathcal{A} .

Phase 2. The adversary can make more decryption queries as in Phase 1.

Guess. Finally, \mathcal{A} outputs a guess $b \in \{0, 1\}$. If $b = b^*$, the challenger \mathcal{C} outputs 1, else outputs 0.

Advantage. \mathcal{A} 's advantage is defined as $\text{Adv}_{\mathcal{TB}\mathcal{E}, \mathcal{A}}^{\text{ind-stag-cca}}(1^n) \stackrel{\text{def}}{=} |\Pr[b = b^*] - \frac{1}{2}|$.

Definition 5.1 (IND-sTag-CCA) We say that a TBE scheme $\mathcal{TB}\mathcal{E}$ is IND-sTag-CCA secure if for any PPT adversary \mathcal{A} , its advantage $\text{Adv}_{\mathcal{TB}\mathcal{E}, \mathcal{A}}^{\text{ind-stag-cca}}(1^n)$ is negligible in n .

For our convenience, we will use the following corollary, which is essentially a q -fold⁸ (transposed) version of [Theorem 4.1](#) with $q' = n$ and 2 bits of linear leakage (rather than $O(\log n)$ bits of arbitrary leakage) per copy. Following [36], the leakage is crucial for the CCA security proof.

Corollary 5.1 Let n be a security parameter and let $0 < \mu < 1/2$ be any constant. Assume that the computational LPN $_{\mu, n}$ problem is $2^{\omega(n^{\frac{1}{2}})}$ -hard (for any superconstant hidden by $\omega(\cdot)$). Then, for every $\mu_1 = \Omega(\log n/n)$ and $\lambda = \Theta(\log^2 n)$ such that $2\lambda \leq \mathbf{H}_\infty(\tilde{\mathcal{B}}_{\mu_1}^n)$, and every $q = \text{poly}(n)$, we have

$$\left((\mathbf{S}_0^\top \mathbf{e}, \mathbf{E}_0^\top \mathbf{s}), \mathbf{e}, \mathbf{s}, \mathbf{A}, \mathbf{S}_0^\top \mathbf{A} + \mathbf{E}_0^\top \right) \stackrel{c}{\sim} \left((\mathbf{S}_0^\top \mathbf{e}, \mathbf{E}_0^\top \mathbf{s}), \mathbf{e}, \mathbf{s}, \mathbf{A}, U_{q \times n} \right),$$

where the probability is taken over $\mathbf{S}_0 \sim (\tilde{\mathcal{B}}_{\mu_1}^n)^q$, $\mathbf{E}_0 \sim \mathcal{B}_\mu^{n \times q}$, $\mathbf{A} \sim \mathcal{D}_\lambda^{n \times n}$, $U_{q \times n}$, $\mathbf{s} \leftarrow \tilde{\mathcal{B}}_{\mu_1}^n$, $\mathbf{e} \leftarrow \mathcal{B}_\mu^n$ and internal random coins of the distinguisher.

⁸Please do not confuse q' with q , where q' is the number of samples in LPN (see [Theorem 4.1](#)) and is set to n (for a square matrix), and q is the number of parallel repetitions of LPN on independent secrets and noise vectors.

5.2 Our Construction

Our construction is built upon previous works in [44, 36]. A couple of modifications are made to adapt the ideas of [44, 36], which seems necessary due to the absence of a meaningful knapsack version for our LPN (with poly-log entropy and non-uniform matrix). Let n be the security parameter, let $\alpha > 0$, $0 < \mu \leq 1/10$ be any constants, let $\mu_1 = \alpha \log n/n$, $\beta = (\frac{1}{2} - \frac{1}{n^{3\alpha}})$, $\gamma = (\frac{1}{2} - \frac{1}{2n^{3\alpha/2}})$ and choose $\lambda = \Theta(\log^2 n)$ such that $2\lambda \leq \mathbf{H}_\infty(\tilde{\mathcal{B}}_{\mu_1}^n)$. Let the plaintext-space $\mathcal{M} = \{0, 1\}^n$, and let $\mathbf{G} \in \{0, 1\}^{q \times n}$ and $\mathbf{G}_2 \in \{0, 1\}^{\ell \times n}$ be the generator matrices that can correct at least βq and $2\mu\ell$ bit flipping errors in the codeword respectively, where $q = O(n^{6\alpha+1})$, $\ell = O(n)$ and we refer to [22] and [33] for explicit constructions of the two codes respectively. Let the tag-space $\mathcal{T} = \mathbb{F}_{2^n}$. We use a matrix representation $\mathbf{H}_t \in \{0, 1\}^{n \times n}$ for finite field elements $t \in \mathbb{F}_{2^n}$ [14, 10, 36] such that $\mathbf{H}_0 = \mathbf{0}$, \mathbf{H}_t is invertible for any $t \neq 0$, and $\mathbf{H}_{t_1} + \mathbf{H}_{t_2} = \mathbf{H}_{t_1+t_2}$. Our TBE scheme $\mathcal{TB}\mathcal{E}$ is defined as follows:

KeyGen(1^n): Given a security parameter n , first uniformly choose matrices $\mathbf{A} \xleftarrow{\$} \mathcal{D}_\lambda^{n \times n}$, $\mathbf{C} \xleftarrow{\$} \mathcal{D}_\lambda^{\ell \times n}$, $\mathbf{S}_0, \mathbf{S}_1 \xleftarrow{\$} (\tilde{\mathcal{B}}_{\mu_1}^n)^q$ and $\mathbf{E}_0, \mathbf{E}_1 \xleftarrow{\$} \mathcal{B}_\mu^{n \times q}$. Then, compute $\mathbf{B}_0 = \mathbf{S}_0^\top \mathbf{A} + \mathbf{E}_0^\top$, $\mathbf{B}_1 = \mathbf{S}_1^\top \mathbf{A} + \mathbf{E}_1^\top \in \{0, 1\}^{q \times n}$, and set $(pk, sk) = ((\mathbf{A}, \mathbf{B}_0, \mathbf{B}_1, \mathbf{C}), (\mathbf{S}_0, \mathbf{S}_1))$.

Enc(pk, t, m): Given the public key $pk = (\mathbf{A}, \mathbf{B}_0, \mathbf{B}_1, \mathbf{C})$, a tag $t \in \mathbb{F}_{2^n}$, and a plaintext $m \in \{0, 1\}^n$, randomly choose

$$\mathbf{s} \xleftarrow{\$} \tilde{\mathcal{B}}_{\mu_1}^n, \mathbf{e}_1 \xleftarrow{\$} \mathcal{B}_\mu^n, \mathbf{e}_2 \xleftarrow{\$} \mathcal{B}_\mu^\ell, \mathbf{S}'_0, \mathbf{S}'_1 \xleftarrow{\$} (\tilde{\mathcal{B}}_{\mu_1}^n)^q, \mathbf{E}'_0, \mathbf{E}'_1 \xleftarrow{\$} \mathcal{B}_\mu^{n \times q}$$

and define

$$\begin{aligned} \mathbf{c} &:= \mathbf{A}\mathbf{s} + \mathbf{e}_1 && \in \{0, 1\}^n \\ \mathbf{c}_0 &:= (\mathbf{G}\mathbf{H}_t + \mathbf{B}_0)\mathbf{s} + (\mathbf{S}'_0)^\top \mathbf{e}_1 - (\mathbf{E}'_0)^\top \mathbf{s} && \in \{0, 1\}^q \\ \mathbf{c}_1 &:= (\mathbf{G}\mathbf{H}_t + \mathbf{B}_1)\mathbf{s} + (\mathbf{S}'_1)^\top \mathbf{e}_1 - (\mathbf{E}'_1)^\top \mathbf{s} && \in \{0, 1\}^q \\ \mathbf{c}_2 &:= \mathbf{C}\mathbf{s} + \mathbf{e}_2 + \mathbf{G}_2\mathbf{m} && \in \{0, 1\}^\ell. \end{aligned}$$

Finally, return the ciphertext $C = (\mathbf{c}, \mathbf{c}_0, \mathbf{c}_1, \mathbf{c}_2)$.

Dec(sk, t, C): Given the secret key $sk = (\mathbf{S}_0, \mathbf{S}_1)$, tag $t \in \mathbb{F}_{2^n}$ and ciphertext $C = (\mathbf{c}, \mathbf{c}_0, \mathbf{c}_1, \mathbf{c}_2)$, first compute

$$\tilde{\mathbf{c}}_0 := \mathbf{c}_0 - \mathbf{S}_0^\top \mathbf{c} = \mathbf{G}\mathbf{H}_t \mathbf{s} + (\mathbf{S}'_0 - \mathbf{S}_0)^\top \mathbf{e}_1 + (\mathbf{E}_0 - \mathbf{E}'_0)^\top \mathbf{s}.$$

Then, reconstruct $\mathbf{b} = \mathbf{H}_t \mathbf{s}$ from the error $(\mathbf{S}'_0 - \mathbf{S}_0)^\top \mathbf{e}_1 + (\mathbf{E}_0 - \mathbf{E}'_0)^\top \mathbf{s}$ by using the error correction property of \mathbf{G} , and compute $\mathbf{s} = \mathbf{H}_t^{-1} \mathbf{b}$. If it holds that

$$\left| \underbrace{\mathbf{c} - \mathbf{A}\mathbf{s}}_{=\mathbf{e}_1} \right| \leq 2\mu n \wedge \left| \underbrace{\mathbf{c}_0 - (\mathbf{G}\mathbf{H}_t + \mathbf{B}_0)\mathbf{s}}_{=(\mathbf{S}'_0)^\top \mathbf{e}_1 - (\mathbf{E}'_0)^\top \mathbf{s}} \right| \leq \gamma q \wedge \left| \underbrace{\mathbf{c}_1 - (\mathbf{G}\mathbf{H}_t + \mathbf{B}_1)\mathbf{s}}_{=(\mathbf{S}'_1)^\top \mathbf{e}_1 - (\mathbf{E}'_1)^\top \mathbf{s}} \right| \leq \gamma q$$

then reconstruct \mathbf{m} from $\mathbf{c}_2 - \mathbf{C}\mathbf{s} = \mathbf{G}_2\mathbf{m} + \mathbf{e}_2$ by using the error correction property of \mathbf{G}_2 , else let $\mathbf{m} = \perp$. Finally, return the decrypted result \mathbf{m} .

Remark 5.1 As one can see, the matrix \mathbf{S}_1 in the secret key $sk = (\mathbf{S}_0, \mathbf{S}_1)$ can also be used to decrypt the ciphertext, i.e., compute $\tilde{\mathbf{c}}_1 := \mathbf{c}_1 - \mathbf{S}_1^\top \mathbf{c} = \mathbf{G}\mathbf{H}_t \mathbf{s} + (\mathbf{S}'_1 - \mathbf{S}_1)^\top \mathbf{e}_1 + (\mathbf{E}_1 - \mathbf{E}'_1)^\top \mathbf{s}$ and recover \mathbf{s} from $\tilde{\mathbf{c}}_1$ by using the error correction property of \mathbf{G} . Moreover, the check condition

$$|\mathbf{c} - \mathbf{A}\mathbf{s}| \leq 2\mu n \wedge |\mathbf{c}_0 - (\mathbf{G}\mathbf{H}_t + \mathbf{B}_0)\mathbf{s}| \leq \gamma q \wedge |\mathbf{c}_1 - (\mathbf{G}\mathbf{H}_t + \mathbf{B}_1)\mathbf{s}| \leq \gamma q$$

guarantees that the decryption results are the same when we use either \mathbf{S}_0 or \mathbf{S}_1 in the decryption. This fact seems not necessary for the correctness, but it is very important for the security proof. Looking ahead, it allows us to switch the “exact decryption key” between \mathbf{S}_0 and \mathbf{S}_1 .

5.2.1 Correctness and Equivalence of the Secret Keys $\mathbf{S}_0, \mathbf{S}_1$.

In the following, we show that for appropriate choice of parameters, the above scheme $\mathcal{TB}\mathcal{E}$ is correct, and has the property that both \mathbf{S}_0 and \mathbf{S}_1 are equivalent in terms of decryption.

- The correctness of the scheme requires the following:
 1. $|(\mathbf{S}'_0 - \mathbf{S}_0)^\top \mathbf{e}_1 + (\mathbf{E}_0 - \mathbf{E}'_0)^\top \mathbf{s}| \leq \beta q$ (to let \mathbf{G} reconstruct \mathbf{s} from $\tilde{\mathbf{c}}_0$).
 2. $|\mathbf{c} - \mathbf{A}\mathbf{s}| \leq 2\mu n \wedge |\mathbf{c}_0 - (\mathbf{G}\mathbf{H}_t + \mathbf{B}_0)\mathbf{s}| \leq \gamma q \wedge |\mathbf{c}_1 - (\mathbf{G}\mathbf{H}_t + \mathbf{B}_1)\mathbf{s}| \leq \gamma q$.
 3. $|\mathbf{e}_2| \leq 2\mu\ell$ (such that \mathbf{G}_2 can reconstruct \mathbf{m} from $\mathbf{c}_2 - \mathbf{C}\mathbf{s} = \mathbf{G}_2\mathbf{m} + \mathbf{e}_2$).
- For obtaining CCA security, we also need to show that \mathbf{S}_0 and \mathbf{S}_1 have the same decryption ability except with negligible probability, namely,
 1. If $|\mathbf{c} - \mathbf{A}\mathbf{s}| \leq 2\mu n \wedge |\mathbf{c}_0 - (\mathbf{G}\mathbf{H}_t + \mathbf{B}_0)\mathbf{s}| \leq \gamma q$, then \mathbf{G} can reconstruct \mathbf{s} from a code within bounded error $|(\mathbf{S}'_0 - \mathbf{S}_0)\mathbf{e}_1 + (\mathbf{E}_0 - \mathbf{E}'_0)\mathbf{s}| \leq \beta q$.
 2. If $|\mathbf{c} - \mathbf{A}\mathbf{s}| \leq 2\mu n \wedge |\mathbf{c}_1 - (\mathbf{G}\mathbf{H}_t + \mathbf{B}_1)\mathbf{s}| \leq \gamma q$, then \mathbf{G} can reconstruct \mathbf{s} from a code within bounded error $|(\mathbf{S}'_1 - \mathbf{S}_1)\mathbf{e}_1 + (\mathbf{E}_1 - \mathbf{E}'_1)\mathbf{s}| \leq \beta q$.

It suffices to show that each Hamming weight constraint above holds (with overwhelming probability) individually and thus polynomially many of them hold simultaneously (with overwhelming probability as well) by [Lemma 2.2](#). First, Chernoff bound guarantees that $\Pr[|\mathbf{e}_1| \leq 2\mu n] = 1 - 2^{-\Omega(n)}$ and $\Pr[|\mathbf{e}_2| \leq 2\mu\ell] = 1 - 2^{-\Omega(\ell)}$. Second, for $i \in \{0, 1\}$ the bound $|(\mathbf{S}'_i)^\top \mathbf{e}_1 - (\mathbf{E}'_i)^\top \mathbf{s}| \leq \gamma q$ is ensured by [Lemma 4.2](#) and we further bound $|(\mathbf{S}'_i - \mathbf{S}_i)\mathbf{e}_1 + (\mathbf{E}_i - \mathbf{E}'_i)\mathbf{s}| \leq \beta q$ with [Lemma 5.1](#) below (proof similar to [Lemma 4.2](#) and thus deferred to [Appendix B](#)).

Lemma 5.1 *For constants $\alpha > 0$, $0 < \mu \leq 1/10$ and $\mu_1 = \alpha \log n/n$, let \mathbf{S} and \mathbf{S}' be i.i.d. to $(\tilde{\mathcal{B}}_{\mu_1}^n)^q$, \mathbf{E} and \mathbf{E}' be i.i.d. to $\mathcal{B}_{\mu}^{n \times q}$, $\mathbf{s} \sim \mathcal{B}_{\mu_1}^n$ and $\mathbf{e} \sim \mathcal{B}_{\mu}^n$. Then,*

$$\Pr \left[|(\mathbf{S}' - \mathbf{S})^\top \mathbf{e} + (\mathbf{E} - \mathbf{E}')^\top \mathbf{s}| \leq \left(\frac{1}{2} - \frac{1}{n^{3\alpha}} \right) q \right] \geq 1 - 2^{-\Omega(n^{-6\alpha}q)} .$$

5.2.2 Security of the TBE Scheme.

We now show that under the LPN assumption, the above scheme $\mathcal{TB}\mathcal{E}$ is IND-sTag-CCA secure in the standard model.

Theorem 5.1 *Assume that the decisional $\text{LPN}_{\mu, n}$ problem is $2^{\omega(n^{\frac{1}{2}})}$ -hard for any constant $0 < \mu \leq 1/10$, then our TBE scheme $\mathcal{TB}\mathcal{E}$ is IND-sTag-CCA secure.*

Proof. Let \mathcal{A} be any PPT adversary that can attack our TBE scheme $\mathcal{TB}\mathcal{E}$ with advantage ϵ . We show that ϵ must be negligible in n . We continue the proof by using a sequence of games, where the first game is the real IND-sTag-CCA security game, while the last is a random game in which the challenge ciphertext is independent from the choices of the challenge plaintexts. Since any PPT adversary \mathcal{A} 's advantage in a random game is exactly 0, the security of $\mathcal{TB}\mathcal{E}$ can be established by showing that \mathcal{A} 's advantage in any two consecutive games are negligibly close.

Game 0. The challenger \mathcal{C} honestly runs the adversary \mathcal{A} with the security parameter n , and obtains a target tag \mathbf{t}^* from \mathcal{A} . Then, it simulates the IND-sTag-CCA security game for \mathcal{A} as follows:

KeyGen. First uniformly choose matrices $\mathbf{A} \xleftarrow{\$} \mathcal{D}_{\lambda}^{n \times n}$, $\mathbf{C} \xleftarrow{\$} \mathcal{D}_{\lambda}^{\ell \times n}$, $\mathbf{S}_0, \mathbf{S}_1 \xleftarrow{\$} (\tilde{\mathcal{B}}_{\mu_1}^n)^q$ and $\mathbf{E}_0, \mathbf{E}_1 \xleftarrow{\$} \mathcal{B}_{\mu}^{n \times q}$. Then, compute $\mathbf{B}_0 = \mathbf{S}_0^\top \mathbf{A} + \mathbf{E}_0^\top$, $\mathbf{B}_1 = \mathbf{S}_1^\top \mathbf{A} + \mathbf{E}_1^\top \in \{0, 1\}^{q \times n}$. Finally, \mathcal{C} sends $pk = (\mathbf{A}, \mathbf{B}_0, \mathbf{B}_1, \mathbf{C})$ to the adversary \mathcal{A} , and keeps $sk = (\mathbf{S}_0, \mathbf{S}_1)$ to itself.

Phase 1. After receiving a decryption query $(\mathbf{t}, (\mathbf{c}, \mathbf{c}_0, \mathbf{c}_1, \mathbf{c}_2))$ from the adversary \mathcal{A} , the challenger \mathcal{C} directly returns \perp to \mathcal{A} if $\mathbf{t} = \mathbf{t}^*$. Otherwise, it first computes

$$\tilde{\mathbf{c}}_0 := \mathbf{c}_0 - \mathbf{S}_0^\top \mathbf{c} = \mathbf{G}\mathbf{H}_t \mathbf{s} + (\mathbf{S}'_0 - \mathbf{S}_0)^\top \mathbf{e}_1 + (\mathbf{E}_0 - \mathbf{E}'_0)^\top \mathbf{s}.$$

Then, it reconstruct $\mathbf{b} = \mathbf{H}_t \mathbf{s}$ from the error $(\mathbf{S}'_0 - \mathbf{S}_0)^\top \mathbf{e}_1 + (\mathbf{E}_0 - \mathbf{E}'_0)^\top \mathbf{s}$ by using the error correction property of \mathbf{G} , and compute $\mathbf{s} = \mathbf{H}_t^{-1} \mathbf{b}$. If

$$|\mathbf{c} - \mathbf{A}\mathbf{s}| \leq 2\mu n \wedge |\mathbf{c}_0 - (\mathbf{G}\mathbf{H}_t + \mathbf{B}_0)\mathbf{s}| \leq \gamma q \wedge |\mathbf{c}_1 - (\mathbf{G}\mathbf{H}_t + \mathbf{B}_1)\mathbf{s}| \leq \gamma q$$

is true, reconstruct M from $\mathbf{c}_2 - \mathbf{C}\mathbf{s} = \mathbf{G}_2 \mathbf{m} + \mathbf{e}_2$ by using the error correction property of \mathbf{G}_2 , else let $\mathbf{m} = \perp$. Finally, return the decrypted result \mathbf{m} to the adversary \mathcal{A} .

Challenge. After receiving two equal length plaintexts $\mathbf{m}_0, \mathbf{m}_1 \in \mathcal{M}$ from the adversary \mathcal{A} , the challenger \mathcal{C} first randomly chooses a bit $b^* \xleftarrow{\$} \{0, 1\}$, and

$$\mathbf{s} \xleftarrow{\$} \tilde{\mathcal{B}}_{\mu_1}^n, \mathbf{e}_1 \xleftarrow{\$} \mathcal{B}_{\mu}^n, \mathbf{e}_2 \xleftarrow{\$} \mathcal{B}_{\mu}^{\ell}, \mathbf{S}'_0, \mathbf{S}'_1 \xleftarrow{\$} (\tilde{\mathcal{B}}_{\mu_1}^n)^q, \mathbf{E}'_0, \mathbf{E}'_1 \xleftarrow{\$} \mathcal{B}_{\mu}^{n \times q}$$

Then, it defines

$$\begin{aligned} \mathbf{c}^* &:= \mathbf{A}\mathbf{s} + \mathbf{e}_1 && \in \{0, 1\}^n \\ \mathbf{c}_0^* &:= (\mathbf{G}\mathbf{H}_{t^*} + \mathbf{B}_0)\mathbf{s} + (\mathbf{S}'_0)^\top \mathbf{e}_1 - (\mathbf{E}'_0)^\top \mathbf{s} && \in \{0, 1\}^q \\ \mathbf{c}_1^* &:= (\mathbf{G}\mathbf{H}_{t^*} + \mathbf{B}_1)\mathbf{s} + (\mathbf{S}'_1)^\top \mathbf{e}_1 - (\mathbf{E}'_1)^\top \mathbf{s} && \in \{0, 1\}^q \\ \mathbf{c}_2^* &:= \mathbf{C}\mathbf{s} + \mathbf{e}_2 + \mathbf{G}_2 \mathbf{m}_{b^*} && \in \{0, 1\}^{\ell}, \end{aligned}$$

and returns the challenge ciphertext $(\mathbf{c}^*, \mathbf{c}_0^*, \mathbf{c}_1^*, \mathbf{c}_2^*)$ to the adversary \mathcal{A} .

Phase 2. The adversary can adaptively make more decryption queries, and the challenger \mathcal{C} responds as in Phase 1.

Guess. Finally, \mathcal{A} outputs a guess $b \in \{0, 1\}$. If $b = b^*$, the challenger \mathcal{C} outputs 1, else outputs 0.

EVENT. Let F_i be the event that \mathcal{C} outputs 1 in Game i for $i \in \{0, 1, \dots, 6\}$.

Lemma 5.2 $|\Pr[F_0] - \frac{1}{2}| = \epsilon$.

Proof. This lemma immediately follows the fact that \mathcal{C} honestly simulates the attack environment for \mathcal{A} , and only outputs 1 if and only if $b = b^*$. \square

Game 1. This game is identical to Game 0 except that the challenger \mathcal{C} changes the key generation phase as follows:

KeyGen. First uniformly choose matrices $\mathbf{A} \xleftarrow{\$} \mathcal{D}_{\lambda}^{n \times n}$, $\mathbf{C} \xleftarrow{\$} \mathcal{D}_{\lambda}^{\ell \times n}$, $\mathbf{S}_0, \mathbf{S}_1 \xleftarrow{\$} (\tilde{\mathcal{B}}_{\mu_1}^n)^q$, $\mathbf{E}_0, \mathbf{E}_1 \xleftarrow{\$} \mathcal{B}_{\mu}^{n \times q}$, and $\mathbf{B}'_1 \xleftarrow{\$} \{0, 1\}^{q \times n}$. Then, compute $\mathbf{B}_0 = \mathbf{S}_0^\top \mathbf{A} + \mathbf{E}_0^\top$, $\mathbf{B}_1 = \mathbf{S}_1^\top \mathbf{A} + \mathbf{E}_1^\top \in \{0, 1\}^{q \times n}$. Finally, \mathcal{C} sends $pk = (\mathbf{A}, \mathbf{B}_0, \mathbf{B}'_1, \mathbf{C})$ to the adversary \mathcal{A} , and keeps $sk = (\mathbf{S}_0, \mathbf{S}_1)$ to itself.

Lemma 5.3 *If the decisional LPN $_{\mu, n}$ problem is $2^{\omega(n^{\frac{1}{2}})}$ -hard, then we have $|\Pr[F_1] - \Pr[F_0]| \leq \text{negl}(n)$.*

Proof. Since the only difference between Game 0 and Game 1 is that \mathcal{C} replaces $\mathbf{B}_1 = \mathbf{S}_1^\top \mathbf{A} + \mathbf{E}_1^\top \in \{0, 1\}^{q \times n}$ in Game 0 with a randomly chosen $\mathbf{B}'_1 \xleftarrow{\$} \{0, 1\}^{q \times n}$ in Game 1. we have that Game 0 and Game 1 are computationally indistinguishable for any PPT adversary \mathcal{A} by our assumption and [Corollary 5.1](#). This means that $|\Pr[F_1] - \Pr[F_0]| \leq \text{negl}(n)$ holds. \square

Game 2. This game is identical to Game 1 except that the challenger \mathcal{C} changes the key generation phase as follows:

KeyGen. First uniformly choose matrices $\mathbf{A} \xleftarrow{\$} \mathcal{D}_\lambda^{n \times n}$, $\mathbf{C} \xleftarrow{\$} \mathcal{D}_\lambda^{\ell \times n}$, $\mathbf{S}_0, \mathbf{S}_1 \xleftarrow{\$} (\tilde{\mathcal{B}}_{\mu_1}^n)^q$, $\mathbf{E}_0, \mathbf{E}_1 \xleftarrow{\$} \mathcal{B}_\mu^{n \times q}$, and $\mathbf{B}_1'' \xleftarrow{\$} \{0, 1\}^{q \times n}$. Then, compute $\mathbf{B}_0 = \mathbf{S}_0^\top \mathbf{A} + \mathbf{E}_0^\top$, $\mathbf{B}_1 = \mathbf{S}_1^\top \mathbf{A} + \mathbf{E}_1^\top \in \{0, 1\}^{q \times n}$ and $\mathbf{B}_1' = \mathbf{B}_1'' - \mathbf{G}\mathbf{H}_{t^*}$. Finally, \mathcal{C} sends $pk = (\mathbf{A}, \mathbf{B}_0, \mathbf{B}_1', \mathbf{C})$ to the adversary \mathcal{A} , and keeps $sk = (\mathbf{S}_0, \mathbf{S}_1)$ to itself.

Challenge. After receiving two equal length plaintexts $\mathbf{m}_0, \mathbf{m}_1 \in \mathcal{M}$ from the adversary \mathcal{A} , the challenger \mathcal{C} first randomly chooses a bit $b^* \xleftarrow{\$} \{0, 1\}$, and

$$\mathbf{s} \xleftarrow{\$} \tilde{\mathcal{B}}_{\mu_1}^n, \mathbf{e}_1 \xleftarrow{\$} \mathcal{B}_\mu^n, \mathbf{e}_2 \xleftarrow{\$} \mathcal{B}_\mu^\ell, \mathbf{S}'_0, \mathbf{S}'_1 \xleftarrow{\$} (\tilde{\mathcal{B}}_{\mu_1}^n)^q, \mathbf{E}'_0, \mathbf{E}'_1 \xleftarrow{\$} \mathcal{B}_\mu^{n \times q}$$

Then, it defines

$$\begin{aligned} \mathbf{c}^* &:= \mathbf{A}\mathbf{s} + \mathbf{e}_1 && \in \{0, 1\}^n \\ \mathbf{c}_0^* &:= (\mathbf{G}\mathbf{H}_{t^*} + \mathbf{B}_0)\mathbf{s} + (\mathbf{S}'_0)^\top \mathbf{e}_1 - (\mathbf{E}'_0)^\top \mathbf{s} && \in \{0, 1\}^q \\ \mathbf{c}_1^* &:= (\mathbf{G}\mathbf{H}_{t^*} + \mathbf{B}_1)\mathbf{s} + (\mathbf{S}'_1)^\top \mathbf{e}_1 - (\mathbf{E}'_1)^\top \mathbf{s} && \in \{0, 1\}^q \\ \mathbf{c}_2^* &:= \mathbf{C}\mathbf{s} + \mathbf{e}_2 + \mathbf{G}_2\mathbf{m}_{b^*} && \in \{0, 1\}^\ell, \end{aligned}$$

and returns the challenge ciphertext $(\mathbf{c}^*, \mathbf{c}_0^*, \mathbf{c}_1^*, \mathbf{c}_2^*)$ to the adversary \mathcal{A} .

Lemma 5.4 $\Pr[F_2] = \Pr[F_1]$.

Proof. Because of $\mathbf{B}_1'' \xleftarrow{\$} \{0, 1\}^{q \times n}$, we have that $\mathbf{B}_1' = \mathbf{B}_1'' - \mathbf{G}\mathbf{H}_{t^*}$ is also uniformly distributed over $\{0, 1\}^{q \times n}$. This means that the public key in Game 2 has the same distribution as that in Game 1. In addition, since $\mathbf{S}_1 \xleftarrow{\$} (\tilde{\mathcal{B}}_{\mu_1}^n)^q$ and $\mathbf{E}_1 \xleftarrow{\$} \mathcal{B}_\mu^{n \times q}$ are chosen from the same distribution as \mathbf{S}'_1 and \mathbf{E}'_1 respectively. By the fact that $\mathbf{B}_1 = \mathbf{S}_1^\top \mathbf{A} + \mathbf{E}_1^\top \in \{0, 1\}^{q \times n}$ is not included in the public key $pk = (\mathbf{A}, \mathbf{B}_0, \mathbf{B}_1', \mathbf{C})$ (and thus \mathcal{A} has no information about \mathbf{S}_1 and \mathbf{E}_1 before the challenge phase), we have that the challenge ciphertext in Game 2 also has the same distribution as that in Game 1. In all, Game 2 is identical to Game 1 in the adversary's view. Thus, we have $\Pr[F_2] = \Pr[F_1]$. \square

Game 3. This game is identical to Game 2 except that the challenger \mathcal{C} changes the key generation phase as follows:

KeyGen. First uniformly choose matrices $\mathbf{A} \xleftarrow{\$} \mathcal{D}_\lambda^{n \times n}$, $\mathbf{C} \xleftarrow{\$} \mathcal{D}_\lambda^{\ell \times n}$, $\mathbf{S}_0, \mathbf{S}_1 \xleftarrow{\$} (\tilde{\mathcal{B}}_{\mu_1}^n)^q$, and $\mathbf{E}_0, \mathbf{E}_1 \xleftarrow{\$} \mathcal{B}_\mu^{n \times q}$. Then, compute $\mathbf{B}_0 = \mathbf{S}_0^\top \mathbf{A} + \mathbf{E}_0^\top$, $\mathbf{B}_1 = \mathbf{S}_1^\top \mathbf{A} + \mathbf{E}_1^\top \in \{0, 1\}^{q \times n}$ and $\mathbf{B}_1' = \mathbf{B}_1 - \mathbf{G}\mathbf{H}_{t^*}$. Finally, \mathcal{C} sends $pk = (\mathbf{A}, \mathbf{B}_0, \mathbf{B}_1', \mathbf{C})$ to the adversary \mathcal{A} , and keeps $sk = (\mathbf{S}_0, \mathbf{S}_1)$ to itself.

Lemma 5.5 *If the decisional LPN $_{\mu, n}$ problem is $2^{\omega(n^{\frac{1}{2}})}$ -hard, then $|\Pr[F_3] - \Pr[F_2]| \leq \text{negl}(n)$.*

Proof. Since the only difference between Game 2 and Game 3 is that \mathcal{C} replaces the randomly chosen $\mathbf{B}_1'' \xleftarrow{\$} \{0, 1\}^{q \times n}$ in Game 2 with $\mathbf{B}_1 = \mathbf{S}_1^\top \mathbf{A} + \mathbf{E}_1^\top \in \{0, 1\}^{q \times n}$ in Game 3, by our assumption and [Corollary 5.1](#) we have that Game 2 and Game 3 are computationally indistinguishable for any PPT adversary \mathcal{A} seeing $(\mathbf{S}_1^\top \mathbf{e}_1, \mathbf{E}_1^\top \mathbf{s})$ in the challenge ciphertext. This means that $|\Pr[F_3] - \Pr[F_2]| \leq \text{negl}(n)$ holds. \square

Remark 5.2 *Note that for the challenge ciphertext $(\mathbf{c}, \mathbf{c}_0^*, \mathbf{c}_1^*, \mathbf{c}_2^*)$ in Game 3, we have that $\mathbf{c}_1^* := (\mathbf{G}\mathbf{H}_{t^*} + \mathbf{B}_1')\mathbf{s} + \mathbf{S}_1^\top \mathbf{e}_1 - \mathbf{E}_1^\top \mathbf{s} = \mathbf{S}_1^\top \mathbf{c}$.*

Game 4. This game is identical to Game 3 except that the challenger \mathcal{C} answers the decryption queries by using \mathbf{S}_1 instead of \mathbf{S}_0 .

Lemma 5.6 $|\Pr[F_4] - \Pr[F_3]| \leq \text{negl}(n)$.

Proof. This lemma directly follows from the fact that both \mathbf{S}_0 and \mathbf{S}_1 have equivalent decryption ability except with negligible probability. \square

Game 5. This game is identical to Game 4 except that the challenger \mathcal{C} changes the key generation phase and the challenge phase as follows:

KeyGen. First uniformly choose matrices $\mathbf{A} \xleftarrow{\$} \mathcal{D}_\lambda^{n \times n}$, $\mathbf{C} \xleftarrow{\$} \mathcal{D}_\lambda^{\ell \times n}$, $\mathbf{S}_0, \mathbf{S}_1 \xleftarrow{\$} (\tilde{\mathcal{B}}_{\mu_1}^n)^q$, and $\mathbf{E}_0, \mathbf{E}_1 \xleftarrow{\$} \mathcal{B}_\mu^{n \times q}$. Then, compute $\mathbf{B}_0 = \mathbf{S}_0^\top \mathbf{A} + \mathbf{E}_0^\top$, $\mathbf{B}_1 = \mathbf{S}_1^\top \mathbf{A} + \mathbf{E}_1^\top \in \{0, 1\}^{q \times n}$, $\mathbf{B}'_0 = \mathbf{B}_0 - \mathbf{G}\mathbf{H}_{t^*}$ and $\mathbf{B}'_1 = \mathbf{B}_1 - \mathbf{G}\mathbf{H}_{t^*}$. Finally, \mathcal{C} sends $pk = (\mathbf{A}, \mathbf{B}'_0, \mathbf{B}'_1, \mathbf{C})$ to the adversary \mathcal{A} , and keeps $sk = (\mathbf{S}_0, \mathbf{S}_1)$ to itself.

Challenge. After receiving two equal length plaintexts $\mathbf{m}_0, \mathbf{m}_1 \in \mathcal{M}$ from the adversary \mathcal{A} , the challenger \mathcal{C} first randomly chooses a bit $b^* \xleftarrow{\$} \{0, 1\}$, and $\mathbf{s} \xleftarrow{\$} \tilde{\mathcal{B}}_{\mu_1}^n$, $\mathbf{e}_1 \xleftarrow{\$} \mathcal{B}_\mu^n$ and $\mathbf{e}_2 \xleftarrow{\$} \mathcal{B}_\mu^\ell$. Then, it defines

$$\begin{aligned} \mathbf{c}^* &:= \mathbf{A}\mathbf{s} + \mathbf{e}_1 && \in \{0, 1\}^n \\ \mathbf{c}_0^* &:= (\mathbf{G}\mathbf{H}_{t^*} + \mathbf{B}'_0)\mathbf{s} + \mathbf{S}_0^\top \mathbf{e}_1 - \mathbf{E}_0^\top \mathbf{s} = \mathbf{S}_0^\top \mathbf{c}^* && \in \{0, 1\}^q \\ \mathbf{c}_1^* &:= (\mathbf{G}\mathbf{H}_{t^*} + \mathbf{B}'_1)\mathbf{s} + \mathbf{S}_1^\top \mathbf{e}_1 - \mathbf{E}_1^\top \mathbf{s} = \mathbf{S}_1^\top \mathbf{c}^* && \in \{0, 1\}^q \\ \mathbf{c}_2^* &:= \mathbf{C}\mathbf{s} + \mathbf{e}_2 + \mathbf{G}_2\mathbf{m}_{b^*} && \in \{0, 1\}^\ell, \end{aligned}$$

and returns the challenge ciphertext $(\mathbf{c}, \mathbf{c}_0^*, \mathbf{c}_1^*, \mathbf{c}_2^*)$ to the adversary \mathcal{A} .

Lemma 5.7 *If the decisional LPN $_{\mu, n}$ problem is $2^{\omega(n^{\frac{1}{2}})}$ -hard, then we have that $|\Pr[F_5] - \Pr[F_4]| \leq \text{negl}(n)$.*

Proof. One can easily show this lemma holds by using similar proofs from [Lemma 5.2](#) to [Lemma 5.6](#). We omit the details. \square

Game 6. This game is identical to Game 5 except that the challenger \mathcal{C} changes the challenge phase as follows:

Challenge. After receiving two equal length plaintexts $\mathbf{m}_0, \mathbf{m}_1 \in \mathcal{M}$ from the adversary \mathcal{A} , the challenger \mathcal{C} first randomly chooses $b^* \xleftarrow{\$} \{0, 1\}$, $\mathbf{u} \xleftarrow{\$} \{0, 1\}^n$ and $\mathbf{v} \xleftarrow{\$} \{0, 1\}^\ell$. Then, it defines

$$\begin{aligned} \mathbf{c}^* &:= \mathbf{u} && \in \{0, 1\}^n \\ \mathbf{c}_0^* &:= \mathbf{S}_0 \mathbf{c}^* && \in \{0, 1\}^q \\ \mathbf{c}_1^* &:= \mathbf{S}_1 \mathbf{c}^* && \in \{0, 1\}^q \\ \mathbf{c}_2^* &:= \mathbf{v} + \mathbf{G}_2 \mathbf{m}_{b^*} && \in \{0, 1\}^\ell, \end{aligned}$$

and returns the challenge ciphertext $(\mathbf{c}, \mathbf{c}_0^*, \mathbf{c}_1^*, \mathbf{c}_2^*)$ to the adversary \mathcal{A} .

Lemma 5.8 *If the decisional LPN $_{\mu, n}$ problem is $2^{\omega(n^{\frac{1}{2}})}$ -hard, then we have that $|\Pr[F_6] - \Pr[F_5]| \leq \text{negl}(n)$.*

Proof. Since the only difference between Game 5 and Game 6 is that \mathcal{C} replaces $\mathbf{c}^* = \mathbf{A}\mathbf{s} + \mathbf{e}_1$ and $\mathbf{c}_2^* = \mathbf{C}\mathbf{s} + \mathbf{e}_2 + \mathbf{G}_2\mathbf{m}_{b^*}$ in Game 5 with $\mathbf{c}^* := \mathbf{u}$ and $\mathbf{c}_2^* := \mathbf{v} + \mathbf{G}_2\mathbf{m}_{b^*}$ in Game 6, where $\mathbf{u} \xleftarrow{\$} \{0, 1\}^n$ and $\mathbf{v} \xleftarrow{\$} \{0, 1\}^\ell$, by our assumption and [Corollary 5.1](#) we have that Game 5 and Game 6 are computationally indistinguishable for any PPT adversary \mathcal{A} . Obviously, we have that $|\Pr[F_6] - \Pr[F_5]| \leq \text{negl}(n)$ holds. \square

Lemma 5.9 $\Pr[F_6] = \frac{1}{2}$.

Proof. This claim follows from the fact that the challenge ciphertext $(\mathbf{c}, \mathbf{c}_0^*, \mathbf{c}_1^*, \mathbf{c}_2^*)$ in Game 6 perfectly hides the information of \mathbf{m}_{b^*} . \square

In all, by Lemma 5.2 \sim Lemma 5.9, we have that $\epsilon = |\Pr[F_0] - \frac{1}{2}| \leq \text{negl}(n)$. This completes the proof of Theorem 5.1. \square

Acknowledgments

Yu Yu was supported by the National Basic Research Program of China Grant No. 2013CB338004, and the National Natural Science Foundation of China Grant (Nos. 61472249, 61572192, 61572149), Shanghai excellent academic leader funds (No. 16XD1400200) and International Science & Technology Cooperation & Exchange Projects of Shaanxi Province (2016KW-038). Jiang Zhang is supported by the National Basic Research Program of China under Grant No. 2013CB338003 and the National Natural Science Foundation of China under Grant No. U1536205.

References

- [1] Related work on LPN-based authentication schemes. <http://www.ecrypt.eu.org/lightweight/index.php/HB>.
- [2] Michael Alekhnovich. More on average case vs approximation complexity. In *44th Annual Symposium on Foundations of Computer Science*, pages 298–307, Cambridge, Massachusetts, October 2003. IEEE.
- [3] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *Advances in Cryptology - CRYPTO 2009*, pages 595–618, 2009.
- [4] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Cryptography with constant input locality. In *Advances in Cryptology - CRYPTO 2007*, pages 92–110, 2007. Full version available at <http://www.eng.tau.ac.il/~bennyap/pubs/input-locality-full-revised-1.pdf>.
- [5] Anja Becker, Antoine Joux, Alexander May, and Alexander Meurer. Decoding random binary linear codes in $2^{n/20}$: How $1 + 1 = 0$ improves information set decoding. In *Advances in Cryptology - EUROCRYPT 2012*, pages 520–536, 2012.
- [6] Elwyn Berlekamp, Robert J. McEliece, and H.C.A. van Tilborg. On the inherent intractability of certain coding problems. *IEEE Transactions on Information Theory*, 24(3):384–386, 1978.
- [7] Daniel J. Bernstein, Tanja Lange, and Christiane Peters. Smaller decoding exponents: Ball-collision decoding. In *Advances in Cryptology - CRYPTO 2011*, pages 743–760, 2011.
- [8] Avrim Blum, Merrick L. Furst, Michael J. Kearns, and Richard J. Lipton. Cryptographic primitives based on hard learning problems. In Douglas R. Stinson, editor, *Advances in Cryptology—CRYPTO '93*, volume 773 of *LNCS*, pages 278–291. Springer-Verlag, 22–26 August 1993.
- [9] Avrim Blum, Adam Kalai, and Hal Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *Journal of the ACM*, 50(4):506–519, 2003.

- [10] Dan Boneh, Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. *SIAM J. Comput.*, 36(5):1301–1328, December 2006.
- [11] Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. In Christian Cachin and Jan Camenisch, editors, *Advances in Cryptology – EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 207–222. Springer Berlin / Heidelberg, 2004.
- [12] Anne Canteaut and Florent Chabaud. A new algorithm for finding minimum-weight words in a linear code: Application to mceliece’s cryptosystem and to narrow-sense BCH codes of length 511. *IEEE Transactions on Information Theory*, 44(1):367–378, 1998.
- [13] David Cash, Eike Kiltz, and Stefano Tessaro. Two-round man-in-the-middle security from LPN. In *Proceedings of the 13th Theory of Cryptography (TCC 2016-A)*, pages 225–248, 2016.
- [14] Ronald Cramer and Ivan Damgård. On the amortized complexity of zero-knowledge protocols. In Shai Halevi, editor, *Advances in Cryptology – CRYPTO 2009*, volume 5677 of *Lecture Notes in Computer Science*, pages 177–191. Springer Berlin Heidelberg, 2009.
- [15] Ivan Damgård and Sunoo Park. How practical is public-key encryption based on lpn and ring-lpn? Cryptology ePrint Archive, Report 2012/699, 2012. <http://eprint.iacr.org/2012/699>.
- [16] Bernardo David, Rafael Dowsley, and Anderson C. A. Nascimento. Universally composable oblivious transfer based on a variant of LPN. In *Proceedings of the 13th International Conference on Cryptology and Network Security (CANS 2014)*, pages 143–158, 2014.
- [17] Yevgeniy Dodis, Yael Tauman Kalai, and Shachar Lovett. On cryptography with auxiliary input. In Michael Mitzenmacher, editor, *STOC*, pages 621–630. ACM, 2009.
- [18] Yevgeniy Dodis, Eike Kiltz, Krzysztof Pietrzak, and Daniel Wichs. Message authentication, revisited. In *Proceedings of the 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT 2012)*, pages 355–374, 2012.
- [19] Nico Döttling. Low noise lpn: Kdm secure public key encryption and sample amplification. In *Public-Key Cryptography–PKC 2015*, pages 604–626. Springer, 2015.
- [20] Nico Döttling, Jörn Müller-Quade, and Anderson C. A. Nascimento. IND-CCA secure cryptography based on a variant of the LPN problem. In *Advances in Cryptology – ASIACRYPT 2012*, pages 485–503, 2012.
- [21] Vitaly Feldman, Parikshit Gopalan, Subhash Khot, and Ashok Kumar Ponnuswami. New results for learning noisy parities and halfspaces. In *47th Symposium on Foundations of Computer Science*, pages 563–574, Berkeley, CA, USA, October 21–24 2006. IEEE.
- [22] David Forney. *Concatenated Codes*. MIT Press, 1966.
- [23] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, pages 197–206, Victoria, BC, Canada, 17–20 May 2008. ACM.
- [24] Y. Gertner, S. Kannan, T. Malkin, O. Reingold, and M. Viswanathan. The relationship between public key encryption and oblivious transfer. In *Foundations of Computer Science, 2000. Proceedings. 41st Annual Symposium on*, pages 325–335, 2000.

- [25] Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In Johnson [31], pages 25–32.
- [26] Shafi Goldwasser, Yael Kalai, Chris Peikert, and Vinod Vaikuntanathan. Robustness of the learning with errors assumption. In *Innovations in Theoretical Computer Science, ITCS'10*, pages 230–240. Tsinghua University Press, 2010.
- [27] Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. *Concrete Mathematics: A Foundation for Computer Science*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2nd edition, 1994.
- [28] J. Håstad, R. Impagliazzo, L.A. Levin, and M. Luby. Construction of pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.
- [29] Nicholas J. Hopper and Manuel Blum. Secure human identification protocols. In *Proceedings of the 7th International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT 2001)*, pages 52–66, 2001.
- [30] Abhishek Jain, Stephan Krenn, Krzysztof Pietrzak, and Aris Tentes. Commitments and efficient zero-knowledge proofs from learning parity with noise. In *Advances in Cryptology – ASIACRYPT 2012*, pages 663–680, 2012.
- [31] D. S. Johnson, editor. *Proceedings of the Twenty First Annual ACM Symposium on Theory of Computing*, Seattle, Washington, 15–17 May 1989.
- [32] Ari Juels and Stephen A. Weis. Authenticating pervasive devices with human protocols. In Victor Shoup, editor, *Advances in Cryptology—CRYPTO 2005*, volume 3621 of *LNCS*, pages 293–308. Springer-Verlag, 14–18 August 2005.
- [33] Jorn Justesen. A class of constructive asymptotically good algebraic codes. *IEEE Trans. Info. Theory*, 18(5):652–656, 1972.
- [34] Jonathan Katz and Ji Sun Shin. Parallel and concurrent security of the hb and hb⁺ protocols. In Serge Vaudenay, editor, *Advances in Cryptology—EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 73–87. Springer-Verlag, 2006.
- [35] Eike Kiltz. Chosen-ciphertext security from tag-based encryption. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography*, volume 3876 of *Lecture Notes in Computer Science*, pages 581–600. Springer Berlin Heidelberg, 2006.
- [36] Eike Kiltz, Daniel Masny, and Krzysztof Pietrzak. Simple chosen-ciphertext security from low-noise lpn. In Hugo Krawczyk, editor, *Public-Key Cryptography PKC 2014*, volume 8383 of *Lecture Notes in Computer Science*, pages 1–18. Springer Berlin Heidelberg, 2014.
- [37] Eike Kiltz, Krzysztof Pietrzak, David Cash, Abhishek Jain, and Daniele Venturi. Efficient authentication from hard learning problems. In *Proceedings of the 30th Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT 2011)*, pages 7–26, 2011.
- [38] Paul Kirchner. Improved generalized birthday attack. Cryptology ePrint Archive, Report 2011/377, 2011. <http://eprint.iacr.org/2011/377>.
- [39] Eric Leveil and Pierre-Alain Fouque. An improved lpn algorithm. In *Proceedings of the 5th Conference on Security and Cryptography for Networks (SCN 2006)*, pages 348–359, 2006.

- [40] Vadim Lyubashevsky. The parity problem in the presence of noise, decoding random linear codes, and the subset sum problem. In *Proceedings of the 9th International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM 2005)*, pages 378–389, 2005.
- [41] Vadim Lyubashevsky and Daniel Masny. Man-in-the-middle secure authentication schemes from lpn and weak prfs. In *Advances in Cryptology - CRYPTO 2013*, pages 308–325, 2013.
- [42] Alexander May, Alexander Meurer, and Enrico Thomae. Decoding random linear codes in $\tilde{O}(2^{0.054n})$. In *Proceedings of the 17th International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT 2011)*, pages 107–124, 2011.
- [43] Daniele Micciancio and Petros Mol. Pseudorandom knapsacks and the sample complexity of lwe search-to-decision reductions. In Phillip Rogaway, editor, *Advances in Cryptology - CRYPTO 2011*, volume 6841 of *Lecture Notes in Computer Science*, pages 465–484. Springer Berlin Heidelberg, 2011.
- [44] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology - EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 700–718. Springer Berlin Heidelberg, 2012.
- [45] Moni Naor and Moti Yung. Universal one-way hash functions and their cryptographic applications. In Johnson [31], pages 33–43.
- [46] Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *Proceedings of the Twenty Second Annual ACM Symposium on Theory of Computing*, Baltimore, Maryland, 14–16 May 1990.
- [47] Krzysztof Pietrzak. Cryptography from learning parity with noise. In *Proceedings of the Theory and Practice of Computer Science (SOFTSEM 2012)*, pages 99–114, 2012.
- [48] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *STOC*, pages 84–93. ACM, 2005.
- [49] Alon Rosen and Gil Segev. Chosen-ciphertext security via correlated products. In Omer Reingold, editor, *Theory of Cryptography*, volume 5444 of *Lecture Notes in Computer Science*, pages 419–436. Springer Berlin / Heidelberg, 2009.
- [50] Amit Sahai. Non-malleable non-interactive zero-knowledge and adaptive chosen-ciphertext security. In *40th Annual Symposium on Foundations of Computer Science*, New York, October 1999. IEEE.
- [51] Jacques Stern. A method for finding codewords of small weight. In *Coding Theory and Applications, 3rd International Colloquium*, pages 106–113, 1988.
- [52] Yu Yu. The LPN problem with auxiliary input. withdrawn, see historical versions at <http://eprint.iacr.org/2009/467>.

A Definitions and Security Notions

A.1 Symmetric-Key Encryption Schemes with Auxiliary Input

Definition A.1 (Symmetric-key encryption schemes) *A symmetric-key encryption scheme Π is a tuple $(\text{KeyGen}, \text{Enc}, \text{Dec})$ with message space \mathcal{M} , such that*

- $\text{KeyGen}(1^n)$ is a PPT algorithm that takes a security-parameter 1^n and outputs a symmetric key k .
- $\text{Enc}_k(m)$ is a PPT algorithm that encrypts a message $m \in \mathcal{M}$ under key k and outputs a ciphertext c .
- $\text{Dec}_k(c)$ is a deterministic polynomial-time algorithm that decrypts a ciphertext c using key k and outputs a plaintext m .

Definition A.2 (Correctness) We say that a symmetric-key encryption scheme $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ is correct, if it holds for every plaintext $m \in \mathcal{M}$ that

$$\Pr_{k \leftarrow \text{KeyGen}(1^n)} [\text{Dec}_k(\text{Enc}_k(m)) \neq m] = \text{negl}(n) .$$

Definition A.3 (IND-CPA/IND-CCA SKE w.r.t. auxiliary input) For $X \in \{\text{CPA}, \text{CCA}\}$, a symmetric-key encryption scheme $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ is IND- X secure w.r.t. sub-exponentially hard-to-invert auxiliary input if there exists a constant $0 < \alpha < 1$ such that for any PPT adversary \mathcal{A} , any $2^{-\Omega(n^\alpha)}$ -hard-to-invert function f

$$\Pr [\text{SKE}_{\Pi, f, \mathcal{A}}^X(1^n, \alpha) = 1] \leq \frac{1}{2} + \text{negl}(n) ,$$

where $\text{SKE}_{\Pi, f, \mathcal{A}}^{\text{cpa}}(1^n, \alpha)$ is the IND-CPA indistinguishability experiment defined as below:

1. On $k \leftarrow \text{KeyGen}(1^n)$, the adversary takes as input $1^n, f(k)$, and is given oracle access to Enc_k . Then, he outputs a pair of messages m_0 and m_1 of the same length.
2. A random bit $b \xleftarrow{\$} \{0, 1\}$ is sampled, and then a challenge ciphertext $c \leftarrow \text{Enc}_k(m_b)$ is computed and given to \mathcal{A} .
3. \mathcal{A} continues to have oracle access to Enc_k and finally outputs $b' \in \{0, 1\}$.
4. The experiment outputs 1 if $b' = b$, and 0 otherwise.

and $\text{SKE}_{\Pi, f, \mathcal{A}}^{\text{cca}}(1^n, \alpha)$ is the IND-CCA indistinguishability experiment defined as below:

1. On $k \leftarrow \text{KeyGen}(1^n)$, the adversary takes as input $1^n, f(k)$, and is given oracle access to Enc_k and Dec_k . Then, he outputs a pair of messages m_0 and m_1 of the same length.
2. A random bit $b \xleftarrow{\$} \{0, 1\}$ is sampled, and then a challenge ciphertext $c \leftarrow \text{Enc}_k(m_b)$ is computed and given to \mathcal{A} .
3. \mathcal{A} continues to have oracle access to Enc_k and Dec_k (with the exception that decryption for challenge ciphertext is not allowed) and finally outputs $b' \in \{0, 1\}$.
4. The experiment outputs 1 if $b' = b$, and 0 otherwise.

A.2 Public-Key Encryption Schemes

Definition A.4 (Public-key encryption schemes) A public key encryption scheme Π is a tuple $(\text{KeyGen}, \text{Enc}, \text{Dec})$ with message space \mathcal{M} , such that

- $\text{KeyGen}(1^n)$ is a PPT algorithm that takes a security-parameter 1^n and outputs a pair of public and private keys (pk, sk) .
- $\text{Enc}_{pk}(m)$ is a PPT algorithm that encrypts message $m \in \mathcal{M}$ under public key pk and outputs a ciphertext c .

- $\text{Dec}_{sk}(c)$ is a deterministic polynomial-time algorithm that decrypts a ciphertext c using secret key sk and outputs a plaintext m (or \perp).

Definition A.5 (Correctness) We say that a public-key encryption scheme $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ is correct, if it holds for every plaintext $m \in \mathcal{M}$ that

$$\Pr_{(pk, sk) \leftarrow \text{KeyGen}(1^n)} [\text{Dec}_{sk}(\text{Enc}_{pk}(m)) \neq m] = \text{negl}(n) .$$

Definition A.6 (IND-CPA/IND-CCA PKE) For $X \in \{\text{CPA}, \text{CCA}\}$, a public-key encryption scheme $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ is IND- X secure if for any PPT adversary \mathcal{A}

$$\Pr[\text{PKE}_{\Pi, \mathcal{A}}^X(1^n) = 1] \leq \frac{1}{2} + \text{negl}(n) ,$$

where $\text{PKE}_{\Pi, \mathcal{A}}^{\text{cpa}}(1^n)$ is the IND-CPA indistinguishability experiment defined as below:

1. On $(pk, sk) \leftarrow \text{KeyGen}(1^n)$, the adversary takes as input 1^n and pk . Then, he outputs a pair of messages m_0 and m_1 of the same length.
2. A random bit $b \xleftarrow{\$} \{0, 1\}$ is sampled, and then a challenge ciphertext $c \leftarrow \text{Enc}_{pk}(m_b)$ is computed and given to \mathcal{A} .
3. \mathcal{A} continues his computation and finally outputs $b' \in \{0, 1\}$.
4. The experiment outputs 1 if $b' = b$, and 0 otherwise.

and $\text{PKE}_{\Pi, \mathcal{A}}^{\text{cca}}(1^n)$ is the IND-CCA indistinguishability experiment defined as below:

1. On $(pk, sk) \leftarrow \text{KeyGen}(1^n)$, the adversary takes as input 1^n and pk , and is given oracle access to Dec_{sk} . Then, he outputs a pair of messages m_0 and m_1 of the same length.
2. A random bit $b \xleftarrow{\$} \{0, 1\}$ is sampled, and then a challenge ciphertext $c \leftarrow \text{Enc}_{pk}(m_b)$ is computed and given to \mathcal{A} .
3. \mathcal{A} continues to have oracle access to Dec_{sk} (with the exception that decryption for challenge ciphertext is not allowed) and finally outputs $b' \in \{0, 1\}$.
4. The experiment outputs 1 if $b' = b$, and 0 otherwise.

B Facts, Lemmas, Inequalities and Proofs Omitted

Proof of Fact 4.

$$\begin{aligned} & \mu \log(1/\mu) \\ & < \left(\mathbf{H}(\mu) = \mu \log(1/\mu) + (1 - \mu) \log(1/(1 - \mu)) \right) \\ & = \mu \log(1/\mu) + (1 - \mu) \log\left(1 + \frac{\mu}{1 - \mu}\right) \\ & = \mu \log(1/\mu) + (1 - \mu) \frac{\ln\left(1 + \frac{\mu}{1 - \mu}\right)}{\ln 2} \\ & \leq \mu \log(1/\mu) + \frac{\mu}{\ln 2} < \mu(\log(1/\mu) + \frac{3}{2}) , \end{aligned}$$

where the first inequality is due to $(1 - \mu) \log(1/(1 - \mu)) > 0$, the second one follows from the elementary inequality $\ln(1 + x) \leq x$ for any $x > 0$, and the last inequality is simply $\ln 2 > 2/3$.

□

Lemma B.1 (Sample-preserving reduction) *For the same assumptions and notations as in the proof of [Theorem 4.1](#), we have*

$$\begin{aligned} & (f(\mathbf{x}, \mathbf{e}; Z), Z, \mathbf{A}, \mathbf{A} \cdot \mathbf{y} + \mathbf{e}, \mathbf{r}^\top, \mathbf{r}^\top \cdot \mathbf{y}) \stackrel{\text{c}}{\sim} (f(\mathbf{x}, \mathbf{e}; Z), Z, \mathbf{A}, \mathbf{A} \cdot \mathbf{y} + \mathbf{e}, \mathbf{r}^\top, U_1) \\ \Rightarrow & (f(\mathbf{x}, \mathbf{e}; Z), Z, \mathbf{A}, \mathbf{A} \cdot \mathbf{y} + \mathbf{e}) \stackrel{\text{c}}{\sim} (f(\mathbf{x}, \mathbf{e}; Z), Z, \mathbf{A}, U_{q'}) . \end{aligned}$$

Proof. Assume for contradiction that there exists a polynomial $p(\cdot)$ and a PPT distinguisher \mathcal{D} such that

$$\Pr[\mathcal{D}(f(\mathbf{x}, \mathbf{e}; Z), Z, \mathbf{A}, \mathbf{A} \cdot \mathbf{y} + \mathbf{e}) = 0] - \Pr[\mathcal{D}(f(\mathbf{x}, \mathbf{e}; Z), Z, \mathbf{A}, U_{q'}) = 0] \geq 1/p(n)$$

for infinitely many n 's and we recall that $\mathbf{y}, \mathbf{r} \sim U_\lambda$. Given input $(z_1, z, \mathbf{A}, \mathbf{A} \cdot \mathbf{y} + \mathbf{e}, \mathbf{r}^\top)$, we use an efficient \mathcal{D}' (which invokes \mathcal{D}) to predict the Goldreich-Levin hardcore bit $\mathbf{r}^\top \cdot \mathbf{y}$ with non-negligible probability (and thus a contradiction to the assumption). \mathcal{D}' chooses a random $\mathbf{u} \stackrel{\$}{\leftarrow} \{0, 1\}^{q'}$, computes a new $q' \times n$ Boolean matrix $\tilde{\mathbf{A}} = \mathbf{A} - \mathbf{u} \cdot \mathbf{r}^\top$, applies \mathcal{D} on $(z_1, z, \tilde{\mathbf{A}}, \mathbf{A} \cdot \mathbf{y} + \mathbf{e})$ and outputs his answer. Note that $\tilde{\mathbf{A}} \sim U_{q' \times n}$ and $\mathbf{A} \cdot \mathbf{y} + \mathbf{e} = \tilde{\mathbf{A}} \cdot \mathbf{y} + \mathbf{e} + \mathbf{u} \cdot \mathbf{r}^\top \cdot \mathbf{y}$. Therefore, when $\mathbf{r}^\top \cdot \mathbf{y} = 0$ we have $(z_1, z, \tilde{\mathbf{A}}, \mathbf{A} \cdot \mathbf{y} + \mathbf{e})$ follows $(f(\mathbf{x}, \mathbf{e}; Z), Z, \tilde{\mathbf{A}}, \tilde{\mathbf{A}} \cdot \mathbf{y} + \mathbf{e})$ and for $\mathbf{r}^\top \cdot \mathbf{y} = 1$ it is distributed according to $(f(\mathbf{x}, \mathbf{e}; Z), Z, \tilde{\mathbf{A}}, U_{q'})$.

$$\begin{aligned} & \Pr[\mathcal{D}'(f(\mathbf{x}, \mathbf{e}; Z), Z, \tilde{\mathbf{A}}, \tilde{\mathbf{A}} \cdot \mathbf{y} + \mathbf{e}, \mathbf{r}^\top) = \mathbf{r}^\top \cdot \mathbf{y}] \\ = & \Pr[\mathbf{r}^\top \cdot \mathbf{y} = 0] \cdot \Pr[\mathcal{D}'(f(\mathbf{x}, \mathbf{e}; Z), Z, \tilde{\mathbf{A}}, \tilde{\mathbf{A}} \cdot \mathbf{y} + \mathbf{e}, \mathbf{r}^\top) = 0 \mid \mathbf{r}^\top \cdot \mathbf{y} = 0] \\ & + \Pr[\mathbf{r}^\top \cdot \mathbf{y} = 1] \cdot \Pr[\mathcal{D}'(f(\mathbf{x}, \mathbf{e}; Z), Z, \tilde{\mathbf{A}}, \tilde{\mathbf{A}} \cdot \mathbf{y} + \mathbf{e}, \mathbf{r}^\top) = 1 \mid \mathbf{r}^\top \cdot \mathbf{y} = 1] \\ = & \frac{1}{2} (\Pr[\mathcal{D}(f(\mathbf{x}, \mathbf{e}; Z), Z, \mathbf{A}, \mathbf{A} \cdot \mathbf{y} + \mathbf{e}) = 0] + 1 - \Pr[\mathcal{D}(f(\mathbf{x}, \mathbf{e}; Z), Z, \mathbf{A}, U_{q'}) = 0]) \\ \geq & \frac{1}{2} + \frac{1}{2p(n)} , \end{aligned}$$

which completes the proof. \square

Proof of [Lemma 5.1](#). Consider $|(\mathbf{S}'_0 - \mathbf{S}_0)^\top \mathbf{e} + (\mathbf{E}_0 - \mathbf{E}'_0)^\top \mathbf{s}|$ conditioned on any $|\mathbf{e}| \leq 1.01\mu n$ (except for a $2^{-\Omega(n)}$ -fraction) and $|\mathbf{s}| \leq 2\mu n$. We have by [Lemma 4.3](#) and [Lemma 4.4](#) that $\mathbf{S}_0^\top \mathbf{e}, \mathbf{S}'_0^\top \mathbf{e}$ are i.i.d. to $\mathcal{B}_{\delta_1}^q$, and $\mathbf{E}_0^\top \mathbf{s}, \mathbf{E}'_0^\top \mathbf{s}$ are i.i.d. to $\mathcal{B}_{\delta_2}^q$, where $\delta_1 \leq 1/2 - n^{-\alpha/2}$ and $\delta_2 \leq 1/2 - n^{-\alpha}/2$. Thus, $((\mathbf{S}'_0 - \mathbf{S}_0)^\top \mathbf{e} + (\mathbf{E}_0 - \mathbf{E}'_0)^\top \mathbf{s})$ follows \mathcal{B}_δ^q for $\delta \leq 1/2 - 2n^{-3\alpha}$ by the Piling-up lemma, and then we complete the proof with [Lemma 2.3](#). \square

Lemma B.2 (Flattening Shannon entropy) *For any $n \in \mathbb{N}$, $0 < \mu < 1/2$ and any constant $0 < \Delta < 1$, there exists some random variable $W \in \{0, 1\}^n$ such that $\mathbf{H}_\infty(W) \geq (1 - \Delta)n\mathbf{H}(\mu)$ and $\text{SD}(\mathcal{B}_\mu^n, W) \leq 2^{-\Omega(\mu n)}$.*

Proof of [Lemma B.2](#). Recall that $\mathbf{H}(\mu) \stackrel{\text{def}}{=} \mu \log(1/\mu) + (1 - \mu) \log(1/(1 - \mu))$ equals to $\mathbf{H}_1(\mathcal{B}_\mu)$. We define the sample-entropy of $e \in \{0, 1\}^n$ with respect to \mathcal{B}_μ^n as

$$\mathbf{H}_{\mathcal{B}_\mu^n}(e) \stackrel{\text{def}}{=} \log(1/\Pr[\mathcal{B}_\mu^n = e])$$

and let $\mathcal{E} \stackrel{\text{def}}{=} \{e \in \{0, 1\}^n : \mathbf{H}_{\mathcal{B}_\mu^n}(e) < (1 - \Delta)n\mathbf{H}(\mu)\}$. Write $\mathcal{B}_\mu^n = (E_1, \dots, E_n)$, and for each $1 \leq i \leq n$ define

$$\xi_i \stackrel{\text{def}}{=} \begin{cases} 1, & \text{if } E_i = 1 \\ \frac{\log(\frac{1}{1-\mu})}{\log(\frac{1}{\mu})}, & \text{if } E_i = 0 \end{cases}$$

and thus we have that ξ_1, \dots, ξ_n are i.i.d. over $\{\frac{\log(1/(1-\mu))}{\log(1/\mu)}, 1\}$, each of expectation $\mathbf{H}(\mu)/\log(1/\mu)$.

$$\begin{aligned} & \Pr[\mathcal{B}_\mu^n \in \mathcal{E}] \\ = & \Pr\left[\sum_{i=1}^n \xi_i < (1 - \Delta) \cdot \frac{n\mathbf{H}(\mu)}{\log(1/\mu)}\right] \\ < & \exp^{-\frac{\Delta^2 n\mathbf{H}(\mu)}{2\log(1/\mu)}} = 2^{-\Omega(\mu n)} , \end{aligned}$$

where the inequality follows from the Chernoff bound (see [Lemma 2.1](#)) and we recall [Fact 4](#). \square

C Improving the Applications Considered in [\[17\]](#)

We recall the definitions about symmetric-key encryption schemes with auxiliary input, correctness and CPA/CCA security introduced in [\[17\]](#) and reproduce them in [Appendix A.1](#).

Our CPA encryption scheme $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ is quite straightforward (and almost the same as the one in [\[17\]](#)). For constants μ and α , let $\lambda = n^\alpha$, $q' = \text{poly}(n)$ (as defined in [Theorem 3.1](#)) with poly to be decided below. The key generation algorithm KeyGen takes as input 1^n and outputs a random $\mathbf{k} \xleftarrow{\$} \{0, 1\}^n$ and a public random matrix $\mathbf{V} \xleftarrow{\$} \{0, 1\}^{\lambda \times n}$ (the basis of the random subspace). The encryption algorithm

$$\text{Enc}_{\mathbf{k}}(\mathbf{m}) = (\mathbf{A}, (\mathbf{A} \cdot \mathbf{V}) \cdot \mathbf{k} + \mathbf{e} + \mathbf{G}\mathbf{m})$$

where $\mathbf{A} \sim U_{q \times \lambda}$, $\mathbf{e} \sim \mathcal{B}_\mu^q$, $\mathbf{G} \in \{0, 1\}^{q \times n}$ is a generator matrix for a decodable code that efficiently corrects up to an arbitrary $\Theta(\alpha)$ -fraction of bit flipping errors. For example, each message bit m_i is encoded to $m_i \cdots m_i \in \{0, 1\}^\delta$ for a repetition of $\delta = \omega(\log n)$ times and thus $q = n\delta$. The decoder takes as input the noisy codeword $\mathbf{G}\mathbf{m} + \mathbf{e}$ and parse it as δ -bit blocks $c_1 \cdots c_n$, where each $c_i = (m_i \cdots m_i + \mathcal{B}_\mu^\delta)$ decodes to $m'_i = 0$ if $|c_i| \leq 2\mu\delta$, $m'_i = 1$ if $|c_i| \geq (1 - 2\mu)\delta$, and $m'_i = \perp$ otherwise. Finally, the decoder outputs \perp if there is any $m'_i = \perp$, or $m'_1 \cdots m'_n$ otherwise. Π is a correct scheme since a decryption error is of probability $2^{-\Omega(\delta)} = \text{negl}(n)$ by a Chernoff bound.

Theorem C.1 *Assume that the decisional (constant-noise) LPN problem is hard, then the above symmetric encryption scheme $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ is IND-CPA secure w.r.t. sub-exponentially hard-to-invert auxiliary input.*

The CPA security is straightforward as it follows from [Theorem 3.1](#) that the LPN on auxiliary input is closed under composition to encrypt many messages (see [Remark 3.1](#)). Note that we could have used more optimal codes (e.g., Justesen code [\[33\]](#)) in the above, but we need $\text{Dec}_{\mathbf{k}}$ to output \perp to indicate an unlikely state such that Π can have the unique-key property below.

Lemma C.1 (Key uniqueness) *The above symmetric-key encryption scheme $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ with key space and message space $\mathcal{K} = \mathcal{M} = \{0, 1\}^n$ has the unique key property (as defined in [\[17\]](#)), i.e., if for every $\mathbf{k} \in \mathcal{K}$ and $m \in \mathcal{M}$*

$$\Pr[\exists \mathbf{k}' \in \mathcal{K} : \mathbf{k}' \neq \mathbf{k} \wedge \text{Dec}_{\mathbf{k}'}(\text{Enc}_{\mathbf{k}}(m)) \neq \perp] = \text{negl}(n),$$

where the probability is taken over the coin tosses of $\text{Enc}_{\mathbf{k}}$.

Proof. Parse the $q \times n$ matrix $\mathbf{A}' = \mathbf{A} \cdot \mathbf{V}$ as n ($\delta \times n$) submatrix $\mathbf{A}'_1, \dots, \mathbf{A}'_n$ (note $q = \delta n$)

such that the noisy codeword for each message bit m_i is $\mathbf{A}'_i \mathbf{k} + m_i \cdots m_i + \mathbf{E}_i$, where $\mathbf{E}_i \sim \mathcal{B}_\mu^\delta$.

$$\begin{aligned}
& \Pr[\exists \mathbf{k}' \in \mathcal{K} : \mathbf{k}' \neq \mathbf{k} \wedge \mathcal{D}_{\mathbf{k}'}(\mathcal{E}_{\mathbf{k}}(m)) \neq \perp] \\
\leq & \sum_{\mathbf{k}' \neq \mathbf{k}} \Pr[\mathcal{D}_{\mathbf{k}'}(\mathcal{E}_{\mathbf{k}}(m)) \neq \perp] \\
= & \sum_{\mathbf{k}' \neq \mathbf{k}} \prod_{i=1}^n \Pr[|\mathbf{A}'_i(\mathbf{k} + \mathbf{k}') + E_i| < 2\mu\delta \vee |\mathbf{A}'_i(\mathbf{k} + \mathbf{k}') + E_i| > (1 - 2\mu)\delta] \\
\leq & 2^n \prod_{i=1}^n \Pr \left[|\mathbf{A}'_i(\mathbf{k} + \mathbf{k}') + E_i| < 2\mu\delta \vee |\mathbf{A}'_i(\mathbf{k} + \mathbf{k}') + E_i| > (1 - 2\mu)\delta \mid \mathbf{k}' \neq \mathbf{k} \right] \\
\leq & 2^n \prod_{i=1}^n \left(\Pr \left[\left| |\mathbf{A}'_i(\mathbf{k} + \mathbf{k}')| - \frac{1}{2}\delta \right| > 2\mu\delta \mid \mathbf{k}' \neq \mathbf{k} \right] + \Pr[|\mathbf{E}_i| \geq (\frac{1}{2} - 4\mu)\delta] \right) \\
= & 2^n (2^{-\Omega(\delta)})^n + 2^{-\lambda} \\
= & \text{negl}(n)
\end{aligned}$$

where the first inequality is a union bound, $\mathbf{A}'_i \sim (U_{\delta \times \lambda} \times U_{\lambda \times n})$ implies that for any $\mathbf{k} + \mathbf{k}' \neq 0^n$, $U_{\lambda \times n}(\mathbf{k} + \mathbf{k}') \sim U_\lambda$ and thus $\mathbf{A}'_i(\mathbf{k} + \mathbf{k}')$ is $2^{-\lambda}$ -close to U_δ , and the last inequality follows from the Chernoff bound. \square

In addition, Π also enjoys the key hiding property against hard-to-invert auxiliary input, whose proof is simply a hybrid argument.

Lemma C.2 (Key hiding) *The above symmetric-key encryption scheme $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ with key space and message space $\mathcal{K} = \mathcal{M} = \{0, 1\}^n$ has the key hiding property with sub-exponentially hard-to-invert auxiliary input (as defined in [17]), i.e., if for every constant $0 < \alpha < 1$, $\lambda = n^\alpha$ every $\mathbf{m} \in \mathcal{M}$, and every $2^{-2\lambda}$ hard-to-invert f , we have*

$$(\text{Enc}_{\mathbf{k}}(\mathbf{m}; \mathbf{A}), \mathbf{V}, f(\mathbf{k})) \stackrel{c}{\sim} (\text{Enc}_{\mathbf{k}'}(\mathbf{m}; \mathbf{A}), \mathbf{V}, f(\mathbf{k}'))$$

where $\mathbf{k}, \mathbf{k}' \stackrel{\$}{\leftarrow} \{0, 1\}^n$.

Next, we can use the general transformation in [17] that converts any symmetric encryption scheme Π that is CPA secure w.r.t. (sub-)exponentially hard-to-invert auxiliary input into one that is CCA secure w.r.t. (sub-)exponentially hard-to-invert auxiliary input. We reproduce the following from [17]. The idea (originating in [45]) is to add a component to the encryption with the additional property “ciphertext unforgeable” (and thus render decryption oracle useless): given oracle access to the encryption oracle $\text{Enc}_{\mathbf{k}}$, it is infeasible to produce a valid ciphertext not returned by this oracle. This implies CCA security. To this end, we need additional tools: CCA-secure public-key encryption scheme $\widetilde{\text{Enc}}_{pk}$ (without auxiliary input), a universal one-way hash function (UOWHF) $h : \{0, 1\}^n \rightarrow \{0, 1\}^{n^\beta}$ (for constant $\beta < \alpha$) and a simulation-sound non-interactive zero-knowledge (NIZK) proof system for NP [50], all implied by trapdoor permutations [45, 46, 50]. The proposed CCA secure scheme Π' consists of the original secret key \mathbf{k} of our CPA secure Π and public parameters: a hashed value $z = h(k)$, a public-key pk (of the standard PKE without auxiliary input) and a common reference string σ for the NIZK proof system. Note that we allow the leakage function f to hardwire (z, pk, σ) , but CRS and the public key pk are independent of \mathbf{k} , and z is of length n^β for $\beta < \alpha$, and thus f remains more than $2^{-1.99n^\alpha}$ -hard-to-invert even when conditioned on these public parameters.

The encryption and decryption algorithms of Π' work as follows: to encrypt \mathbf{m} , compute $c = \text{Enc}_{\mathbf{k}}(\mathbf{m})$, $d = \widetilde{\text{Enc}}_{pk}(c, \mathbf{k})$ and produce $C = (c, d, \pi)$ as ciphertext, where π is a NIZK proof for the statement “ (pk, z, c, d) satisfies the claim that there exist secret \mathbf{k} , message \mathbf{m} and randomness r , such that $z = h(\mathbf{k})$, $\text{Dec}_{\mathbf{k}}$ decrypts c to \mathbf{m} , and d is the encryption of the pair (c, \mathbf{k}) using randomness r ”. To decrypt $C' = (c', d', \pi')$ using $(\mathbf{k}, z, pk, \sigma)$, first check the validity of the proof π' , outputs \perp if invalid, otherwise output $\text{Dec}_{\mathbf{k}}(c')$.

Theorem C.2 ([17]) *If Π is CPA-secure w.r.t. (sub-)exponentially hard-to-invert auxiliary input (which is implied by constant-noise LPN), and assume that $\widetilde{\text{Enc}}_{pk}$ is the encryption algorithm of a CCA-secure public-key encryption scheme, h is a UOWHF, and π is a simulation-sound NIZK (all implied by trapdoor permutations), then the above scheme Π' is a CCA-secure PKE w.r.t. (sub-)exponentially hard-to-invert auxiliary input. In particular, one can build such a scheme assuming the existence of trapdoor permutations and the constant-noise LPN assumption.*

MORE APPLICATIONS. [17] further pointed out that CPA-secure symmetric-key encryption scheme w.r.t. (sub-)exponentially hard-to-invert auxiliary input further implies average-case obfuscators for point functions w.r.t. (sub-)exponentially hard-to-invert auxiliary input and reusable randomness extractor, and that CCA-secure SKE counterpart implies robust extractors. We refer the full version of [17] for the further details. To conclude, with our technique all the applications introduced in [17] are now based on standard constant-noise LPN with improved efficiency (by reducing the noise level from polynomially close to uniform to constant) and tighter security against sub-exponentially (rather than exponentially) hard-to-invert auxiliary input.